Send a message   ➤

Bloomberg Law

The Power of the Prompt

A guide to artificial intelligence and the profession of law, from an introduction to the technology to practical guides for in-house counsel and law firms.

Note from the Editor

What is the power of the AI prompt? Is the legal industry ready for AI? The answers depend on whom you ask.

Some in the legal industry have been deploying artificial intelligence for the better part of a decade to parse data and query documents.

The explosion of interest in powerful tools like ChatGPT has changed the conversation entirely, leading some to ban use of the latest technology for fear of the potential risks.

And somewhere in the vast middle between the longtime users and the skeptics is everyone else, expressing varying degrees of enthusiasm and caution for the amazing impression that AI technology has made.

That spectrum of embrace for AI is playing out across industries all over the globe. But in law, the pain is particularly acute.

Attorneys at law firms are already on the front lines of advising clients in those industries how to handle matters in a range of practice areas as governments move to regulate at different speeds.

General counsel are collaborating across their companies with business partners, crafting use guidelines and advising on just where the risk may lie.

Even as they do so, both in–house counsel and law firm attorneys must grapple with how potent AI tools will affect their own jobs.

Once ChatGPT passed the bar, the world began asking if "the prompt" and what the model spits out would have the power to replace lawyers. The answer is a qualified "no." But just how much of a hybrid profession—one part AI, one part human—law will one day be is a more interesting question.

In this special report, we explore the conversation around AI in the legal industry as a whole—from law firms to in–house counsel to legal operations. If your AI knowledge is base level, start with our glossary of terms; if you're looking for more practical content, we have a list of checklists specific to practice areas.

For industry and legal news on AI, sign up for a daily newsletter from Bloomberg Law. For practical guidance, including deeper analysis and tools, visit Bloomberg Law's In Focus: AI page.

Thanks for reading, and as always, we welcome your feedback.

Rachael Daigle
Editor, Bloomberg Law
rdaigle@bloombergindustry.com

Table of Contents

Glossary of Terms

Algorithm

A step–by–step process used to solve a problem.

Machine Learning

The process of feeding data into computer algorithms so they get more refined and sophisticated over time.

Natural Language Processing

The branch of AI that helps computers to understand, process, and generate speech and text the way a human would.

Chatbots

These are products that can hold advanced, human–like conversations with people about anything from historical trivia to lists of creative recipes using a watermelon.

Deep Learning

The most common form of AI, in which software is taught to classify something such as a video or a loan application from a very large set of labeled data.

Generative AI

This refers to the production of entirely new creative works—pictures, music, text, poetry—from simple prompts after AI is trained on vast quantities of pre–existing material.

Large Language Models

The backbone of natural language processing that can summarize and generate text using information from all over the internet. Perhaps the most well–known is OpenAI's GPT–4.

Hallucinations

The phenomenon by which AI chatbots may confidently provide false information (sometimes ludicrously so) in response to a prompt.

Adapted from "A Cheat Sheet to AI Buzzwords and Their Meanings: QuickTake," Bloomberg News.

Credit: SOPA Images

Corporate AI Legal Policies Race to Keep Up With Technology

By Isabel Gottlieb

When ChatGPT burst onto the scene last year, in–house lawyers had to scramble to figure out how to govern the use of new generative AI tools, and decide who would take charge of those decisions.

Topping their concerns: protecting confidential business and customer data, and establishing human backstops to safeguard against the technology's propensity to "hallucinate," or spit out wrong information.

Artificial intelligence isn't new. But generative AI—tools trained on oceans of content to produce original text—created ripples of panic among legal departments when ChatGPT debuted, because its full legal implications were both far–reaching and not entirely clear. And with public–facing platforms, the tool is easily accessible to employees.

"Generative AI is the first thing that can violate all our policies at once," said Dan Felz, a partner at Alston & Bird in Atlanta, referring to a company's perspective.

AI Oversight

As the technology evolves and the legal implications multiply—and with regulation on the horizon in multiple jurisdictions—companies should have a person or team dedicated to AI governance and compliance, said Amber Ezell, policy counsel at the Future of Privacy Forum. The group this summer published a checklist to help companies write their own generative AI policies.

That role often falls to the chief privacy officer, Ezell said. But while AI is privacy adjacent, it also encompasses other issues.

Bloomberg Law

Toyota Motor North America has established an AI oversight group that includes experts in IP, data privacy, cybersecurity, research and development, and more to evaluate internal requests to use generative AI on a case–by–case basis, said Gunnar Heinisch, managing counsel.

The team is "continually trying to evaluate what the risks look like versus what the benefits are for our business" as new issues and use cases arise, Heinisch said.

"Meanwhile, in the background, we're trying to establish what our principles and framework look like—so, dealing with the ad hoc questions and then trying to establish what that framework looks like, with a long–term regulatory picture in mind," he added.

Salesforce, the San Francisco–based enterprise software giant, has been using AI for years, said Paula Goldman, chief ethical and humane use officer at the company. While that meant addressing ethical concerns from the start, she noted, generative AI has raised new questions.

The company recently released a new AI acceptable use policy, Goldman said.

"We know that this is very early days in generative AI, that it's advancing very quickly, and that things will change," she said. "We may need to adapt our approach, but we'd rather put a stake in the ground and help our customers understand what we think is the answer to some of these very complicated questions right now."
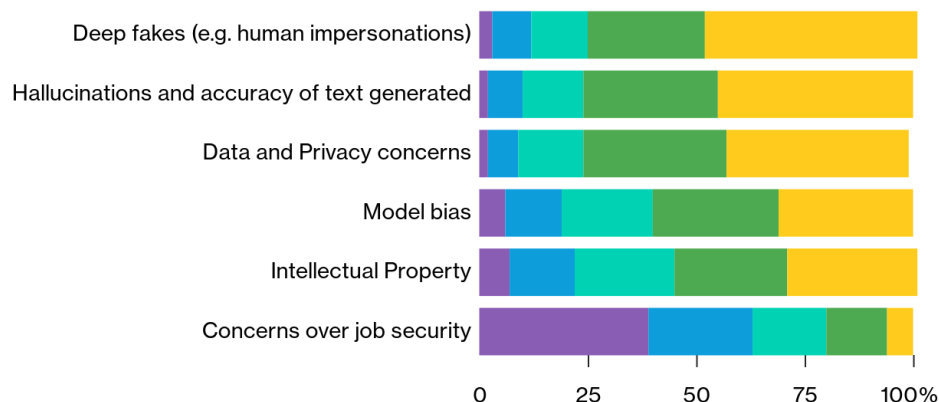
> *" Generative AI is the first thing that can violate all our policies at once."*
>
> –-Dan Felz, Partner at Alston & Bird in Atlanta

Concerns Around Generative AI are High for Lawyers
'Please select your level of concern with the following potential issues...'

