Using Generative AI in Legal Departments: How to Guard Against the Risks Without Stifling Innovation

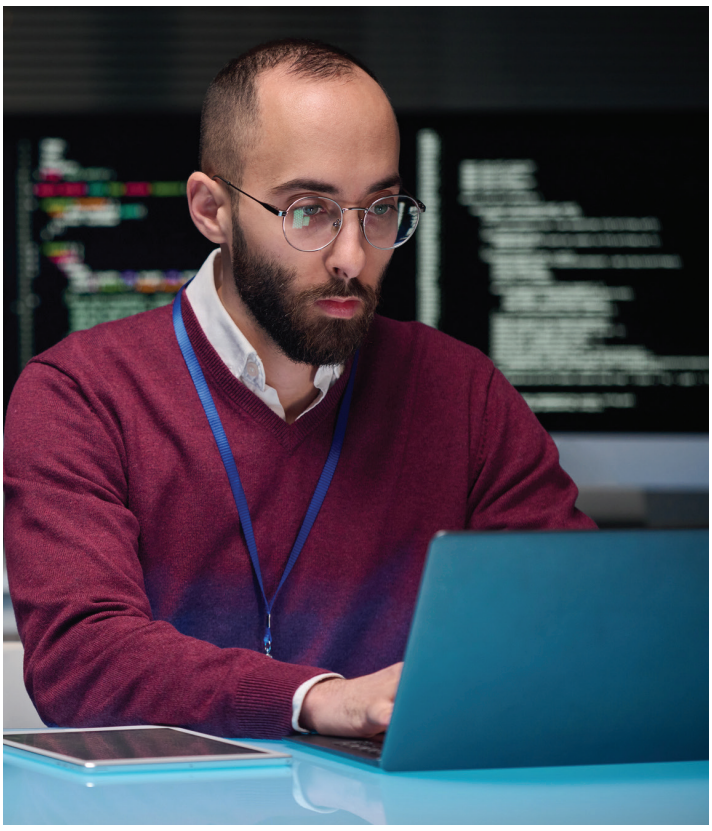

As in-house legal departments face more pressure than ever to be nimble and efficient while balancing ballooning workloads and tightening budgets, the capabilities of generative artificial intelligence (AI) technology have leapfrogged. But is AI the savior of overworked legal teams, a dangerous liability, or something in between? The answer, of course, depends on how it's used.

The evolution of AI technology heralds an exciting new chapter for businesses looking to streamline operations, but it also carries legal, ethical, and compliance risks that even its creators haven't fully untangled. For California in-house counsel, the challenge isn't just keeping up. It's striking the right balance between harnessing AI's power and protecting their company from costly missteps.

This practical guide to using AI includes explainers and best practices for legal departments to develop smart AI policies, mitigate risk, and ensure compliance without hindering progress.

Key technical concepts

It's impossible to anticipate all the legal risks of generative AI without a thorough understanding of how it operates. Generative AI isn't like traditional software, which relies on rigid, predefined rules. Instead, it generates outputs by identifying patterns in vast datasets, making predictions that are both creative and contextually relevant. However, this method also makes the results unpredictable and inconsistent — triggering legal, ethical, and compliance concerns.

Here are the most important concepts to be familiar with:

- **Generative AI:** AI technology that uses a large language model to produce new content that did not previously exist. Well-known examples include chatbots such as ChatGPT, Claude, and Gemini.
- **Large language models (LLMs):** Gigantic software systems that organize and respond to natural language inputs with natural language outputs (or other forms, such as code).
- **Training process:** LLMs are trained on vast datasets to simulate language knowledge. LLMs do not learn language or facts; they simply ingest a large amount of information, whether accurate or inaccurate.
- **Tokens:** When LLMs process these large amounts of data, they break this text into words and word fragments — known as tokens — and map relationships between them. For example, the LLM might see the word "apple" appear in various contexts alongside words like "fruit," "orchard," and "crisp." Meanwhile, it might also encounter "Apple" with words like "technology," "iPhone" and "Steve Jobs." By analyzing these patterns, the LLM learns to distinguish between the fruit and the company based on context.

- **Probabilistic vs. deterministic:** Unlike traditional programming, LLMs make probabilistic guesses about language and generate different responses to the same query. Unlike a calculator that will always return the same answer to, say, 2+2, LLMs do not have a set pathway to reach a single "correct" response. This is because they are making educated guesses and generating new content every time.
- **Hallucinations:** Because LLMs predict text based on patterns rather than retrieving verified facts, they can confidently produce responses that sound plausible but are entirely incorrect. For example, an LLM might invent a legal case that doesn't exist or attribute a quote to the wrong person.

The biggest AI risks for legal professionals

Generative AI's nondeterministic nature puts in-house counsel in a tricky position as they balance business objectives with legal requirements. This technology doesn't follow a fixed set of rules or paths to a definitive answer; instead, it generates predictions based on patterns in large datasets. In a field where precision and consistency are non-negotiable and mistakes can have significant consequences, this inherent unpredictability is problematic.

