FOLEY
FOLEY & LARDNER LLP

Navigating the AI Frontier:
Legal and Operational Insights Into Generative AI

FALL 2023

FORWARD

In our ever-changing technological environment, artificial intelligence (AI) is beginning to exert its influence on numerous sectors, revolutionizing our approach to work. As businesses and organizations worldwide adopt generative AI technologies like ChatGPT, they are faced with the promise of unprecedented opportunities, challenges, and legal complexities.

Welcome to Navigating the AI Frontier: Legal and Operational Insights Into Generative AI. In this eBook, we take a deep dive into the world of generative AI, exploring its transformative power and the legal and operational risks it introduces. We will delve into various topics, from the impact of generative AI on dealmaking processes to its implications in venture capital, intellectual property, and even employment law.

In each chapter, we will explore the intricacies of generative AI, its implications, and how businesses can navigate this landscape while remaining compliant with the law and fostering responsible AI adoption. Whether you are a legal professional, an AI enthusiast, or simply curious about the future of technology, this eBook aims to share the knowledge and insights needed to thrive in the AI frontier.

Louis Lehot
Partner
Silicon Valley

TABLE OF CONTENTS

Deep Dive into Generative AI and What Will Drive Tomorrow

Published March 2023 by Foley & Lardner LLP

AUTHOR

Louis Lehot | llehot@foley.com

2023 has seen an explosion of activity in the artificial intelligence (AI) space that can only be referred to as a "hype-cycle." There is a growing consumer and enterprise appetite for this technology, and more and more companies are jumping on board. And with all this buzz there is an even greater interest on the part of investors, who are looking to position themselves ahead of each new wave of development. More regulation is also being pushed to protect intellectual property and privacy.

When seeing a hype-cycle, we ask: how'd we get to this point and where are we going?

AI is the development of computer systems performing tasks that typically require human intelligence such as visual perception, speech recognition, decision-making, and translation between languages. It's a combination of computer science techniques as applied to datasets to enable problem solving. Building on that, generative AI is a term describing algorithms that can create new content, including audio, code, images, text, simulations, and video. The key to generative AI is extensive large language modeling (LLM) that can analyze and learn from natural language interfaces, partnered with media that can procure natural language inputs, with modeling and infrastructure that can analyze inputs to create generative outputs.

The phrase was coined in December 1997 when two computer scientists, Sepp Hochreiter and Jurgen Schmidhuber, invented long short-term memory networks, which improve memory capacity in neural networks to allow for pattern recognition in training data. In 2012 Alex Krizhevsky gave us AlexNet, a convolutional neural network trained on graphical processing units, breaking 75% accuracy in identifying images from a manually tagged database. While Google has led research and development in the field for over

a decade, focused on a product called "BERT," OpenAI formed and released its first generative pre-trained transformer in June 2018 called GPT, ushering in a new era of LLMs. A year later, Microsoft invested US$1 billion into OpenAI, launching the modern-day AI arms race and culminating in a new text-to-image model called Dall-E and chatbot ChaptGPT.

NVIDIA, Google, Microsoft, Meta, and even government agencies are leading this arms race and are locked in battle, each deploying billions of dollars to win. Microsoft announced a multiyear, multibillion-dollar investment in OpenAI and exclusive integration with its Azure system in January 2023. Stability.ai and Amazon have partnered on AWS. Google announced BARD and invested US$300 million in Anthropic.

PitchBook released its report "Vertical Snapshot: Generative AI," examining the venture capital trends, industry overview, and market landscape in this space. Below we look at the AI sub-verticals, technologies, and startups that are getting the most funding and why.

Natural-language user interface (LUI or NLUI) is a type of computer human interface where the user and the system communicate using natural, human language. Most of us use this daily on our phones or other devices e.g., (Siri or Alexa virtual assistants.) The largest areas of investment within this specific segment were chatbots, voicebots, and personal assistants, which captured US$544.9 million in 2022 (59.6% of all dollars invested).

2D media is exactly as it sounds — any artwork that exits in two dimensions, such as paintings, drawings, or prints. AI can be used not only to create 2D media but also to convert 2D media into 3D media. In this area, investors seem to see the most possibility in avatars, video generation, and editing, grabbing 37.7% and 40.8% of total dollars invested, respectively.

The AI Core and biotech categories also brought in impressive numbers. PitchBook found that AI Core, which includes foundation model developers and infrastructure for model development, raised US$5 billion between 2018 and 2022. Biotech startups utilizing generative techniques have also piqued the interest of venture capital investors, with US$1.6 billion invested during the same period.

One further fascinating finding of the PItchBook report is the predicted growth in generative AI due to the incredible enterprise applications of this technology. Natural language interfaces will likely be the primary catalyst of this growth. In fact, PitchBook "expects the market at a 32.0% CAGR to reach [US]$98.1 billion by 2026."

As we look forward, the costs of foundation model training are dropping. Custom hardware and accelerated software tools are making it cheaper to train new LLMs. Price points are accessible to startups. Meanwhile, chief information officers at larger corporations are pushing for digital transformation and AI adoption across the enterprise. According to a report published by MIT in September 2022, only IT, supply chain, and finance departments were gaining widespread adoption of AI, and only at the 40% level. We believe the launch of ChatGPT 4 is driving adoption deeper into those functions and powering expansion into sales, marketing, product development, and human resources.

This space has fantastic potential, and it will be interesting to see how it ultimately impacts industries across the board and becomes an even more significant part of our daily lives. One area that will be critical to watch is how lawmakers approach regulation and what kind of implications those regulations will have on this rapidly evolving technology.

What are the guard rails of intellectual property? Who owns the content output if the inputs were copyrighted? Who is responsible for copyright infringement? What if the output is wrong and it drives tortious conduct? Who will bear the risk of liability? Will it be insured? How can you protect your image, likeness, speech, and images, and how can you enforce them?

It is a brave new world; we may only be in the early innings of a new technology revolution. For startups looking for funding, we believe the keys to success will lie in foundational technology combined with unique algorithms that have demonstrated utility and traction in at least one vertical but with applications across verticals. While the banking crisis may have interrupted funding flows for a limited period, the hype-cycle in generative AI has momentum with staying power.

Largest areas of investment (2022)

chatbots, voicebots, and personal assistants

$544.9M

IP Lawyer vs. ChatGPT: Top 10 Legal Issues of Using Generative AI at Work

Published March 2023 by Foley & Lardner LLP

AUTHORS

Kate Wegrzyn | kwegrzyn@foley.com
Shabbi Khan | skhan@foley.com
ChatGPT

My mom says you should read science-fiction if you want to know what the future will look like. In her 50+ years as an avid reader of the genre, she has seen things once considered preposterous materialize into reality. Her assessment of ChatGPT and other generative AI is that it is coming for us all.

I'm more of a historical fiction person myself, so my take on generative AI is a bit more measured – it is an impressive tool that will become an integral part of our everyday existence, much the same as the internet did in the 1990s and early 2000s.

About a month ago, my colleague Shabbi Khan and I were discussing how the use of this tool in the workplace will continue to expand. While in agreement that we should write a blog on the legal issues surrounding the use of generative AI (like ChatGPT) at work, we lamented that it would take quite a bit of time to accomplish. This is a challenge for anyone, but particularly for lawyers where the product you sell is an hour broken into six minute increments.

I mused that we should just ask ChatGPT to write the blog post for us. We had a little laugh. From there, this 'man vs. machine' experiment was born. I generously offered to take the task of prompting ChatGPT to write the blog post while Shabbi was tasked with drafting it the "normal" way.

Author's note: ChatGPT is a mouthful. It must have been named by a computer programmer and not a marketer. From here on, I am going to call ChatGPT "Cathy" after "Chatty Cathy," the 1960s pull-string doll that likewise was a technological marvel for its time.

The results of this experiment were unsurprising:

Efficiency: +1 Point for the Machine. My time commitment to this experiment, a mere 11 minutes (or .2 non-billable hours, if you will), paled in comparison to that of Shabbi, who spent 10 non-billable hours over 8 days researching, pondering, and then finally drafting and editing the article. To no one's surprise, Cathy wins this point.

Bias: +1 Point for the Human. I award this point to Shabbi because Cathy did not mention that she hallucinates (that is, she makes up responses sometimes when she doesn't know the answer). It's funny – one would think that this lack of self-awareness would have been the human trait, but not so in this case. Perhaps this deserves extra weight because the fact that Cathy is so confident, and "doesn't know what she doesn't know," could easily lull one who relies on her into a false sense that her result is more accurate than it actually is.

Now, Shabbi is an IP lawyer and his list of 10 issues is mostly centered on IP topics. But I forgave this in my scoring because when Shabbi and I decided to undertake this project, we both agreed that to do it fully, we would need to ask many of our colleagues across various practice areas to weigh in (e.g., labor and employment, data privacy). It was determined that doing so would slow the process down so much that ChatGPT would be outdated technology by the time we had completed the blog post. Perhaps this demonstrates

that Cathy deserves another point, but she was already awarded a point in efficiency and, frankly, I feel humans need a finger on the scale at the moment.

Readability: +1 Point for the Machine. Cathy's responses were snappy, short, and easy to read. But that readability came at the expense of depth. See next bullet.

Effectiveness: +1 Point for the Human. Shabbi's article took a much deeper dive into the topics that he raised than Cathy's did. She also repeated some of the same issues more than once in slightly different ways to round out a list of ten. For this reason, Shabbi gets the point.

Winner: With man and machine each scoring two of the available four points on my totally made-up scoring system, we have a tied ballgame.

Takeaways: My takeaways from this experiment, as well as with the content of the articles themselves, is that there may be a place for using Cathy at work but the boundaries of when its use is appropriate are still being established. For now, here are some practical tips for using Cathy (and generative AI in general) to be more efficient while avoiding trouble:

1. Figure out how to prompt it in a way that gives the best result.

2. Use it for appropriate projects – blog posts on 'man vs. machine' style experimentation is a good example of such an appropriate project.

3. Do not feed it confidential information – the user cannot control what it does with that information.

4. Always verify what it gives you is accurate – trust, but verify.

5. Do diligence to ensure that the response it gives you is not plagiarized.

6. Include appropriate notices and disclaimers about the item being produced with ChatGPT. I expect someday this will be like a Prop65 warning – it'll be on everything and consequently barely noticed by the reader.

The bottom line is that the way everyone works is likely about to change at a breakneck speed, and no one is perfectly clear what that means from a legal perspective. The laws will play catch up to the technology. In the meantime, here are some of the legal issues that Shabbi and Cathy identified with respect to using generative AI in the workplace:

Top 10 Legal Issues of Using Generative AI at Work

Author:	Human (Shabbi Khan)	ChatGPT-3.5 (Prompted by Kate Wegrzyn)
Total Time:	10 hours	11 minutes
Work Done:	Brainstorming issues (1 hour) Web-based research (3 hours) Drafting article (4 hours) Editing article (2 hours)	Running various prompts (6 minutes) Reading article and adding bold headings (5 minutes) Note that my initial prompt was "Write a 1500 word blog article on the following topic: What are the top legal issues in using generative AI at work?" I moved to the Top 10 list, because the 1500 words seemed arbitrarily limiting.
Result:	Although various versions of generative AI models have been available to the public for the past few years, it was the release of ChatGPT that has gotten everyone's attention. Just two months after its launch, ChatGPT reached 100 million monthly users, making it the fastest growing software application in history.	Generative Artificial Intelligence (AI) has the potential to revolutionize the way businesses operate, automate repetitive tasks, and create new products and designs. However, with any new technology, there are also legal implications to consider. Organizations using generative AI at work must be aware of the top legal issues that arise

| **Author: Human (Shabbi Khan)** | **Author: ChatGPT-3.5 (Prompted by Kate Wegrzyn)** |

Generative AI refers to artificial intelligence that can generate new content, such as text or images. Generative AI software include complex machine learning models that have been trained on massive amounts of publicly available data, such as websites, images, videos, among others. To generate text, the software uses these machine learning models to predict the next word based on previous word sequences until the complete passage of text is generated. Similarly, to generate images, the software uses machine learning models to predict the next part of an image based on other images that include similar portions until a complete image is generated.

Many professionals have started using generative AI technologies at their workplaces, for example, to generate blog articles, draft emails, generate graphics, with some even daring to use ChatGPT as an automated chatbot. The professionals, however, are often using these generative AI technologies without truly understanding the business and legal risks associated with the use of generative AI.

The following are the top ten legal issues that companies and their employees using these generative AI tools should be aware of:

Unauthorized use of Copyrighted Data to Train Generative AI Models

Generative AI models are trained on massive amounts of publicly available data including copyrighted data. There are concerns that the use of copyrighted data to train these machine learning models violates the copyrights of the copyright owners. Courts in the U.S. have not yet weighed in on whether generative AI systems use of copyrighted data is permissible under the fair use doctrine. That said, other jurisdictions don't have such permissive fair use exceptions as the U.S. and so therefore, there may be exposure to generative AI companies in other jurisdictions outside the U.S. Because it is uncertain whether generative AI companies will be allowed to use these models in the future if they are found liable for copyright infringement, there is concern about end users or companies that have produced output using such models and whether liability can extend to such end users or companies.

Does the Output Violate Copyright Laws?

Separate from the liability that may be imposed on generative AI companies that have trained these generative

when using this technology. In this article, we will discuss the top 10 legal issues that organizations face when using generative AI at work.

Intellectual Property Rights

One of the primary legal issues in using generative AI at work is intellectual property rights. When generative AI is used to create new products or designs, the question of who owns the rights to these creations becomes a concern. In many cases, the output generated by generative AI may be considered a "joint work," with both the AI system and the human users sharing the rights. However, the specific ownership and distribution of rights can be complex and can depend on the circumstances of each case.