Legend: Not concerned at all | Slightly concerned | Somewhat concerned | Moderately concerned | Extremely concerned

- Deep fakes (e.g. human impersonations)
- Hallucinations and accuracy of text generated
- Data and Privacy concerns
- Model bias
- Intellectual Property
- Concerns over job security

(x-axis: 0, 25, 50, 75, 100%)

Source: Bloomberg Law's Second State of Practice Survey 2023, conducted from June 29, 2023 to July 21, 2023. Respondents were asked to select one level of concern for each potential issue regarding Gen AI. Percentages rounded to nearest whole number.

Bloomberg Law

The conversation about responsible use of the technology will continue as laws evolve, she added.

Creating Policies

The first appearance of ChatGPT was, "All hands on deck! Fire! We need to put some policy in place immediately," said Katelyn Canning, head of legal at Ocrolus, a fintech startup with AI products.

In a perfect world, Canning said, she would have stopped internal use of the technology while figuring out its implications and writing a policy.

"It's such a great tool that you have to balance between the reality of, people are going to use this, so it's better to get some guidelines out on paper," she said, "just so nothing absolutely crazy happens."

Some companies banned internal use of the technology. In February, a group of investment banks prohibited employee use of ChatGPT.

Others have no policies in place at all yet—but that's a dwindling group, Ezell said.

Many others allow their employees to use generative AI, she said, but they establish safeguards—like tracking its use and requiring approval.

"I think the reason why companies initially didn't have generative AI policies wasn't because they were complacent or because they didn't necessarily want to do anything about it," Ezell said. "I think that it came up so fast that companies have been trying to play catch-up."

According to a McKinsey Global Institute survey, among respondents who said their organizations have adopted AI, only 21% said the organizations had policies governing employee use of generative AI. The survey data was collected in April and included respondents across regions, industries, and company sizes, McKinsey said.

For companies creating new policies from scratch, or updating their policies as the technology evolves, generative AI raises a host of potential legal pitfalls, including security, data privacy, employment, and copyright law concerns.

As companies wait for targeted AI regulation that's under discussion in the EU, Canada, and other jurisdictions, they're looking to the questions regulators are asking, said Caitlin Fennessy, vice president and chief knowledge officer at the International Association of Privacy Professionals. Those questions are "serving as the rubric for organizations crafting AI governance policies," she added.

"At this stage, organizations are leveraging a combination of frameworks and existing rulebooks for privacy and anti-discrimination laws to craft AI governance programs," Fennessy said.

" We're trying to establish what our principles and framework look like—so, dealing with the ad hoc questions and then trying to establish what that framework looks like, with a long-term regulatory picture in mind."

--Gunnar Heinisch, Managing Counsel at Toyota Motor North America

What's a 'Hard No?'

At the top of most corporate counsels' concerns about the technology is a security or data privacy breach.

Bloomberg Law

If an employee puts sensitive information—such as customer data or confidential business information—into a generative AI platform that isn't secure, the platform could offer up the information somewhere else. It could also be incorporated into the training data the platform operator uses to hone its model—the information that "teaches" the model—which could effectively make it public.

But as companies seek to "fine-tune" AI models—train them with firm- and industry-specific data to obtain maximum utility—the thorny question of how to safeguard secrets will remain at the forefront.

Inaccuracy is also a major concern. Generative AI models have a propensity to hallucinate, or produce incorrect answers.

Companies must be careful to not allow unfettered, un-reviewed use, without checks and balances, said Kyle Fath, a partner at Squire Patton Boggs in Los Angeles, who focuses on data privacy and IP.

A "hard no" would be using generative AI without internal governance or safeguards in place, he said, because humans need to check that the information is factually accurate and not biased, and doesn't infringe on copyrights.

Risks and Guardrails

Using generative AI for HR functions—like sorting job applications or measuring performance—risks violating existing civil rights law, the US Equal Employment Opportunity Commission has warned.

The AI model could discriminate against candidates or employees based on race or sex, if it's been trained on data that is itself biased.

Recent guidance from the EEOC is consistent with what employment lawyers had been advising their clients, said David Schwartz, global head of the labor and employment law group at Skadden Arps in New York. Some jurisdictions have already enacted their own AI employment laws—such as New York City's new requirement that employers subject AI hiring tools to an independent audit checking for bias.

There's also already regulatory attention on privacy issues in the US and EU, Fath said.

Employee use of generative AI also puts companies at risk of intellectual property law violations. Models that pull data from third-party sources to train their algorithms have already sparked lawsuits against AI providers by celebrities and authors.

"It's probably not outside of the realm of possibility that those suits could start to trickle down to users of those tools," beyond just targeting the platforms, Fath said.

Companies are looking closely at whether their current privacy and terms of use policies allow them to touch customer or client data with generative AI, he added.

State AI Laws & Regulations

Click to learn more.

Keep track of state laws and regulations related to artificial intelligence with this state-by-state look at codified statutes, legislative proposals, codified regulations, and rulemakings.

Bloomberg Law subscription required.

PRACTITIONER INSIGHT

How Employers and In-House Counsel Can Adapt to AI Workplace Use

By Alexandra Stathopoulos   Orrick Herrington & Sutcliffe

Business leaders and their legal departments are navigating a new workplace reconfigured by artificial intelligence. It's impacting hiring, employer decision–making, and how employees do their jobs. In response, state and local governments are enacting a patchwork of AI–related laws, producing a dizzying landscape of regulation in the absence of a uniform federal approach.

These state and local laws have far–reaching consequences, impacting transparency and accountability of AI systems, potential bias and discrimination in AI algorithms, data privacy and protection, workforce displacement, compliance and reporting obligations, and the ethical use of AI. And, existing legal frameworks related to intellectual property may also have profound implications for employee use of generative AI.

There is an urgency for employers and in–house counsel to prepare for these fast–changing legal requirements to ensure they can quickly adapt as new laws are enacted. Here is a roadmap to prepare.

Form an AI Working Group

To set the groundwork for a robust AI compliance program that addresses different laws and regulations across multiple locations, employers should establish a cross–functional AI working group to identify and track all the ways the company uses AI.

