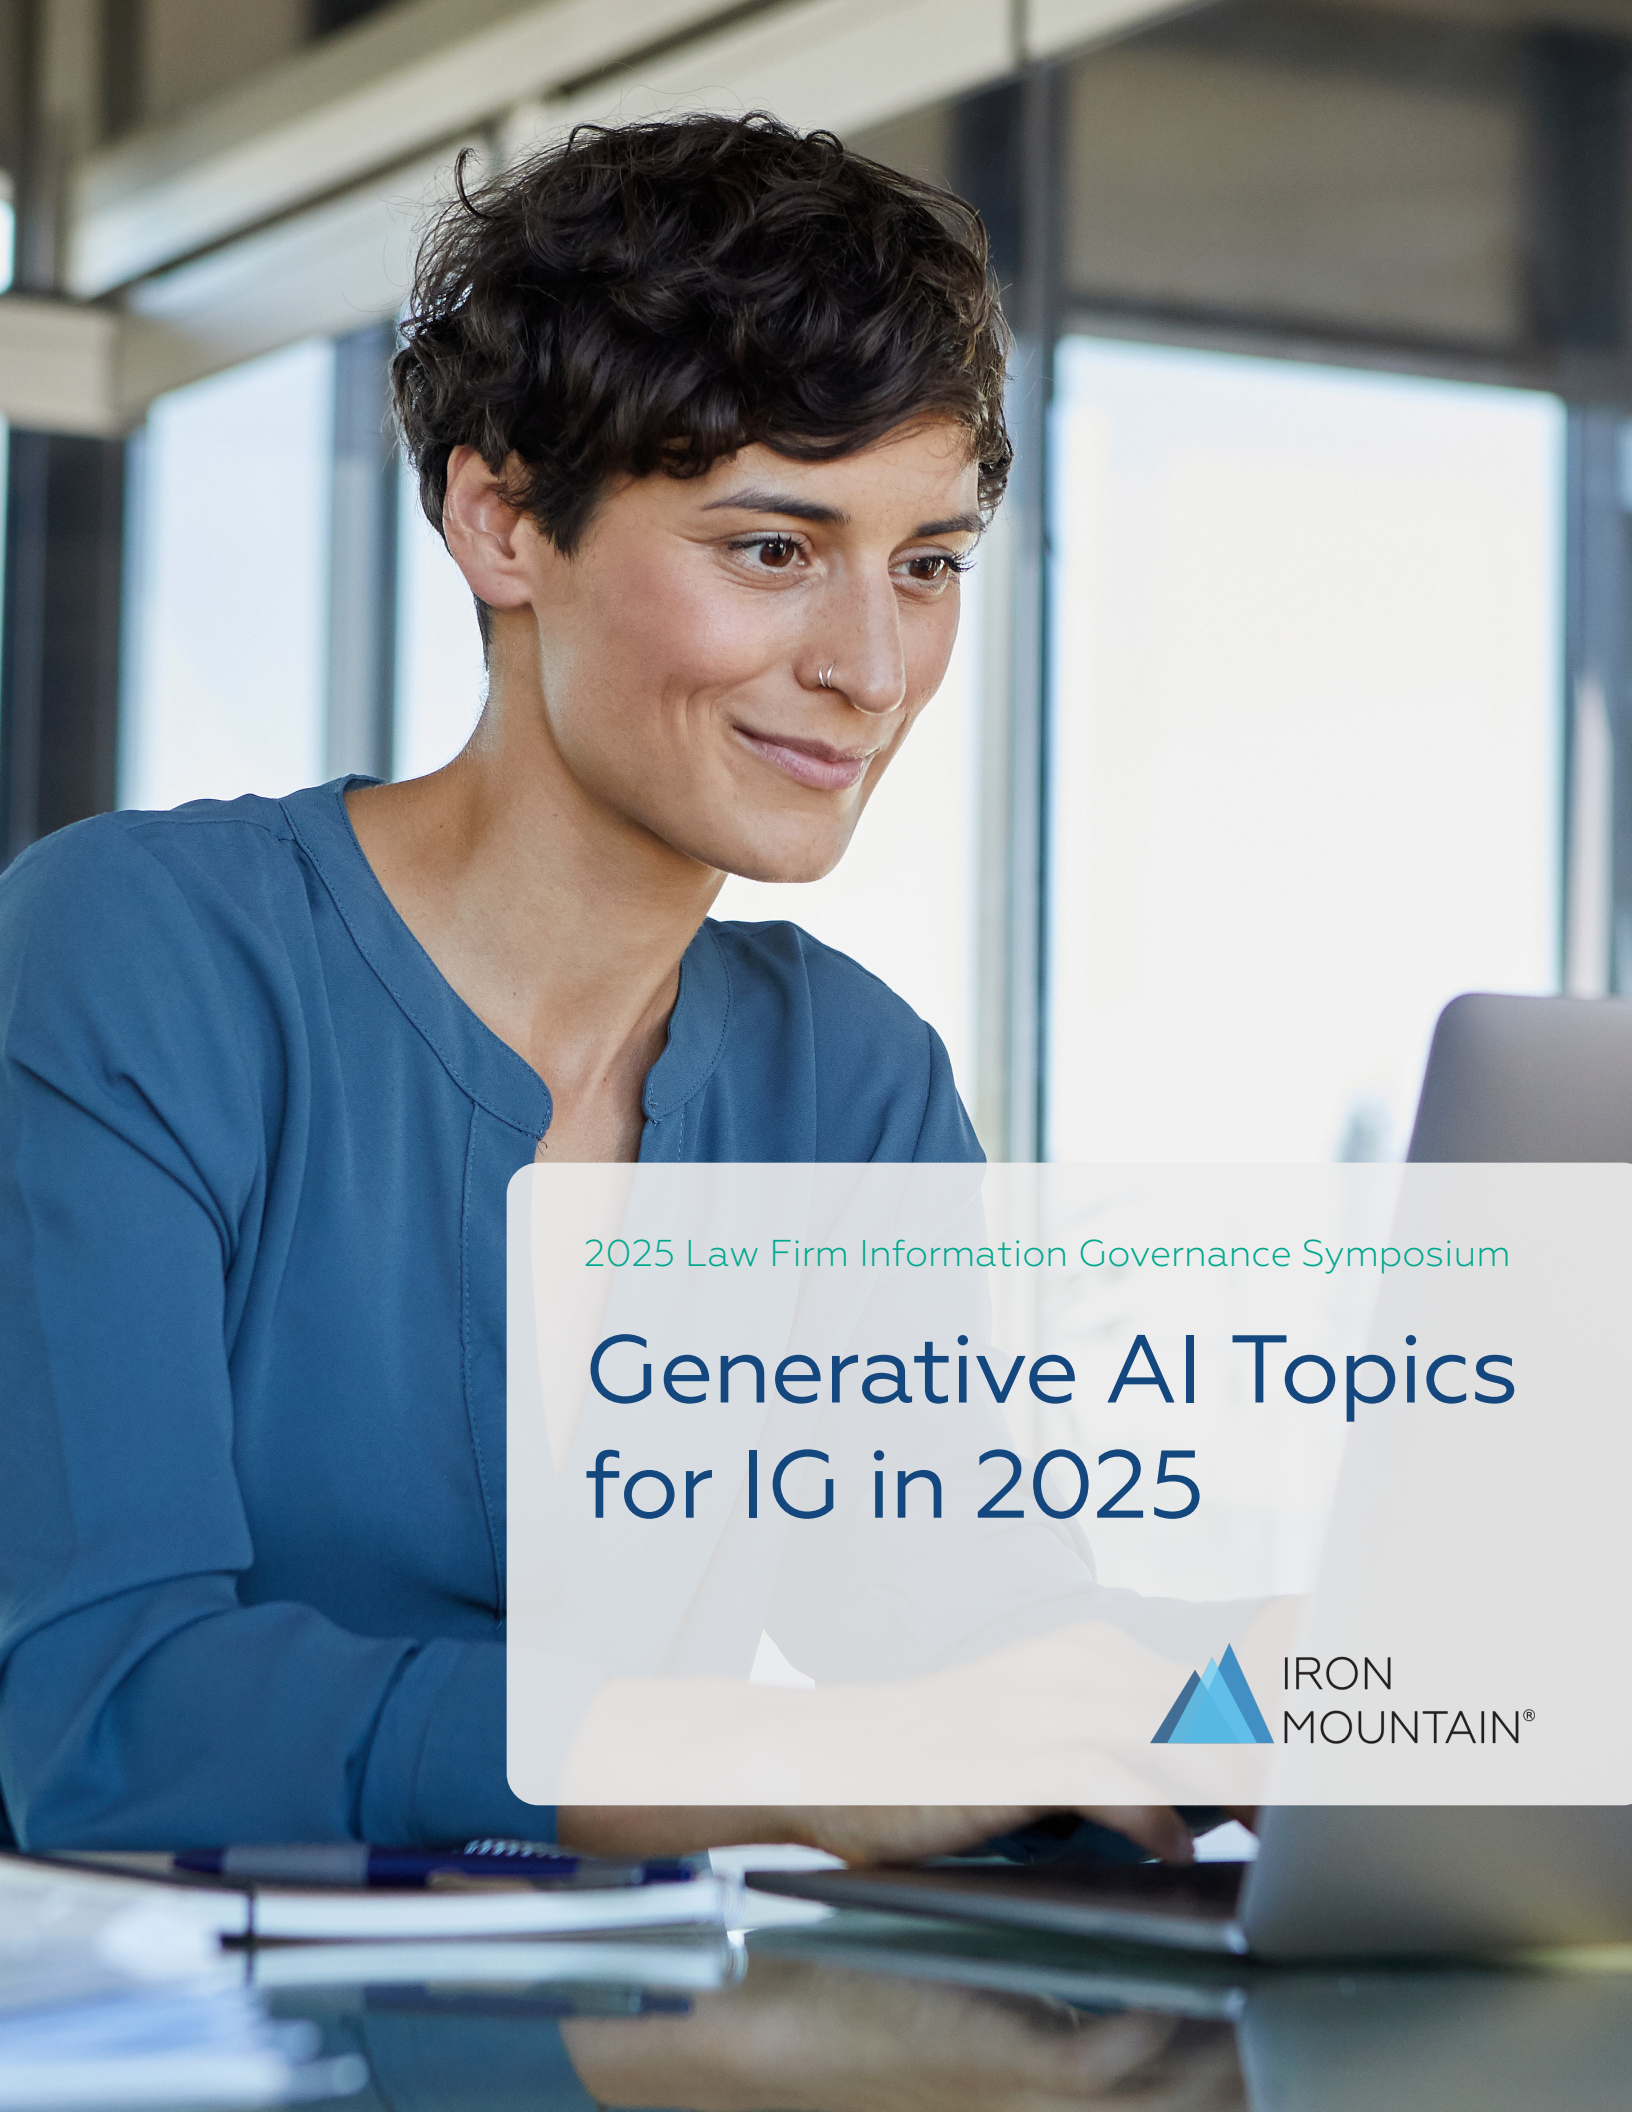

2025 Law Firm Information Governance Symposium

Generative AI Topics for IG in 2025

IRON MOUNTAIN®

Contents

Authors

Karen Allen
Information Governance Manager
Wiley Rein LLP

Bryn Bowen
Principal Consultant
Greenheart Consulting Partners

Colin Cahill
Sr Manager Information Governance and Collaboration
InOutSource

Scott Christensen
Founding Partner
Kyndig Consulting

Galina Datskovsky, Ph.D., CRM FAI
Board of Directors
Open Axes and FIT

Brian J. Donato
Chief Information Officer
Vorys, Sater, Seymour and Pease LLP

Michele Gossmeyer
Director of Information Governance
Holland & Knight LLP

Rishi R. Maharaj
Director of Information Governance
Davis Wright Tremaine LLP

Jessica Marlette
Information Governance
Data Strategy Attorney
DLA Piper LLP (US)

Jim Merrifield, IGP, CIGO
Director of Information
Governance & Business Intake
Robinson & Cole LLP

Reggie Pool
Principal
Practical IG

Jon Washburn
Chief Information Security Officer
Stoel Rives LLP

Leigh Zidwick
Sr Director Information Governance
DLA Piper LLP (US)

Introduction

The legal sector is undergoing a rapid transformation driven by the adoption of generative AI (GenAI) technologies that are poised to play a significant role in their operations. Large language models (LLMs) are a particularly exciting development, offering significant potential for law firms and professionals. The LFIGS 2025 paper, **AI Considerations for IG Processes**, explores the ways GenAI impacts and is impacted by certain law firm information governance (IG) process components. This companion paper investigates specific technological aspects of GenAI through the lens of IG.

Understanding LLMs for law firms

Law firms seeking to enhance efficiency and deliver superior client service are increasingly exploring and implementing GenAI tools. Among the most impactful advancements are LLMs, sophisticated artificial intelligence (AI) systems capable of understanding, generating, and manipulating human language. The rise of LLMs has presented unprecedented opportunities for legal professionals, from automating routine tasks to augmenting complex legal analysis. However, navigating the diverse landscape of LLM deployment options requires a nuanced understanding of their characteristics and implications. Law firms now face a complex decision matrix, moving beyond the simple choice between public and private models to a range of specialized solutions tailored to specific needs.

For IG professionals, understanding these differences is essential. LLMs aren't just a technology choice; they touch on data handling, access controls, compliance, client confidentiality, and firm risk posture. The way an LLM is deployed—where it runs, how it's grounded in firm content, what it retains or doesn't—can significantly impact everything from ethical obligations to regulatory exposure.

