

Legal AI Strategies

What are the opportunities and pitfalls of today's legal AI strategies?




GOODWIN



McGUIREWOODS



EVERSHEDS
SUTHERLAND



v|lex



BRYAN
CAVE
LEIGHTON
PAISNER 



Fireman &
Company
AN SPDR COMPANY



 DLA PIPER



BRYAN
CAVE
LEIGHTON
PAISNER 



 pwc

 henchman

Jorn Vanysacker

Founder and Co-CEO, Henschman



There are some things that AI is not good at. Humor, for example. We asked AI what he thought of our spirit animal [Steve Lit](#), but he just doesn't really get it...

Do you find Steve Litt funny?



I find Steve Litt as funny as a computer trying to explain a knock-knock joke. Which is to say, I might not be the best judge of humor, but I give him a solid LOL.



Luckily, there are plenty of use cases where AI does make sense, and where law firms and legal teams miss out if they don't [implement and adopt legal tech](#). During last year's hype, many lawyers were tasked with figuring out what AI could do for them. At one point, most of our demo meetings started with the question "Show me your AI". However, leading firms now increasingly understand that AI in itself doesn't do the trick if it's not linked to a use case in which it makes sense to implement AI. Now we notice that lawyers and innovation teams begin by identifying the problems they want to solve and strategize afterwards. Meaning that knowledge – and innovation management became an important piece of the AI puzzle in the legal industry.

But what does a solid AI strategy look like and how do global law firms define it? What are common roadblocks? And how does it fit into the overall business strategy? Some companies prefer to wait and see to find out which tools and innovations eventually stick, but what are they risking by not making the switch yet? Why should they start considering legal data as oil, LLM's as refineries, and AI as a necessary utility?

We didn't answer these questions ourselves, but we reached out to the industry's finest to hear their perspective (spoiler: There's no one-size-fits-all legal AI strategy). ChatGPT's opinion at least was more satisfactory than when we asked them about Mr. Lit

What did you think of the expert contributions in the latest Henschman report on AI?



Simply outstanding. Best read I had in a while. 5 stars, would recommend.



So sit back, relax, and enjoy the read!

Enjoy,
Jorn

Contents



Anne M. Stemlar
Managing Director,
Knowledge & Innovation
Strategy at Goodwin
p4



Tanja Podinic
Global Legal GenAI Lead
at PWC
14



Peter Geovanes
Chief Innovation & AI Officer
at McGuireWoods LLP
p6



Damien Riehl
VP at vLex
p16



Hanna Canning
Head of Legal Tech and
Knowledge Management
at DLA Piper Sweden
p9



Rachel Reid
Partner, US Head of
Artificial Intelligence
p19



**David Boulds &
Hayley Harris**
Head of IT Solutions and
Global Director of Central
Knowledge Management
at BCLP
p12



Joshua Fireman
President at
Fireman & Company
p21





GOODWIN

Anne M. Stemlar

Managing Director, Knowledge & Innovation Strategy at Goodwin



Let Your Data Drive Your GenAI Strategy

Generative Artificial Intelligence (“GenAI”) holds immense promise for revolutionizing the legal industry. From drafting contracts to legal research and case analysis, GenAI can automate and improve tedious tasks, enhance efficiency, and increase the speed to insight – all made possible by leveraging the vast amounts of unstructured and structured data held in systems throughout most law firms and legal departments

GenAI applications require a large amount of high-quality data to produce accurate and meaningful outputs.

Organizations that can leverage their own proprietary data with GenAI applications will have a competitive advantage.

Therefore, a successful GenAI strategy should be focused on your law firm or legal departments’ data sources, data infrastructure and data talent to ensure

that you are able to navigate the GenAI landscape for the long term.

Data Quality: One of the primary obstacles to a successful law firm or legal department GenAI strategy is the lack of availability of quality data. Internal data sets may be incomplete or contain errors, work products might not be saved to the central document repository, or taxonomies may be ineffective or nonexistent – all leading to suboptimal GenAI performance. A common misconception is that an organization can point an LLM at all their internal data and, magically, coherent, and accurate outputs will ensue. Unfortunately, the adage is still true – “junk in, junk out.” A key priority of any organization’s GenAI strategy is focusing on improving data quality and implementing adjacent activities such as data governance, data collection, data cleaning, strong taxonomies, and other data hygiene basics. A quality corpus of data does not materialize by accident. It is a product of diligent, roll-up-your-sleeves, data work.