Data Privacy

Another legal issue that organizations face when using generative AI at work is data privacy. AI systems require vast amounts of data to train and improve their performance, and this data can often be sensitive or personal in nature. Organizations must ensure that they comply with all relevant data privacy laws and regulations, including the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA).

Employment Law

AI systems can automate many tasks that were previously performed by humans, and this can have a significant impact on the employment landscape. Organizations must ensure that they comply with all relevant employment laws and regulations, including fair selection processes for which tasks will be automated and providing appropriate training and support to affected employees.

Liability

When an AI system generates a product or service, the question of who is responsible for any harm that may arise from the use of that product

Author: Human (Shabbi Khan)	Author: ChatGPT-3.5 (Prompted by Kate Wegrzyn)
AI models, there are concerns that if the output of a generative AI system is too similar to a copyrighted work, the output may violate copyright laws and therefore, expose the entity that distributes or publishes the output to copyright infringement claims. The biggest risk here is that because users are not aware of all of the copyrighted works out there, a user will not know how similar the output is to a copyrighted work and therefore, unwittingly publish the output (likely passing the work as their own). Doing so will not give rise to copyright infringement claims but it can also result in significant damage to the user's reputation.	or service becomes a concern. Organizations must ensure that they have appropriate liability insurance in place to cover any potential claims.

<div style="display:flex">
<div style="width:50%">

AI models, there are concerns that if the output of a generative AI system is too similar to a copyrighted work, the output may violate copyright laws and therefore, expose the entity that distributes or publishes the output to copyright infringement claims. The biggest risk here is that because users are not aware of all of the copyrighted works out there, a user will not know how similar the output is to a copyrighted work and therefore, unwittingly publish the output (likely passing the work as their own). Doing so will not give rise to copyright infringement claims but it can also result in significant damage to the user's reputation.

Confidentiality of Input of Prompts

Prompts are the queries a user inputs into a generative AI system to generate an output. These prompts can be used by the generative AI software to improve their models and for other purposes. As such, care should be taken to avoid sharing confidential or sensitive information as the generative AI system can incorporate the prompts to generate outputs for other users. Many generative AI systems mention using third party contractors to review both the input and the output for safety and moderation purposes, which means that submitting confidential information as a query may result in liability stemming from a breach of confidentiality.

Data Ownership of the Prompt and the Output

Users of generative AI software should review the Terms of Use of the generative AI system to understand the ownership rights associated with the input or prompt and the output generated by the AI system. The user should understand what rights, if any, the generative AI system has in the input and how the AI system may use it. Similarly, the user should understand what rights, if any, the generative AI software and the user have in the output, and what restrictions the user may have on the output. Because the generative AI system may have rights to the output data, it is possible that the generative AI system may reproduce the same output for another user. This can result in claims for copyright infringement and plagiarism claims so users should proceed with caution when dealing with outputs generated from the AI system.

Authorship

Who is the author of the output of a generative AI system? Is it only the person who input the prompt or is it only the

</div>
<div style="width:50%">

or service becomes a concern. Organizations must ensure that they have appropriate liability insurance in place to cover any potential claims.

Discrimination

Organizations must also be aware of the potential discrimination risks that may arise from the use of generative AI. AI systems must be designed to be inclusive and avoid biases that could impact specific groups. Regular monitoring and audits should be conducted to ensure that the AI systems are working fairly and effectively.

Fairness

The use of generative AI must also be fair and transparent. Organizations must ensure that the AI systems are not making decisions that are discriminatory or biased and must be able to demonstrate how decisions were made if necessary.

Contract Law

Organizations must also be aware of contract law when using generative AI at work. Contracts must clearly define the rights and responsibilities of the AI system, the human users, and any other parties involved in the creation of the output.

Regulatory Compliance

Organizations must ensure that they comply with all relevant regulations, including industry-specific regulations, such as financial services regulations or healthcare regulations.

Cybersecurity

Cybersecurity is a major concern when using generative AI at work. Organizations must implement appropriate security measures to protect the data, including encryption, access controls, and monitoring systems.

International Law

Organizations must also be aware of international law when using generative AI at work. Different countries have different

</div>
</div>

Author: Human (Shabbi Khan)	Author: ChatGPT-3.5 (Prompted by Kate Wegrzyn)

generative AI model that produced the output or is it a combination? OpenAI suggests users mention the output was generated in part using its generative AI models. For instance, OpenAI has provided some stock language that a user may use to describe the creative process, which states the following:

"The author generated this text in part with GPT-3, OpenAI's large-scale language-generation model. Upon generating draft language, the author reviewed, edited, and revised the language to their own liking and takes ultimate responsibility for the content of this publication."

Failure to accurately list the authorship may result in potential liability. In particular, it may violate the Terms of Use of certain generative AI systems, some of which require that the author may not represent that the output from the generative AI software was human-generated when it is not. And because the same output may be generated for another user, failing to represent that generative AI was used to generate the content could result in claims of misrepresentation once detected.

Seeking Copyright Protection on the Generative AI Content

Copyright protection for AI-generated works vary from country to country. For example, in the United States, copyright laws do not protect works created solely by a computer but works in which an individual can demonstrate substantial human involvement may qualify for copyright protection. In the United Kingdom, works generated completely by a computer can be protected. In the European Union, things are less clear and it allows for human creativity to be expressed through an AI system. Without copyright protection on certain works of art, companies may not be able to enforce their rights over others in the case of blatant copying. This may be important for media companies or individuals that need copyright protection on the works that they generate. Accordingly, it is important for companies to understand the risks of not having copyright protection on such works.

Bias in Outputs

Users may consider using the generative AI software for various use cases. For example, using ChatGPT as a chatbot or for evaluating resumes of candidates or for creative writing. In each of these use cases, the output

laws and regulations regarding the use of AI, and organizations must ensure that they comply with the laws and regulations of each jurisdiction in which they operate.

Conclusion

Generative AI has the potential to transform the way businesses operate, automate repetitive tasks, and create new products and designs. However, organizations must be aware of the legal implications of using generative AI at work.

Author: Human (Shabbi Khan)	Author: ChatGPT-3.5 (Prompted by Kate Wegrzyn)

from the generative AI software may exhibit bias that if left undetected, may result in discriminatory behavior. Discriminatory or offensive behavior can result in negative public relations crises, litigation exposure, and civil penalties. Accordingly, users and companies should monitor the use of generative AI to ensure that the systems are not exhibiting bias.

Factual Inaccuracies in Outputs

Generative AI systems are great for automatically generating content. However, the content that is generated may not be factually correct. One of the biggest challenges with generative AI systems is the concept of hallucinations – which is the generative AI's ability to make up information that seems true but is not. The risk of hallucinations increases when the generative AI model is asked to output larger amounts of information. Users and companies should ensure that the AI generated content they use is factually correct to avoid reaching incorrect conclusions or taking actions that are based on the factually incorrect information. Failure to do so may result in exposure to the company if others rely on the information to take ill-advised actions. For instance, a generative AI chatbot in a medical setting may hallucinate and tell the patient to take a medication to which the patient may have a severe allergic reaction.

Algorithmic Disgorgement

Several generative AI software companies are facing legal challenges due to copyright infringement claims, and it is unclear what liability, or penalty, these companies may face. A more recent enforcement mechanism, referred to as algorithmic disgorgement, has been used by the Federal Trade Commission (FTC) to penalize companies that have illegitimately collected data and used that data to train an algorithm by requiring a company to either delete the data they collected, or delete the algorithm or models the company developed using the data. The threat of algorithmic disgorgement can extend to the output generated from these models and as such, users may be at risk of losing their work product that was generated using generative AI models that have violated other entities' rights.

Indemnification

Various generative AI software require users to indemnify the generative AI software as part of the Terms of Use.

Author: Human (Shabbi Khan)	Author: ChatGPT-3.5 (Prompted by Kate Wegrzyn)
As such, copyright owners may consider suing a user of the generative AI software for use of the generative AI software that has been trained using the copyright owner's copyrighted data. This risk of exposure becomes higher when using generative AI models to generate an image as output that is substantially similar to copyrighted works of a particular visual artist or if the output inserts a watermark or other insignia indicating that the model was trained using copyrighted data of the visual artist or image source. As companies and employees begin to adopt generative AI solutions in their workplace, they should understand the risks associated with the adoption of generative AI and establish policies that will help mitigate such risks.	

Context

It is not lost on us that in this experiment we, being lawyers, used legal content as our test case – specifically asking Cathy to spot and write about the legal issues surrounding the use of generative AI. As lawyers, we are especially concerned about providing accurate content for both our client work and, as here, where our analysis is offered for general information and not considered legal advice. But accuracy is critically important in almost any field.

Thus, as noted above, even before Shabbi would post his work for public consumption, we said he would normally consult with more knowledgeable attorneys (which he didn't do here.) The same is true for anyone who relies on Cathy. The content she generates should not be used on its own, and should be checked and vetted before anyone relies on it or offers it up to someone else as authoritative. Needless to say, the content generated by this experiment is merely an illustration and is not intended to be relied upon for its substantive accuracy.
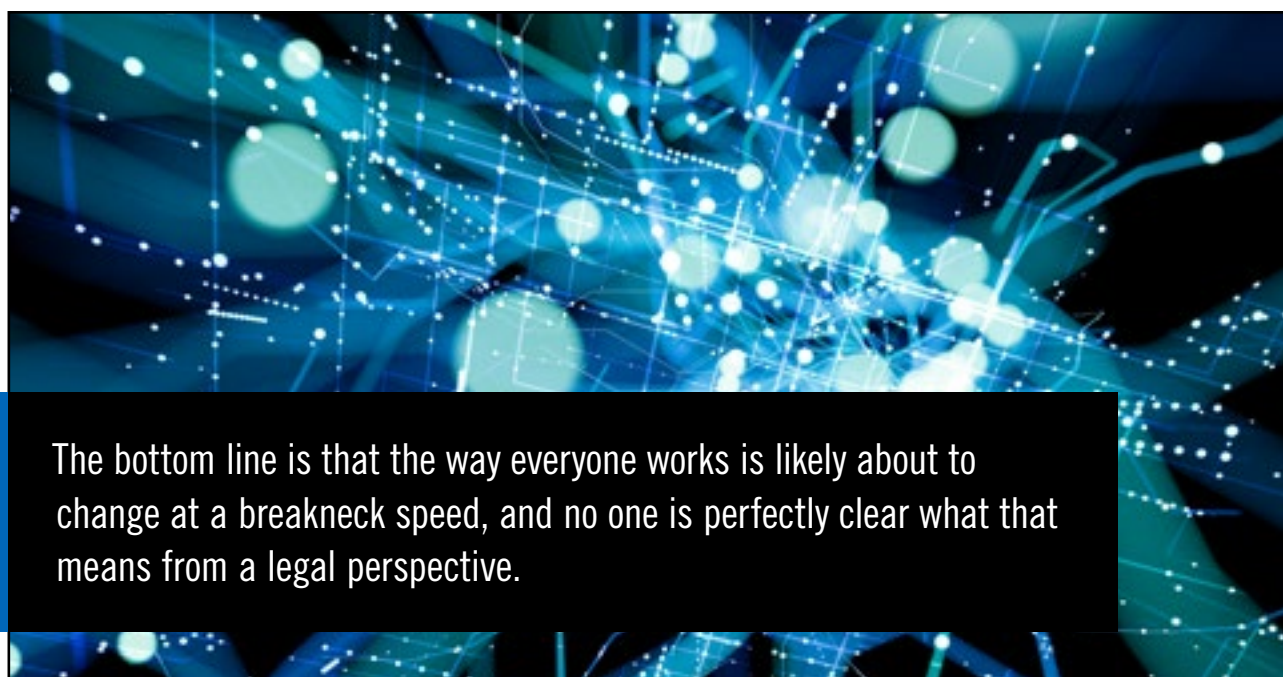
The bottom line is that the way everyone works is likely about to change at a breakneck speed, and no one is perfectly clear what that means from a legal perspective.

Key Legal and Operational Risks for Enterprise AI

Published April 2023 by Legaltech News

AUTHOR

Louis Lehot | llehot@foley.com

The ability of machines to learn and improve without explicit instructions has the potential to revolutionize many industries, but businesses that use AI must be aware of the legal and operational risks that come with it.

The adoption of artificial intelligence tools in the enterprise is set to accelerate as venture capitalists and large corporations alike deploy billions of dollars toward creating and releasing new foundation models. Following ChatGPT's rapid adoption, arguably the fastest adoption of a new technology application in the history of science, there is no turning back. Fasten your seatbelts and prepare for the disruption. The ability of machines to learn and improve without explicit instructions has the potential to revolutionize many industries, from health care to finance and up the courtroom steps and ultimately into the judge's chambers. However, we all know that with great power comes great responsibility, and businesses that use AI must be aware of the legal and operational risks that come with it.

Potential for AI in the Enterprise

As evidenced by the billions of dollars in new capital being deployed to OpenAI, Anthropic, Stability AI, and other startups, the opportunities for new businesses to take large data pools and monetize them within new foundation models that provide premium services, will disrupt consumer and enterprise applications forever. Creators of intellectual property have the opportunity to unlock untapped revenue streams. Developers and enterprises will pay for these premium services, potentially on a volumetric basis. At last, enterprises will be able to consume advanced analytics scalably.

Public markets reward those enterprises out in front (e.g., Nvidia) and punish those who look to be replaced by ChatGPT and its progeny (e.g., Buzzfeed).