We'll begin by breaking down the different types of LLMs through a legal use case lens. Not all "private" models are created equal. Some implementations—such as Microsoft 365 Copilot—combine public model infrastructure with private data handling in ways that require thoughtful evaluation. Similarly, open-source models introduce valuable flexibility but come with their own operational responsibilities.

This overview is designed to help legal and IG teams ask the right questions, align stakeholders, and evaluate LLM options with clarity, not just curiosity.

Public LLMs

Public LLMs are general-purpose models made available by companies like OpenAI, Google, and Anthropic. They are typically accessed through the cloud via an API. These models are powerful and cost-effective, making them a popular starting point. However, because they process data in external environments under terms that may not address best practice security concerns, they can raise concerns around confidentiality—particularly when handling client-sensitive materials. To address this, some firms use a method called "fine-tuning," which means training the model further on firm-specific information (like internal contracts or case law) to improve relevance.

What is fine-tuning?

Fine-tuning is the process of taking a pre-trained LLM (which already understands general language patterns) and training it further on a smaller, focused dataset—often unique to a specific organization or field. This helps the model adapt to a particular tone, vocabulary, or workflow. For law firms, that might mean using internal case files, contracts, or legal memos to make the model more accurate and relevant for legal tasks. Think of it as customizing a general AI to speak your firm's language.

Private LLMs

Private LLMs are single-tenant or customer-controlled deployments (on-prem, VPC, or dedicated cloud tenancy) where the firm can control the data and model lifecycle. They may be a tool the firm hosts itself, or a tool a provider deploys to a logically isolated instance in the firm's subscription. This setup gives the firm more control over how the model behaves and where the data lives, which is especially important when security and compliance are priorities. These models can also be fine-tuned using internal legal data, which helps tailor their outputs to the firm's specific language, tone, and practice areas. However, private models typically require more resources—both technical and financial—to manage.

Hybrid LLMs

A hybrid approach blends public and private models to get the best of both worlds. For example, a firm might use a public LLM for general tasks like drafting or research but switch to a private model when dealing with sensitive content such as client documents. This setup requires careful and smart configuration so that sensitive data is always handled securely. A good example of this model is Microsoft 365 Copilot, which uses public LLMs for language processing and keeps prompts and responses inside Microsoft 365 with tenant-scoped access controls; neither prompts nor responses are used to train foundation models.

