JND

CONSIDERATIONS FOR AN AI GOVERNANCE FRAMEWORK IN LEGAL ENVIRONMENTS

DEREK DRAGOTTA, SENIOR VICE PRESIDENT, JND LEGAL ADMINISTRATION
BEN SEXTON, SENIOR VICE PRESIDENT - INNOVATION & STRATEGY, JND eDISCOVERY

Recently, law firms, corporate counsel, privacy experts, and Government employees have started seeking guidance for bringing AI into their organizations, better and more safely. While every organization and use case is different, we all share the common basic need to govern our use of technology to protect the privacy and confidentiality of our data. This piece provides some practical thoughts for how to do so in the rapidly evolving AI environment.

Setting the Stage

AI is already transforming the legal practice across the globe. The most popular use cases include legal research fact finding, case strategy, drafting and review automation. Lawyers and technologists are rapidly exploring new ways to automate and enhance their work with AI.

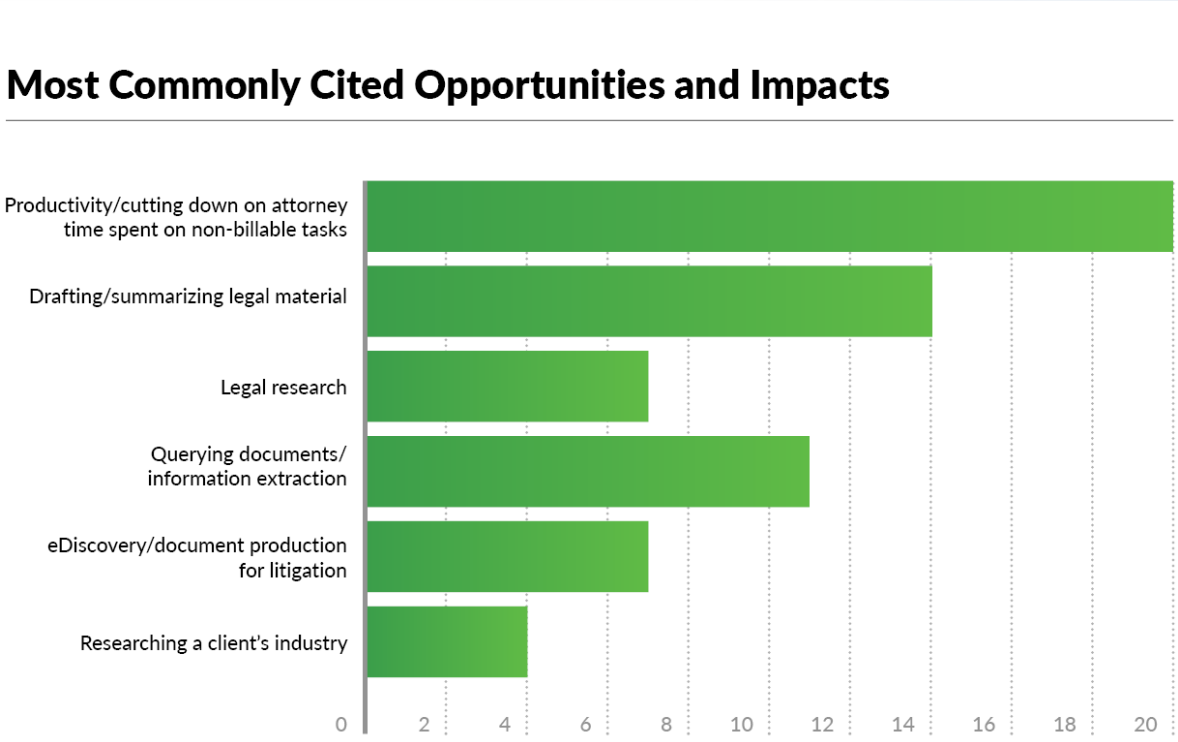

(The American Lawyer)

With that said, a "go fast and break things" approach can lead to disaster. So, how can you, as a stakeholder, advance your organization into the AI era without creating new risk for your business or firm?