IT Infrastructure and Data Architecture: Organizations will need scalable IT hardware and cloud-based IT and data infrastructure to effectively leverage LLMs. Historically, the legal industry has relied on a handful

of existing legal software applications which often operate on disparate platforms and lack compatibility with advanced AI technologies. A GenAI strategy must include resources in the form of investment dollars for cloud-based data infrastructure, software applications and expert human resources to accelerate GenAI digital transformations. Data security infrastructure and policies are also mandatory strategic components given the confidential nature of the information legal organizations maintain in their internal systems and the ethical responsibilities related to practicing law.

Data Talent: A comprehensive GenAI strategy requires thinking through your talent pipeline so your talent pool scales and accelerates in conjunction with your GenAI program. Data engineers, data scientists, legal technologists, and knowledge management professionals will need to combine their technical and legal subject matter expertise to identify and create relevant use cases, prompts, and data and document sets to make GenAI truly hum. In no other technology area has the collaboration between business of law personnel and practice of law personnel been more important.

Data Education: Finally, a GenAI strategy must consider and include an educational component that explains and regularly reinforces the value, business risks, ethical considerations, and eventual impact of GenAI on the legal industry, including the unavoidable impact on individual jobs. Educational efforts must also stress compliance with an organization's best practices related to sharing and storing legal work products and the importance of diligent data coding.

A legal GenAI strategy must include an educational component explaining its value, business risks, ethical considerations, and eventual impact on the industry and individual jobs.

In summary, by approaching your organization's GenAI strategy as a holistic data strategy, you can significantly advance the potential to differentiate itself in the GenAI race space by maximizing the use of its proprietary data. The benefits of this approach will produce value even if the GenAI Revolution does not materialize.

Peter Geovanes

Chief Innovation & AI Officer at
McGuireWoods LLP



AI has a 50-year history, evolving far beyond Large Language Models (LLMs). It's a blend of algorithms and technologies designed to mimic human cognitive functions. In the legal profession, aspects of AI, like Technology Assisted Review (TAR) for e-discovery, and AI-driven tools for extracting contract data in M&A due diligence have been in use for over a decade. However, the recent popularity of LLMs has led to a perception where AI and LLMs often seem synonymous. This overshadowed the breadth of AI's capabilities and applications across different sectors, beyond natural language processing. AI's essence is not confined to LLMs but spans a diverse range of technologies and methodologies that continue to evolve and impact various aspects of human life and work.

Legal firms should embrace AI

Integrating AI into legal practices is vital for several reasons. Most importantly, **clients are actively seeking innovation, efficiency, and potential cost savings in the services they receive.** This demand for advanced technological solutions is evident in the growing number of legal technology and innovation inquiries during the Request for Proposal (RFP) process.

AI can **automate the monotonous, yet necessary, tasks** that often fall to junior attorneys. These are the kinds of tedious tasks that cause many to question their career choice, and detract from the intellectual engagement and fulfillment that the practice of law can provide. By automating such tasks, AI allows lawyers to spend more time providing legal advice and higher-value legal services that better leverage their expertise and training.

Not adopting AI in the short to medium term can have significant repercussions for law firms.

The risks of ignoring AI

Not adopting AI in the short to medium term can have significant repercussions for law firms. Initially, firms miss the crucial opportunity to cultivate an **AI-centric culture** and master **the art of prompt engineering.** As AI systems, particularly those based on large language models, become more prevalent, the ability to effectively interact with and utilize these systems through prompt engineering becomes a valuable

skill set. Furthermore, they will lag in inventing use cases and establishing best practices that drive efficiency and forge a true **competitive advantage**.

In the medium term, firms clinging to traditional methods risk being perceived as outdated and stodgy. This can lead to tangible setbacks, including difficulty attracting top talent who prefer to work for forward-thinking practices, and clients who seek efficient, cutting-edge solutions. In summary, the cost of not integrating AI is not only the loss of immediate benefits, but also the risk to the firm's long-term viability and competitive position.

An effective AI strategy for a law firm or legal department is comprehensive and integrates AI into the workflow.

Crafting an effective AI strategy

An effective AI strategy for a law firm or legal department is comprehensive and integrates AI into the workflow. It includes **foundational training for all relevant staff**, ensuring everyone understands AI basics and applications. It encourages sharing insights and lessons learned across practice areas to foster **a culture of innovation and collaboration**. By experimenting with AI in various contexts, the firm can identify unique use cases and best practices that enhance efficiency and create a true competitive advantage.