Law Firms Wooing AI Experts

As clients demand more for less, law firms are hiring growing numbers of staff who've studied technology—not tort law—to try and stand out from their rivals.

Many firms are now putting together teams to figure out how to use generative AI across all their practice areas, and they're seeking out data scientists and software engineers.

■ Law Firms Wooing AI Experts as Clients Demand 'More for Less'

With the involvement of in-house counsel, this group should gather information about existing AI systems or ones the company plans to deploy, assess their purpose and potential impact on employees, customers, and other stakeholders, consider whether there are any AI tools the company should license, and whether to restrict or limit employees' use of other AI tools.

The working group should also assess the degree that employees are independently deciding to use AI tools in the course and scope of their employment, and how they are using them. This initial step will help employers understand the specific functionalities and capabilities of each AI system, including any data inputs, algorithms, and decision-making processes, and will create quicker responses to new AI laws and regulations.

Review, Update Policies and Practices

Conduct periodic reviews of AI-related policies and practices to identify any gaps or inconsistencies with applicable laws and regulations, with an eye towards changes on the horizon. This review may include:

Policies and practices directly related to AI procurement, development, deployment, and monitoring processes. Contracts with third-party AI vendors, paying particular attention to data privacy and security provisions, and any indemnification clauses. Policies and practices impacted by use of AI, such as confidentiality and IP assignment agreements, recruitment and hiring, job descriptions, codes of conduct, and data collection, storage, and retention.

Assessment of whether it is appropriate for employees to use generative AI tools in light of the need for freedom to exploit commercially and ownership of any company assets created with the assistance of such tools. Insurance policies, to assess whether there is coverage for potential AI-related liabilities. Employers and in-house counsel may need to create bespoke AI policies or addenda to comply with the unique laws and regulations in each jurisdiction.

☑ **Form an AI Working Group**

☑ **Review, Update Policies and Practices**

☑ **Assess Data Collection, Storage, Usage**

☑ **Provide Employee Training and Awareness**

☑ **Stay Informed on Regulations**

Assess Data Collection, Storage, Usage

Review data collection, storage, and usage practices in relation to AI systems. Applicable laws and regulations may impose restrictions on the collection, use and storage of certain types of data, such as biometric data—e.g., data gleaned from "face recognition" technology), sensitive personal information, or information relating to protected class characteristics.

Employers and their third-party vendors should obtain appropriate consent from applicants and employees for the collection and use of their data in relation to AI systems where required by applicable law (and there are many jurisdiction-specific nuances). In addition, review employee notices or disclosures to ensure they are clear and provide sufficient information about how the AI systems will be used, describe the data collected, and how any automated decisions are made and reviewed.

Employers should also consider carrying out a privacy impact assessment to help determine any risks associated with the use of AI systems and identify ways to mitigate any identified risk.

Supervised by counsel, employers may also want to conduct periodic privileged reviews of AI algorithms to test for potential bias, both in terms of input data and output results. Taking proactive steps to ensure that the algorithms avoid discrimination or bias against protected groups is likely to put an employer in a better position to comply with the forthcoming wave of legislation focusing on this issue.

Provide Employee Training and Awareness

Employees play a critical role in the use of AI in the workplace, and training and awareness programs can help an employer stay on top of changing laws and regulations. Employers should provide comprehensive training and awareness programs to educate employees on their responsibilities and obligations related to the use of AI systems.

This may include training on when it is appropriate to use generative AI tools in the performance of duties, ethical use of AI, understanding the limitations and potential biases of AI algorithms, being aware of "hallucinations" (false AI outputs), steps to take to preserve intellectual property rights in company assets created using AI tools, promoting transparency and accountability in using AI for decision–making, and educating employees on the nuances of applicable laws and regulations that apply to the use of AI.

To reinforce efforts to establish transparency and accountability, employers should consider developing an incident response plan and training managers and HR on how to address any issues or complaints related to the use of AI, including procedures for escalating potential issues to in–house legal counsel, and investigating and addressing any allegations of discrimination, bias, or other potential legal or policy violations.

Stay Informed on Regulations

Make sure to follow the evolving legal landscape. This includes monitoring proposed and enacted laws at the federal, state, and local levels that specifically address use of AI in the workplace. Employers should keep a close eye on legislative developments, including proposed bills, amendments, and regulatory guidance, and seek legal counsel to interpret and analyze the implications for their specific industry and operations in each jurisdiction.

With over 20 states enacting or developing new laws that will shape AI workplace regulation, more changes are inevitable. But these basic approaches ensure your team remains ahead of the curve as AI becomes further ingrained in the workplace.

AI Moves to the Boardroom

The newest member of the corporate boardroom may add just the impersonal touch companies are looking for.

Some companies are using generative artificial intelligence in corporate governance decisions–from sales and marketing issues to finance and legal services. Others are treading slowly due to privacy and security concerns and because, according to recent surveys, most directors aren't knowledgeable about using AI tools. Read more:

- Companies Grapple With Limits in Bringing AI into the Boardroom

In-House AI Use, Planning Quickly Outpacing Law Firms

By Stephanie Pacheco

The legal industry has long been derided as averse to technological change, but the advance of generative artificial intelligence appears to be causing a changing of the tide, based on recent survey data from Bloomberg Law.

Bloomberg Law's most recent State of Practice Survey asked over 450 law firm and in–house practicing attorneys about their familiarity with and views on generative AI and how their organizations have responded to recent developments in the technology.

When attorneys were asked about generative AI in the spring survey, they exhibited an overarching lack of familiarity with the rapidly emerging technology and a high level of uncertainty about the legal industry's practical uses for it. And although uncertainty still lingers, the new data from the summer version of the survey suggest that the profession is quickly moving toward a broader understanding of the technology.

Attorneys Dive Right In

Attorneys' experience with generative AI has passed a tipping point since the spring, according to the data.

Sixty percent of attorneys reported this summer that they've used generative AI. That's an almost complete flip from our spring survey, in which 63% of respondents said they had no experience using the technology. The ubiquity of the topic has likely played a role in the rise of attorneys reporting that they have at least tested the capabilities of generative AI, but the increase in those using it for work is rather interesting.

The percentage of attorneys who reported using generative AI for professional

<div style="background:black;color:white;">

Dive Deeper

Read more from Bloomberg Law analysts:

- Three Considerations for Attorneys Using Generative AI

- Are AI Chatbots a Product or a Service?

- Seeing Is Not Believing– Authenticating 'Deepfakes'

- What Lenders Should Know About AI and Algorithmic Bias

- New Threats, Same Rules for Finance Generative AI

- Proposed SEC AI Rules Could Deter Finance From Tech

Bloomberg Law subscription required.