Safety First

In late 2022, soon after ChatGPT was publicly released, many organizations issued company-wide edicts banning the use of AI. This prohibition was, undoubtedly, the sensible thing to do for any organization without formal policies and procedures governing the use of AI. Run a few quick internet searches and you'll see why.

Among other concerns, by default, most Large Language Models (LLMs) train their models using user inputs. That means that if an associate uses ChatGPT to summarize a deposition, the contents of the deposition, and the summary, may be available to other users who prompt the LLM in the future. There are straightforward ways to prevent this (e.g., most enterprise-level AIs allow you to disable usage of prompts as training data), but not without the proper controls in place to prevent it.

Competitive Pressures

Since its nearly across-the-board ban in 2022, pressure to use AI has started to build. Organizations are pushing law firms to find new ways to cut costs using AI, and software providers like Relativity are making better "fit for use" products. Practitioners are also working to develop safe and defensible workflows to enhance their product or service offerings. Interestingly, the largest firms seem to be leading the charge for adoption, likely because they're feeling the most competitive pressures from clients. According to a Lexis survey, as of January 2024, "53% of Am Law 200 firms have already purchased "Legal AI" tools (Gen AI developed specifically for the legal industry), compared to just 20% of non-Am Law 200 firms with 50 or more attorneys". The survey also found that 43% of AMLaw 200 firms have a dedicated budget for AI in 2024.

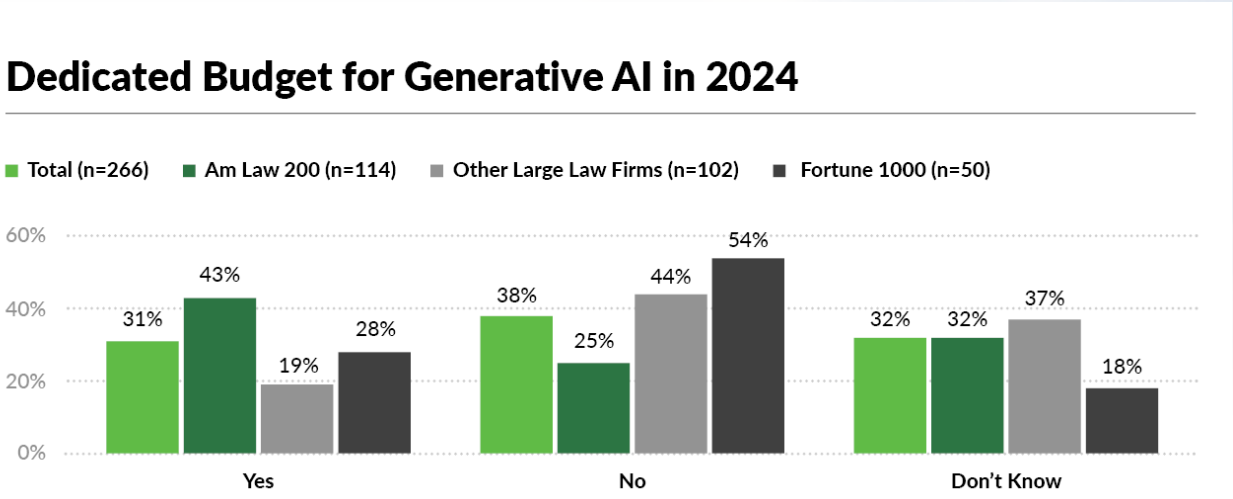

(LexisNexis)

Most organizations didn't intend for their "ban" to last forever, but in order to move forward, they understandably need a well-defined plan with checks and balances to ensure they're not creating unbridled risks. Hopefully what follows will serve as a head start for those endeavoring down that path.

Build a Gate and Establish Your Gatekeepers

Whether it be to manage organizational or technological changes, chart a course for security initiatives, or for something more specific to your individual organization, you need a team and policies in place to securely manage the adoption and use of AI. Change is constant and with AI it's fast. The following steps can serve as guardrails to ensure you're keeping up with the technology safely.