A recent study published by MIT suggests that the deployment of AI in the enterprise is rapidly accelerating across functions and departments, landing and expanding at greater velocities. A recent Pitchbook Data report estimated the value of the AI and machine learning market at US$197.5 billion in end user spending in 2022 and forecasts that spending will double by 2025.

Amid the hype-cycle, in a somewhat stunning development an unexpected group of bedfellows across academia, the entrepreneur community, "big tech" and old economy corporate executives self-styled as the "Future of Life Institute" came together to publish an open letter and policy recommendations demanding that governments around the world mandate a moratorium for all AI labs to pause, for six months at least, the training of AI systems more powerful than GPT-4. The moratorium was demanded to protect against the profound risks to society and humanity in light of the lack of planning for or management of machines that no one — not even their creators — can understand, predict, or reliably control. The "pause" is demanded so that AI labs and independent experts can jointly develop and implement shared safety protocols for advanced AI design and development that would be rigorously audited and overseen by independent outside experts, ensuring safety beyond a reasonable doubt. Simply put, the open letter demands that humanity design AI governance to ensure that humans control machines rather than ceding control to machines.

While the open letter has yet to lead to governmental action, entrepreneurs, executives, and investors would do well to consider resolving ownership questions regarding underlying intellectual property and mitigate the obvious legal and operational risks to AI development and deployment as they design the roadmap forward.

## Fundamental Questions of IP Ownership

Artificial intelligence engines are receiving terabytes of data in the form of text, video, audio, and images (the "inputs"), running large language models and algorithms on this data, and then generating responses to queries (the "outputs"). In other words, given that the inputs inform the outputs, the debate is raging about who owns the intellectual property created by the AI engines. Who is the creator? Is it the author of the original content that the AI engine used to train itself, or is it the engine's designer that created the outputs? Another question before the courts is whether the outputs can benefit from copyright laws at all, given that there is no human creator. Does the AI training process infringe copyrights in other works? Do AI outputs infringe copyrights in other works? Is the process of scraping copyrighted data and using

it to train AI engines that create outputs a "fair use" that is protected under U.S. copyright laws? Lawsuits from authors and artists demanding compensation are dropping all over the United States. The consequences of losing could be significant.

Legal Risks of AI

Another of AI's most significant legal risks is the potential for bias. AI systems are as good as the data they are trained on. If that data is biased, the AI system will also be biased. This can lead to outcomes that violate anti-discrimination laws. For example, an AI hiring system trained on historical data that reflects biased hiring practices may perpetuate that bias and result in discrimination against certain groups.

Another legal risk of AI is the potential for violating privacy laws. AI systems often require access to large amounts of data. If that data includes personal information, businesses must comply with relevant privacy laws, such as the General Data Protection Regulation (GDPR) in the European Union or the California Consumer Privacy Act (CCPA) in the United States.
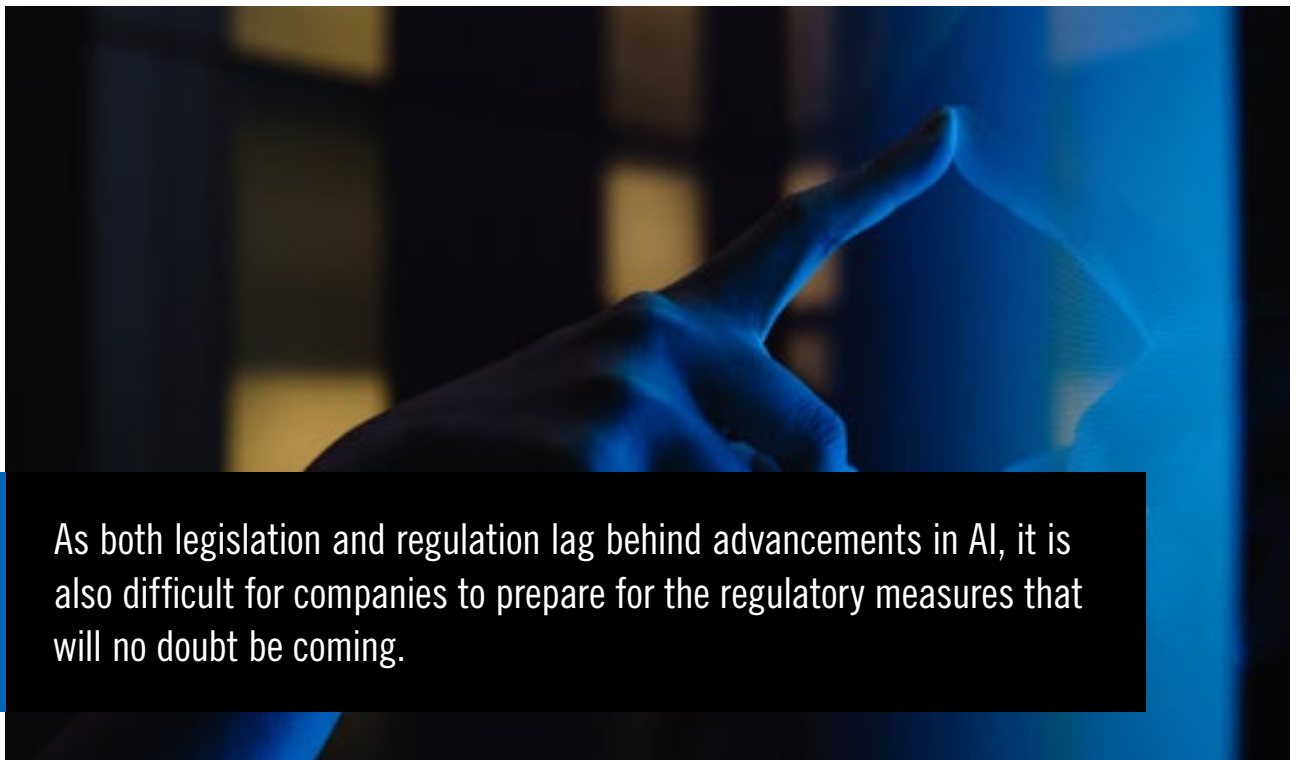
As both legislation and regulation lag behind advancements in AI, it is also difficult for companies to prepare for the regulatory measures that will no doubt be coming.

> AI systems can be complex and challenging to understand, and they may produce surprising results that can harm a business's operations or reputation.

As both legislation and regulation lag behind advancements in AI, it is also difficult for companies to prepare for the regulatory measures that will no doubt be coming. These will, of course, have significant legal concerns as they begin to roll out over time. It will be important that companies developing and utilizing AI are prepared to make some substantial shifts if necessary.

Operational Risks of AI

In addition to legal risks, the operational risks associated with AI must also be considered. One of the biggest is the potential for unexpected outcomes. AI systems can be complex and challenging to understand, and they may produce surprising results that can harm a business's operations or reputation. Some very public examples of well-known AI systems making poor recommendations or diagnoses resulted in significant consequences.

There is also the risk of user error. AI systems must still be operated and ultimately controlled by humans, so the risk of user error will always exist. The results can range from an embarrassing situation for a company to substantial harm to reputation, business operations, or customers.

Another operational risk is the potential for cyberattacks. AI systems can be vulnerable to attacks by hackers who may attempt to manipulate the system for their gain or to disrupt the business's operations. This can lead to data breaches, financial losses, or other types of damage.
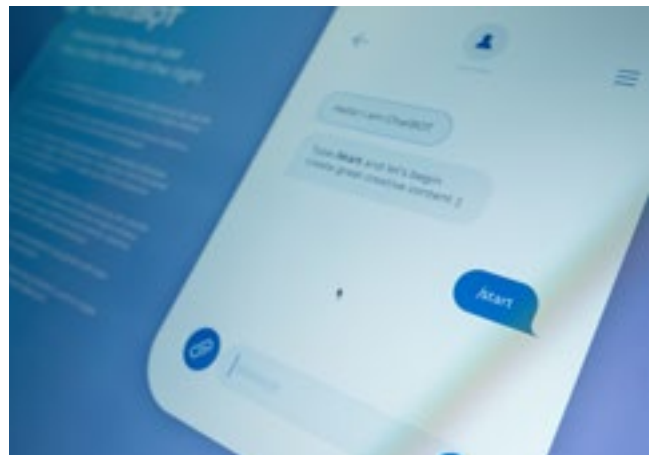
Mitigating Risks

Businesses should take several steps to mitigate AI's legal and operational risks. First, they should ensure that their AI systems are designed with fairness, accountability, and transparency in mind. This may involve developing guidelines for the collection and use of data, testing the system for bias, and providing explanations for the decisions made by the system.

Second, businesses should ensure they comply with relevant laws and regulations, such as anti-discrimination and privacy laws. This may involve conducting regular audits of the AI system to ensure compliance and providing training to employees on the legal and ethical implications of AI.

Third, businesses should develop robust cybersecurity measures to protect their AI systems from cyberattacks. This may involve encryption and multi-factor authentication, conducting regular vulnerability assessments, and training employees on cybersecurity best practices.

Finally, consideration should be given to licensing inputs and compensating artists and authors for outputs. Watch this space. It could change who are the winners and who are the losers.

Businesses that use AI must be aware of the legal and operational risks that come with it. By designing their AI systems with fairness, accountability, and transparency in mind, complying with relevant laws and regulations, developing robust cybersecurity measures, and fairly compensating creators, businesses can mitigate these risks and reap the benefits of this powerful technology.

My Employees Are Using ChatGPT. What Now?

Published July 2023 by Foley & Lardner LLP

AUTHORS

Natasha Allen | nallen@foley.com

Chanley Howell | chowell@foley.com

David Kantaros | dkantaros@foley.com

Shabbi Khan | skhan@foley.com

Carrie Hoffman | choffman@foley.com

Michael Ryan | mryan@foley.com

In the rapidly evolving world of artificial intelligence (AI), one development stands out for its transformative potential: the rise of generative AI tools. Many major technology companies are building the large language models (LLMs) that power these tools, training them on billions of inputs. But among its peers, OpenAI's ChatGPT has emerged as a game-changer; becoming the fastest web platform to reach 100 million users. This milestone is not just a testament to the tool's capabilities but also a clear indication that generative AI is here to stay.

The Proliferation of ChatGPT in the Workplace

ChatGPT is a LLM that has been fine-tuned to be useable as a general purpose chatbot. The current base models are OpenAI's GPT-3.5 and GPT-4 LLMs. ChatGPT understands and responds to natural language prompts and is beginning to find its way into various professional settings. From small startups to multinational corporations, employees across the spectrum are leveraging this tool to enhance their productivity and streamline their workflows.

The applications of generative AI tools in the workplace are diverse. They are being used to draft content, generate documents, conduct fact-checking and research, and even write software code. This widespread use, while enhancing productivity, also brings with it a host of potential risks that organizations need to address. The integration of AI into the workplace is not a simple plug-and-play scenario; it requires careful consideration and strategic planning.

The Need to Provide Guidance to Employees

Given the potential risks associated with the use of ChatGPT and similar tools, it's crucial for companies to provide guidance to their employees. This guidance can take the form of a formal policy or a more general best practice guide that leans on existing information security policies implemented by the company.

While a few large companies may have the resources to build their own internal LLMs, most companies do not. For these companies, adopting the use of third-party generative AI tools safely and with guidance can help them remain competitive against companies with access to their own LLMs. However, this doesn't mean that smaller companies should rush to adopt these tools without due diligence. The potential risks and challenges associated with the use of AI tools like ChatGPT should be carefully evaluated and mitigated.

Understanding the Risks

The use of ChatGPT and other LLMs in the workplace can pose several risks, including:

Confidentiality: Sharing confidential company or client information with generative AI systems may violate contractual obligations or expose trade secrets to public disclosure. This is a significant concern, especially for companies that handle sensitive data. Employees need to be aware of the potential risks associated with sharing confidential information with AI tools. In addition, most generative AI tools are cloud-based or software-as-a-service, meaning that data is being sent to a third-party service provider. If companies provide confidential information to these third-party platforms and the information is then fed back into the model for training purposes, companies could lose trade secret protection for that information.

Personal Data and Privacy Violations: Sharing personal information about customers, clients, or employees with generative AI systems can create privacy risks. This is particularly relevant in the context of stringent data protection regulations such as the General Data Protection Regulation (GDPR) in the European Union and the increasing number of state data privacy laws in the United States, including the California Consumer Privacy Act (CCPA). Further, if private information ever enters the training data of a LLM, e.g., the information is scraped from a data leak, it is feasible that a malicious actor can extract the private information from the model.

Quality Control: Generative AI tools, while remarkable in their capabilities, are susceptible to producing erroneous outputs, leading to potential quality control issues. This propensity to inaccuracies may be further exacerbated by the phenomenon known as automation bias, where users over-rely on the outputs of these AI tools, often without questioning their accuracy. Despite the substantial enhancement of productivity that AI tools offer, they are not impervious to faults. The dangerous part lies in the fact that generative AI tools can produce incorrect results in a very convincing manner, mimicking human-like generation and causing users to trust in their authenticity.

As these generative AI tools evolve over time, they will likely improve in accuracy and exhibit fewer 'hallucinations' or false creations, making it increasingly challenging to detect incorrect information. Don't be fooled – use of the term 'hallucinations' is simply a euphemism for 'mistakes'. This reduced detectability, coupled with automation bias, escalates the likelihood that individuals will become more prone to accept outputs without thorough scrutiny, thereby increasing the potential for misinformation to be propagated or erroneous decisions to be made. Companies in particular (and by extension, their employees) need to be vigilant and cannot absolve themselves of responsibilities and liabilities by solely relying on the use of generative AI tools. It is incumbent upon them to ensure that any information generated and acted upon is as accurate and reliable as possible.