</div>

purposes has more than doubled in just three months—an impressive adoption rate for a profession that is known to be change–averse. And even though fewer than one–sixth of attorneys (15%) have now used generative AI for work, how these attorneys are implementing this technology into their workflows suggests that professional use is likely to keep growing.

Attorneys Becoming More Familiar With Generative AI
'Which of the following best describes your use of generative AI?'

- Never heard of this
- Heard of this, but have no experience using
- Have used generative AI, but only for personal use or to simply test the capabilites
- Have used this in a professional setting or to perform a work task

Source: Bloomberg Law's 2023 State of Practice Surveys, conducted from March 21, 2023 to April 14, 2023 (Spring 2023), and June 29, 2023 to July 21, 2023 (Summer 2023). Respondents were asked to select one of the answers above. Percentages rounded to nearest whole number.

Bloomberg Law

For example, drafting communication—such as emails or letters to opposing counsel—and drafting legal documents were the two most common use cases reported by the lawyers that have already started implementing generative AI–powered tools into their workflows (66% and 53%, respectively). Both tasks often require a lot of time from attorneys, which means many may be enticed by generative AI's ability to reduce that initial workload and create time for more complex tasks—assuming their organization allows the use.

In-House Legal Departments Swim Ahead

The surveys also asked attorneys to identify how their organization has responded to developments in generative AI, and the results show a dramatic spring–to–summer increase in activity, revealing two key insights: (1) the industry at large is getting better at how organizations communicate with their lawyers about the technology, and (2) in–house attorneys are seeing more drastic AI–related changes in their organizations than law firm attorneys are.

From the first survey to the second, the number of attorneys who reported they are "not sure" of what their organization has done in response to developments in generative AI dropped by more than half. A drop this notable indicates that
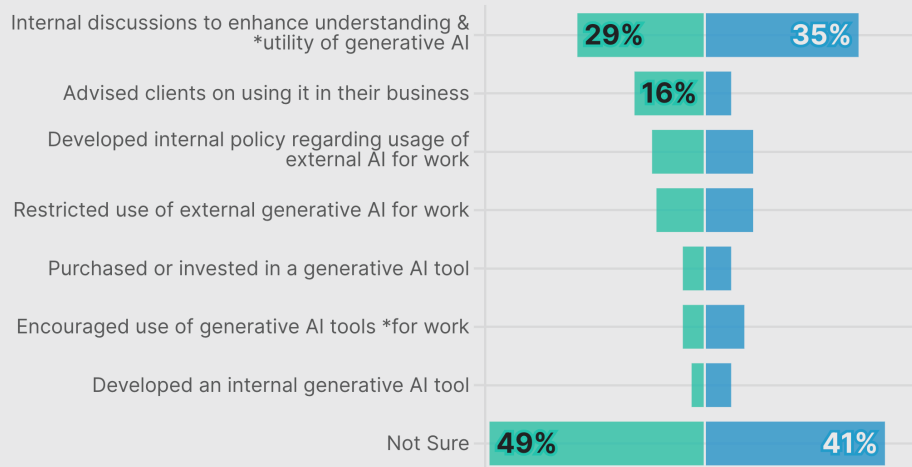
INTERACTIVE: Click to learn more.

Legal Industry Ramps Up Response to Generative AI
'Which of the following has your organization done in response to generative AI?'

Spring 2023 | Summer 2023

◼ Law Firm ◼ In-House

Internal discussions to enhance understanding & *utility of generative AI	29%	35%
Advised clients on using it in their business	16%	
Developed internal policy regarding usage of external AI for work		
Restricted use of external generative AI for work		
Purchased or invested in a generative AI tool		
Encouraged use of generative AI tools *for work		
Developed an internal generative AI tool		
Not Sure	49%	41%

Source: Bloomberg Law's 2023 State of Practice Surveys, conducted from March 21, 2023 to April 14, 2023 (Spring 2023), and June 29, 2023 to July 21, 2023 (Summer 2023). • Respondents were asked to select all that apply. If respondents selected "Not Sure" they were unable to select additional answers. The response "Other" was surveyed but not included in this graphic.
Note: Clicking legend icons will change chart view. "*" indicates slight wording change from State of Practice 1 to 2. "+" indicates new option in State of Practice 2.

Bloomberg Law

organizations have enhanced communication with their attorneys greatly in the last three months.

But the changes from spring to summer have not been equal when it comes to the responses from law firm vs. in-house attorneys.

In the spring, law firms and legal departments were on a similar playing field in terms of their responses to generative AI. Roughly the same percentages of law firm and in-house attorneys reported that their organizations had developed certain policies or restricted the use of generative AI. The largest gaps were in the categories of client advising, which appeared to be more popular in law firms (16% of attorneys) than legal departments (6%), and in internal discussions about AI, which were more common within legal departments (35%, compared to 29%).

Those internal discussions could have contributed to the rapid developments that in-house lawyers saw in this summer survey.

Both groups of attorneys reported much greater activity on the AI front in the summer, but the changes among in-house legal departments outpaced the changes in law firms in every response option across the board. For example, the development of internal policies on generative AI rose from 11% to 50% in corporate legal

departments, compared to an increase from 12% to 36% in law firms. And while advising clients on AI use was cited by 31% of attorneys in both groups in the summer, that 31% represented a five-fold increase for in-house counsel, but less than double for law firm lawyers.

Another look at the "not sure" totals reveals the speed of change—and increase in lawyer awareness—among in-house legal departments. In the spring survey, "not sure" was the most-selected option by in-house attorneys but in the summer, it was the least-selected. By contrast, "not sure" responses only dropped to fifth on the list of options for law firm respondents, leaving it as one of the top-five selected responses in the summer survey.

Such dramatic shifts could be a sign that corporate legal departments are just more innovative than law firms, at least in regard to generative AI. Or, alternatively, the difference in clientele necessitates a more rapid understanding of technological developments for in-house legal departments who are likely advising their corporations on the potential of widespread company adoption and utilization.

But although in-house attorneys are reporting more rapid changes, the survey data clearly show that both groups are moving in the right direction with generative AI.

Related video

Click on the video to watch.

Corporate Legal Departments Start Embracing AI, Cautiously

By Isabel Gottlieb

Generative artificial intelligence is on the verge of transforming how corporate legal departments deal with routine functions—even as it creates new obstacles that must be navigated.