1. **Establish an AI Governance Board:** This is the group who will approve how and where AI can be used and govern the general direction your organization will take with its use. Think of it like the House or Senate but with less tweeting and actually getting things accomplished. This group should meet regularly and be comprised of staff from a multitude of departments. For this function to be effective, you'll need stakeholders from Information Technology, Information Security, Legal, Operations, Finance, Third Party Management, etc. Once formed, this group should be accountable for the overall approach your organization takes with AI.

2. **Create an AI Governance Policy:** This policy establishes an avenue through which AI is adopted. You don't need to get bogged down in the weeds here. This should be simple, succinct, and formally state that the AI Governance Board is the body through which any and all AI solutions must be reviewed and approved. All other AI software and/or solutions are strictly prohibited. This body should also be responsible for approving AI Usage. Ensure that references to this are included in any acceptable use policies you require your staff to review and attest to. You cannot hold people accountable if they do not know the rules. To that end, ensure your policies have enforcement clauses so you can take action when they are not followed.
3. **Create an AI Usage Policy:** Aside from formally approving AI software and solutions, the governance body should also approve how this AI can be used. For example, you may want to use AI to help respond to RFPs, but your AI Governance policy will require the proposal team to review and revise, where needed, any content drafted by the AI. For each AI use case you adopt, include relevant details in policy and procedural documentation to govern how your staff use the technology.

AI Software Procurement

You likely already have policies and procedures in place that govern third party risk management and those should apply to any AI software and technology providers. As a matter of fact, all third-party software should undergo a security and privacy review if it's transmitting, processing, or storing, sensitive data. AI, however, presents new privacy and security considerations that should be assessed prior to any usage. While most AI systems utilize similar 'back-end' technology (e.g., LLMs such as, OpenAI or Llama), how these technologies use your information, and where it's stored, for example, can vary. Depending on software/solution provider, you can control most of this, but it may not be secure-by-design, with the default settings configured to favor the provider, not the user.

While there are many things to consider when vetting any software or solution, here are five to get you started. You'll notice that these are largely the same as what you're used to today.

1. **Data Privacy and Security:** Confirm the AI solution has robust security controls in place to protect sensitive and confidential information. This includes end-to-end encryption of data in transit, secure data storage, data segregation between clients or uses, and regular security assessments. Ensure that the AI provider has clear policies on data handling and storage, and that these policies comply with relevant privacy laws and corporate policy.
2. **Data Usage and Retention Policies:** Check the provider's policies on data usage, particularly whether the data will be used to train AI models. It's crucial to ensure that your data is not used to improve the AI without explicit consent. Also, review the data retention policies to ensure that data is not stored longer than necessary and that there is a clear procedure for data deletion upon request or after the end of the retention period.
3. **Compliance with Legal and Ethical Standards:** Ensure that the AI software complies with statutory and regulatory requirements, legal standards, and ethical guidelines relevant to the industry. This includes considerations around the unauthorized practice of law and ensuring that the AI tool does not make decisions for your staff, but instead supports them in making informed decisions. The provider should also have mechanisms to address bias and ensure fairness in AI outputs.
4. **Auditability and Transparency:** The AI system should be transparent and auditable. This means you should have insights into how decisions are made by the AI (explainability) and the ability to audit the use and decision-making process of the AI to ensure compliance with legal requirements and corporate policies. Transparency in AI processes helps in understanding and trusting the AI outputs and in demonstrating compliance to opposing parties and regulators.
5. **Vendor Reputation and Reliability:** Research the provider's reputation in the market. Consider their track record with other law firms or legal departments, their expertise in legal technology, and their financial stability. It's also beneficial to assess the level of customer support they provide, including training for your team and technical support for troubleshooting. Additionally, consider how critical this service is to your organization. Review SLAs and, if necessary, retain alternate providers that can be used in the event your primary provider becomes unavailable.