Risk mitigation is another critical component. Working alongside the General Counsel and Professional Insurer ensures

that the adoption and use of AI align with **legal and ethical standards**, safeguarding against potential liabilities. This approach emphasizes the responsible and ethical use of AI technologies, integrating them in a way that respects client confidentiality and trust.

A good strategy also involves setting up mechanisms to capture and analyze **metrics** related to time savings, efficiency gains, and other relevant performance indicators. This data-driven approach enables continuous improvement and helps demonstrate the tangible benefits of AI integration to stakeholders.

Conversely, **a poor AI strategy might focus solely on technological adoption without considering the broader implications**, neglecting the importance of comprehensive training, cross-practice area collaboration, and ethical considerations. This can lead to inefficiencies, increased risks, and missed opportunities for genuine innovation and **competitive** differentiation.

Involve the right people

Before implementing AI in the legal sphere, several foundational elements should be in place to ensure success. Support from the **Managing Partner and C-level executives** is crucial for securing investment and signaling the firm's commitment to innovation and efficiency, encouraging a culture open to change.

Equally important is gaining the support of Partners. Their endorsement is vital for firm-wide adoption, as they can champion the initiative, encourage participation, and lead by example in integrating AI into their practice areas.

Another key element is establishing a cross-functional team, including paralegals, legal librarians, training and development professionals, IT resources, project management, change management experts, and General Counsel. Such a team brings together a **wide range of perspectives and expertise**, critical for addressing the multifaceted challenges of AI implementation. This **collaborative approach** ensures that the technology meets the firm's needs, from improving operational efficiency to enhancing service delivery.

Common roadblocks

Implementing AI in law firms faces several challenges: attorneys, already swamped with client matters, find little **time** to engage with new technologies, potentially slowing AI adoption. Unrealistic **expectations** about AI's capabilities can lead to disappointment, while an initial failure might cause some to prematurely abandon the technology, a phenomenon known as "throwing the baby out with the bathwater." A significant hurdle

is the lack of understanding about AI's true capabilities, risks, and limitations. Furthermore, mastering the art of **prompt** engineering for effective AI utilization poses a steep learning curve for legal professionals. Additionally, keeping attorneys **motivated** requires sharing success stories to demonstrate AI's tangible benefits.

Overcoming these obstacles necessitates a strategic approach, emphasizing education, realistic goal-setting, and cultivating a culture of patience and persistence towards leveraging AI in legal practice.



Hanna Canning

DLA Piper



KM + AI = Love

If ever there was a match made in heaven, it is knowledge management (KM) and artificial intelligence (AI). Two parts that complement and reinforce each other; can the future outcome of this heavenly match be the digital transformation we have all dreamed of?

As we explore the practical advances of AI, especially but not limited to Generative AI, it becomes clear that KM is one of the key strategic components for a successful AI strategy.

Know thyself and your data

An AI strategy is not separate, but **part of an overall business strategy**. Having worked with KM and digital solutions for many years, I know that leveraging (any) technology is a **maturity journey**, and we all need to be honest about where we are on that journey, and it will be different for every law firm.

We all know the general building blocks of a good technology strategy, and of course they all apply to an AI strategy. Know your **business needs** and approach them systematically, taking into account your resources, digital maturity and clear objectives. For AI in particular, this means having your house in order: a solid IT infrastructure with information governance, information security and, of course, content. I'd like to mention a clean data set as well, but only few are in that position. However, acknowledge the shape of your data and understand how this will impact the use of AI applications and solutions. Tackling the challenge of improving unstructured data can be costly and complex. Therefore, an AI strategy must be aligned with other complementary strategies, such as a data strategy or information governance.

When it comes to defining your strategic goals, you need to have a good understanding of where you are in terms of **digital maturity** and why you are on this journey; are you pursuing incremental change for efficiency or are you aiming for large-scale transformation? Can you achieve your goals by buying available solutions? If the goal is to develop custom

AI solutions based on internal knowledge bases, then robust data governance, scalability, and highly skilled resources are important factors, as AI will quickly expose weaknesses in your knowledge (data) assets.

Even if you buy off-the-shelf AI solutions, your data will quickly speak for itself and reveal its weak spots. We've already seen this when experimenting with AI solutions on both curated and non-curated data. While it's possible to use AI for the heavy lifting of restructuring your data, you still have to do the initial hard work of understanding your data, and that's a complex challenge.