Open source LLMs

Open source LLMs deserve a specific callout. These models, such as Meta's LLaMA, are made publicly available to use, modify, and deploy. It's important to note that "open source" refers to how the model is shared and licensed—not how it's hosted. An open source model can be run publicly (e.g., hosted online for general access) or privately (e.g., deployed within a firm's secure systems). This flexibility allows firms to have full visibility into how the model works and the ability to customize it. However, open source options typically require more internal expertise to set up, maintain, and adapt to legal workflows. Fine-tuning an open source LLM with a firm's own data can create a highly customized solution for tasks like e-discovery, legal research, or contract analysis.

Domain-specific LLMs

Some LLMs are designed specifically for the legal field. These models are trained or fine-tuned on legal texts and terminology, making them more attuned to how lawyers think and work. Often delivered as part of a software-as-a-service (SaaS) platform, these models can plug into practice management tools or document management systems, helping streamline tasks like contract review or legal research. Because they are focused on legal use cases, they often deliver better results with minimal configuration.

Edge LLMs

Edge LLMs are smaller models designed to run locally on a device or inside a secure system without relying on the cloud. They are ideal when speed, privacy, or limited connectivity are key concerns—such as reviewing documents on a secure laptop or processing data in a high-security environment. While less powerful than full-scale models, edge LLMs are improving rapidly and offer a practical solution for firms prioritizing control and performance in specific workflows.

Summary of LLM options

LLM Type	Definition	Derived From	Key Characteristics	Actual Product Examples	IG Considerations
Public LLM	General-purpose models accessible via public APIs, trained on open datasets (sources vary by provider)	Publicly trained models	Broad capabilities; privacy concerns unless used with strict controls; fast and flexible	OpenAI ChatGPT (API), Anthropic Claude, Google Gemini	High risk of data leakage; typically not suitable for sensitive or client data without additional safeguards
Private LLM	Custom-trained and/or fully controlled by an organization, often hosted internally	Proprietary models	Maximum control; secure; expensive to maintain; deeply customizable.	In-house model built on Azure ML, hosted LLaMA instance	Full control over data lifecycle and model behavior; enables detailed auditability and compliance tracking
Hybrid LLM	Combines public models with private data or infrastructure to balance security and performance	Public base with private integration	Secure integration with enterprise systems; uses organizational data without training the model	Microsoft 365 Copilot, Azure OpenAI + RAG pipeline	Strong IG alignment via Microsoft Graph for Copilot; maintains tenant boundaries; supports audit and compliance policies
Open-source LLM	Freely available models that can be hosted and customized by firms	Public open-source code	Flexible deployment; firm controls model behavior and data access	LLaMA , Mistral, Falcon	Can be hardened for compliance; requires IG oversight to manage lifecycle, versioning, and usage monitoring
Fine-tuned public LLM	A public model further trained with domain-specific data	Public base model	Higher relevance for niche tasks; cloud or hybrid hosting	GPT-5 fine-tuned via OpenAI API or Azure OpenAI	IG must ensure data used for tuning complies with retention and confidentiality rules; risk of overfitting sensitive data
Fine-tuned private LLM	A privately deployed model enhanced with firm-specific data	Private base model	Highly specialized and secure; supports domain-specific automation and workflows	Custom legal assistant LLM built in a secure VPC	Strongest IG control; data residency, audit, and retention policies fully enforceable
Domain-specific LLM	Purpose-built or fine-tuned models for a particular field, like law	Public or open-source base	Targeted legal reasoning and terminology; often SaaS-based	CoCounsel, Harvey.ai, Spellbook	Evaluate vendor's IG posture (data use, retention, transparency); may require contracts with data use clauses
Edge LLM	Lightweight models deployed on local devicesor secure infrastructure	Public or open-source	Designed for low-latency and high-security use cases; offline or air-gapped environments	Phi-3, DistilBERT for local document processing, Apple Intelligence	Strong on-device IG control; good for regulated environments; must manage version control and local storage hygiene

Choosing the right approach

Selecting the right LLM strategy is ultimately about trade-offs. Firms must weigh their need for data security, performance, cost, and compliance. There's no one-size-fits-all answer. The best path forward aligns with the firm's specific goals, the sensitivity of the work being done, and its ability to manage the technology. Whether it's a public model, a fully private deployment, or something in between, the key is striking a balance between innovation and risk.

Model Type	Pros	Cons	Use Cases
Public LLMs	- Powerful, easy to use - Inexpensive to start	- Data privacy issues - Limited customization	Basic drafting, quick research
Fine-tuned public LLMs	- Better for specific tasks - Backed by well-known providers	- Fine-tuning takes work - Ongoing privacy concerns	Contract analysis, legal research
Private LLMs	- High security - Total control	- Expensive - Needs expertise	Sensitive data, custom apps
Privately fine-tuned LLMs	- Highly tailored - Firm-specific results	- Big investment - Ongoing work	Specialized legal tasks
Hybrid LLMs	- Balances cost and security - Flexible use	- Complex setup - Data routing issues	Blending public/private tasks
Open-source LLMs	- Full control, customizable - Transparent code	- Needs in-house management and expertise	Tailored apps, research
Fine-tuned open-source LLMs	- Specific accuracy - Fully customizable	- Needs high implementation effort and expertise	E-discovery, contract analysis
Domain-specific LLMs	- Made for legal work - Integrates with tools	- Limited to legal tasks - Vendor lock-in	Contract review, case law
Edge LLMs	- Fast, private - Runs locally	- Less powerful - Needs local setup	On-device review, local apps

The reality for most firms

Public LLMs offer broad capabilities and general-purpose functionality, while private LLMs are custom-trained models that are owned or licensed and controlled by a specific organization. While private LLMs offer a powerful advantage, the reality is that developing and maintaining one is a complex and expensive undertaking, making it a viable option only for those willing to make the required investments in resources and expertise.

However, law firms don't necessarily have to choose between using a public and a private LLM. In fact, a strategic and controlled approach that leverages both can offer significant benefits:

- **Complementary strengths:** Public and private LLMs excel in different areas. Public LLMs can be very helpful with initial research, legal trend analysis, and generating basic legal content for public education when maintaining confidentiality is not a requirement. Private LLMs shine in secure, in-depth legal research tailored to the firm's expertise, automating tasks, and mitigating bias. By using both, firms gain a wider range of capabilities.
- **Cost-effectiveness:** Public LLMs are often free or low-cost, making them ideal for tasks that don't require the firm's confidential data. Private LLMs, while more expensive, are more efficient for specific legal workflows, potentially saving time and resources in the long run.
- **Workflow optimization:** Public LLMs can handle initial research around specific matters (e.g., understanding how a specific term is used across multiple industries) and basic tasks (e.g., creation of marketing and training materials), freeing up lawyer time for complex legal analysis and client interaction. Private LLMs can then streamline internal processes like document review and generation.

However, a strategic governance approach is crucial:

- **Data security:** Public LLMs should never be used with sensitive client or firm data. Beyond sensitive client data, whether any client data can be used with public LLMs requires a firm's careful risk assessment and review of any client requirements regarding the use of their data.
- **Verification and accuracy:** Always verify the information generated by any LLM with reliable legal sources, especially because many of these tools may produce outputs that are factually incorrect or misleading, often referred to as a 'hallucination'.
- **Workflow integration:** Clearly define tasks best suited for each LLM to avoid redundancy and ensure a smooth workflow.