Actionable Takeaway

As an actionable takeaway, consider including the following requirements in your privacy and security assessments. Note that many of these are likely already covered in your current risk assessment framework. However, it's worth addressing how these are impacted by AI.

Category	Inquiry to Provider	Description
Encryption	What encryption standards are used to secure scoped data both in transit and at rest?	Data should be encrypted both in transit and at rest. This prevents unauthorized access during data transfer and when it is stored. Ensure that the AI software uses strong encryption standards, such as AES-256 for data at rest and TLS 1.2 or higher for data in transit.
Access Controls	Describe Access Controls to ensure only authorized personnel can access scoped data.	This should include using role-based access controls (RBAC), and the least privilege principle to minimize access to sensitive data. Multi-factor authentication (MFA) should also be used to secure user access.
Data Segregation	How is data segregated between clients?	In multi-tenant environments ensure that data segregation is strictly enforced. This prevents data leakage between different clients or projects and makes data easier to destroy when it's no longer needed. Ask the provider about their data isolation practices and the technologies used to segregate data.
Data Retention	Does the AI retain data (e.g., prompts)?	The software should be configured to prevent the AI from retaining any text or data submitted through the prompt or otherwise. Likewise, the AI should be prohibited from using client data for the purpose of training its model.
API Security	How are APIs secured?	If the AI software interacts with other systems via APIs, ensure that these APIs are secure. This involves regular updates, using secure protocols, and implementing rate limiting and authentication on API endpoints to prevent abuse.
Assessments	How often are security assessments performed?	AI Technology is evolving quickly. Conduct regular security assessments, including vulnerability scans and penetration testing, to identify and remediate security weaknesses in the AI system. Also, ensure that the AI software is regularly updated to patch any security vulnerabilities.
Compliance and Certification	Verify that the AI software complies with relevant industry standards and regulations.	Look for certifications like ISO 27001/27002, SOC 2 Type 2, FedRAMP, or specific compliance standards relevant to your industry (e.g., HIPAA for healthcare, GLBA for finance, etc.).

Wrapping Up

Everything is scary until you have a plan and when it comes to AI, you certainly need one. However, as you can see, there's no need to reinvent the wheel. The keys to a robust AI Governance Policy are to involve the proper stakeholders from across your organization, lay your framework, develop your policy, design your use cases, and educate your staff. Hopefully these recommendations give you the "head start" you need to define actionable steps toward a safe adoption of AI in your organization.

Citations

Henry, J. (2024, January 29). We Asked Every Am Law 100 Law Firm How They're Using Gen AI. Here's What We Learned. https://www.law.com/americanlawyer/2024/01/29/we-asked-every-am-law-100-firm-how-theyre-using-gen-ai-heres-what-we-learned/

Simpson, M. (2024, January 31). 2024 Investing in Legal Innovation Survey The Rise of GenAI at Top Firms & Corporations. https://www.lexisnexis.com/pdf/genai-report.pdf

Derek Dragotta

Senior Vice President, JND Legal Administration

Ben Sexton

Senior Vice President - Innovation & Strategy, JND eDiscovery

About JND eDiscovery

JND eDiscovery is an innovative legal technology and eDiscovery services provider supporting law firm, government and corporate clients in the areas of: Generative AI, Litigation Readiness, Identification, Forensic Collection and Preservation, Early Case Assessment and Data Processing, Hosted Review, Managed Review, Analytics and Technology-Assisted Review (TAR), and Production.

JND is a subsidiary of JND Legal Administration (JNDLA), founded by Jennifer Keough (CEO), Neil Zola (Executive Managing Director), and David Isaac (Executive Managing Director). The company's eDiscovery service line was formed in 2016 with the acquisition of Minneapolis-based Alloy Group (now JND eDiscovery). JND eDiscovery's leadership and key personnel have been working side by side for more than 10 years. JND continues to operate out of Minneapolis and is led by EVP of eDiscovery, Scott Lombard and Ben Sexton, SVP of Innovation and Strategy.