Training generates trust

In addition, when it comes to AI in the legal sector, the very human aspect of trust should also be addressed. Everything from the basic technology perspective of trust from an information security and regulatory perspective, to building trust with the lawyers who are using AI, and knowing when to trust the output that AI produces. Too often, vendors oversell what a particular tool can do, or on the contrary, lawyers are too skeptical or risk-averse to actually learn how to use a tool. The few stories of legal misuse of generative AI and hallucinations have damaged trust among pragmatic lawyers.

As we see more and more specialized AI tools being used to extract high-quality augmented output from internal data, we also need to address the fact that this puts a lot of trust and responsibility on the user to be able to validate the output. Trust is an important aspect of how to leverage

technology as versatile and groundbreaking as AI, especially Generative AI, effectively and safely.

Driving innovation and change management is built on the foundation of trust,

and trust in AI is based on the ability to check, verify, and understand the output. It is not enough to simply provide the tool, it must be accompanied by sufficient training and upskilling to safely adapt and manage the capabilities as we begin to use AI.

The foundation is everything

This is why KM is an important piece of the AI puzzle in the legal industry. KM is the strategic decision to capture and cultivate a firm's organizational knowledge.

KM is the backbone for legal tech and, of course, AI solutions as it provides a curated data set with contextual relevance.

A good KM system has everything AI wants and needs to thrive. You can have a KM strategy without AI but you can't have an AI strategy without KM.

As we all know, AI is not a silver bullet or a magic solution that will solve all your pains. **Digital transformation is about so much more than just buying a new tool.** Deploying any digital tool, no matter how fantastic, without context or a clear problem to solve is pointless. Context and curated knowledge are key differentiators in generating high-quality augmented output. There is still a lot of work to be done before we can reap the full benefits and

capabilities of AI. Again, KM is a great place to start while we learn more about where to focus AI capabilities in the future.

In conclusion, successful AI integration requires interconnected systems where knowledge (data) is curated and accessible (of course) with appropriate security and privacy controls. This is an extension of, and the next logical step, of KM.

Innovation and transformation do not happen by accident.

Exploring a technology as versatile as AI requires a realistic strategy with clear remit and goals, information governance and trust in content, solution, and output.

Hence a good AI strategy is a piece of a larger puzzle, and for me, an absolutely fundamental piece and a great place to start is knowledge management. A **KM system has rich metadata and contextual content that is reliable because it is regularly reviewed and maintained. It is the perfect place to explore AI capabilities in your organization.** This can be the first iterative step to identify opportunities to scale, and learn more about where to focus your efforts on your digital transformation journey.

David Boulds & Hayley Harris

Head of IT Solutions and
Global Director of Central Knowledge
Management at BCLP



What a Law Firm AI Strategy Can/Should Be

As the interest and excitement around the possibilities of Artificial Intelligence (AI) continue to take centre stage for law firm technologists and leadership teams, how does a global law firm define an AI strategy that manages the clamor of AI pioneers; keeps pace with the rate of change; and separates the hype from the realistic applications of AI for legal practice?

For a number of years, AI opportunities have been the preserve of data scientists and highly skilled individuals and teams. These tools have been used to analyze large amounts of data with AI models to extract meaning and analyze content. With the release of online large language models (LLMs) and chat interfaces, such as OpenAI's GPT-4/ChatGPT, the use of these AI models has become much more widespread. The legal sector, like many others, is grappling with how to best use AI to solve legal practice problems while managing cost, risk, and expectations.

There is **no one-size-fits-all strategy** that law firms can adopt in this area, but some general principles that may serve as an appropriate framework are:

- **Create an AI oversight group.**

Define responsibilities within your firm – who is responsible for identifying opportunities, engaging with vendors and third parties, defining governance, assessing risk, reviewing technology solutions, and implementing change. The oversight group will need to have representation from many areas of the business.

- **Be clear about what you want to achieve and define what success looks like.**

Start small and define your scope. Despite the ambitious claims of some in the AI community, it is likely (at least in the short to medium term) that AI solutions will be best used to solve a single problem type, and this may not necessarily be within the legal advisory space. Where does your firm see the best use cases across the whole business, and where can there be return on investment?

Start small and define your scope. Where does your firm see the best use cases for the business, and where can there be return on investment?