While AI has been in development for some time, its widespread commercial use is still relatively new, making it difficult to pinpoint all the risks involved. One major hurdle is the overconfidence some companies have in deploying untested or new AI tools in the name of efficiency and innovation, often without fully understanding the technology's capabilities and limitations. While AI promises significant advances, many leaders have an unrealistic sense of what it can actually achieve.

Rapid AI adoption is also expanding in-house counsel's responsibilities. Legal teams are now tasked with negotiating contracts for AI products and services, engaging with AI security professionals, and navigating an expanding web of privacy and legal compliance issues. For in-house counsel, developing proficiency in AI is no longer optional but essential for the long term. Understanding where AI can be used with minimal risk versus high risk will be key to deploying it effectively.

High-risk uses of AI

Some AI applications in legal work pose more significant risks than others, especially when it comes to areas requiring accuracy and precision. Here are some of the highest-risk uses.

- **Legal research:** Asking AI for factual answers can be problematic. While it can provide general information about legal principles or case law, it was not designed to give precise answers to specific legal questions. For example, if you ask a chatbot, "What was the ruling in the 2022 *Smith v. Jones* case?" you're more likely to get a probabilistic guess than a factually accurate response.
- **Brief writing:** Using AI to assist with drafting is useful, but relying on it to create entire arguments from scratch is a high-risk move. While AI can make predictions based on the data it has been trained on, it can't verify the accuracy of the information it generates. Legal-specific models are more refined, but they still operate on educated guesses rather than ensuring the precision of the information.
- **Inputting client information into public chatbots:** AI models learn from the data they process. This adaptive learning means that if confidential client information is entered into a public chatbot, there is a risk that it could be used to train the model. This could inadvertently lead to sensitive information being revealed in a future response to another user outside the company via a phenomenon known as "regurgitation."
- **Unexplained behavior:** Despite the significant advancements in AI, there's still a lot we don't understand about how these models work. Data scientists are still uncovering why some AI models produce certain outputs, and this lack of transparency makes the technology harder to manage, both in terms of effectiveness and in mitigating risks.

Lower-risk uses of AI

Some applications of AI present opportunities for efficiency with minimal downside when used thoughtfully. Here are some of the safer ways to integrate AI into your legal department's workflow.

- **Administrative tasks:** AI can streamline time-consuming, repetitive tasks such as scheduling meetings, summarizing emails, or managing document versioning. Since these applications don't require fact-based accuracy in the same way legal analysis does, the risks are significantly lower.
- **Drafting boilerplate language:** While AI-generated legal arguments are unreliable, using AI to draft standardized contract clauses or template agreements can be a practical and efficient way to save time. Lawyers should still review and refine the output, but AI can provide a solid starting point for routine documents.

- **Internal knowledge management:** AI-powered tools can help legal teams organize and retrieve internal resources more efficiently. Unlike public AI chatbots, private legal AI systems trained on a firm's own data can offer a controlled environment for surfacing relevant information without the risk of exposing sensitive client details.
- **Summarizing documents:** AI excels at condensing lengthy contracts, case law, and other legal texts into digestible summaries. While human review is still necessary to ensure accuracy, this can significantly speed up the process of identifying key provisions and takeaways.
- **Compliance monitoring:** AI can assist legal teams in tracking regulatory changes, flagging potential compliance risks, and monitoring industry developments. By automating updates on evolving legal frameworks, AI can reduce the burden of manual tracking while ensuring lawyers stay informed.

Even in these lower-risk applications, human oversight remains crucial. AI can enhance efficiency, but you must remain actively involved to ensure accuracy, relevance and compliance.

Professional responsibility considerations

Ethical and professional obligations are a cornerstone of the legal profession — and this extends to AI. While not updated this year, the State Bar of California's AI Guidance (2023) remains highly relevant for attorneys navigating AI-related ethical responsibilities. Remember, perfection isn't the goal — being diligent and adaptable is.

Here are the key considerations and risk mitigation strategies for in-house counsel to consider:

Competence

All California attorneys have a duty to stay informed about developments in the law and legal practice, and that includes the benefits and risks of emerging technologies. Now that AI is increasingly embedded in legal tools, business applications and client operations — often without explicit disclosure — it's important to proactively assess whether it's involved in the services or tools your company uses. Under Rule 1.1 of the California Rules of Professional Conduct, lawyers must "keep abreast of the changes in the law and its practice, including the benefits and risks associated with relevant technology."

- **Risk mitigation strategy:** Invest time in AI education and stay current with both legal and technical developments in the field. Develop more than a superficial understanding of its implications.

Confidentiality

LLMs may use inputs to continually train and improve, potentially leading to the regurgitation of confidential information in responses to other users. This creates significant risks for maintaining client confidentiality under Rule 1.6 of the California Rules of Professional Conduct and Business and Professions Code section 6068.

- **Risk mitigation strategy:** Be extremely careful about the client information you input into AI systems. Consider using closed models or vendors with clear data protection policies. Review the AI policies of your vendors and ensure they align with your confidentiality obligations.