Bias and Discrimination: Generative AI systems can produce biased and discriminatory results. As LLMs are trained on data available on the internet, they are capable of repeating biases found there. If companies rely on generative AI tools, they need to ensure that they do not engage in any biased or discriminatory actions based on the use of these tools. These risks are particularly present and acute in connection with employee recruiting, screening, and hiring.
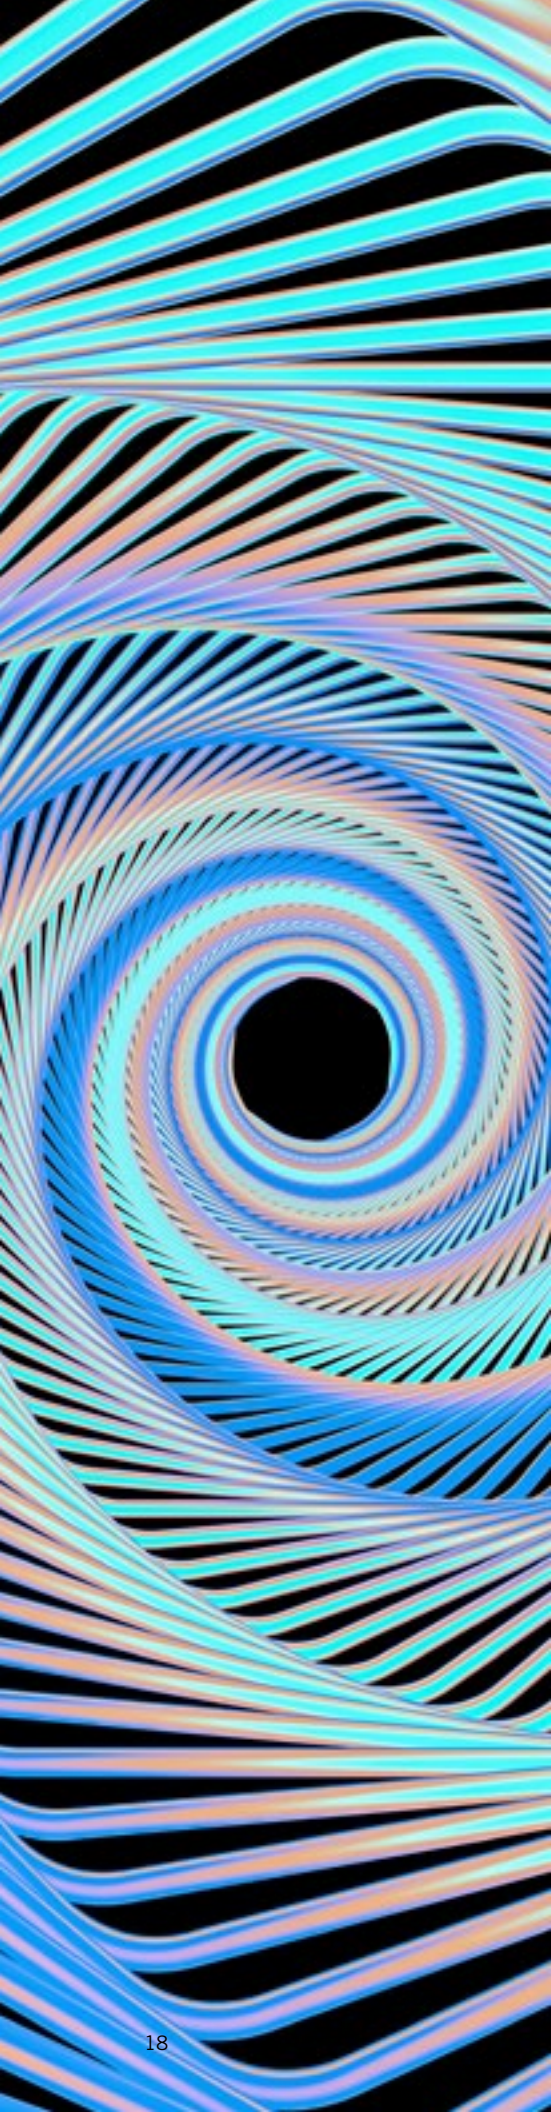
Product Liability: Generative AI tools can be used in product research, design, development, and manufacturing phases. Products may be physical (e.g., construction products) or software-based (such as autonomous driving technology). If a product or system powered by AI makes a decision that harms a user, it could result in claims and liability for all actors and organizations in the "chain" of the AI development and use.

Intellectual Property Ownership: The use of generative AI systems raises complex IP issues, including whether documents or code generated by generative AI systems are entitled to legal protection or whether the company can be held liable for using the output of generative AI systems. The U.S. Copyright Office has stated it "will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author." This is a grey area in the law that is yet to be fully resolved.

Separately, there is still some uncertainty around the ownership of data generated by generative AI tools. The terms of use of certain tools mention that the user owns the data. However, owners of the data that was used to train the LLM may also have certain ownership claims.

Misrepresentation: Claiming output is human-generated when it's not can lead to consumer protection claims or other public relations concerns depending on its use. Companies should be aware that there is a risk of unfair or deceptive practice claims under state or federal law if they are incorrectly using AI tools. Accordingly, transparency is key when using AI tools. In addition, social media is notorious for calling out content that was created using generative AI but not disclosed, which at a minimum would be a blight on the user's reputation and credibility.

Insurance Coverage: Depending on the policy, insurance may not provide coverage for liability resulting from the use of generative AI tools. As generative AI tools become more integrated into business operations and automate more functions, the possibility of adverse events rises, bringing about additional exposure to companies. It is possible these tools fall outside the scope of existing policies

Future Requirements: There may be future requirements to clearly identify AI-generated content if a company needs to make representations in a transaction (such as a sale or financing) or in connection with a commercial agreement with a vendor or a customer. Other regulations relating to transparency, consent, and notice will likely be enacted.

Potential Employment Discrimination: The use of generative AI systems may adversely affect the performance of individuals who are not using it relative to their peers. This could potentially lead to employment decisions stemming from the use or non-use of AI tools, which may have a discriminatory or adverse impact on a protected class of individuals (e.g., persons 40 and over). Companies need to ensure that the use of AI tools does not create an unlevel playing field within the workplace or potential bias in employment decision-making.

Addressing the Risks

Companies should provide employees guidance on how to use generative AI tools responsibly, including how to avoid the risks associated with their use. The guidance can come in the form of a new generative AI acceptable use policy or something less formal, such as a best practices guide. In addition, companies should consider additional mechanisms including systems for monitoring internal use and providing procedures or mechanisms for reporting inadvertent sharing of confidential information with the generative AI systems.

> If companies decide to directly provide the results of generative AI tools to their clients and customers without human review, they must understand the higher level of risk it entails as these tools can produce incorrect or even offensive results.

Confidentiality and Privacy: In light of potential data breaches and the privacy risks associated with sharing personal data, the company's generative AI acceptable use policy or guidance should clearly define what data can and can›t be shared with AI tools, with a particular emphasis on protecting sensitive company, client, and employee information. This approach helps ensure contractual obligations and trade secrets are maintained while complying with data protection regulations such as GDPR. Further, if companies use generative AI tools offered by third-party service providers, they must do so while abiding by contractual obligations to their clients, vendors, employees, etc. and confirm that they are authorized to share the information with the platform.

Employees should be asked to opt out of having any data be used for machine learning training if such an option is available. For certain service providers, users may opt out by sending an email or filling out a form provided by the vendor.

Companies should also consider providing a reporting mechanism through which employees can report to management if confidential or sensitive data was inadvertently shared with a generative AI tool. This reporting mechanism could involve informing a manager or sending an email to a designated email address managed by the information security team.

Quality Control and Factual Inaccuracies: Due to the potential quality control issues arising from inaccuracies in AI-generated content, the company's acceptable use policy or guidance should ask that employees review any output generated for accuracy and correctness. Employees should apply their expertise and exercise sound judgment on how best to use the output. In addition, if companies decide to directly provide the results of generative AI tools to their clients and customers without human review, they must understand the higher level of risk it entails as these tools can produce incorrect or even offensive results.

Intellectual Property Risks: Because copyright laws vary across jurisdictions, companies may not have copyright protection over the data output from a generative AI solution. Accordingly, others may be able to copy the data without any risk of copyright infringement. Separately, there are also concerns that owners of the data used to train a generative AI tool may have certain ownership claims on the output generated. Companies should be judicious about the use of such content and understand that they may not have clear title of ownership or copyright protection on the generated content.

Misrepresentation: To avoid misrepresentation claims and potential public relations issues, companies need to be transparent about their use of AI tools. In a world where consumers value authenticity, revealing that content was generated using generative AI is not just ethical, but could be appreciated by the audience. Certain jurisdictions already have notice provisions that require informing consumers about the use of a chatbot or other automated bots, and additional regulations will likely be passed that may require providing notice of the use of generative AI tools.

Companies should be reminded that taking a very conservative approach to the use of generative AI tools may create an environment where employees may still use these tools but on personal devices and outside the purview of the company.

Insurance: With the risks associated with AI usage, insurance policies need to evolve in tandem. Companies must engage in discussions with insurers and insurance brokers to ensure coverage extends to potential liabilities resulting from AI use.

Regulations: Given the fast-paced evolution of AI regulations, it's important for businesses to stay ahead of the curve. Anticipating and preparing for potential regulatory changes will keep the company compliant and limit future disruptions.

Employment: To ensure fair play within the workplace, companies need to monitor the use of AI tools like ChatGPT. Measures should be implemented to ensure that the performance evaluation of individuals not using AI tools is not adversely affected and consider training for those employees who are not early adopters of generative AI tools, thereby mitigating potential discrimination. Moreover, the company's generative AI acceptable use policy or guidance should consider restricting the employees use of personal email accounts to log into AI tools for work-related activities.

More generally, each company should evaluate these risks and assess what types of measures they should take when implementing generative AI acceptable use policies. Companies should be reminded that taking a very conservative approach to the use of generative AI tools may create an environment where employees may still use these tools but on personal devices and outside the purview of the company. Therefore, a more nuanced approach that understands the use cases where employees can use generative AI and determines what conditions and circumstances are allowed is advised.

The Time to Act is Now

The growth of ChatGPT and other generative AI tools is not slowing down. As these tools become more integrated into our professional and personal lives, the importance of addressing the associated risks becomes increasingly critical. Companies must be proactive in understanding these risks and implementing strategies to mitigate them.

In the face of potential future regulations and the evolving nature of AI technologies, it's essential for organizations to stay informed and adaptable. As we navigate this new landscape, the key to harnessing the power of generative AI tools lies in striking a balance between leveraging their potential and managing their risks.

The rise of ChatGPT and other generative AI tools presents both opportunities and challenges. By providing clear guidance to employees, transparency with consumers, understanding the associated risks, and implementing robust policies, companies can navigate this new landscape responsibly and effectively.

Click here to download our cheat sheet with guidance on addressing specific areas of risk.

The authors gratefully acknowledge the contribution of Mathew Cha, a student at UC Berkeley School of Law and 2023 summer associate at Foley & Lardner LLP.

How Artificial Intelligence is Disrupting the Dealmaking Process

Published September 2023 by Westlaw

AUTHORS

Louis Lehot | llehot@foley.com
Eric Chow | echow@foley.com

The use of artificial intelligence (AI) and machine learning (ML) technologies has exploded across the board. According to EY's July 2023 CEO Outlook Pulse,1 almost half of the CEOs who participated are focusing capital allocation on these technologies, with 43% responding they have already fully integrated AI-driven product or service changes into their capital allocation process and are actively investing in AI-driven innovation.

As more businesses adopt these technologies, it has become clear that there are countless business applications for AI and ML. This includes their ability to transform the merger and acquisition process.

Negotiating a deal and running due diligence in a merger or acquisition has always been complex, time-consuming, and resource intensive. However, AI and ML are changing the game and making the process much more efficient on many levels. These transformative technologies can streamline deal processes with their ability to better analyze data and improve outcomes, particularly regarding due diligence and contract analysis.

In any transaction, participants exchange long lists of questions, requests, and spreadsheets demanding documents, data, and information, generally referred to as due diligence request lists or DDRLs. This has traditionally been a cumbersome process, involving mass amounts of digital and paper files and teams of bankers, lawyers, accountants, and analysts who must review the documents, glean information, negotiate agreements, produce disclosure schedules, and allocate risks and rewards.

This largely human and manual process of digesting the information could be more efficient, as well as the source of many budgets exceeded, details missed, questions never asked, and opportunities for synergies lost.

However, with the introduction of AI in the dealmaking landscape, the due diligence process is becoming more efficient as it is able to analyze what can be overwhelming amounts of data in a fraction of the time.

Tools like Robotic Process Automation (RPA) can provide tremendous potential for time and cost savings as they automate data analysis and can quickly flag anything that would require further review. AI tools also provide predictive modeling tools, allowing companies to examine various scenarios and what kinds of risks might be involved.

Companies such as Deloitte are launching products to assist with the due diligence process. Their iDeal product uses AI and ML to organize and tag the huge amounts of data involved. It can also learn from corrections made by humans, making it more accurate the more it is used.

Numerous companies are also launching AI products designed to improve the contract review process. Take, for example, Kira Systems, a machine learning contract and document review software that can scan and analyze vast amounts of data sources and contracts. This kind of software allows users to analyze complex contracts and documents with a much higher level of efficiency, saving time and money and allowing for identifying and mitigating risks.