By harnessing AI to do the grunt work on tasks like reviewing simple contracts, in–house counsel can reduce the time and money needed to perform those tasks to a fraction of what would be needed by humans.

But generative AI is new, and the industry is navigating critical questions around the security and quality of data.

"We are at a pivotal moment" in the evolution of legal operations technology, said Mary O'Carroll, the former president of the Corporate Legal Operations Consortium and former head of legal operations at Google. O'Carroll is now the chief community officer at Ironclad, a contract management technology company that's building generative AI into its products.

A recent increased focus on legal operations technology has been driving investment, she added, "and AI is just this rocket propelling all of us."

'Pilot Mode'

Many in–house departments are already using automation tools to streamline tasks like contract management.

For example, companies are automating "self-serve" non-disclosure agreements to easily swap in language reflecting different state laws, said Brian McGovern, executive director of corporate legal and claims at Mitratech, an enterprise software company that sells legal workflow and automation products.

These tools are useful for "high volume, low complexity tasks," he said. "NDAs a couple of years ago were taking a month to negotiate, now they're getting done in five minutes."

Generative AI moves beyond those automation tools because it isn't just operating according to specific instructions from humans. It can interpret language and produce original text–such as red-lining or suggesting edits based on its interpretation of other contracts a company has negotiated.

Using AI, legal departments can carry out certain work as much as 20 times faster, O'Carroll said.

Currently, faced with a crowded market of AI and automated legal technology vendors, companies are still weighing and testing their options, said Sean Monahan, a senior director at Harbor, a legal technology and operations services company.

"The vast, vast majority of legal departments are in pilot mode or evaluation mode, much more than application mode," he said.

Legal technology companies, including Ironclad, said many of their customers are already using AI tools.

Early adopters are likely companies that are already well advanced in adopting technology into their legal operations, O'Carroll said.

> ## QuickTake on AI Regulation
>
> As Congress debates whether and how to impose binding regulation on AI, President Joe Biden has called on leading AI companies to meet voluntary transparency and security standards as a first step. For a look at what Biden has done and where things stand in Congress, as well as what regulation currently exists and what's happening at a state level, read this explainer:
>
> ■ Regulate AI: Here's What That Might Mean in the US

Security and Accuracy

But adoption of the technology brings its own set of concerns.

Legal departments worry that sensitive company or customer data could be compromised by feeding it into a third-party platform, or that they could run afoul of data privacy regulations in the US or abroad. There's also a risk that proprietary information could become part of the platform's training model, if the service is built directly on an open platform and the information isn't walled off.

Companies could also run into intellectual property law issues. When information is fed into a large language model, Monahan said, the question is: "Who owns it?"

Because generative AI can "hallucinate"–turn out incorrect answers–humans need to check the work. That could mean hiring dedicated staff, Monahan said.

Training Data

The AI can only "learn" from the data it can access—but privacy and security constraints limit the amount of real-world data, such as actual contracts, the tool can be fed.

Finding good training data "is a major problem" for training the models, said Jason Boehmig, CEO and co-founder of Ironclad.

When generative AI first burst onto the scene, Monahan said, the industry was saying, "If we can point this at our stuff, it's going to be awesome."

But, he added, "What are you pointing? What stuff? Are you sure it's yours?" and, "Are you sure it's not your client's?"

There's publicly accessible data—contracts in SEC filings, for example—but "that would be like training a car dataset on cars that wound up in the impound lot," Boehmig said. "It's highly not representative."

Bad training data can also create bias, a potentially expensive problem.

Once there's bias in the model, it's difficult to un-corrupt it without millions of dollars of training, said Jerry Ting, founder and CEO of Evisort, a contract management technology company that focuses on AI.

Companies can train models themselves, but that can be costly, Monahan said.

"Fine-tuning" the data going into the model, and its responses, is a challenge, he added.

What's Next?

Currently, sentiment toward the new technology runs "a wide gamut" across the legal ops industry, and even within companies, Ting said.

He said he expects to see most companies within a year make an initial investment in automated negotiation and automating the manual tracking of when contracts expire. In five years, Ting said, "the legal industry will be completely changed. But I think it's going to take some time."

The speed at which generative AI is evolving is one reason to be hesitant for now, said Ed Sohn, head of solutions at Factor, a legal services company.

"We believe it has insane potential, really a massive opportunity," Sohn said. But, he added, "we think it's really important—because of the speed at which it's developing—that we're cautious to not over-build around the present state of the art because what's really awesome in June is outmoded in July."

Given the pace of change, "obsolescence" of the new technologies is a major concern for legal departments considering investing in them, Monahan said.

" The legal industry will be completely changed. But I think it's going to take some time. "

--Jerry Ting, Founder/
CEO at Evisort

For now, the industry is grappling with questions kicked up by the new technology.

"This came on us really fast," McGovern said. "We haven't had time to figure out all the things that are going to go wrong."

As adoption moves forward across the industry, "the laggard won't be the technology," he said. It'll be the law, the precedent, and attitudes toward it, he said.

Like any new, shiny object, AI is the center of attention now, said Brittany Leonard, general counsel at Civix, a public-sector technology and services provider. Once the hype dies down, it's likely to become comfortable and useful for in-house counsel, she said.

Meanwhile, the technology is also evolving rapidly.

"Stuff I thought might be achievable in my career got achieved last month," Boehmig said.

Bloomberg Law competes in this market and sells AI-based tools that provide contract solutions.

Related video

Click on the video to watch.

The Artificial Intelligence Dilemma: Can Law...

Watch later   Share

\\AI generated

\\AI generated

\\AI generated

Bloomberg Law News

Watch on ▶ YouTube

Big Law's AI Challenge Drives New Conversation About Training

By Sam Skolnik

Michael Best & Friedrich blocks its lawyers and staff from using ChatGPT on the job.

The law firm made the decision to ban internal use of open–source generative AI tools after consulting some of its clients, according to Sarah Alt, Michael Best's chief process and AI officer. The move also followed feedback from the firm's insurers, she said.

A key question the firm needed to answer, Alt said, is could Michael Best use the power of such publicly accessible programs while making sure it's "a safer place" for staff and clients? The answer was "no," she said, especially given that many client companies are taking a cautious mindset about ChatGPT because of data concerns.

Other major firms are taking different a different tack, allowing limited use of the technology on the job. Some are opting to build or buy their own AI tools.

What most agree on is that lawyers and staff need to be trained on the technology. Fast.

"Developing that technical competence has become increasingly important in the legal profession," said Jeffrey Chivers, who teaches an AI course at Yale Law School. Law firms are "error–intolerant" businesses, which means attorneys need to understand the technology that they—and their clients—are using.

