

# AI SMARTS FOR THE LEGAL SECTOR

*Preparing Your People  
and Content for AI*

---

BY GABRIEL KARAWANI

**2**023. What a year! Artificial Intelligence (AI) has upended industries worldwide, and the legal sector is no exception. Whether you are a believer or a skeptic, whether you are an early adopter or a laggard, it is hard to avoid being sucked into the vortex of questions, expectations, and concerns.

The fact is that a shift has happened. Technology is now allowing us to do or at least imagine how we might be able to do things that we couldn't before.

Whether we are going to go through a few rounds of resetting expectations, we will inevitably see an increased use of AI, both as specific tools and within tools that we use daily. Consequently, this will – sooner or later – shine a light on our enterprise content's readiness for AI. For many organizations, this is creating a here-and-now desire to get ready for AI. The risk of doing nothing and falling behind is too significant to ignore. Equally, the opportunity to get ahead is too attractive to dismiss with conservative skepticism.

In parallel, users need answers, guidance, and training, whether jumping in with both feet or just dipping their toes.

### **Start by Focusing on Your Cases**

Organizations are now faced with many choices. Wait or go early. Build in-house or buy third party. Give access to everyone or only specialists. Expose all or just small sets of content to AI.

To determine the right approach and have a higher chance of success, firms should start with a clear definition of what they want to do with AI and how it will impact their business. It is, after all, technology, and all technology should be purchased and implemented with specific intent.

It is certain that Generative AI will not be suitable for all tasks or use cases and in most cases, firms will not

let it anywhere near anything with a hint of client-facing advisory work.

The best way to start is to identify core use cases that deliver value through improved productivity, growth or supporting new business models. Prioritize those that meet as many of the following criteria as possible:

- There are clear and limited datasets that can be associated with the use case.
- The associated data is low-sensitivity and does not contain PII (personally identifiable information).
- The use case does not involve many different groups of people, particularly complex cross-organizational processes.
- The use case is conceptually simple and can be modelled through manual steps.
- If implemented successfully the use case will have high visibility and impact.

### **People-Centric Actions**

People will be at the heart of successful outcomes. The explosive adoption of Generative AI in the mainstream of technology usage has caused many myths, misunderstandings, and false expectations. It has also given rise to a whole new set of AI experts perpetuating said myths, misunderstandings, and false expectations.

Onboarding AI – particularly Generative AI – as a technical capability within a firm therefore needs a bit more thinking than your average technology roll-out. In some cases, you will need to move some users back to “square one” and reset their thinking, before moving forward as a whole.

Equally, it is vitally important to address the user

community's desire for using AI tools, sooner rather than later, to avoid the proliferation of shadow AI.

We recommend starting with a primary and firm-wide internal communication and education initiative, provided alongside a safe and secure, organizationally controlled AI interface where users can experiment with zero risk.

Educate and allow users to experiment with:

- Basic prompts to understand a GPT LLM's limitations.
- Prompts known to provide incorrect answers.
- Similar prompts to understand the nature of LLM responses being "non-deterministic" (i.e., you might not get the same answer every time).
- Prompts where the LLM will be leveraging an authoritative source of your data as part of generating a response.

As part of the education, provide clear guidance and communication to further people's understanding of Generative AI, such as ChatGPT, and how/why it can sometimes be prone to answer incorrectly (with great confidence as well!) and how users can avoid exposure to risks caused by, among other things:

- lack of transparency,
- accuracy and bias of models,
- intellectual property (IP) issues,
- sustainability concerns.

Like end-users being responsible for the content in their emails, firms are likely to put the onus on end-users to check any work drafted by Generative AI. For this and

many other reasons, organizations will need to implement policies and controls to detect the use of data in prompts, biased/inaccurate outputs, and so forth.

But rolling out AI is also about getting the basics right. It can be met with or create profound cultural resistance – sometimes for good reasons. Unsurprisingly, given the amount of misinformation about AI, it can foster anxiety, nervousness and fear.

Effective change management and communication strategies are necessary to ensure smooth adoption and overcome user resistance or fear related to implementing any AI use case to end-users.

Depending on the use cases chosen as high priorities, the change programs can – and should – be highly adaptable, and you can learn from each of these before moving on to the next one.

But when it comes down to real business and for AI to be effective, knowledge management teams – and subject matter experts such as PSLs – are the perfect home for honing AI-related skills and ensuring the truth and validity of the firm's data is both protected and leveraged in what we refer to as the AI Sweet Spot (more about this below). The KM function has, in effect, never been more important or relevant than it is now.

We recommend creating a cross-functional and cross-practice Center of Excellence (CoE) within your firm to build and disseminate the expertise and experience gathered around Generative AI. The CoE can ensure the firm maximizes the potential of AI through centralized governance and strategy, helping to bring together the right people, clarifying the objectives, owning the use cases, executing the implementation, and measuring progress.