Diligence and accuracy

Lawyers cannot blindly rely on AI-generated outputs or vendor assurances without independent verification. LLMs can confidently produce fabricated information, such as nonexistent case citations or inaccurate legal analysis. As data scientists have noted, some AI model behaviors remain unexplained, further complicating their reliability.

- **Risk mitigation strategy:** Verify all AI-generated content against reliable sources, especially for legal research or brief writing.

Supervision obligations

Rules 5.1 and 5.2 establish responsibilities for both supervisory and subordinate lawyers. Supervisors must ensure compliance with professional rules, while subordinates cannot blindly follow directions involving AI without understanding the technology.

- **Risk mitigation strategy:** Establish clear policies on appropriate AI use and ensure proper training at all levels.

Vetting your AI vendors: critical questions to ask about AI tools

Relying on assumptions or vague assurances about a vendor's AI tool's security, training data, and compliance with legal and ethical standards is risky. In-house counsel must independently assess whether an AI system aligns with their firm's professional responsibilities, data security requirements, and regulatory obligations.

Before adopting an AI-powered tool via your company or a vendor, consider these critical questions:

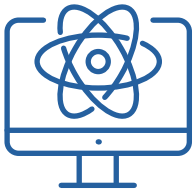

Training data:

- What is the AI model trained on?
- Were the sources used to train the model lawfully obtained?
- Are they reputable and appropriate for legal work?
- Does the training take into consideration inherent biases, and are there safeguards in place to ensure those do not affect results?
- Is the system continuously training on new user inputs? If so, does that create risks of confidential data exposure?

Data protection, control and security:

- Who can access the tool and its outputs?
- Does the vendor have safeguards to prevent unauthorized access?
- Can other users see or benefit from your company's data?
- Does the AI tool's data storage and processing impact attorney-client privilege?

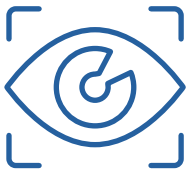

Ongoing AI development:

- Is the model evolving in a way that aligns with legal compliance?
- Does the vendor disclose updates and changes to how the AI system operates?
- Are there clear policies on managing AI-generated outputs that might be inaccurate or misleading?

Document vendor responses to these questions and ensure contractual protections address data privacy, confidentiality and compliance risks.

Developing an AI policy: key considerations for legal teams

A well-defined AI policy is essential for ensuring compliance, protecting client confidentiality, and maintaining the integrity of legal work. Here are the key elements to consider when developing a policy for your organization:

Define acceptable uses

Establish clear guidelines on when and how AI tools can be used. Determine whether AI is appropriate for tasks such as legal research, contract analysis, document review, or internal communications. Consider whether its use should be restricted to nonsubstantive tasks or require human verification for legal conclusions.

Ensure compliance with legal and ethical standards

AI policies must align with California's data privacy laws, industry-specific regulations, and professional responsibility rules. This includes adherence to attorney-client privilege obligations, compliance with the California Consumer Privacy Act (CCPA) as well as any other international laws that might apply, and safeguarding against AI-generated errors that could compromise legal work.

Set boundaries on data input

Specify what types of data can and cannot be input into AI tools. Prohibit the entry of confidential client information, privileged communications, or proprietary business data unless using a secure, closed system with strict safeguards. Ensure employees understand the risks of inadvertently exposing sensitive information.

Training and accountability

Educate all employees, contractors, and other stakeholders on AI's limitations, potential biases, and risks of misinformation. Establish training programs to ensure responsible use and require verification of AI-generated outputs to prevent reliance on inaccurate or fabricated information. Conduct training regularly, especially after making policy updates and if there are any important regulatory changes. Clearly define oversight responsibilities to ensure compliance with the firm's AI policy.

By proactively developing a comprehensive AI policy, legal teams can harness AI's benefits while minimizing risks and maintaining professional responsibility. Revisit this policy frequently (at least every three months) as AI technology continues to change rapidly.

CEB®

Making AI work for your in-house legal department

As one of the most exciting — and disruptive — developments of our time, the challenge of integrating AI into company operations can feel overwhelming. In-house counsel is under increasing pressure to embrace these tools, balancing innovation with the weight of ethical and regulatory obligations. But adopting AI responsibly doesn't have to be a massive undertaking, nor does it have to be a significant source of liability.

The key is structured policies, continuous education, and a willingness to adapt. AI's risks and benefits will shift over time — some aspects will become safer, while new challenges will emerge. Staying ahead requires some level of comfort with uncertainty and a readiness to revise your approach as technology evolves. By implementing thoughtful guardrails now and committing to ongoing oversight, you can harness AI's potential without compromising professional responsibilities.

*This guide distills insights from a recent CEB webinar, "Managing AI Risks and Ensuring Responsible Use: An In-House Counsel Discussion," featuring Jeffer Mangels Butler & Mitchell partner and cybersecurity/privacy group chair Michael Gold and Google's former in-house AI general counsel, King & Spalding partner Benjamin Softness. For a more detailed discussion, the webinar recording is available at **ceb.com/webinars**.*