Firms are trying to quickly devise comprehensive internal training regimens, even as the technology advances. They want to avoid pitfalls that range from from exposing sensitive client data to inaccurate information, like the AI–generated court brief littered with wholly fabricated court decisions that recently landed two New York lawyers in hot water.

Some firms are offering in-person seminars taught by outside specialists and video-learning presentations from law firm training platforms.

Orrick Herrington & Sutcliffe has partnered with tech training provider AltaClaro to develop a "prompt engineering" curriculum being rolled out across the firm, said Daryl Shetterly, managing director of Orrick Analytics. Prompt engineering involves structuring and refining the queries that produce AI-generated content.

The training includes instruction on how generative AI programs work, Shetterly said in an email. It also "highlights the importance of data privacy, confidentiality, checking accuracy, and other issues folks are likely to encounter."

As firms wrestle with the new tech, several are in the process of developing policies that aim to lessen the risks generative AI tools pose.

Early this year, for example, BakerHostetler issued a firmwide directive advising its staff not to use large language models combined with any client data, said Katherine Lowry, the firm's chief information officer. Data ingested by the public ChatGPT and similar tools isn't kept private and could be breached, according to the firm.

Firm leaders are "deep in discussions" over a broader policy regarding chatbots, Lowry said.

"Firms are looking for candidates comfortable and knowledgeable in this space."

--Stacy Leeds,
ASU Law Dean

Law Schools

US law schools are also working to meet student demand for up-to-date courses about the fast-evolving tech.

"Students are hungry for it, and professors are aware of it," said Dyane O'Leary, chair of the Association of American Law Schools' section on technology, law, and legal education.

O'Leary will be teaching a new winter intercession course—Generative AI for Lawyers—at Suffolk University Law School. Elective classes are one thing, but O'Leary says there are many students they won't reach.

The "huge trick," according to O'Leary: "Is this going to be integrated into the existing law school curriculum?"

Arizona State University's law school gained attention in July when it announced that it will allow prospective students to use chatbots to help prepare their applications.

"Firms are looking for candidates comfortable and knowledgeable in this space," ASU Law dean Stacy Leeds said in an interview.

Rob Chesnut

GOOD
COUNSEL

Bloomberg Law

Rob Chesnut is a former general counsel and Justice Department prosecutor. He writes Good Counsel, a regular column for Bloomberg Law on in-house, corporate, and legal ethics issues.

COMMENTARY

Is AI Friend or Foe of General Counsel? It Just Might Be Both.

By Rob Chesnut

Like many of you, I've been wrestling with the question of what in-house counsel should be aware of regarding AI. So I decided to go to the expert on the subject: ChatGPT.

Predictably, Chat GPT created a good list of nine key issues for me to consider. There was only one problem with the list: It omitted three obvious and critical issues.

To be fair, the ChatGPT list covers a lot of ground, and it's a breezy, high-level starting point for lawyers thinking through the laundry list of issues for the first time. But it's hardly genius.

Take the first point. "Stay informed about the evolving legal and regulatory landscape surrounding AI. Familiarize yourself with laws and regulations related to data protection, privacy, intellectual property, consumer protection and discrimination."

Hey Renee, you've got this one. Get back to me next week with answers. Make sure you cover all the global implications, too.

Or take its advice on data protection: "Ensure that data collection, storage, sharing, and usage comply with relevant laws and regulations."

That simple statement is a full-time job for a whole team of lawyers and engineers in many global companies. And "consider implementing appropriate data protection policies, data retention policies, and data breach response plans."

Consider it? If you don't already have these things in place, you're already way behind.

We asked ChatGPT: "What should in-house counsel be aware of regarding AI?" This is a condensed list of the nine points it offered as an answer.

1. Stay informed about the evolving legal and regulatory landscape surrounding AI.

2. Ensure that data collection, storage, sharing, and usage comply with relevant laws and regulations. Consider implementing appropriate data protection policies, data retention practices, and data breach response plans.

3. Develop strategies for patent filing, copyright protection, and IP licensing agreements.

4. Identify and mitigate biases in AI algorithms and data sources.

5. Assess potential product liability and safety risks.

6. Review and negotiate contracts with AI vendors, suppliers, and partners.

7. Consider the potential impact on employees, consumers, and society as a whole.

8. Establish mechanisms to monitor and ensure ongoing compliance with legal and regulatory requirements related to AI.

9. Be prepared to handle litigation or alternative dispute resolution processes.

IP considerations, bias and discrimination, product liability and safety … all good points. The list goes on.

But what about the big issues? The ones that aren't on the list. I'm not a supercomputer, and I don't have terabytes of data at my fingertips, but here's a quick crack at what I'd add to ChatGPT's nine–point list.

10 AI may give you a completely wrong answer. And when it's wrong, it's not just wrong. It's often confidently wrong, lulling you into a false sense of security with answers that often lack any qualifying language.

The fact that it makes mistakes is one thing. But even more interesting is the fact that, in creating an otherwise exhaustive list about what issues in–house lawyers need to know about artificial intelligence, the AI omits the fact that it can make terrible mistakes. Mistakes that—if blindly relied upon—can damage brands, ruin careers, and hurt people. AI likes to give strong opinions and demonstrate confidence in its answers. Is it possible that AI doesn't like to admit that it might be fallible? That sounds almost … human.

It's easy to be wowed by something called "artificial intelligence." It's the bright shiny object that everyone is chasing right now. Companies are tripping over themselves to introduce the next AI version of their product, law firms are now AI driven, and I'm betting that when the IPO market opens up sometime in the next year, every company is going to claim that its products and services are powered by AI.

But it can botch simple math problems. It can miss basic geography questions and generate code that doesn't work. It can create a legal brief with citations that are completely made up, and it can misidentify facial features in a manner that can lead police to wrongfully accuse someone of a crime. So are you going to rely on ChatGPT to write your next contract, draft a brief, or handle your next quarterly filing? I don't think so.

11 While over–reliance on AI might be a dangerous siren song, ignoring AI is probably just as dangerous. AI isn't perfect. But it's impressive, and it's getting better, fast. Your lawyers need to be exploring how AI can be used to double check work, look for missed issues in briefs or clauses in contracts, find issues in vendor reviews, or review your security policies. If you aren't exploring how AI can make your legal department better, you are probably missing opportunities to work better, faster, and cheaper.