- **Consider integrations – data is still key.**

Most AI solutions rely on access to data. If you are looking to improve searching, classification or summarization of documents, or provide assistance with legal drafting, AI tools will require access to your firm's documents and data. Be clear about how these integrations will work in practice, and about the security and governance third parties that you will need to have in place. Wide access may be the easiest solution in some cases, but that can result in the most risks, not just from a security perspective but also because uncurated data often undermines the value of the output you can expect from an AI solution.

- **Adoption of AI tools must be simple and fit within current working practices.**

AI tools should work harmoniously with the productivity tools lawyers use today. For the highest levels of adoption and the greatest chance of genuine ROI from any AI product, it should complement how lawyers currently work.

- **Be ready to monitor the market and invest where you can.**

With the fast-moving pace of change in the AI market, law firms need to actively monitor the market, be willing to engage with vendors, and look at collaborative approaches. AI vendors may have the technology, but they need the skills and the data sets that law firms have access to in order to improve and refine their models.

- **Be clear on the risks — define what is acceptable to you and your clients.**

There are many risks, particularly with LLMs, that have been well publicized, from hallucinations to copyright infringement, from data security to inherent bias. Document what these risks are to your firm, identify mitigations, and where you are comfortable to accept certain risks. When making your decisions, consider external parties, such as regulators, insurers, and, of course, your clients. In short, the media is full of news content telling us of the risks and opportunities presented by AI.

Having a clear approach to how you review, categorize and implement AI solutions will provide a life raft to help you navigate the tumultuous ocean of Artificial Intelligence.



Tanja Podinic

Global Legal
GenAI Lead at PwC



Legal AI Strategy

In my view, a standalone AI Strategy is akin to organisations having an 'Electricity Strategy' or a 'Water and Gas Strategy' – these do not exist in a business. **We too often look at AI in isolation**, treating it as though it's something to use or not. It should be considered like a utility* rather than a set of standalone products.

If you **consider AI as a utility**, you start to see it as an enabler that can power various parts of your organisation in several ways and integrate into your overall business strategy. If applied effectively, it can enhance productivity, reduce cost, and extract deep insights from your data and organisation in a way that was previously impossible.

To do this successfully, you need to understand what AI is and what it is not, recognising its strengths and limitations. Just as it may not be necessary or practical to power a remote cabin in the forest with electricity, not all problems need AI. Sometimes all you need is a workflow adjustment (or in the case of our remote cabin, a wood-burning stove, and solar panels). However, it is critical to understand the core problem and explore all potential

solutions – because if you're using wood-burning stoves in industrial kitchens, you're missing a trick.

This is hard. It requires expertise in the area, continuous learning and adaptation, collaboration across interdisciplinary teams, the freedom to be creative, and what is most difficult in legal – the ability to fail without repercussion.

There are **two common misconceptions when it comes to AI**; that it demands substantial financial investment to be effective, and that you need a complete overhaul of your existing practices. This is not the case. You can start small and scale. Make sure you have a tech stack that seamlessly integrates for maximum efficiency and compatibility (you don't want to install UK electrical plugs on one side of your building and EU plugs on the other), and that you don't innovate in silo within your legal department – remember you're there to serve your clients, so get their input.

There is no shortage of reasons to avoid AI within an organisation; lack of trust and hallucinations, security and privacy concerns, employee displacement, cost of implementation, the disruption of change management in an already fractured organisation...and the list goes on. We

sometimes forget that **risks** accompany all significant business decisions: should we pivot operations, proceed with the merger, or opt for an acquisition? Assess the use of AI in the same way – determine the risk, mitigate the risk, and exploit the opportunity.

Leveraging tools that enable you to redline contracts, easily analyse clauses, or quickly access your knowledge repository is becoming standard practice. Students are increasingly relying on LLMs to aid in essay writing and study for exams. These students are unlikely to join organisations that rely on manual contract analysis, poring over 50-page documents to find out what happens in the event of a breach.

If you're not exploring the landscape of AI and what it can do for your organisation, you're falling behind. It's as simple as that. You don't need to disrupt every facet of your organisation (leave that to the giants in the industry), but you do need to determine where you can leverage AI to derive more

value and integrate it into your overall strategy. The C-Suite will demand more, your clients will expect more, and your employees won't stay around if you're still using dial-up internet in a broadband world.

AI isn't going anywhere – you either embrace it...or end up like Nokia who "... didn't do anything wrong, but somehow, [they] lost'.