Spellbook, Syntheia, and IronClad are other examples of companies focused more specifically on contract review. While these tools do not replace human judgment and experience, they help significantly speed up what is traditionally a prolonged process and help identify red flags that can turn into major issues down the road.

The use and benefits of AI in transactions go beyond just due diligence and contract review, extending into target identification and valuation — two areas that can often present difficulties.

Target identification requires extensive research and detailed analysis of, again, what can be mass amounts of data. In order to better match targets and buyers, AI tools can compound numerous data sets to pinpoint patterns that humans cannot otherwise recognize, enhancing the target identification process.

The use and benefits of AI in transactions go beyond just due diligence and contract review, extending into target identification and valuation.

It is important to remember that AI tools do not take into account corporate value compatibilities and visions of the buyer, but they can help narrow down available targets and provide insights. They can also analyze market trends, competitor performances, and customer behavior to identify the impact of a target and a company's most suitable acquisition targets. An essential part of this process is the technology's ability to assess the risks of a transaction or the risk indication of a potential target.

Determining a "fair" valuation can be one of the most difficult parts of the transaction process, but the predictive modeling techniques we discussed earlier can also assess the target's future potential and forecast financial performance more accurately. AI algorithms can analyze historical financial data, market trends, and macroeconomic factors to project potential growth scenarios.

AI also has the capability to analyze thousands of previous valuations and learn from them, putting that knowledge to use in the valuation process. This data-driven valuation approach, combined with the human intuition of dealmakers, helps negotiate better deals and avoid overvaluation or undervaluation of the target company.

Potential Issues and Concerns

Despite AI's numerous advantages to the deal landscape, as with any new and developing technology, some issues and concerns must be considered.

Data Privacy and Security: Because of the vast amounts of data processed by AI tools, there are real concerns about data privacy and security breaches when used in the deal context. Sensitive and private information is shared throughout a transaction process, and the last thing parties want is for their data to be compromised. Companies should ensure robust data protection measures to prevent unauthorized access and potential misuse of sensitive information.

Lack of Human Judgment: AI lacks human judgment and intuition. It is important to bear in mind that while it is a valuable tool, AI does not make human involvement obsolete or irrelevant, especially when evaluating qualitative factors and understanding the nuances that can impact a deal's success.

Regulatory Compliance: Relying heavily on AI may inadvertently lead to non-compliance with various regulatory frameworks governing transactions. Similar to the lack of human judgment, some areas of AI technology still need to be fully developed enough to rely on entirely. Compliance requirements may easily be overlooked by AI technology.

AI is an Investment: AI is still in the very early innings, and everyone may not be able to take advantage of it right away. To some organizations, AI means they may have to go out of their way to invest time and money into technology and possibly even personnel. Companies should be mindful of the time and resources required to implement AI tools and carefully weigh the risks and rewards.

AI will undoubtedly continue to transform the way deals are conducted, and companies are already jumping in to launch new tools to automate this process even more. Its potential benefits are vast, enabling people to make more informed and efficient decisions. However, it is vital to strike a balance between leveraging AI capabilities and incorporating human expertise to address serious concerns that exist.

AI is not yet at a place where the technology can replace all human interactions and procedures within a deal, and AI users should be cognizant of the issues and concerns presented above. But with proper safeguards in place, it can serve as a powerful ally in creating more informed and ultimately more successful deals.

> With the introduction of AI in the dealmaking landscape, the due diligence process is becoming more efficient as it is able to analyze what can be overwhelming amounts of data in a fraction of the time.

Retraining the Acquisition Model: How to Approach the Risks and Rewards of Deals With Artificial Intelligence Targets

AUTHORS

Natasha Allen | nallen@foley.com

David Kantaros | dkantaros@foley.com

Shabbi Khan | skhan@foley.com

Avi Ginsberg | aginsberg@foley.com

Graham MacEwan | gmacewan@foley.com

Published February 2023 by Foley & Lardner LLP

In recent years, companies specializing in artificial intelligence (AI) technologies have been increasingly coveted acquisition targets. With the AI field transforming our approaches to key issues – from climate change to cancer treatment – and pushing the boundaries of human capabilities through technologies such as self-driving cars and advanced space travel, it has the potential to dramatically alter life for humankind. Accordingly, it is no surprise that AI has emerged as a prime sector for M&A activity. As of October 2022, AI deals had surpassed $16.9 billion in aggregate value for the year and experts expect deal activity only to grow. Of particular interest are generative AI companies – those that produce algorithms able to create new content based on data inputs. Generative AI's potential capability to revolutionize productivity has led to a considerable increase in investment activity and company valuations since 2020.

As active as the overall market has been, engaging the AI field does not come risk-free to the would-be acquirer. As the field has exploded in recent years, some of its familiar participants have warned of the unknowns and harms associated with the technology and its capabilities, as well as the companies developing it. Beyond the caution with which a traditional M&A transaction is treated, those seeking to purchase an AI company must approach the deal clear-eyed and deploy the resources and expertise to mitigate the associated risks.

The Unique Due Diligence Review

Conducting due diligence on a target in the AI space may be considerably more difficult than on conventional technology companies. Buy-side advisors investigating a traditional technology company can readily review its proprietary software – usually at the heart of the transaction – for standard intellectual property concerns; for example, advisors will assess the company's use of open source code and whether the underlying IP has been properly assigned to the company. The review of an AI company's chief product can prove far more nebulous, as the company's value is often derived from its datasets and proprietary models which absorb and analyze information.

Accordingly, when conducting diligence on an AI company, advisors are encouraged to expand and revise the scope of the review. They should review the company's rights to its models, data, and "outputs" therefrom – a more difficult and nuanced investigation for the uninitiated. Buy-side advisors must draw on expertise in data ownership and privacy, as well as knowledge of the specifics of the AI field, to probe the company for risks to which a traditional technology target might not be as vulnerable. They must undertake a comprehensive review exceeding that of the traditional IP diligence investigation, both to assess the merits of a proposed transaction and to adjust the terms governing it appropriately.

Although this additional diligence may be cumbersome for both the target and buyer – and adds an additional layer of costs to the transaction – completing it can help the buyer better evaluate the AI capabilities of the target. Through this process a buyer can get a better

understanding of various risks to which the target is exposed. For instance, many startup AI companies claim to have a fully automated system when in reality a lot of their processes are still manual; performing the additional diligence can help a buyer understand how far along the target's AI technology actually is. In addition, through diligence, a buyer can determine whether the target's AI technology is capable of handling real-world data or whether it has only been fine-tuned for a particular customer of the target. Diligence can also uncover risks associated with a target's data management and compliance procedures, as well as potential regulatory risks to which the target may be exposed based on the type of data the target collects, processes, and stores.

To help guide your thinking, below is a non-exhaustive list of diligence questions that could be asked of acquisition targets working in the AI space during the evaluation process:

1. Are you using AI, Machine Learning (ML), or any technologies that a reasonable consumer might think are AI or ML?

2. Do you engage in any type of processing that uses models or predictive analysis?

3. Do you engage in tracking and/or modeling of the real world (through the use of cameras, sensors, etc.)?

4. What are the sources of your training data?

   a. Please share all license agreements regarding training data.

   b. Do you scrape any websites/resources to obtain the data?

5. What steps have you taken to ensure adequate data usage rights for all data (first and third party) that you use for development purposes (including to train models, optimization, benchmarking, and debugging)?

6. What steps have you taken to determine which regulations apply to your processing of data?

7. What analysis has been performed to ensure your processing of data is compliant with applicable regulations?

   a. Has a data processing impact assessment been performed?

   b. Has a risk analysis been performed?

   c. Please provide all analysis and an overview of determinations.

8. Does your system store or use biometric data? If so, how have you verified compliance with all applicable regulations (e.g., the Illinois Biometric Privacy Act)?

9. How do you test and/or validate your models?

10. How do you detect and correct for bias in your products and services (including in training data, models, and third-party software components)?

> Generative AI's potential capability to revolutionize productivity has led to a considerable increase in investment activity and company valuations since 2020.

Drafting the Purchase Agreement

The standard form purchase agreement, including its representations and warranties, may not adequately address the risks involved in AI. But the representations and warranties contained in a definitive agreement will describe how the target company uses its AI assets and shift risks associated with those assets. Alongside the diligence review, these provisions (if drafted properly) will provide the buyer with the necessary information to understand more fully what it is acquiring.

As there are risks inherent and unique to the AI space not always encountered in traditional technology fields, the representations and warranties should be tailored specifically to the target company and its assets. The buyer's legal advisors should be careful not to resort to generic, off-the-shelf language for a standard acquisition in the tech space. For instance, the buyer's advisors should draft the representations such that risk is shifted to the seller with regards to the target's rights to use the outputs associated with its AI models. By identifying such specific risks associated with AI and drafting the representations to allocate liability appropriately, a buyer investing in such AI technology can feel increasingly comfortable.

The definitions of the purchase agreement must also be crafted to the specifics of the field. The definitions pertaining to AI and the company's products should be broad enough to capture the various techniques employed to create the AI. They should also adequately address the target's specific niche in AI – a field of remarkable variety and expanse. As these definitions will govern the scope and applicability of the provisions contained in the agreement, it is crucial to tailor them to the specifics of the industry and target.

Additional Thoughts

The active nature of the acquisition market for AI companies itself presents risks to both buyer and seller. Such activity has pressured companies to emphasize rapid growth, which can come at the cost of the long-term health and stability the company as well as its attractiveness as a target for acquisition. A potential buyer would do well to ensure the company has scaled appropriately and efficiently.

And, as significant as the targets' AI products are to many of the recent acquisitions, these transactions are often driven instead by the desire to purchase talent—namely, the targets' coveted AI researchers and engineers. This approach can be fraught with risks, from the bloating of the transaction value to downstream concerns of misalignment and retention. While any M&A transaction can raise questions of target compatibility, such risks are likely heightened by the particularities of the AI field.

Further, the unique, nascent nature of the field lends itself to a host of additional headaches for the would-be buyer—some of which perhaps escape the purview of a traditional due diligence review—from issues regarding the product's quality and consistency to its future legal and ethical stature. While the AI field presents remarkable opportunity in both the narrow and broadest senses, buyers would do well to approach the transaction with the caution and sophistication the field warrants.

Action in Venture Capital is Focused on Early Stage and AI Business

Published February 2023 by Foley & Lardner LLP

AUTHORS

Louis Lehot | llehot@foley.com

Lyman Thai | lthai@foley.com

Many startups have seen their valuations drop in 2023 due to a myriad of issues on a global scale. That being said, the valuations of U.S.-based startups remain strong, particularly for some specific segments. It should come as no surprise that startups at the angel and seed stages are leading the way, with median seed valuations growing by 16.7% to US$10.5 million.

Prior analysis of data shows that while early-stage startups did maintain an overall healthy level of deal activity in 2022, Q3 and Q4 were less robust than Q1 and Q2. There was a downward trajectory from Q1 to Q4 as the market continued to soften, and the early-stage median value fell by 33.3% from the beginning of 2022 until the end of the year.

These very early-stage companies have been able to weather the storm much better than those at later stages, and it's not just for U.S.-based startups. We are seeing this on a global scale. In this current climate, investors are taking bets on these early-stage companies and scaling back on investments later in the startup life cycle. So, it makes sense that this would be reflected in higher valuations at the seed and angel stages.

> These trends will likely continue at least through the rest of 2023, as economic conditions have not dramatically changed since last year.

There continues to be a more difficult road for later-stage companies as they are faced with lower valuations and a lack of exit options, mainly because of a freeze in IPOs. Some findings in this area include:

- Median and average public listing valuations fell to US$214.0 million and US$604.0 million, respectively—their lowest points since 2016.

- Acquisitions remained relatively resilient, with median acquisition valuations declining by just 3.5%.

- Public exit valuation step-ups tumbled as well, falling 32% to 1.05x, offering little upside for investors who participated in the most recent funding round.

These trends will likely continue at least through the rest of 2023, as economic conditions have not dramatically changed since last year. The interest on the investor side seems to be leaning towards these very early-stage companies, but they will not stay at that stage forever and will need to move on to additional rounds of funding down the road. Innovative founders need to give themselves a longer runway so they have more time for conditions to improve before their next investment round.

Meanwhile, for those raising capital in the meantime, we can share anecdotally that we are seeing businesses lead their pitch with how AI is integrated into their product roadmap, its impact on AI infrastructure, and how AI accelerates their penetration of vertical markets.

The Most Critical Factors for AI Legal Compliance: Transparency & Explainability

Published September 2023 by Foley & Lardner LLP

- As artificial intelligence (AI) and generative AI (GAI) continue to evolve and become integral to business operations, businesses must be mindful of the risks associated with deploying AI solutions.

- Although there is not yet a comprehensive law governing AI, regulators have tools to hold businesses accountable. They are focused on transparent and explainable AI solutions to ensure that consumers and key stakeholders understand how these systems operate and make decisions.

- Because regulators are focused on these requirements, businesses should develop AI solutions aligned with industry standards, such as leveraging the NIST framework.

AUTHORS

Chanley Howell | chowell@foley.com
Lauren Hudon | lhudon@foley.com

What Is Transparency and Explainability?