12 If you keep working on #11, and AI keeps advancing to the point where it makes fewer mistakes than humans, in ten years you might not need to worry about issue #10 anymore. And that's a sobering thought for millions of people who make a living practicing law around the world.

Tools to Get Organized

Need a checklist to help you identify the various risks AI may pose to your company or your clients? Check out the following resources from Bloomberg Law:

- Addressing privacy, security, and other data risks for AI

- Laws employers that use AI need to know

- AI considerations for employer policies

- M&A due diligence checklist for AI

Bloomberg Law subscription required.

Artificial Intelligence Newsletter

SIGN UP

Bloomberg Law | Bloomberg Tax | Bloomberg Government

Avoid the Pitfalls of AI with Multidisciplinary Teams, Robust Vendor Review

*A conversation with Bloomberg Law analyst **Stephanie Pacheco**.*

You've spent a lot of time thinking about AI and the legal industry. What's the one thing that's really surprised you?

There's still that cautionary approach that legal always has when they're adopting technology, but you can see that there's also a sense of urgency with the entire industry.

I think there's an overall industry approach that this is something we do need to understand, and we can't be the last people to get it, which you normally see in legal.

And yet some law firms and companies don't plan to use it. How long can that last? Is there a point at which they fall behind?

I think they're already behind, honestly. This is something that is going to impact not just the legal industry, but every market, every industry—and very quickly.

You have to understand it not only to use it in your own department or your firm to increase your own efficiency, but you're going to have clients that are asking about it and there are going to be novel legal issues coming out of these developments.

If you don't know anything about it: One, you run the risk of trying to advise clients on issues you don't understand, and then you run the risk of violating ethics codes and committing malpractice.

Or, two: You run the risk of losing business because these are the next wave of legal issues we're going to see roll out.

If you don't know anything about it or you're not taking the time to learn, you're going to just kind of fall out. It's not that AI will replace the lawyers. It's that those who understand it will replace those who don't.

Are legal departments and their outside counsel on the same wavelength when it comes to how they should be using AI?

I think it really falls back to one main issue: billing.

How is outside counsel billing for their time and what are they doing to justify that bill? I think for a lot of in-house attorneys and general counsels, they're not going to want to pay for a firm that is just using generative AI to do everything. But they're also not going to want to pay the cost for firms that aren't keeping up with the technology.

What you'll see is a lot of work will remain in-house at first draft, and then they'll confirm with outside counsel.

It's going to be a lot less of that initial back and forth…starting with a blank page and

an open problem that you go initially to outside counsel to get insight on.

They're just not going to want to pay for that lengthy bill for a firm's associates to spend weeks developing a motion or what have you.

It sounds like legal departments have a lot of leverage here.

I think so. Legal departments have been wanting to push for alternative fee arrangements for a while, and this technology is going to be that leverage for them to change the way legal billing is done, especially from an in–house to outside counsel perspective.

Some law firms say they're trying to build the AI–assisted firm of the future. But nobody really has the expertise for that yet—everybody is learning as they go. Who is really going to provide that service? They have to turn to somebody, right?

The legal industry at large—especially law firms—needs to lean into the idea and understand the importance of multidisciplinary teams.

This is something I wrote about last year: how the legal industry needs to make room for data experts. I think that holds true now.

The legal industry needs to welcome in data scientists, data analysts, data engineers— all of these people that have the technological expertise to help them create the models they need. It's going to be imperative for the industry, especially law firms, to lean into outside expertise.

What could blindside a general counsel or a chief legal officer on AI?

Understanding the source of the data and where it's coming from is a super important step that you can't skip in adopting these technologies, because that's where your product starts.

Understand what's going into your data pool, where you're pulling from, what happens to the data you're putting into the model. If it's a prompt–type tech like ChatGPT, where is your data going?

You need to understand what open sources you might be accidentally contributing to that you don't know about.

Read the terms and conditions of all of your vendors, especially as companies begin merging together and adopting or working with AI models. It's really imperative that GC's stay on top of that.

It might require a multilevel examination of terms and conditions.

Really look at the level of review that you're doing on your terms and conditions with your vendors and make sure that if they are taking steps to work with AI companies or AI entities, you understand what that means for your data.

And that's not just vendors within the legal department but the company as a whole?

It's your entire company.

'Supercool': OpenAI's New Top Lawyer Preps to Face Novel Issues

By Brian Baxter

OpenAI Inc. has a new general counsel and head of litigation as the company behind ChatGPT faces a spate of lawsuits and increased regulatory scrutiny.

Che Chang, a former OpenAI deputy general counsel, took over in July as the company's top lawyer. OpenAI the same month hired Renny Hwang from Alphabet Inc.'s Google to oversee litigation.

"It's an exciting time, both energizing and exhausting, as there's a lot of stuff to cover and we still have a relatively small team compared to the amount of surface area and issues that we face," Chang said in an interview. "We're tackling issues of first impression everywhere, which is super interesting and supercool."

Generative artificial intelligence programs like OpenAI's ChatGPT, which has conversational applications that allow it to create content via text, have opened the door to new legal issues.

Thousands of authors, including the comedian Sarah Silverman, have hit OpenAI and others with copyright claims for allegedly using their works as training material for such chatbots.

Chang replaced Jason Kwon, who took on the new role of chief strategy officer. Chang now handles the day–to–day management of OpenAI's legal group, which he said currently has about 20 lawyers on staff.

The company is looking for lawyers "across all disciplines," he said.

Hwang, who confirmed he joined OpenAI in July as a deputy general counsel and head of litigation, is among the more prominent legal additions. The former Latham & Watkins associate spent nearly 14 years at Google, where he was most recently a senior legal director and global head of litigation.

His new role puts him in the hot seat, as a growing number of lawsuits have been filed in recent weeks against OpenAI.

'Branding Issue'

"AI has a bit of a branding issue," Chang acknowledged in an interview from OpenAI's London office. "There is a lot of work required—not just on the legal, regulatory, or policy end—but as a society that we're trying to figure out."

Chang said he has "high confidence" that "sophisticated legal jobs" won't be directly affected by the technology. Instead, more rote tasks better suited to having a robot take a first crack at something are where ChatGPT is best deployed, Chang said.

OpenAI's legal group uses ChatGPT to rewrite, simplify, and summarize "chunks of text," usually to make them sound "less like legalese," Chang said. "It's very good at that and it's a simple thing that's not risky." Lawyers still vet anything generated by the artificial intelligence program, Chang said.

Chang has used ChatGPT for drafting job descriptions and potential outlines for talks and legal presentations, he said.