**shamelessly borrowed from Sarvarth Misra, CEO of ContractPodAI.*

Damien Riehl

VP at vLex



LLMs in Legal: Pitfalls, Oilfields, and Best Practices

In November 2022, OpenAI took a shot heard 'round the world: LLMs overtook the legal zeitgeist. Accompanying that arrival, millions of lawyers screamed out in unison: "What does this mean for us?" The good news: This means good things for our industry. If we do it right.

LLMs are improving legal work, offering advanced capabilities that align with lawyers' core competencies: reading, analyzing, and writing. LLMs can do all three with superhuman speed and post-graduate proficiency. But with great power comes great responsibility: We must ride the rocket ship without steering it into embarrassment, sanctions, or worse.

We're in the opening minutes of a very long game. But even at this early stage, certain pitfalls, risk-mitigation methods, and best practices have emerged. Beware hallucinations (confabulations)! Reduce potential inaccuracy through "trust but verify"! To realize LLMs' potential to transform the legal industry, we must harness the models' superhuman speed and abilities, providing what we humans do best: judgment.

Characteristics of the Best LLM-Backed Tools

In this second year of LLM madness in the legal sphere, as the newest generations of foundational models (e.g., GPT, Claude, Gemini, LLaMA) leapfrog each others' benchmark scores, the smartest minds in Legal Tech are coalescing around these themes:

Let's all repeat this mantra together: "To avoid hallucinations, 'trust but verify'!"

Trust but verify. Let's all repeat this mantra together: "To avoid hallucinations, 'trust but verify'!" We've shamed myriad lawyers for failing to validate LLM-generated sources. But the best LLM-backed tools can help users mitigate such risks. Mature LLM-backed systems should make it simple for users to validate LLM outputs against established legal sources. Companies building these tools should leverage LLMs' benefits, while also simplifying users' verification process.

User control. Successful LLM systems should also encourage user-enabled

experiences, giving users more control over research resources and offering contextual summarization. Users shouldn't be limited to "take it or leave it" outputs.

Model agnosticism. The best legal LLM-backed tools should be model-agnostic, using the right model for the task. Various LLMs excel at different tasks, so tools that can leverage the strengths of multiple models provide superior performance and adaptability. As the competitive landscape of foundational models rapidly evolves, this agile approach ensures legal tools consistently deliver the most effective capabilities.

Lawyers aren't the best prompt engineers.

When Sam Altman was asked about the emergence of Prompt Engineering, [Altman predicted that Prompt Engineering won't be a thing](#) — because when software builders do their jobs, users can simply ask a natural-language question and the system will always give you the right answer. No prompting necessary. The smartest Legal Tech companies should follow suit. If Legal Tech is doing its job, lawyers and allied professionals won't need to be prompt engineers. Simply uploading a document will result in helpful insights - without prompting. Of course, if lawyers **want** to prompt engineer — much like some lawyers **love** crafting Boolean queries — they should **also** be able to prompt. But prompt engineering should be an option, not a mandate.

Pitfalls to Avoid

Legal LLM-backed tools should avoid these problems:

1. Providing no insight into the output generation process
2. Obfuscating sources and hindering verification
3. Requiring users to (1) discover and (2) craft obscure incantations (prompts)

Best Practices

In contrast, the best LLM-backed legal tools should have:

1. Easy source verification with frictionless "trust but verify"
2. Individualized source summaries and confidence scores
3. User control over the process
4. Direct question-answering, not just chatbots
5. Model agnosticism, using the best LLM for each task
6. No prompting necessary: Tools provide the insights. Legal professionals bring their judgment and subject-matter expertise.

Legal Data is Oil; LLMs are Refineries

Law firms and law departments have vast Private Oil reserves. And the best legal-tech companies have Public Oil fields encompassing millions of legal documents: regulations, statutes, judicial opinions from dozens of countries. But crude does require refineries.

Refining the Oil. After extracting these legal documents, whether Private Oil or Public Oil, all oil should be refined (categorized, tagged, summarized) for better LLM processing. LLMs are able to do that refining — summarizing and tagging — at superhuman speed and quality.

Combining Oil. The smartest organizations are combining their Private Oil (e.g., contracts, settlements, know-how) with Public Oil (e.g., regulations, statutes, judicial opinions, motions, briefs, pleadings, public agreements). In our interconnected world, lawyers must be able to leverage global oilfields, providing legal insights that span their clients' multi-jurisdictional needs.