Transparency refers to the capability of a user to understand how the AI model is being built, how the data is being used and processed, and how the data is affecting internal weights and biases. In other words, transparency enables users to understand "what happened" in the system.[1] For example, if the AI solution hallucinates, a user should be able to identify what caused the hallucination. Further, transparency allows users to act when an AI solution generates incorrect content or could otherwise lead to negative consequences.[2] Because transparent AI systems enable users to better understand how it creates content, it increases user trust and confidence in its capabilities.[3]

Explainability refers to understanding "how" the GAI system made a decision and makes it easier for others to describe a model, its expected impact, and potential biases. Like transparency, an explainable GAI model increases users' trust in the model's outputs.[4] There are various types of explainability, including global and local-level explainability.[5] Global-level explainability is a general understanding of the algorithm's behavior at a high level.[6] Conversely, local-level explainability is understanding how the algorithm targets individuals, such as making decisions based on an individual's credit or job application.[7]

Transparency and explainability are related concepts but are not interchangeable terms. Transparency provides a user with information about how a solution makes decisions, which allows for external auditing and evaluation. Conversely, explainability offers a user a rational justification for why a solution made a particular decision in a way humans can understand. Transparency and explainability are necessary components of AI governance and increase the credibility and trustworthiness of AI solutions.

Using Transparency & Explainability to Manage Risk

A significant focus surrounding transparent and explainable AI solutions is managing risk. Many AI solutions are inherently opaque, akin to a black box. AI solutions such as GAI models are trained on many inputs, with the most significant models having a corpus size in the trillions. These inputs are then used to determine connections within a neural network, which enable the model to find patterns in the data, make predictions, and generate content. The opaque nature of some AI solutions can make it difficult to manage risk and fully understand how data is used and processed, even by the machine learning researchers that build them. However, the inability to understand an AI solution is not a sufficient defense to legal liability.

> **Although there is not yet a comprehensive law governing AI, the European Union (EU) recently voted to approve the Artificial Intelligence Act to establish rules surrounding explainability and transparency in AI applications.**

To effectively manage risk with opaque systems, businesses need to document design decisions and training data, the structure of the model, its intended use cases, and how and when deployment, post-deployment, or end-user choices were made and by whom. Further, businesses should consider having a policy that notifies a human operator when a potential or actual negative outcome caused by a GAI system is detected.[8] Finally, companies must ensure that using data does not infringe on copyrighted works.[9]

Transparency & Explainability with Third-Party Solutions

Another challenge for businesses deploying AI solutions is managing third-party risk. AI as a service will continue to grow, and companies must be mindful of the risks associated with engaging third parties to develop or operate their AI solutions.[10] Contracts with third parties should include governance structures, risk assessment frameworks, monitoring and auditing protocols, and technical safeguards.[11] Further, businesses should have policies and procedures for using third-party solutions, evaluation criteria, and technical safeguards. Companies should proceed with caution if a third party is not transparent about the risk metrics or methodologies used to develop or train the AI solution, as the lack of transparency presents a considerable level of risk.

Transparency & Explainability Regulatory Landscape

Although there is not yet a comprehensive law governing AI, the European Union (EU) recently voted to approve the Artificial Intelligence Act to establish rules surrounding explainability and transparency in AI applications. The Artificial Intelligence Act has global-level explainability requirements. The Act requires technical documentation of an AI system including, but not limited to, general and detailed descriptions of the AI system, detailed information about monitoring, functioning, and control of the AI system, a detailed description of the risk management system, and a description of any change made to the system over its lifecycle.[12]

Further, the EU General Data Protection Regulation (GDPR) contains transparency requirements related to automated decision-making. Controllers that engage in automated decision-making must comply with the GDPR's transparency requirements. Under the GDPR, automated decision-making means the ability to decide about a data subject based solely on automated means without human involvement.[13] The GDPR prohibits solely automated decision-making that produces legal or similarly significant effects unless the decision is (i) necessary for the performance of a contract, (ii) authorized by EU or member state law, or (iii) is based on the data subject's explicit consent.[14] Under the GDPR, controllers must process personal data in a transparent manner, which includes the data subject's right to receive information about the controller's identity and the nature of the processing, whether or not their personal data is being processed, and if so, the nature of the purposes of that processing, and any personal data breach that is likely to result in a high risk to their rights and freedoms.[15]

Beyond the EU, federal regulators in the United States. also want businesses to provide documentation related to their AI solutions. For example, the U.S. Federal Trade Commission (FTC) recently opened an investigation into OpenAI, the creator of ChatGPT, for potential violations of Section 5 of the FTC Act.[16] In short, the FTC is requiring OpenAI to provide detailed descriptions of each large language model (LLM) product, the data used to train their LLMs, the policies and procedures followed to assess the risk and safety of new LLM products, how they prevent personal information from being included in the training data

for any LLMs, and how their LLM products generate information or statements about individuals.[17]

State regulators are also focused on transparent and explainable AI solutions. New York City passed Local Law 144, enacted on May 6, 2023. The Rule governs the use of Automated Employment Decision Tools (AEDT) and makes it unlawful for an employer or employment agency to use an AEDT to screen candidates and employees unless (i) the tool has undergone a bias audit no more than one year before its use, (ii) a summary of the most recent bias audit is made publicly available, and (iii) notice of the AEDT use and an opportunity to request an alternative selection process is provided to each candidate and employee who resides in New York City.[18]

Recent regulatory actions demonstrate the importance of transparent and explainable AI solutions. In 2019, consumer complaints surfaced regarding Apple Card's creditworthiness decisions. Consumers claimed that Apple Card's creditworthiness determinations violated the federal Equal Credit Opportunity Act because women received substantially lower credit limits than men.[19] Consumers further alleged that Apple relied on algorithms and machine learning that Apple employees could not explain.[20] The New York Department of Financial Services (NYDFS) opened an investigation into Apple's underwriting data and ultimately determined that there was no evidence of discrimination.[21] As part of the investigation, Apple and its partner bank had to provide their policies related to creditworthiness determinations and underwriting data.[22] The bank provided its policies and explained its creditworthiness decisions for each consumer who complained.[23] The bank identified the factors they rely on to make creditworthiness determinations, such as credit score, indebtedness, income credit utilization, and missed payments, among other credit history elements.[24] Further, the NYDFS found that based on these factors, men and women with similar credit histories received similar credit limits.[25]

Beyond New York, California also intends to regulate AI. Earlier this month, the California Privacy Protection Agency (CPPA) released draft regulations related to Risk Assessments under the CPRA.[26] Businesses that use artificial intelligence or automated decision-making technologies will be subject to additional risk assessment requirements. These new requirements are focused on transparency and explainability of AI solutions. For example, the draft regulations provide that businesses that use automated decision-making technology will need to provide plain language explanations of how the business evaluates their use of the automated decision-making technology for validity, reliability, and fairness.[27] The draft regulations also require businesses to identify any third parties that provide software or other technological components for their automated decision-making technology.[28] Further, businesses that make their artificial intelligence or automated decision-making

Navigating the AI Frontier: Legal and Operational Insights Into Generative AI

available to other businesses must provide "facts necessary" for those recipient businesses to conduct their own risk assessments.[29] The draft regulations also provide that businesses must identify the degree and details of any human involvement in the use of their automated decision-making technology and whether the human can influence how the business uses the outputs generated by the technology.[30] Although the draft regulations are subject to change, businesses using automated decision-making technology or artificial intelligence should monitor these developments and ensure that they have policies and procedures in place to comply with these future requirements under the CPRA.

Recent class actions regarding training data further illustrate the importance of transparent and explainable AI solutions. Businesses that operate large language models (LLMs) must proceed cautiously as copyright owners are bringing copyright infringement claims related to LLM training data.[31] Because LLMs are trained on massive amounts of text from various sources, some text is likely subject to copyright protection. When faced with copyright infringement claims, a business must demonstrate that their AI solutions did not infringe on copyrighted works, such as proving that the solution was not trained on a specific work. Furthermore, because such claims are relatively new, it is unclear how arguments will shake out in court and whether a business can attempt to invoke trade secret protection related to how their LLM is trained. As such, companies need to know what data their LLMs are pulling from and have policies and procedures in place to ensure their model is not infringing on copyrighted works.

Minimizing the Risk of Potential Liability

Businesses that operate AI solutions face significant liability if their solutions are not transparent and explainable. Explainable and transparent AI solutions not only build trust and allow stakeholders to have confidence in their outputs but can help minimize the risk of potential liability if they can clearly explain how their AI solutions work, how it is trained, and why they made certain predictions or decisions. Businesses also need policies and procedures to protect privacy and security and make necessary disclosures to consumers related to how the AI solution operates and uses consumer data.

Endnotes

1 "Artificial Intelligence Risk Management Framework (AI RMF 1.0)" National Institute of Standards and Technology, 1 January 2023 at 17.
2 *Id.* at 16.
3 *Id.* at 15.
4 *Id.* at 16.
5 "What is AI Transparency?" Holistic AI, 6 February 2023.; CERRE Think Tank "Towards and EU Regulatory Framework for AI Explainability [Video]" YouTube
6 *Id.*
7 *Id.*
8 *Id.* at 15.
9 *Id.* at 16.
10 *Id.* at 5.
11 *Id.*
12 "Annex IV, Technical Documentation referred to in Article 11(1)." European Commission, 21 April 2023.
13 "Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679" at 8, 23. European Commission, 22 August 2018.
14 GDPR Article 22(2)(a)-(c); see also "Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679" at 21-23. European Commission, 22 August 2018.
15 See GDPR Articles 13-15, 34.
16 "Federal Trade Commission ("FTC") Civil Investigative Demand ("CID") Schedule FTC File No. 232-3044." United States Federal Trade Commission. Accessed 18 September 2023.
17 *Id.*
18 "Use of Automated Employment Decision Tools." New York City Department of Consumer and Worker Protection, Accessed 18 September 2023.
19 "Report on Apple Card Investigation." at 4. New York State Department of Financial Services, March 2021
20 *Id.*
21 *Id.* at 5-7.
22 *Id.*
23 *Id.* at 7.
24 *Id.*
25 *Id.*
26 See "Draft Risk Assessment Regulations for California Privacy Protection Agency September 8, 2023 Board Meeting." California Privacy Protection Agency, 8 September 2023.
27 *Id.* at 13-15.
28 *Id.* at 14.
29 *Id.* at 16.
30 *Id.* at 15.
31 Marathe, I. "In OpenAI Copyright Lawsuits, Discovery Complications Likely to Take Center Stage | National Law Journal." The National Law Journal, 21 July 2023.

FTC Issues Guidance on AI-Powered Products

Published March 2023 by Foley & Lardner LLP

AUTHORS

Louis Lehot | llehot@foley.com

Natasha Allen | nallen@foley.com

AI is booming and it's fair to say it has become the latest catchphrase with OpenAI's ChatGPT and other emerging technologies, bringing AI systems into our everyday conversations and lives. In the past few months, generative AI has dominated mainstream consciousness and this technology is moving fast.

As with most significant technological developments, regulation typically needs to catch up to advancement and adoption. So, while AI becomes more and more popular and lawmakers play catch-up on AI technology, the U.S. Chamber of Commerce recently released a report on generative artificial intelligence, calling on lawmakers to create some regulations around the ballooning technology. Meanwhile, the FTC has recently cautioned marketers not to make unsubstantiated claims about AI-powered products. The FTC has warned companies in the past with some frequency about its AI concerns, particularly issues relating to discrimination. In its most recent guidance, the FTC explicitly focuses on advertising and advising companies to be transparent about how their AI products work and what the technology can do.

According to the agency, companies relying on AI may be subject to FTC enforcement if they've exaggerated the claims of their products, overpromised or underdelivered, and not correctly accounted for reasonable risks to consumers. The agency is closely watching how developments with this technology play out and will use its enforcement authority to penalize conduct it views as unfair or deceptive.

In recent years, there has been an increase in the number of products and services that use artificial intelligence-powered technologies. These products and services can perform or enable new tasks or results. Some customers have been hesitant to trust these services with the potential consequences that could come with using these technologies.

However, much evidence suggests that these technologies are coming into play soon. Some service providers have used artificial intelligence algorithms to create products and services that are entirely customizable. This means that end-users can ensure that the product or service best serves their needs and interests.

End-users could use these technologies to personalize their experiences or to create things that are only possible because the end-user has given information to the service provider. This is important because it means the end-user is cared for and can be more hands-on with the product. This can make the end-user more likely to be satisfied with the product.

Finally, some may use AI technologies to make money. For example, an end-user might use AI to prepare product reviews — which may be genuine, fake, or a mix of the two. We can only wait to see what this trend will lead to.

The FTC's recent warning and increased FTC enforcement action surrounding AI signals a seriousness about regulating this technology. There is an increasing push for Congress to enact regulations around AI as well. Companies operating in this space must understand how their practices could lead to FTC scrutiny and what kind of regulatory hurdles could be coming.

Keys to AI IP Strategy: A Comprehensive Framework

Published September 2023 by Foley & Lardner LLP

As artificial intelligence (AI) technologies have advanced over time, strategies for patenting them have evolved in kind. Most recently, the promise of generative AI to help leverage the knowledge contained in unstructured data has captured global attention, spurring novel advances from traditional high-tech companies and startups.

This article outlines a framework for creating valuable patents for protecting AI technologies, as part of a series that covers topics including why to invest in patents for AI, how to overcome the biggest challenges in patenting AI, and business factors to consider in building your patent portfolio. The framework highlights several complementary approaches for claiming the inventions underlying any particular AI technology, and answers the following questions that characterize useful patent claims:

1. does the claim **prevent competitors** from performing AI functionality that provides a **competitive advantage** in your market

2. is the claim **resistant** to **invalidity** challenges

3. can **infringement** of the claim be **detected**

AUTHORS

Shabbi Khan | skhan@foley.com
Nikhil Pradhan | npradhan@foley.com

At a high level, most (if not all) AI technologies can be broken down into the inputs provided to an AI algorithm, such as a machine learning algorithm, how the algorithm processes the inputs (which may often be seen as a black box), and the outputs of the algorithm. For example, a system for designing new drugs could include an algorithm trained on examples of drugs that meet existing criteria, so that at runtime the algorithm can generate candidate drugs expected to meet specific criteria of interest. In another example, a customer service chatbot could fine tune a GPT-type large language model based on positive examples of effective conversations (as well as negative examples of ineffective conversations.)