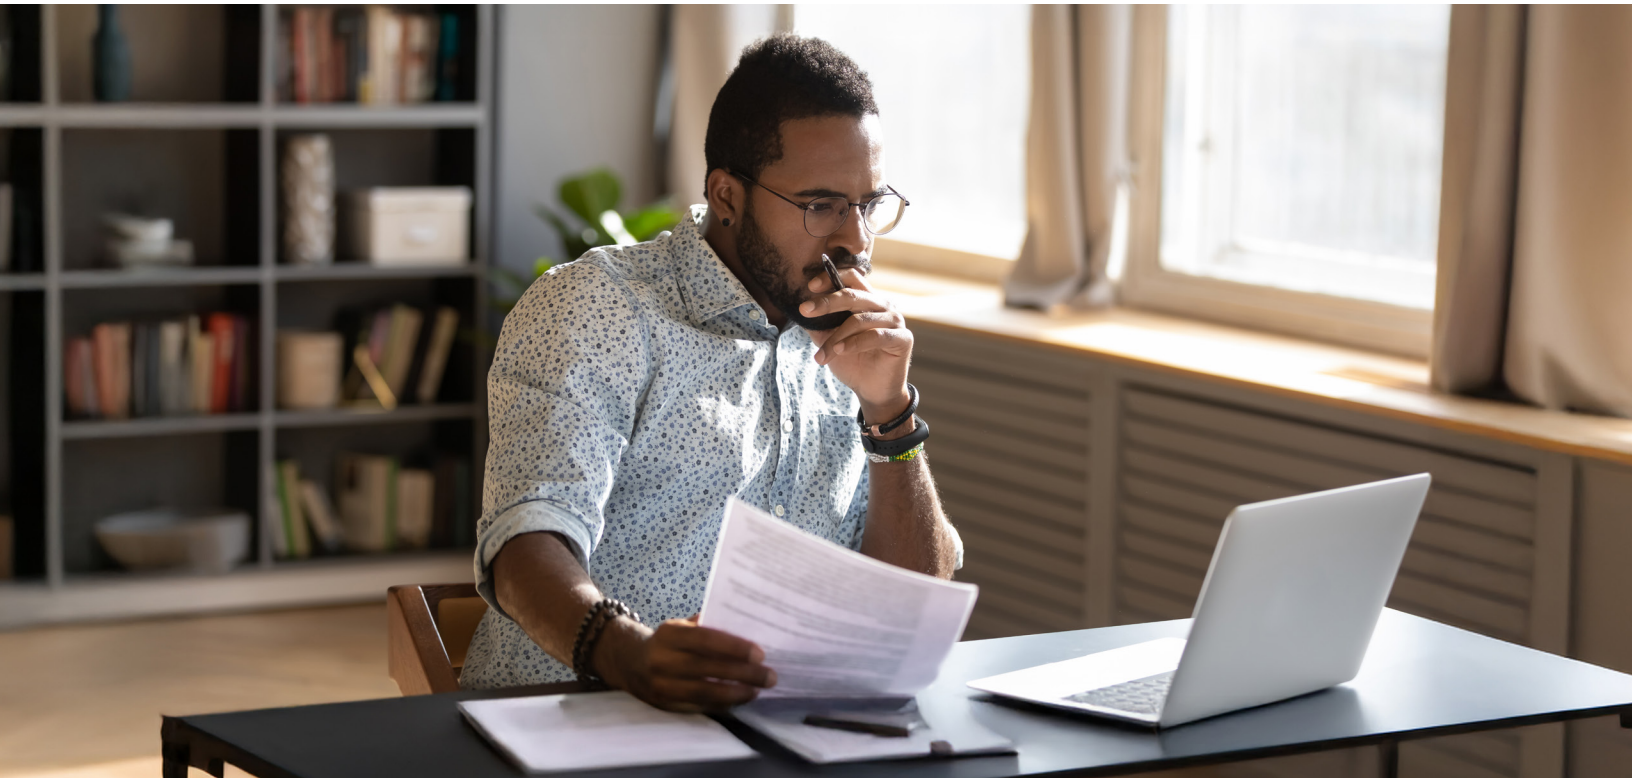

Most law firms will likely leverage the capabilities offered by third-party LLM products integrated into existing legal software. This provides a more accessible and cost-effective way to benefit from private LLMs, particularly for firms already utilizing Microsoft 365 products like Copilot.

Main challenge	How to address	Information governance impact
Data security and privacy risks	* Implement data anonymization techniques. * Enforce strict access controls/respect ethical walls. * Regularly monitor and audit the LLM environment. * Develop a comprehensive incident response plan.	* Defines data minimization practices (anonymization). * Establishes access control policies. * Requires data security audits and reporting. * Defines data breach response procedures.
Biased training data	* Conduct a thorough data assessment. * Ensure training data reflects the diversity of the data set. * Continuously monitor outputs for bias.	* Requires data quality checks for bias. * Defines data collection and selection processes. * Establishes data monitoring and remediation procedures.
Overreliance on AI	* Develop clear guidelines for lawyer interaction with LLM. * Emphasize critical thinking and independent judgment. * Require lawyers to explain reasoning when relying on LLM outputs.	* Defines protocols for human oversight of AI outputs. * May require changes to lawyer training and evaluation processes. * Establishes record-keeping practices for AI-assisted work.
Cost and implementation challenges	* Conduct a cost-benefit analysis. * Consider third-party LLM products. * Develop a phased implementation plan. * Invest in training for lawyers and staff.	* Requires justification for LLM development or adoption. * May impact data security practices depending on third-party solutions. * Defines change management procedures for workflow integration. * Establishes training needs and documentation for LLM use.
Explainability and transparency	* Train LLM on data with explanations. * Develop mechanisms for tracing LLM reasoning. * Provide clear explanations of limitations and any existing or potential biases.	* Defines data quality standards for training data. * May require specific data lineage practices. * Establishes disclosure requirements for LLM limitations.

By strategically leveraging both public and private LLMs, law firms can gain a competitive edge, optimize workflows, and ultimately deliver better service to their clients.

Understanding Microsoft Copilot

Microsoft Copilot is an advanced AI assistant developed based on OpenAI's GPT-4 series of LLMs. With the potential to revolutionize the workplace, Copilot aims to enhance productivity, efficiency, and creativity across various Microsoft applications.

Copilot for M365 is a private LLM, as it is restricted to information within the firm's tenancy, whereas the general public-facing Copilot operates with Microsoft applications and data but is not confined to tenant-specific content. Copilot aims to employ AI to boost users' productivity and creativity in various domains and tasks. It can generate natural language, code, data insights, and more, tailored to the user's data inputs and preferences.

Copilot can be integrated into all of Microsoft's products, including Microsoft 365, Azure, Power Platform, Dynamics 365, and GitHub. Copilot leverages Microsoft's expertise in AI research and development, along with its commitment to privacy and security, to deliver high-quality and reliable solutions that empower users to accomplish more with less effort.

AI is not merely a feature of Microsoft's Copilot—it has become a core technology integral to all Microsoft operations. Bernstein Research estimates that AI solutions currently impact over 42% of Microsoft's revenue and have the potential to double Microsoft Cloud revenue. They foresee that Copilot solutions will transform Microsoft and possibly the software and cloud computing industries.

Much like a copilot is vital for flying a large aircraft, Microsoft's plans for Copilot-based solutions are expected to become the standard method for powering its applications.