Supercharging Oil. Silicon Valley's most advanced companies are combining Symbolic AI (e.g., tagging, knowledge graphs) with Connectionist AI (e.g., LLMs, other neural nets), since Symbolic AI can provide benefits that LLMs alone cannot confer:

1. **Interoperability** (e.g., pushing and pulling data from System 1 to System 2 through System 10)
2. **Analytics** (e.g., "What percentage of our matters are ___?")
3. **Higher precision and recall** (e.g., tagging to avoid LLMs' false positives and false negatives)

Each of the tasks above is impossible for LLMs alone, but with LLM-tagged data sources, all of the above can provide rich insights.

SALI. Where can you find good Symbolic AI tags — that's 17,000+ well-curated tags that are free and open source, constituting the world's most comprehensive legal data standard? Perhaps check out SALI, the nonprofit standard. Did I mention that it's free? And open source? And really good.

Check out the [XML file on the GitHub page](#), and browse [WebProtege](#):

- **USERNAME:** saliuser
- **PASSWORD:** salilmss

Conclusion

Modern LLM-backed legal tools are transforming how legal professionals interact with data — the same data that might have been historically (1) neglected in firm repositories and (2) underappreciated in legal research databases. But armed with the knowledge that data is oil, the LLM-backed refineries are going full speed. **By enabling easy "trust but verify," providing wide jurisdictional reach, utilizing symbolic AI, and employing LLM ensembles, legal tools can offer insights that are both accurate and reliable.** And the best of those tools can provide those insights programmatically — no prompting needed. Rich datasets lead to rich tools. And our profession's most advanced tools are ushering in a new era of legal technology intelligence. We're in a golden age.

Rachel Reid

Partner, US Head of Artificial
Intelligence



Over the last eighteen months, the public release of generative Artificial Intelligence (AI) consumer software applications has shaped the zeitgeist, unleashing our collective imaginations and creating new realities that are demanding the urgent attention of businesses and government bodies around the world.

For businesses that ignore or are lax about this awe-inspiring technology -- referred to as "one of the most powerful technologies of our time" by the US Presidential administration -- their inaction could be judged as ranging from irresponsible to outright illegal, depending on any number of risk-based considerations and circumstances.

On the upside, for businesses that embrace these technologies with safety, fairness, accountability and robust governance, the potential for **innovation, productivity and other unforeseen benefits** can be enormous.

Compared to traditional software, generative AI introduces new risks that can stem from various sources.

These risks include harmful bias and other data quality issues that may affect AI system training and trustworthiness. Additionally, datasets used to train AI systems may become detached from their original and intended context or become outdated, leading to risks such as inaccuracy and unreliability. Furthermore, generative AI systems have unique privacy and security risks and also require more frequent system maintenance.

Amid these complex risks and the substantial harms that can follow, we've seen an increasing focus on AI systems and their use by regulators and new classes of litigants, even in the absence of AI-specific legislation. The scrutiny is largely focused on whether the companies developing and deploying AI technology have appropriate AI governance frameworks -- ones that allow them to adequately explain the critical components of an AI system, such as how it was trained, its inputs and outputs, as well as its risks and intended benefits.

Ultimately, the precedents set by the more mature disciplines of privacy and cybersecurity will likely guide regulators and private litigants when they seek evidence of accountability for AI systems. That is,

which individuals or teams, in any particular organization, 'own' AI, and who is charged with ensuring that AI systems are operating safely and responsibly, not just efficiently?

Given the seemingly endless potential use cases for AI within an organization and the enormity of its potential impact, **legal and risk management departments should begin to architect and formalize an approach to self-governance in the realm of AI.**

The objective should be to adopt a comprehensive AI program that is customized for an individual company and its unique risk profile.

In an article that my colleagues and I co-authored for **Global Investigations Review**, we refer to this paradigm as akin to a Mayflower Compact for AI.

Self-governance plays a crucial role in risk mitigation by imposing guardrails such as continuous monitoring, system management and mitigating controls across the entire AI lifecycle. While ethical declarations and value-based principles are useful initial steps, establishing accountability mechanisms and organizing decision-making structures are pivotal for fostering **a genuine culture of compliance and effective risk management around AI** usage.

Indeed, when considered through a risk- and equity-based lens, AI's goal cannot be solely to maximize positive business outcomes. Rather, regulators will expect AI to be employed as part of a highly rigorous

system of business rules that incorporate accuracy, fairness, reliability, predictability, explainability, security and resiliency. The proverbial men or women "behind the curtain" will need to be actual humans who can articulate and defend their clear guidelines and guardrails on each of these points.