As shown below, patent claims can generally be prepared to cover at least three categories of inventions for AI technologies that map to the input-processing-output breakdown.

Input	Processing	Output
Training Data Generation: Where a machine learning model is used, what steps are taken to make the training data more useful -- which results in a more useful model -- such as filtering the data, reserving data for validation, automating the generation of training data or incorporating user feedback?	**Model Training**: Where a machine learning model is used, how is the model being trained and what is the value-add generated by using a trained model as opposed to other approaches, such as rule-based or engineered approaches?	**Insights** and **Actions**: What is the outcome of using the AI solution? What patterns or new knowledge are identified? How is a physical machine, such as a medical device, industrial automation robot, or autonomous vehicle controlled using information generated by the AI solution?

The framework outlined below applies this breakdown in a manner suited for **machine learning** technologies that train models to generate useful **insights** and **actions**. It can also be readily adjusted to address AI solutions that do not necessarily use machine learning models, such as heuristics/rule based engines:

Machine Learning Patent Protection Framework

Category	Why	How	Detectability
Insights/ Actions	This is the real world result and value add	How does the action improve the underlying technology? Are you faster/more accurate/able to respond to situations before they occur?	Focus on inputs and outputs while varying the granularity of claims from the processing components up to the system in which the AI processing is implemented Ideally can treat the algorithm itself as a black box
Model Training	This is where you enable the technical improvements that make your system have a competitive advantage	Does training the model enable your system to perform functions it could not previously? Does training the model enable your system to be faster or more accurate than other ML approaches, human approaches or rules-based approaches?	Focus on input data including the source of the data and the combination of input data parameters Capture the most likely types of models that could be used using varying claim scope
Training Data Generation and Pre-Processing	Remedy for "garbage in, garbage out" problem Automation of the pre-processing to make implementing AI solutions faster/more efficient/ use fewer computational resources	Emphasize computer steps used to improve the training data Identify unique approaches for manipulating baseline data into data that may be expressed in different forms relevant to the technology area, such as filtering, changing dimensionality, automated annotation, identifying a subset of parameters that are more significant than others for the model, etc.	Focus on input data including the source of the data and the combination of input data parameters Capture the right types and scope of data, sift through to get the optimal amount of data, and properly label the data to teach your system the correct way for insights/ actions

> Patent claims can generally be prepared to cover at least three categories of inventions for AI technologies that map to the input-processing-output breakdown.

Typically, patent strategies for AI technologies focus on one or more of these aspects, depending on where the innovation lies and/or where the most value is created for end users. Companies should keep in mind that as the AI patent landscape becomes more crowded, more nuance will be necessary to properly craft applications and claims. In the case of generative AI technologies including large language models and image generation models (e.g., diffusion models), invention harvesting and application preparation efforts should focus on extracting as much information as possible regarding how the technology provides benefits that could not be achieved simply by using the baseline, publicly available models. For example, understanding how the technology is made so that it can utilize previously challenging unstructured data, how it ensures that outputs are accurate (i.e., solve the hallucination problem), and how it creates benefits that are not possible with more traditional AI will provide significant dividends for building valuable patent portfolios.

By starting from the framework shown above and incorporating innovation-specific nuances, companies can identify multiple ways for patenting your AI technology, resulting in a stronger patent portfolio that will be easier to enforce and more difficult for competitors to design around.

Intelligent Hiring: AI Recruiting Tools Get EEOC Scrutiny

Published March 2023 by Foley & Lardner LLP

The U.S. Equal Employment Opportunity Commission (EEOC) has released its [Draft Strategic Enforcement Plan for 2023-2027 (SEP)](#), outlining the agency's goals and priorities for enforcing federal employment anti-discrimination laws. One key area of focus in the plan is employers' use of artificial intelligence (AI) in hiring and recruitment.

The EEOC recognizes that the use of AI in hiring and recruitment is on the rise and is an incredible tool, aiding employers in sifting through hundreds or thousands of resumes to identify key candidates with qualities that match the skill set and experience required for a particular job. However, as we have addressed [time](#) and [again](#), the EEOC's draft SEP also highlights the potential to perpetuate and exacerbate discriminatory practices against certain groups, such as women, people of color, and individuals with disabilities. This is because AI algorithms are only as unbiased as the data that is fed into them. And, if that data reflects historical patterns of discrimination and/or inputs from users with specific biases (conscious or unconscious), then the AI system will likely perpetuate those patterns, risking disparate and unlawful impact discrimination against certain protected groups.

To address this issue, the draft SEP calls for increased scrutiny of AI tools used in hiring and recruitment, and for employers to ensure that their use of AI is non-discriminatory. This includes conducting regular audits and testing of AI tools to detect and correct any biases as well as providing transparency to job applicants about the use of AI in the hiring process.

Additionally, the EEOC is urging employers to prioritize diversity, equity, and inclusion in their hiring practices and to consider the potential impact of AI on underrepresented groups, especially those

AUTHOR

John Litchfield | jlitchfield@foley.com

> The key takeaway from the draft SEP is the EEOC's new focus on rooting out biases in recruitment technology.

with mental or physical disabilities. This may involve creating diverse hiring teams, conducting focus groups and outreach to underrepresented communities, and incorporating alternative hiring methods to supplement or moderate use of AI tools.

The key takeaway from the draft SEP is the EEOC's new focus on rooting out biases in recruitment technology. The obvious advantages of using such tools — including removing individualized decision-making about specific applicants or employees to limit liability of disparate treatment claims — is tempered by the magnified risk of disparate impact liability and the EEOC's coming scrutiny of AI recruitment practices.

Overall, the EEOC's draft SEP highlights the growing importance of addressing the potential for discriminatory outcomes in the use of AI in hiring and recruitment. By taking proactive steps to ensure that their AI tools are non-discriminatory and prioritizing diversity and inclusion, employers can help to create a more equitable and inclusive workforce.

Using AI to Help Hire?

Published June 2023 by Foley & Lardner LLP

It is expected that AI either already is or will become a routine part of recruiting and hiring, perhaps eventually even helping to manage performance. The use of AI is likely to help ease the administrative burden to screen applications and even conducting first interviews. But it is likely to also increase risk for disparate impact claims.

The EEOC has issued guidance on just this point. The new guidance relies on the 1978 Uniform Guidelines on Employee Selection Procedures, which generally require validity studies be done to ensure that characteristics being measured are reliable indicators of successful performance in the job.

As a quick refresher, disparate impact occurs when a facially neutral employment practice (such as using AI to screen applications to determine who meets minimum standards) results in disproportionately screening out people of a certain protected category (such as people forty or older) or a combination of protected categories (for example, Hispanic women).

A first question — and related takeaway — is how do you determine when an impact is disproportionate? An easy common approach is the 80% rule (also called the four-fifths rule). Let's assume your applicant group has 100 men and 100 women. AI says 50 of the men (50%) and 25 of the women (25%) are qualified. Compare the selection rate of the protected group to the unprotected group — 25/50 — and the difference is 50%. Because the ratio is under 80%, it indicates there is a statistically disparate impact.

But, this may not be enough to conclude there is actually a disparate impact. The 80% rule is not appropriate for every scenario, and other types of tests may be better suited for the circumstances. Even if the 80% rule is appropriate, the disparate impact may be

AUTHORS

Dabney Ware | dware@foley.com

legally justified if the company can prove the selection tool is job-related and consistent with business necessities. This means that, if your job action doesn't satisfy the 80% rule, you need to do more or different testing or provide more justification as to why, despite this outcome, your testing methods are necessary.

Alternatively, being in compliance with the 80% rule does not automatically mean there is no legal exposure. If the job description, or another part of the recruiting process, discourages people based on a protected characteristic, the screening tool may pass the 80% rule test, but overall still poses a legal problem or can be subject to a legal challenge.

Plus, employers cannot simply rely on the vendor's representations about testing in order to escape liability from a disparate impact challenge. The legal risk does not simply move to the vendor, and the employer can still be liable. No matter the scenario, proper use of AI in screening and hiring will involve time and money to justify its use and outcomes.

> No matter the scenario, proper use of AI in screening and hiring will involve time and money to justify its use and outcomes.

> Consider the advantages of using AI in screening and hiring and compare that advantage to the extra investment that may be involved if the AI tool is challenged.

A related problem may be records retention. A recent case alleging improper use of an AI screening tool developed by HireVue, is based on a job application from 2021. In a Washington Post article from 2019, HireVue said a standard thirty-minute assessment may provide up to 500,000 data points, although it's not clear whether this is the same type of tool at issue in the recent lawsuit. That's a lot of data to save for two years or longer until the case is concluded. All of this resource investment is in addition to making sure the substance of the AI tool is proper. Assuming all of the data about the impact on selection is retained, organizations also need to remember what the AI process looked like two years ago. In an area of such rapid change, that could be a big challenge.

In short, companies need to pay attention and keep paying attention. As the news is telling us, AI is rapidly changing. Consider the advantages of using AI in screening and hiring and compare that advantage to the extra investment that may be involved if the AI tool is challenged.

Contributors

Natasha Allen

Partner
+1 650 251 1112 | nallen@foley.com
Silicon Valley

Natasha Allen is a strategic advisor, supporting companies in all stages of growth in complex decision making across a broad range of corporate matters. She is a partner with the firm, serving as Co-Chair for Artificial Intelligence within the Innovative Technology sector, Co-Chair of the Venture Capital Committee, and is a member of the Venture Capital, M&A, and Transactions Practices. She also serves as Pro Bono chair for the firm's Silicon Valley office.

Natasha has a significant background advising startups, emerging companies, and late stage companies on debt and equity financing, venture capital financing, commercial matters, and general corporate matters including formations, corporate governance, structuring, and share issuances across a variety of industries including Artificial Intelligence, FinTech, robotics, and life sciences. Natasha's robust transactional experience includes mergers, acquisitions, and divestitures, including domestic and cross-border buy-side and sell-side transactions in Augmented/Virtual Reality, software, cybersecurity, telecommunications, and healthcare industries. She also regularly advises Canadian startups looking to establish themselves in the United States on structure, financing, and other issues as they enter this market.

Prior to joining Foley, Natasha was a founding partner at a corporate and transactional law firm, as well as managing counsel at a technology enabled law firm where she acted as external general counsel for clients. She draws on the depth of experience she has developed in a combination of law firms, business consulting firms, and startups in both Canada and the United States to help her clients navigate complicated transactions in the United States and abroad.

Eric Chow

Partner
+1 650 251 1120 | echow@foley.com
Silicon Valley

Eric Chow is a business lawyer representing public and private companies in a broad range of transactions, including mergers and acquisitions, investments, joint ventures, as well as other corporate transactions, with a focus on the Northern California market. He is a partner with the firm and a member of the Transactions Practice.

Having lived on four continents, Eric not only understands the culture and business environment in these countries but also speaks and writes Cantonese, Mandarin Chinese, and conversational Portuguese. He has a passion for cross-border work, in particular U.S./Greater China region transactions and U.S./European transactions, bridges the cultural gaps and differences in professional practices with his language skills and prior experience working in a big law firm's Hong Kong office.

Being Silicon Valley trained, Eric's experience naturally gravitates towards the innovative technology sector, where he regularly advices technology company clients on all kinds of legal matters typically encountered by Silicon Valley style venture backed companies. He also works with clients across a wide variety of other industries, including pharmaceuticals, medical devices, real estate, manufacturing, energy, and beyond.

Contributors

Carrie Hoffman

Partner
+1 202 999 4262 | choffman@foley.com
Dallas

Carrie Hoffman represents and counsels major employers nationwide in all areas of labor and employment law across a wide range of industries including retail and hospitality. She serves as chair of Foley's Labor & Employment practice, and she is highly regarded for her experience with wage and hour issues, as well as employment discrimination and retaliation claims. She regularly reviews and drafts employment agreements – such as covenants not to compete – and advises clients on a wide variety of labor and employment issues such as:

- Workplace safety
- Workplace harassment
- Union avoidance and election campaigns
- Leaves of absence
- Hiring, termination and severance policies and procedures
- Litigation avoidance

Carrie has defended clients throughout the United States in major employment litigation at both the federal and state court level including class and collective actions. She also represents employers in investigations and audits by agencies such as the Equal Employment Opportunity Commission, Texas Workforce Commission, the National Labor Relations Board and the Department of Labor.

Clients have a high regard for Carrie's knowledge of wage and hour laws as well as her practical approach to problem solving and direct manner.

Contributors

Chanley Howell

Partner
+1 904 359 8745 | chowell@foley.com
Jacksonville

Chanley T. Howell is a partner and intellectual property lawyer with Foley & Lardner LLP, where his practice focuses on a broad range of technology law matters. He is a member of the firm's Technology Transactions, Cybersecurity, and Privacy Practice and the Sports, Health Care and Automotive Industry Teams.