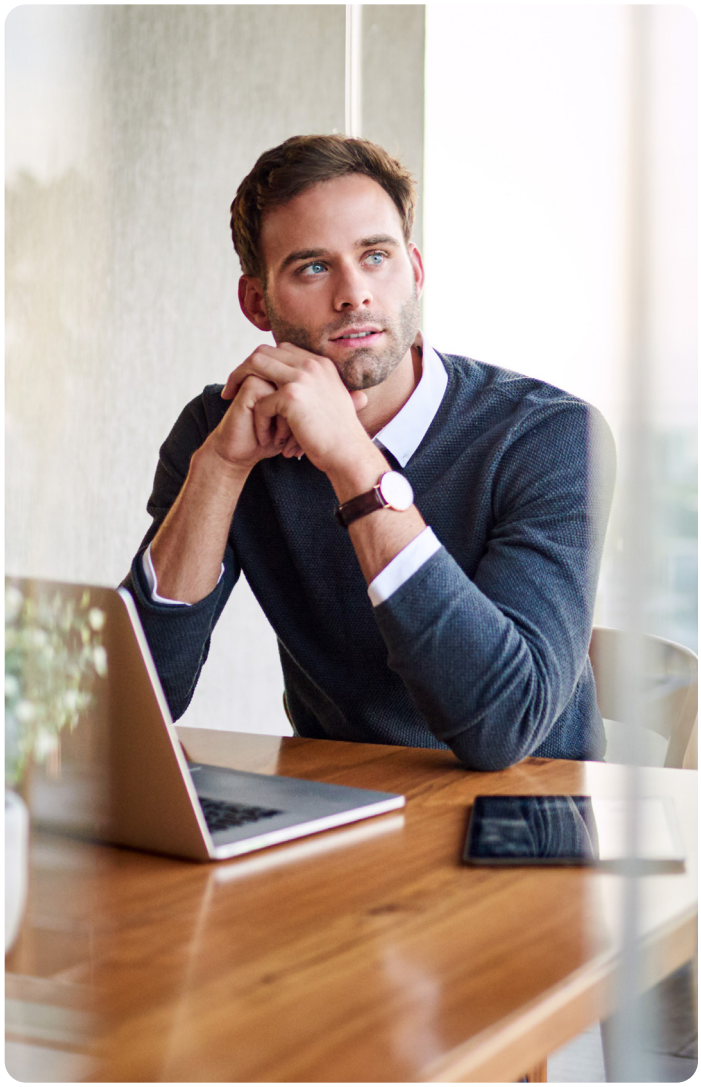

Microsoft Azure AI and Copilot

- Azure AI is a set of cloud-based services and tools that enable developers and data scientists to build intelligent applications using Microsoft's Azure Machine Learning and Azure Cognitive Services. Copilot runs on the Azure cloud infrastructure, providing the necessary computing power and scalability to handle complex AI tasks across a vast customer base.
- In essence, Azure AI provides the underlying technology and infrastructure that powers Copilot's intelligent capabilities, allowing it to understand user requests, generate relevant responses, and perform a wide range of tasks, leveraging the power of Azure Machine Learning and Azure Cognitive Services to understand the context and intent of the user, generate natural language and data responses, and perform tasks such as formatting, summarizing, translating, or scheduling.

Many Microsoft solutions, such as Windows 11, Edge browser, Office, Microsoft 365, and Dynamics 365, have integrated AI capabilities through their Copilot features. Microsoft also created a comprehensive AI development platform and tech stack (Copilot AI Stack) to help customers and partners design their own AI-enabled solutions.

Copilot security and privacy

Microsoft has provided clarity around the data privacy issues associated with using GenAI on Azure. Microsoft assures that your data is secure, private, and never shared with anyone or even used to train LLMs. Microsoft Red Teams (security testing teams) have run millions of scenarios to try to break the security built into the Copilot AI stack or get the system to perform unacceptable behavior, and the company is still learning (and training). Open-source offerings are likely not to have similar enterprise-class privacy and security features and may require more in-house engineering to reach the same posture.

According to Microsoft, Azure OpenAI instances are isolated from other customers, and your data is not used to train or enrich the foundation AI models. In other words, customer data is fully protected, which is key for customers to really leverage AI. For the most updated information, please refer to **Microsoft's website**.

Microsoft Copilot vs. Microsoft Copilot for M365

Microsoft Copilot and Microsoft Copilot for M365 are advanced AI-powered assistants developed by Microsoft to enhance productivity and streamline tasks across various applications. While both share a common foundation, they serve different purposes and audiences, offering distinct functionalities and integrations.

The following table summarizes these similarities and differences:

	Microsoft Copilot	Microsoft Copilot for M365
Similarities	**AI foundation:** Both Copilot and Copilot for M365 are built on the same underlying AI technology, leveraging LLMs to understand and respond to user prompts. **Natural language processing:** Both tools can process and understand natural language, enabling users to interact with them in a conversational manner. **Task automation:** Both can automate various tasks, such as summarizing text, generating content, and answering questions.	
Target audience	Primarily designed for individual home users, offering general assistance with tasks like creating documents, summarizing text, or answering questions.	Tailored for enterprise professionals, aiming to improve productivity within the Microsoft 365 suite of applications.
Functionality and integration	Offers a more general-purpose AI assistant experience, with features like web search, document creation, and basic task management.	Integrates seamlessly with Microsoft 365 apps like Word, Excel, PowerPoint, and Teams, providing context-specific assistance for tasks related to those applications.
Examples	**Document creation:** Ask Copilot to write a summary of a news article or create a draft email. **Information retrieval:** Query Copilot for general knowledge or search the web for specific information.	**Document enhancement:** In Word, ask Copilot to rewrite a sentence, suggest alternative phrasing, or summarize a document. **Data analysis:** In Excel, request Copilot to analyze data, create visualizations, or identify trends. **Meeting assistance:** In Teams, use Copilot to generate meeting summaries, action items, or suggest follow-up tasks.
Privacy		Offers a more specialized experience tailored to the needs of enterprise professionals and limits its source content to the firm's Microsoft tenant.

Understanding where Copilot for Microsoft 365 fits in the LLM landscape

As law firms assess how to bring generative AI into their environments, Copilot for Microsoft 365 is often one of the first tools under consideration. It's already embedded in a platform many firms rely on daily, and it introduces LLM capabilities in a way that feels accessible and familiar. Because of that, it serves as a practical example for understanding how different LLM deployment models—private, hybrid, or something in between—can take shape in a real-world legal context.

While Copilot uses a commercial large language model hosted in the cloud, its configuration within Microsoft 365 is designed to align with enterprise security, access controls, and compliance requirements. Depending on how a firm defines its thresholds for data control and system ownership, Copilot can reasonably be viewed as exhibiting characteristics of both **private** and **hybrid** LLM models.

Private-like behavior through scoped access and security controls

Despite this foundation in a public LLM, Copilot offers characteristics that resemble a private model. For instance, it personalizes responses by referencing firm-specific data stored within Microsoft 365 applications like SharePoint, OneDrive, and Outlook. This context-aware capability allows Copilot to deliver highly relevant insights without directly modifying the underlying model. Furthermore, because data processing occurs securely within the organization's Microsoft 365 tenant, sensitive information remains protected, reducing the risk of data leakage.

Despite these safeguards, Copilot is not a fully private LLM in the traditional sense. Copilot relies on Microsoft's managed Azure OpenAI Service environment, which provides a secure but shared infrastructure.

Hybrid architecture: Public model + private context = RAG-like behavior

Copilot is a commercial model combined with organizational context and data to generate more relevant outputs.

This includes:

Component	Function
Hosted LLM	Provides the language generation capabilities
Graph API / Directory Data	Supplies context about users, permissions, and relationships
Firm's M365 content	Acts as real-time reference material for prompts and responses
Middleware orchestration	Manages routing of prompts, grounding, formatting, and permissions checks

This structure operates similarly to a retrieval-augmented generation (RAG) setup. **The model doesn't store or learn from internal firm data—it pulls in relevant content at the time of the prompt, based on current access and context.** This reduces hallucination risk and allows for tailored, document-aware responses without requiring direct data ingestion.

What is RAG?

Retrieval-augmented generation (RAG) is an AI architecture that combines a language model with a search or retrieval system. Instead of relying solely on the model's internal training data, RAG pipelines dynamically "look up" relevant information from an external source—like a firm's document management system or Microsoft 365 environment—at the time of the query.

This approach helps ensure responses are **grounded in real, up-to-date content**, improving accuracy and reducing the risk of hallucinations. It's especially useful in legal settings where precision and context matter.

