



Morae

moraeglobal.com

WHITE PAPER



David Malkinson

More Signal, Less Noise: Refined Knowledge is Key to Generative AI Success



Generative AI

As the breathless hype around Generative AI turns (thankfully) into a more grounded realism, the focus has now shifted to managing the technology as a human aid and working within the limits of AI platforms.

Few doubt that GenAI is here to stay, and I believe the clearest long term use case is the ability to deliver knowledge at a speed and ease that hasn’t been possible before. Today’s Large Language Models (LLMs) are not even close to replacing human judgement, nor applying legal rules to even vaguely nuanced situations, but LLMs will enable lawyers to leave behind some routine tasks and take on higher value, more complex ones. Searching, filtering, and delivering knowledge onto the page is one such task that will be dramatically improved.

The legal profession relies on judgement, precision, and factual accuracy – not exactly what GenAI is best known for. Yes, you can improve accuracy with long form prompts, but making the AI return results only from trusted, cited sources is emerging as this year’s key theme.



OpenAI’s Andrej Karpathy put it this way in his recent [State of GPT](#) talk:

“LLMs don’t want to succeed. They want to imitate training sets...”

Legal work product must be trustworthy and unambiguous. Automated output must be based on trusted checklists, firm best practice, and maintained, curated knowledge content.

Knowledge Curation – Generative AI’s Foundations

Firms that don’t have centralised, curated internal knowledge and experience data could find themselves at a significant disadvantage in the near future. Specific, tailored, proprietary datasets significantly enhance the performance of LLMs - Bloomberg realised this very quickly and developed BloombergGPT trained on their own financial data.

Such is the clamour for attention and relevancy that some recent announcements have made a big deal of spinning up what seems to be little more than a branded instance of OpenAI in Azure. But if knowledge is your primary currency, then it’s a huge, missed opportunity to embark on an AI project unless you’re putting your own institutional knowledge front and centre.

Determining where this knowledge resides is the first challenge. Pointing AIs at entire document repositories isn’t the answer – this is not only prohibitively expensive, but the tools are not that great (yet) at honouring ethical walls. Most, if not all, products currently on the market limit the upload of documents, and don’t respect ethical walls. While these repositories can be useful to support research, their inherent limitations fall short of the knowledge use case presented by AI to facilitate rapid, ethical access to the right information in accordance with company policies and best practices.

A lot of firms have already done the hard graft. The front runners have curated processes to manage knowhow, best practice, marked up closing sets, and FAQs explicitly identified and wrapped in valuable metadata. These datasets are the best possible source to ground AI models to ensure accurate, reliable results.

Many realise the value of knowledge and data curation initiatives but haven’t executed on it. They still rely on recycling a good-ish example rather than relying on a known template. Not only is this practice inherently risky, but it also complicates the implementation of a need-to-know security model.

Legal departments have it particularly bad, often surviving on a thin gruel of random, unindexed G: drives, and inbox searches. Leavers have years of knowledge obliterated as soon as they’re out the door. New starters are presented with a blank OneDrive as their DMS, then go right ahead and engage counsel to solve the same problem that’s been paid for 20 times already.

Nicola Shaver, CEO of the [LegalTech Hub](#) has been tracking the latest developments in the LegalAI space and notices a distinct shift in focus since the initial hype at the start of the year:

“After ChatGPT and then GPT 4.0 launched, the initial hype was extreme, and we saw firms leap into multi-million dollar contracts very quickly. There now seems to be more of an understanding about the limitations of the technology and the need to put in place a structured data model to get the most out of deploying it. More firms are recognizing the benefits of fine-tuning a large language model with their own internal data and numerous firms are building accordingly. Of course, that only really makes sense for firms that have already done the work to structure and clean their data.”



The Problem with Public LLMs for Legal Work

Whether we choose to acknowledge it or not, chatbots like ChatGPT, Google’s Bard, and Bing Chat have become integral tools in the daily workflows of most professionals. For some time, they have been routinely used for improving emails, proofreading text snippets, summarizing meeting notes, and interpreting language. Even if the firm has turned off access, it’s likely personal devices are still being used to perform these tasks.

It’s well known that AI chatbots are not reliable enough to produce final work product. Data privacy issues bar any client content from being used as prompts. Users must carefully anonymise text, which can be time consuming, and still problematic in many scenarios. With ChatGPT, there’s no way of being sure what the source is for any produced content, resulting in a manual QA every time. By the time you’ve entered the prompt, eyeballed the output, and tweaked the content, it’s usually a slower process than simply picking content from a known template.

Content providers such as LexisNexis and Thomson Reuters have been quick to address this problem. Their trusted content sources and early market access to OpenAI give them a significant advantage. Early adopters of these tools will enjoy productivity gains and be able to use their investment to showcase their AI credentials to clients. However, if/when adoption normalises, there will be little in the way of competitive advantage.

Starting Your AI Initiative in the Right Way

If you’re embarking on a GenAI initiative (who isn’t?), then the first priority is to familiarise staff with the pitfalls of AI, and roll out guardrails in the form of usage policies – with “no client data” front and centre. Depending on your risk appetite, consider standing up a platform that is not subject to the public facing privacy policies. If legal work falls under HIPAA rules, you will need to exercise additional caution.

Second, think critically about the problem to be solved, and make sure that GenAI is the right answer – in many instances AI will slow you down. This is particularly true in some drafting and review tasks that rely on long-form prompts. Often a well-designed workflow or old-school menu-driven input will get the job done quicker.