Chanley represents companies in a variety of technology law areas, such as:

- Data Privacy and Security Compliance – Counsel and advise clients with respect to compliance with data privacy and security laws requirements and obligations, including federal privacy laws (Gramm-Leach-Bliley, HIPAA, telemedicine / telehealth); consumer protection / marketing laws (TCPA, CAN-SPAM, COPPA); FTC privacy requirements, investigations and enforcement actions; state data privacy, security and breach notification laws; artificial intelligence / AI issues; EU Privacy Shield compliance and certification; regulatory risk assessment and compliance audits; privacy and information security policies; records retention, destruction and management policies

- Software Agreements – Software licensing, software support and maintenance; joint development, IT professional and consulting services; software license audits and disputes

- Technology Agreements – IT outsourcing; EHR (electronic health record) system agreements; ERP licensing, acquisition and implementation agreements; telemedicine/telehealth related agreements and policies; artificial intelligence / AI related agreements; patent, copyright and trademark licenses; joint development and marketing agreements; telecommunications agreements; technology acquisition and distribution agreements; IT/IP due diligence for mergers and acquisitions

- Internet – Online and electronic contracts and agreements, ESIGN and electronic signature issues, telemedicine / telehealth, ecommerce and online transactions, domain names, website terms and conditions, privacy policies, social media policies

Contributors

Dave Kantaros

Partner
+1 617 342 4068 | dkantaros@foley.com
Boston

David W. Kantaros serves as Co-Chair for Artificial Intelligence within our Innovative Technology sector and is a member of the Private Equity & Venture Capital and Transactional & Securities Practices. He represents venture capital and private equity funds as well as publicly and privately held corporations in the emerging technology and life science industries.

David's experience includes:

- Representing private equity funds, family offices, and public and private companies in mergers and acquisitions, joint-ventures, public offerings, and control investments.

- Advising venture capital firms, strategic investors, and high net worth individuals and investment groups in private placement activities, including preferred stock and secured and convertible debt financing in technology and life science companies.

- Counseling high-tech, as well as biotechnology, pharmaceutical and medical device companies as to corporate structure, technology protection and transactional matters.

- Advising investors, corporate boards and public and private companies on complex domestic and cross-border transactions including corporate migrations, re-incorporations and joint venture arrangements.

- Structuring and negotiating technology license agreements, including master license agreements, OEM agreements, Beta test agreements, shrinkwrap agreements and strategic relationships involving the sharing of technology.

Shabbi Khan

Partner
+1 617 502 3291 | skhan@foley.com
Boston

Shabbi S. Khan is a partner and intellectual property lawyer with Foley & Lardner LLP. His practice focuses on patent portfolio counseling and management, preparation and prosecution of patent applications, patent infringement analysis, patent validity analysis, and intellectual property due diligence primarily in the fields of artificial intelligence and machine learning, computer software including cloud and SaaS-based technologies, medical devices, digital therapeutics, and artificial intelligence in health care. He is a member of the Electronics Practice Group and the firm's Innovative Tech Sector.

Contributors

Louis Lehot

Partner
+1 650 251 1222 | llehot@foley.com
Silicon Valley

On a very short list of leading corporate lawyers in Silicon Valley and San Francisco, Louis Lehot has handled some of the highest profile matters in the tech, health care, and clean energy spaces. A partner in the firm's Private Equity & Venture Capital, M&A, and Transactions Practices, as well as the Technology, Health Care, Life Sciences, and Energy Industry Teams, he advises entrepreneurs and their management teams, investors, and financial advisors at all stages of growth.

Louis is skilled at taking clients from "garage to global," helping them achieve hyper-growth, go public, and successfully obtain optimal liquidity events. He guides emerging private companies as they secure venture capital financing, prepare for IPO, de-SPAC, or an M&A transaction, and navigate the exit. Clients routinely turn to Louis for his domain experience in public offerings and private placements of equity, equity-linked, and debt securities, mergers, acquisitions, dispositions, spinoffs, strategic investments, and joint ventures, as well as corporate governance and securities law compliance matters. Additionally, Louis regularly represents U.S. and non-U.S. registrants before the SEC, FINRA, NYSE, and NASDAQ.

Louis draws on more than 20 years of global experience and leverages the latest legal technology tools to drive strategies and solutions that make sense. He has been recognized by peer-reviewed industry guides and leading practitioners as one of the most innovative, creative, and versatile corporate counsel. Chambers USA noted that he has considerable experience representing both emerging companies and investors in venture financings, stating, "Louis Lehot is known for the high quality of his advice, his responsiveness and passion for his clients."

Contributors

John Litchfield

Partner
+1 312 832 4538 | jlitchfield@foley.com
Chicago

John Litchfield's primary practice includes counseling clients on complex labor and employment-related matters, including disability accommodations under Titles I and III of the Americans with Disabilities Act and the Rehabilitation Act, sexual harassment and safe workplace compliance, leave management, restrictive covenants, executive employee terminations, reductions in force, wage and hour compliance, and other state and federal employment laws.

John is a member of the firm's Labor & Employment, Business Litigation & Dispute Resolution, and Family Office Practices and the firm's ADA & FMLA Compliance Team. He is also a member of the Environmental, Social, and Corporate Governance (ESG) Team. John is a leader in the firm's Artificial Intelligence (AI) Task Force and for years has been counseling clients on AI risks and opportunities in employee recruitment, retention, and management.

John regularly represents employers in federal and state litigation matters relating disability discrimination, the Family and Medical Leave Act, and harassment and discrimination. John regularly conducts on-site trainings for clients and their employees covering workplace conduct, sexual harassment prevention, diversity and inclusion, and many other topics.

He also has experience in counseling clients on state and local campaign finance regulations, and has successfully represented clients with disputes before the City of Chicago Department of Procurement Services.

John is the chair of Foley's LGBGTQIA+ Affinity Group and serves on the firm's Recruiting Committee.

Contributors

Lyman Thai

Partner
+1 415 438 6436 | lthai@foley.com
San Francisco

Lyman Thai represents fast-growing companies in the technology sector, the entrepreneurs who build them, and the investors who back them. He is a San Francisco and Silicon Valley based partner in the firm's Transactions Practice Group and Venture & Growth Capital Team.

Lyman advises founders, management teams, and boards of directors in a wide range of transactions across all stages of growth. He advises business clients on company formation, founder and employee equity arrangements, fundraising, corporate governance, federal and state securities laws, stock and option administration, intellectual property protection, employment and compensation, commercial contracts, mergers and acquisitions, and other strategic transactions, including equity, debt, and convertible instrument financings, and investments. He also represents venture capital firms and corporate investors in their investments in high-growth companies, including seed capital, venture capital, growth capital, and strategic investments.

Lyman has represented clients across a wide range of industries, including cloud computing, cybersecurity, digital media, enterprise software, financial technology, mobile applications, on-demand services, online travel, social media, software-as-a-service (SaaS), video games, and other emerging businesses and technologies.

In addition to advising clients on general business issues, Lyman leads client teams and coordinates specialists to address licensing, patent and trademark prosecution, tax, employee benefits, employment and general commercial litigation, real estate, privacy, CFIUS/national security, HSR/antitrust, and other regulatory matters.

Prior to Foley, Lyman represented emerging companies and investors from the San Francisco and Palo Alto offices of a global technology law firm headquartered in Silicon Valley. He also previously served as the general counsel for a venture capital-backed legal technology company. In that capacity, Lyman helped the company raise US$65 million in venture capital through multiple rounds, supported the company's growth from 25 to 200 people, negotiated strategic transactions and commercial contracts, and managed business and legal risks across the organization.

Contributors

Kate Wegrzyn

Partner
+1 414 297 5778 | kwegrzyn@foley.com
Milwaukee

Kate Wegrzyn (pronounced "Way Grin") is a partner and business lawyer with Foley & Lardner LLP. Kate is chair of the firm's Supply Chain team. She is also a member of the firm's Manufacturing Sector Advisory Board and is a member of the Food & Beverage Team. Kate is also the Editor in Chief of the LexisNexis-published treatise entitled "International Agency and Distribution Agreements."

Kate's areas of focus include drafting commercial contracts and counseling businesses on supply chain matters, including force majeure issues, pricing, procurement strategies, dealer and customer arrangements, licensing issues, and logistics and transportation. Clients appreciate Kate's practical, business-focused approach to commercial counseling and contract negotiation.

Nikhil Pradhan

Senior Counsel
+1 617 226 3143 | npradhan@foley.com
Boston

Nikhil T. Pradhan is a senior counsel and intellectual property lawyer with Foley & Lardner LLP. He is a member of the Mechanical & Electromechanical Technologies Practice. Nikhil has experience in preparing and prosecuting patent applications across a broad range of technologies, including energy, automotive, aerospace, computer vision, machine learning, artificial intelligence, medical device and surgical methods, and control system technologies. Nikhil provides high-level intellectual property counseling, including competitive strategy, clearance, and invalidity/non-infringement opinions. He also provides counsel for litigation, post-grant proceedings, and transactional matters, including performing due diligence and preparing licensing and development agreements.

Nikhil was a summer associate with Foley in 2013. He also served as a summer associate with Ford Motor Company in 2012.

Prior to joining Foley, Nikhil served as an intellectual property intern with the U.S. Forest Service Patent Program. He completed a judicial externship with the Wisconsin State Court of Appeals. During law school, Nikhil worked as a graduate teaching assistant for the University of Wisconsin – Madison's Chemical Engineering Department. He holds experience as a research intern for the MIT Center for Collective Intelligence (CCI).

Contributors

Michael Ryan

Senior Counsel
+1 713 276 5175 | mryan@foley.com
Houston

Michael Ryan is a senior counsel in the Labor & Employment Practice Group in Foley's Houston and Denver offices. He counsels employers on a wide range of employment matters, including unfair competition and trade secret laws, hiring and firing decisions, EEO and FLSA matters, enforcement of personnel policies, leave issues, accommodations, drug-testing laws, safety regulations, and governmental investigations. He has significant experience representing clients in prosecuting and defending restrictive covenants, trade secret theft claims, discrimination and retaliation claims, and wage-and-hour matters.

Michael partners with clients of all sizes, from multinational corporations to local business and executives. He has particular experience counseling employers and litigating cases involving oil and gas services and production, health care, security, safety, chemical, service, and retail. He also frequently presents at seminars and conducts trainings on various employment topics, including EEO matters and internal and governmental investigations.

Prior to earning his law degree, Michael served as a judicial intern for the Hon. Hilda G. Tagle, U.S. District Court Judge for the Southern District of Texas-Brownsville Division. He also worked in-house as a compliance and employment law clerk for an oilfield services company in Houston.

Dabney Ware

Of Counsel
+1 904 359 8737 | dware@foley.com
Jacksonville

Dabney Ware is of counsel and a litigation lawyer with Foley & Lardner LLP where she focuses her practice in the area of labor and employment law. Ms. Ware has extensive experience both in employment-related litigation and in counseling clients on all aspects of compliance with the myriad of federal, state, and local laws. She is a member of the firm's Labor & Employment Practice and Food & Beverage Industry Team.

Ms. Ware does significant counseling regarding wage and hour, harassment and discrimination, employee counseling and discipline, and workforce reductions, including disparate impact analysis. She also counsels employers on affirmative action obligations under Executive Order 11246, affirmative action plans, AAP audits, and all matters involving the OFCCP.

Contributors

Avi Ginsberg

Associate
+1 617 226 3149 | aginsberg@foley.com
Boston

Avi Ginsberg is a cybersecurity and data privacy attorney with Foley & Lardner LLP based in Boston. He has substantial industry experience in cybersecurity and international business transactions. He regularly helps corporate clients respond to ransomware attacks and data breaches, evaluate privacy risk and compliance strategies for complex data sets, and develop kids' privacy approaches for online games and metaverse experiences.

Avi focuses his practice on nuanced cybersecurity and data privacy issues, including:

- advising clients on data privacy and information security programs,
- cybersecurity risk mitigation and incident response,
- artificial intelligence privacy compliance and reputational risk mitigation,
- website, web application, and mobile application privacy compliance,
- online trust and safety issues, and
- performing diligence in technology transactions.

Avi has over a decade of experience with web application design and programming, cybersecurity, databases, website exploit discovery/penetration testing, automation, bot development, software development, and reverse engineering.

Graham MacEwan

Associate
+1 617 502 3306 | gmacewan@foley.com
Boston

Graham MacEwan is an associate in the Business Law Department with Foley & Lardner LLP. Graham is based in the Boston office where he is a member of the Transactions Practice Group.

During law school, Graham was a summer associate with Foley in the Transactions Practice Group. He also served as a legal extern at Wayfair, where he reviewed company contracts, conducted compliance research, and processed intellectual property infringement matters. Graham also served as a Research Assistant to Professor Cathleen Kaveny, and as a president of the school chapter of the American Constitution Society.

ABOUT FOLEY & LARDNER LLP

Foley & Lardner LLP is a preeminent law firm that stands at the nexus of the energy, health care and life sciences, innovative technology, and manufacturing sectors. We look beyond the law to focus on the constantly evolving demands facing our clients and act as trusted business advisors to deliver creative, practical, and effective solutions. Our 1,100 lawyers across 25 offices worldwide partner on the full range of engagements from corporate counsel to IP work and litigation support, providing our clients with a one-team solution to all their needs. For nearly two centuries, Foley has maintained its commitment to the highest level of innovative legal services and to the stewardship of our people, firm, clients, and the communities we serve.

⊞FOLEY

FOLEY & LARDNER LLP

FOLEY.COM

AUSTIN | BOSTON | BRUSSELS | CHICAGO | DALLAS | DENVER | DETROIT | HOUSTON | JACKSONVILLE | LOS ANGELES | MADISON | MEXICO CITY | MIAMI | MILWAUKEE
NEW YORK | ORLANDO | SACRAMENTO | SALT LAKE CITY | SAN DIEGO | SAN FRANCISCO | SILICON VALLEY | TALLAHASSEE | TAMPA | TOKYO | WASHINGTON, D.C.