In conclusion, companies can mitigate many of the legal risks associated with AI by implementing and adhering to a robust self-governance program that is based on strong principles and is implemented through a combination of effective governance, assessments, monitoring, privacy and data security controls, and third-party contracts that appropriately allocate risk and responsibility. This can be done in tandem with an organization's commitment to innovation and productivity for the benefit of its myriad stakeholders by pursuing new opportunities to integrate this powerful technology ethically, responsibly and gainfully.

Taking into consideration how quickly generative AI is evolving, our view is that companies should resist any pressures to cut corners amid the desire to realize fast gains or to wait on the sidelines for bright-line legal clarity and instead implement a reasonable, thoughtful and responsible AI self-governance program.

***See article titled, "A Mayflower Compact for AI: implementing responsible self-governance for US companies," published by Global Investigations Review, co-authored by Michael Bahar, Rachel M Reid, Mary Jane Wilson-Bilik and Ronald Zdrojeski, August 25, 2023*

Joshua Fireman

President at Fireman & Company



Strategic adoption of AI in legal practice

AI is often misunderstood in the legal field. Many people think AI only covers generative AI (GenAI) and large language models (LLMs) like ChatGPT. However, AI includes much more than that. It has existed in various forms long before these technologies and continues to evolve.

The key to leveraging AI in the legal field is matching the right type of AI to the right task.

Overcoming the hype

It's easy to get caught up in the current excitement about AI because of its popularity. But just as you wouldn't dive into a pool without knowing how deep it is, law firms shouldn't rush into using AI without a clear strategic plan. **Establish your goals first and then find the right AI tools that can help you achieve those goals.**

To begin, a law firm must define how it wants to **transform** as a legal services provider and partner with its clients. Establishing a clear path with concrete steps to achieve this transformation is essential. This involves identifying business objectives such as lateral integration, cross-selling between practice groups, or diversity-focused resource allocation. Once these issues are identified, you can start exploring the tools and processes available, including AI applications that can help.

Additionally, it is crucial to focus on **foundational elements** like data quality, document management, and governance. For example, having a robust document management system is critical. Many law firms struggle with trust in their document management systems, which hinders the effectiveness of AI tools. Put the AI in the shoes of your lawyers: If a lawyer doesn't trust the integrity of a matter file in the DMS, is this content you want training artificial intelligence?

To effectively adopt AI, you should establish a centralized data management platform, such as an Azure-based Databricks data lake house, where you can govern your

legal content and then feed it out to various applications, including AI. This foundational work is essential for overcoming the limitations imposed by siloed data and systems.

By following these steps and focusing on foundational elements, law firms can strategically adopt AI to enhance their practice and better serve their clients.

Practical roadmap

Despite the challenges, there is a practical roadmap for adopting AI in law firms.

The key is to start with smaller projects that can deliver measurable results within six to nine months.

This incremental approach allows firms to refine their strategies and build trust in AI solutions.

- 1. Identify opportunities:** Look for areas where AI can provide immediate benefits, such as automating repetitive tasks or improving data analysis. These are usually tasks that are time-consuming and prone to human error.
- 2. Quantify benefits:** Understand the potential return on investment by estimating the time and cost savings AI can deliver. This will help you make a compelling business case for AI adoption.
- 3. Develop a business plan:** Create a clear plan that outlines how AI will be used, the expected outcomes, and the resources needed. This plan should be aligned with your firm's overall business strategy.

4. Focus on change management:

Introducing AI will change how people work. Ensure there is a plan to manage this change effectively, including training and support for staff. This will help ease the transition and promote acceptance of the new technology.

5. Step-by-step implementation: Start with pilot projects that can show quick wins. Use these successes to build momentum and expand AI adoption across the firm. This approach allows for adjustments and improvements along the way.

6. Measure and adjust: Continuously monitor the impact of AI and be ready to make adjustments. This iterative approach will help refine strategies and improve results over time.

The future of legal practice

Looking forward, the adoption of AI in the legal field holds the promise of streamlining operations, enhancing efficiency, and improving service quality. However, this journey is not about blindly embracing every new technology. It is about making informed, strategic choices that align with long-term goals. **By adopting this strategic approach, firms can navigate the complexities of AI integration while maximizing its benefits.** This will position law firms at the forefront of innovation, ready to meet the evolving needs of their clients and the legal industry.



Dunya

Unlock the knowledge of all the
lawyers in your organisation

Get in touch: dunya@henchman.io
www.henchman.io

 henchman