"People find it useful for inspiration and getting started, it's that first draft part, and then you can refine it yourself to get what you want," Chang said. "Those are popular use cases."

As for potential safety risks involving artificial intelligence, Chang claims that academia, industry, and government need to work together to find answers to difficult questions. OpenAI conducts extensive testing before deploying its products but is open to "iterating" them after release, he said.

OpenAI has recruited about a dozen lawyers so far this year. Most recently the company has recruited for a product counsel for artificial intelligence and a congressional lead for public policy, as well as a lead policy analyst for public policy.

> **" *There is a lot of work required— not just on the legal, regulatory, or policy end—but as a society that we're trying to figure out.* "**
>
> **––Che Chang, General Counsel at OpenAI**

+ Follow **us on LinkedIn**

Bloomberg Law

PRACTITIONER INSIGHT

How to Get Around the AI Trust Paradox in Legal Practice

By Anup Iyer   Moore & Van Allen

In a bustling law firm, trust isn't just a five–letter word—it's the very backbone of collaborative work. As a junior partner said at a recent AI task force, "After working with a four–year associate for a while, in time, I will trust the first draft of their work product enough to not spend too much time reviewing it. With generative AI, I will never develop that level of trust." This captures a deeply entrenched challenge in the world of legal tech—the trust paradox.

Understanding the Sentiment

At the heart of this sentiment lies a complex cocktail of human emotion, experience, and intuition. As lawyers work together over time, they develop a mutual understanding. They get a feel for each other's strengths, weaknesses, and quirks.

When a senior attorney reviews a document crafted by an associate they have collaborated with for years, they see the words on paper, and understand the human thought process behind them. This tapestry of shared experiences, understanding each other's reasoning, and the occasional shared sigh over a late–night coffee builds a level of trust that's hard to replicate.

With generative AI, the narrative changes. Even if the AI consistently produces impeccable drafts, there's no shared journey, no late–night brainstorming sessions, no exchanged glances when a complex case gets a breakthrough.

The AI tool doesn't evolve in its understanding in the way a human associate might. It remains, in essence, an enigma, producing outcomes without the rich tapestry of human experiences. This abstract nature of AI becomes a trust barrier.

Addressing the Challenge

Human–AI Collaboration. One way to bridge this trust gap is through active human–AI collaboration. Instead of letting the AI draft in isolation, attorneys can work with it iteratively.

Feed the AI preliminary thoughts, get a draft, make edits, provide feedback, and then let it refine the document. This iterative process, over time, can instill a greater sense of control and predictability, fostering trust.

Transparent Algorithms. If attorneys understand the how and why behind AI decisions, they may feel more comfortable with the output. Offering transparency into how the AI model works, the data it was trained on, and the logic it employs can alleviate some trust concerns.

Consistent Quality Checks. While a junior partner might trust an associate's draft, they still typically conduct reviews—it's part of the legal due diligence. The same should apply to AI. Periodic quality checks, even if they decrease in intensity over time, can provide assurance of the AI's reliability.

Feedback Loops. Just as human associates learn from feedback, sophisticated AI models can too. By incorporating feedback loops, where the AI learns from the corrections and preferences of the attorney, the system can continuously improve and align more closely with the lawyer's expectations.

Empathy and Relatability. One of the things AI lacks is the human touch, the ability to empathize, and relate. Perhaps the solution doesn't lie in making AI more human but in recognizing and appreciating it for what it is—a tool. Knowing the strengths and limitations of this tool can set realistic expectations and foster trust in its capabilities.

While the trust paradox is real, it isn't insurmountable. With time, understanding, and consistent collaboration, even the chasm between human intuition and AI's algorithms can be bridged.

After all, every good lawyer knows that a seemingly watertight argument can always find room for a counter, given the right perspective.

" Even if the AI consistently produces impeccable drafts, there's no shared journey, no late–night brainstorming sessions, no exchanged glances when a complex case gets a breakthrough."

—Anup Iyer, Senior Associate at Moore & Van Allen

Bloomberg Law
Insights Perspectives From Legal Experts and Thought Leaders

Click on the squares to read more.

Why ChatGPT Won't Put Discovery Counsel Out of a Job

By Courtney Duncil
Moore & Van Allen

Law Firms of the Future Will Be Different in Three Critical Ways

By Barclay T. Blair, Bennett Borden and Danny Tobey
DLA Piper

In-House Counsel Should Jump on Board the AI Train

By Canby Wood
LexFusion

By Richard Coel
WMATA

ChatGPT Doesn't Have Ethical Obligations, But Attorneys Do

By Natalie Pierce and Stephanie Goutos
Gunderson Dettmer

NYC's New AI Bias Law Broadly Impacts Hiring and Requires Audits

By Jonathan Kestenbaum
AMS

ChatGPT Can Help Lawyers Do Their Jobs Today

By Eric Pacifici and Kevin Henderson
SMB Law Group

New Rules for AI in the EU Attempt to Control the Technology

By Kamales Lardi
Lardi & Partner Consulting

AI Systems Automate Efficiency, But Invite Regulatory Scrutiny

By Sarah Hutchins and Robert Botkin
Parker Poe

Section 230 Immunity Isn't a Guarantee in a ChatGPT World

By Kurt Drozd
Sionic Advisors

What the NYSBA AI Task Force Hopes to Achieve for Law Practice

By Richard Lewis
NYSBA

AI 'Authorship' Muddies the Waters of Copyright Law Claims

By Zvi Rosen
Southern Illinois University School of Law

Future of AI in the Tax World Prompts Both Excitement and Fear

By Grant Wardell-Johnson
KPMG

A Governance Strategy Is the Best Path to Head Off Digital Risk

By Sophie Ross
FTI Consulting

New Legal Professionals Should Learn to Use AI to Their Advantage

By Jo-Anne Pugh
BPP University Law School

Why Arizona State Law School Applicants May Use Generative AI

By Stacy Leeds
Arizona State University College of Law

AI-Generated Inventions Need Human Ingenuity and Patents

By Robert Plotkin
Blueshift IP

To contribute readers should contact Insights@BloombergIndustry.com.

About Bloomberg Law

Bloomberg Law helps legal professionals provide counsel with access to action-oriented legal intelligence in a business context. Bloomberg Law delivers a unique combination of Practical Guidance, comprehensive primary and secondary source material, trusted content from Bloomberg Industry Group, news, timesaving practice tools, market data, and business intelligence.

For more information, visit pro.bloomberglaw.com.

Bloomberg Law