Implications for law firms

From a governance and risk perspective, this hybrid model:

- Offers bounded, role-based access to firm content, minimizing exposure.
- Keeps control over where data resides and how it's processed.
- Supports compliance with internal information governance frameworks, ethical barriers, and confidentiality requirements.
- Adapts behavior and style but doesn't, on its own, resolve data-handling risks: fine-tuning ≠ privacy.
- Is in line with the modern reality. Many firms start with RAG and evaluations, using fine-tuning mainly for stable patterns that RAG can't address, such as formatting, style or tone.

This type of architecture represents a middle ground, offering firms the ability to work with advanced generative AI models while maintaining alignment with internal security, privacy, and compliance standards. It also highlights how not all LLM deployments are simply "public or private"—many fall into blended categories that reflect the nuances of legal work and enterprise data handling.

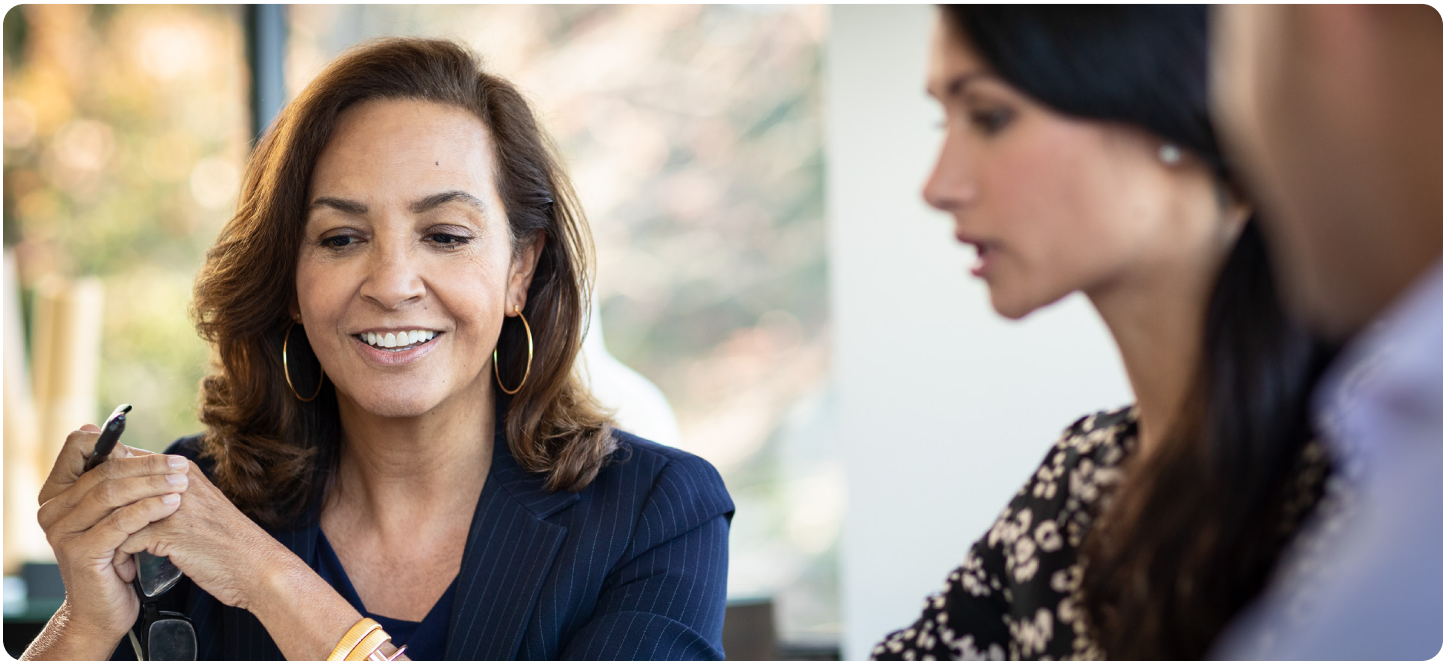

AI and data classification

When dealing with the vast amount of legal content law firms generate, the use of existing AI tools, such as file analysis software, to classify data may make it "safer" to use generative AI on that data. AI can help train itself by using classified source documents.

Document understanding refers to the process of using **artificial intelligence (AI)** and **machine learning (ML)** techniques to extract meaningful insights from text-based documents.

The following are examples of how AI-powered "document understanding" can be applied in this scenario:

Classifying by document type:

AI can categorize legal documents such as contracts, emails, and court filings by identifying keywords and formats.

- **Contracts:** Keywords like "agreement" and "breach" flag a document as a contract. Additionally, AI can recognize sections such as payment terms, confidentiality clauses, and dispute resolution mechanisms, ensuring that all essential components are easily accessible.
- **Emails:** Elements like "To," "From," and "Subject" distinguish emails. AI can further analyze the body text for legal implications, identifying correspondence that pertains to ongoing cases or contract negotiations.
- **Legal filings:** Specific formatting and keywords like "complaint" and "motion" help classify these documents. AI can also detect the type of filing, such as affidavits, pleadings, and briefs, aiding legal professionals in efficiently managing their caseloads.

Extracting key concepts:

Once categorized, AI can identify key concepts within documents.

- **Parties involved:** AI can extract and catalog the names and titles of entities in contracts, such as plaintiffs, defendants, attorneys, and witnesses, which is crucial for understanding the context of legal documents.
- **Obligations:** AI can highlight obligations related to performance standards, deliverables, and compliance requirements, allowing legal teams to monitor adherence and identify potential breaches.
- **Dates and deadlines:** AI can identify important dates for performance or termination, then create timelines and set reminders for critical deadlines, ensuring that no important dates are missed and that all actions are taken promptly.

Recognizing relationships:

AI can find connections across documents, linking email chains with related contracts. By establishing links between various documents, AI can provide a comprehensive view of all related materials, enabling legal professionals to see the big picture. For example, it can correlate contract amendments with email negotiations, ensuring that all modifications are accurately documented and agreed upon.

Refining the data:

AI cleans and enriches data, allowing LLMs to focus on relevant information for a better understanding of legal concepts and language. AI can remove duplicates, correct errors, and fill gaps in data, creating a high-quality dataset. It can also enrich documents with metadata, such as document type, creation date, and involved parties, facilitating more effective search and retrieval. This refined dataset enables LLMs to deliver more precise and nuanced legal analysis.

Overall, AI-powered document understanding creates an organized dataset, helping LLMs provide accurate and insightful legal information.

AI classification risks

While using machine learning (ML) and natural language processing (NLP) AI to prepare content for a legal-specific LLM model has potential benefits, there are several key risks to consider when using AI for document understanding and subsequently training LLMs on that legal content:

- **Biased training data:** AI systems rely on the quality of the data on which they are trained. If the training data contains biases, such as a lack of diverse legal cases or a bias in the classification algorithm, the LLM may inherit those biases and reflect them in its outputs. This could lead to skewed or inaccurate legal analysis.
- **Inaccurate classifications:** AI systems are not infallible, and misclassifications can occur, which may result in the LLM missing important legal information. Overreliance on these systems necessitates thorough quality checks. Misclassification could lead to overlooking critical contract clauses or incorrect sentencing guidelines.

To mitigate these risks, law firms can consider the following steps:

- Curate training data to ensure it is diverse and representative of the legal issues handled by the firm. Constraining data to specific classifications can be overly restrictive and introduce biases.
- Implement robust quality control measures to identify and correct AI misclassifications.
- Choose AI systems that provide transparency into their decision-making processes.
- Maintain strong data security protocols to protect client confidentiality, prohibiting the ingestion of classes of data that are expressly restricted by clients from interacting with LLMs.
- Foster a culture within the firm that emphasizes that AI complements human expertise and does not replace it. All output requires human review.

Overall, AI offers significant benefits for processing and understanding legal content, but it's important to proceed with caution and implement safeguards to avoid potential pitfalls.