Next, ensure you have your data in good shape for future use. That means:

- **Centralize the institutional knowledge** – this content rarely resides solely in the ‘official’ repositories: per-departmental folders buried on someone’s OneDrive, shared team matters, checklists in OneNote are common sources embedded into daily workflows. Deal bibles are a crucial source – make sure they are all discoverable in a central location.
- **Enrich with metadata** – Make sure you have a system to apply metadata to content – a DMS is the most obvious answer here. Labelled or categorised data is essential for training a generative AI model. The quality and quantity of metadata applied to each record makes a huge difference to overall usability. Metadata can range from key fields such as document type, knowledge type, area of law, legislation/regulation etc. to complete legal taxonomies such as [SALI](#).

The Growing LegalTech Ecosystem

There are a number of avenues to explore to light up this knowledge for use with LLMs. The first approach is to pull data from the source system’s API, copy, translate into vector databases stored elsewhere, and then use this for grounding via R.A.G., or use documents for fine-tuning. Some firms will be hesitant to adopt this approach – they will be apprehensive about moving data en-masse between jurisdictions, they will have concerns about respecting client confidentiality, and they lack the appropriate technical resources to execute. The better option, if possible, is to perform this in-situ without moving data.

Knowledge is most often located in the Document Management System (DMS), eDiscovery platforms, CRM, and time/billing/matter management systems. Harnessing AI tools supplied by those vendors will address the security and data movement challenges, reduce cost to implement, and remove potential points of failure.

Announced in August, iManage’s Insight+ comes with a built-in Knowledge Curation and automatic categorisation engine, semantic search, and a chatbot functionality – massively improving search and research tasks.

Alex Smith, Product Director for Search, Knowledge & AI at iManage says iManage Insight+ has been built specifically to help findability, and is part of a multi-phase rollout for iManageAI in the coming months:

“iManage deliberately has a knowledge, search AND AI roadmap because these things are inter-related – this new generative technology is only an enabler to delivering on what’s important - enabling knowledge workers to get their job done with the right trusted information at the right moment in their workflows.



People are very obsessed about “controlling” this AI with degrees in ‘prompt engineering’, but making these services only generate off the information you feed it (textual or structured) and offering structured levers to hone answers (specific metadata) should be firms’ top priority. Some IA before AI needs to be in place, for some this is cleaning up existing repositories, for others introducing new cultures of sharing with a new additional purpose”

Contract drafting and review tools were among the first tools to integrate with LLMs. Products such as [Henchman](#) and several CLM tools gain access to previously negotiated terms in order to really add value. Having immediate access what’s market for any provision, and what’s been agreed previously can speed up drafting considerably and is only possible by having access to trusted source of finalized contracts.

For eDiscovery, summarization and chronology creation are two use cases where GenAI is already making an immediate impact, with live beta testing happening via multiple vendors. Industry experts anticipate that RelativityOne will soon announce their GenAI feature, which is likely to include a review component. Notably, some recent entrants in the eDiscovery AI review market already offer this. Similar to Predictive Coding and Active Learning advances over the last ten years, document review will be redefined, with resources redirected from eyes-on review to focus on meaningful insights from datasets that were previously cost prohibitive to discover.

Cloud First

As seen with recent advancements in the legal technology market, cloud offerings are always prioritised for new features, with on-premises or hosted versions often trailing in development. iManage’s new AI features will only ever be a cloud-only offering.

Relativity’s cloud platform RelativityOne’s development has a monthly release cadence. Relativity on-prem hosting providers only receive a single release a year - considering the pace of development in the space, they will have to choose between waiting for the tech to be released for the own platforms, or biting the bullet and building something themselves.

If there was ever a persuasive reason for firms to shift to the cloud versions of iManage or RelOne, this would be it.



Final Thoughts

The practice of law is focused as much on the historical narrative of past matters than anything else - a lawyer’s job is close to impossible without having access to trusted knowledge and prior work product.

Generative AI has the capability to provide access to this information in a truly transformative way. If the foundations are built correctly, staff will be able to achieve tasks previously out of reach via natural language queries - reducing time-to-value, slashing development costs, and minimizing end-user training.

Despite many firms’ best efforts, physical proximity can’t be relied upon anymore to disseminate knowledge – this is the biggest challenge of the modern workplace, and Generative AI appears to be the technology with the best chance of solving it.

About the author

David Malkinson
Senior Managing Director
Morae

David leads Morae's digital transformation function which helps clients leverage technology to improve productivity, reduce risk and retain a competitive advantage.

David has been working in tech enablement for professional services firms for more than 20 years. He was instrumental in establishing the global consultancy arm of Phoenix from a small four man consultancy team to the 100+ team it was when they joined Morae in 2019.

David has had a unique insight into the different challenges and complex environments that companies must manage, and has a laser focus on delivering business outcomes. He has worked with hundreds of firms, from small-medium enterprises to large conglomerates and top tier international law firms in the strategic advisory, implementation, and management of document management, information governance and contract enablement technologies.

[Connect with David on LinkedIn](#)



Changing
Legal
Together

Morae Global Corporation is trusted worldwide by leading law firms, legal departments, and compliance functions for the delivery of award-winning solutions for digital transformation and legal transformation.

Abu Dhabi
Bangalore
Chicago

Frankfurt
Hong Kong
Houston

London
Maastricht
New York

Sydney
Washington DC
Zurich

© Copyright 2023 Morae Global Corporation

moraeglobal.com