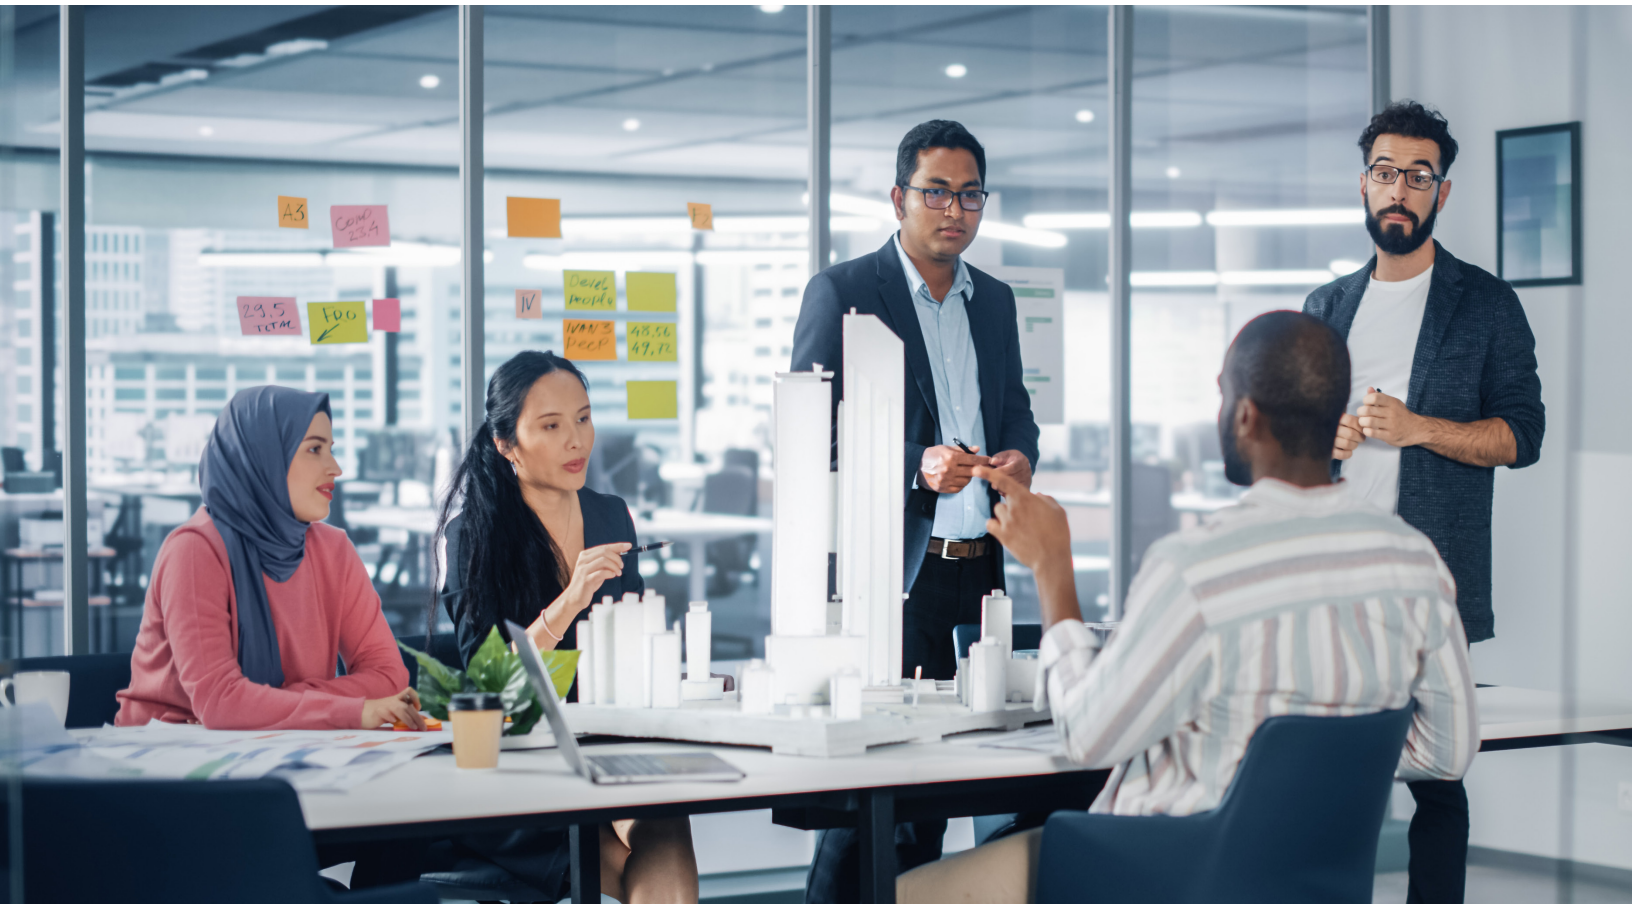

Prompt engineering

Prompt engineering is the art of getting a GenAI tool to provide a desired output. The question you ask to prompt the AI tool matters greatly, and how you evolve the conversation to fine-tune the results is the difference between getting a meaningful response and a superficial attempt at an answer.

An initial prompt might include a simple question such as "How should I go about building a project plan?" A well-engineered prompt includes the question, context about the question, instructions about how to go about solving the question, examples, and other potential modifiers. For example, "Answer the following as though you were an expert project manager: How should I go about managing a project to create a white paper on prompt engineering that includes deliverables such as assigning team members, identifying topics, assigning writers to each topic, and finally assembling a paper. Suggest topics and a potential project plan."

While there is a clear advantage to prompt engineering expertise, there are also questions about how useful it will be in the near future. First, many products tuned specifically for the legal industry (such as Lexis Protégé) limit the size and type of prompts available. Second, as models continue to be refined, it seems likely that the advantage of a great prompt over an average prompt will diminish. Finally, newer models such as ChatGPT-4o can help generate an effective prompt for you while agentic or thinking models such as ChatGPT o1 are, in essence, creating prompts behind the scenes to accomplish the goal of the original prompt. In fact, **there is some research to suggest** that humans are unlikely to stumble upon the most effective prompts that may be generated automatically by LLMs.

Instead, as this **HBR article** suggests, problem formulation is likely to be a much more enduring skill. From the article, "Prompt engineering focuses on crafting the optimal textual input by selecting the appropriate words, phrases, sentence structures, and punctuation. In contrast, problem formulation emphasizes defining the problem by delineating its focus, scope, and boundaries."

Still, prompt engineering will be a vital skill in the near term, especially as firms continue to utilize private, secure versions of public models. Fortunately, there is no lack of resources available for becoming skilled at prompt engineering. For example, OpenAI has an entire section of its documentation dedicated to **prompt engineering**. Microsoft also offers **similar resources**, as does **Anthropic for Claude**.

Finally, the addition of **generative pre-trained transformers** (GPTs), which are self-contained combinations of prompts and documents by OpenAI, means that skilled prompt engineers can create their own useful applications very quickly. Imagine, for example, taking an existing 50-state survey of employment laws and allowing lawyers to ask it questions. With the proper documents and prompts, this can easily be packaged for reuse by anyone within the law firm.

Governance and ethical concerns of prompt engineering

Prompt engineering can be like a nuanced art, requiring skill and intuition to elicit desired responses from AI systems. While it promises better results, it also brings forth governance and ethical concerns that warrant careful consideration.

From a governance standpoint, the foremost concern lies in establishing robust quality control and standardization protocols for prompt engineering. The lack of a uniform framework can lead to inconsistencies in AI-generated outputs, thereby affecting their reliability and utility. Furthermore, transparency in prompt design is pivotal. Opacity in how prompts are engineered could undermine trust and hinder accountability, especially when AI-generated advice leads to adverse outcomes. This necessitates a transparent approach that allows stakeholders to understand and evaluate the methodologies behind prompt construction.

Ethical concerns about transparency and client awareness of AI's role in their representation can be addressed through information governance by developing communication strategies to inform clients about AI use and ensuring that AI decisions are transparent. Finally, if a firm does use GPTs, which are available both from OpenAI and via Microsoft's various AI services, the firm and IG professionals should be concerned about the information embedded in those GPTs, first determining if, and then determining how it should be systematically governed.

Structural bias

Ethically, the concerns are multifaceted and profound. Bias and fairness emerge as critical issues, as the way prompts are structured can inadvertently perpetuate or amplify existing biases within the AI's training data. This can lead to skewed or discriminatory outputs, raising significant ethical red flags, particularly when these systems are used in decision-making processes that impact human lives.

Overreliance

Moreover, the increasing reliance on AI-generated advice, predicated on skillfully engineered prompts, poses the risk of over-trust in AI systems. This dependency could be perilous, and lawyers cannot take AI's suggestions at face value without understanding how the system involved arrived at that advice. Overreliance on AI can lead to a lack of critical thinking and independent judgment. IG can address this by developing clear guidelines for human oversight of AI outputs and ensuring that lawyers explain their reasoning when relying on AI.

Digital divides

Additionally, the skill gap in prompt engineering could lead to a new form of digital divide within law firms, where the effectiveness of AI technologies becomes contingent on one's proficiency in prompt crafting, further entrenching inequalities in performance. This is an old problem, but the aggressive adoption of GenAI in law firms might give it a new, riskier life. To mitigate these risks, IG practitioners could work to establish training programs to equip a wider range of employees with prompt engineering skills. This will help to democratize access to AI-powered tools and reduce the risk of a digital divide.

In light of these concerns, it is imperative to develop a comprehensive framework that addresses both governance and ethical challenges. This would involve not only the establishment of best practices and standards for prompt engineering but also the implementation of mechanisms to ensure fairness, transparency, and accountability in AI-generated outputs.

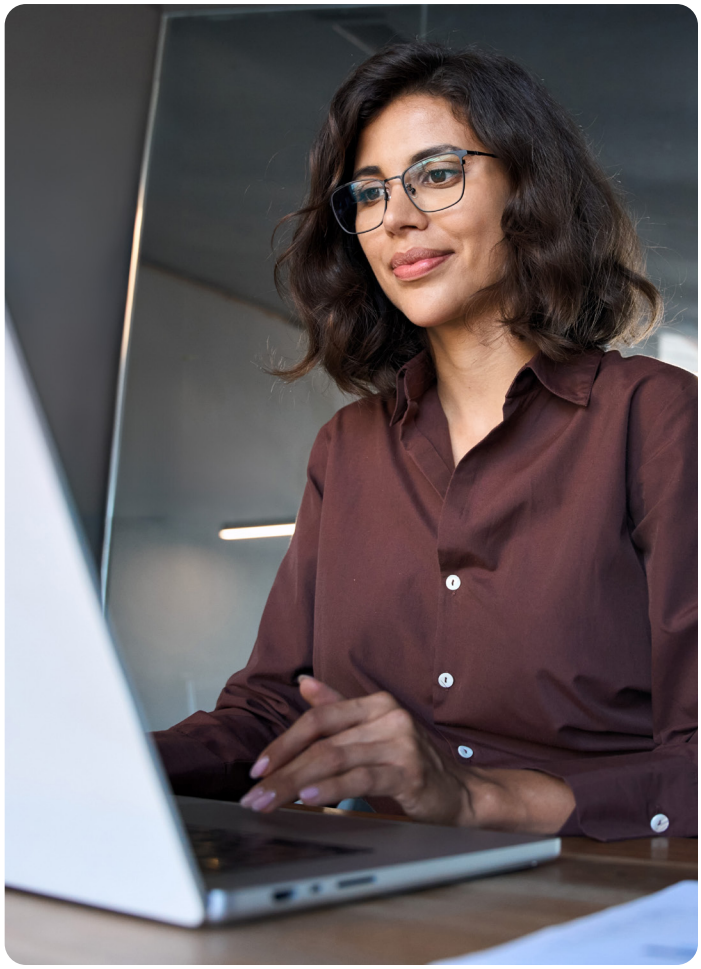

A renewed focus on information governance

The rise of both public and private LLMs presents challenges and opportunities for IG practitioners within law firms. A strong IG program is essential for successful LLM implementation. A solid focus on the fundamental tenets of IG can help guide the implementation of LLMs within a firm.

Challenges

Data governance: Ensuring the quality, security, and privacy of data used to train the LLM requires robust data governance policies and procedures. IG practitioners will need to develop and implement new protocols for data collection, classification, retention, and disposal specific to LLM training data. Practitioners should work to ensure the data used to train the LLM is relevant, accurate, and up to date, leading to more effective models. The benefits of good data governance include reduced costs, more efficient training of models, reduced bias in data and subsequent outputs, and a clearer understanding of the data lineage underlying the LLM outputs. Many providers now train on licensed and public data, while litigating or negotiating access to the rest. This landscape is evolving quickly, so be sure to review the vendor's current data-usage disclosures.

Change management: Implementing and integrating AI solutions such as private LLMs requires effective change management strategies to encourage user adoption and address potential resistance. IG practitioners can play a key role in developing training programs and communication plans to educate lawyers and staff on the benefits and limitations of LLMs.

Opportunities

Strategic role: IG practitioners can play a strategic role in shaping the firm's AI strategy by advocating responsible AI development and use. They can ensure that AI solutions are implemented in a way that aligns with the firm's ethical values and professional obligations.

New skillsets: Developing expertise in AI governance, data ethics, and bias mitigation will be increasingly valuable for IG practitioners. Staying informed about the latest developments in AI and its legal implications will be crucial for effectively managing the risks associated with LLM use.

By proactively addressing the challenges and seizing the opportunities presented by private LLMs, IG practitioners can play a critical role in ensuring the successful and ethical adoption of AI within their firms. They can become trusted advisors, helping lawyers leverage the power of AI while safeguarding the firm's data, mitigating bias, and maintaining the highest ethical standards.

Conclusion

The rapid rise of generative AI models is occurring at a lightning pace. Governance professionals need to stay abreast of the advancements in technology in order to effectively govern the firm's data. Keeping up with your vendors' adoption and integration of AI technology into their products is essential to knowing how the tools are trained, how the data flows through them, and why this knowledge is fundamental to effective governance.

Lawyers should always review AI-generated work to ensure its accuracy and relevance. While there is debate in the industry as to the value and/or relevance of maintaining a recording or evidence that a draft was generated by AI, there is currently no legal requirement or case law indicating that you should. However, under the EU AI Act as an example, deployers of certain high-risk systems must keep logs. While legal drafting and research tools generally aren't classified as "high-risk," firms may still need audit trails for client or regulatory reasons. Logging every prompt that was used in the drafting of a document can be a complex and voluminous undertaking. Just as firms have differing levels of risk tolerance with records retention and disposition programs in general, the same can be said about the level of detail that needs to be maintained and retained when using AI. An argument can be made that the use of AI for drafting is not so distant a practice from lawyers using legal research software such as LexisNexis or Westlaw to draft documents. This may change, however, and thus it is important to remain current on laws, regulations, and case law that may mandate disclosure.

This paper reflects the state of the industry as of the date of publication. Given this rapidly evolving landscape, the goal of LFIGS is to continue revisiting AI technologies in future papers.

IRON MOUNTAIN®

800.899.IRON | ironmountain.com

About Iron Mountain

Iron Mountain Incorporated (NYSE: IRM), founded in 1951, is the global leader for storage and information management services. Trusted by more than 225,000 organizations around the world, and with a real estate network of more than 98 million square feet across more than 1,400 facilities in over 60 countries, Iron Mountain stores and protects billions of information assets, including critical business information, highly sensitive data, and cultural and historical artifacts. Providing solutions that include secure storage, information management, digital transformation, secure destruction, as well as data centers, art storage and logistics, and cloud services, Iron Mountain helps organizations lower cost and risk, comply with regulations, recover from disaster, and enable a more digital way of working. Visit www.ironmountain.com for more information.

XXXX-WP-XXXXXXA