### Content-Centric Actions

As organizations start to experiment with AI interfaces that work with organizational content, outcomes

typically fall short of the end-user expectations. This is mainly caused by one thing in particular: content is not yet ready.

This lack of data readiness stems from an overall – and often historic - lack of data governance. Data is typically over shared, under tagged, lacking in consistent version control, and abundantly duplicated and outdated, to name some problems.

Another practical issue is that the content is not available in a system or form that the AI can work with. For instance, Microsoft 365 Copilot will require high-quality data to be available in SharePoint Online so it can be indexed in the Microsoft Graph to be made available via the Semantic Index.

For those that have worked on enterprise search solutions, you will notice some similarities in these challenges and indeed, many of the same hygiene factors

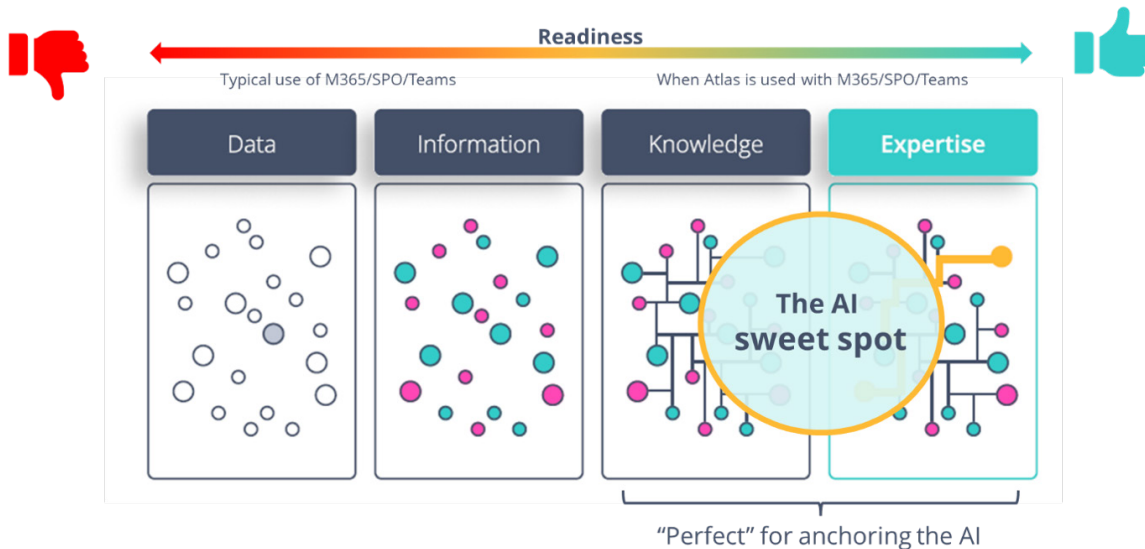
that an efficient and trustworthy search solution requires will need to be in place for AI.

In this context, we talk about an AI Sweet Spot, which is “good content” that can be used for anchoring AI and providing better “grounding” during the prompt process.

At one end of the scale, content is just a binary object (a file) with a little bit of metadata (such as a filename, date and author) which provides little context outside of the actual document content.

At the other end of the scale, a file might be tagged with related sectors, legal subject areas, and legal jurisdiction, turning it into content that can be used as knowledge in the proper context. From that, content can now be connected to create expertise.

This end of the scale, where content exists as knowledge and can be linked together as expertise, is where AI can thrive. This is the AI sweet spot.



The technology architecture behind this dictates that the more that can be inferred from what we know about content, the better we can identify content relevant to the user's prompt. The reason is that the user's prompt is broken down into different "sub prompts", some of which can simply be thought of sub-searches that are happening behind the scenes. This is why the challenges with AI are very similar to those of enterprise search.

But arguably, when – and if – you consider using Generative AI to draft an overview of relevant specialisms for a proposal, article or summary, for instance, the challenges and consequences go deeper. To illustrate this, consider the following prompt (ignoring any system prompt):

"For a tender response to Bank of Laska, provide an overview of our relevant specialisms, each with a header and a 100-word description, based on other tender responses to the same or similar clients."

An experienced human would obviously conduct a few searches and put the answer together with a high degree of confidence.

But what will the Gen AI do? First, we must assume it can access "your data" across your platforms. Second, it must be able to infer similar clients from the content or metadata. Third, it will have to search for the most relevant documents from which it will extract content. And so on.

Just by looking at these three steps, we can quickly identify areas to focus on when getting ready for AI.

1. Consolidate content in a modern content management system, such as Microsoft 365, or connect and streamline content indexing from multiple systems (iManage, Netdocs, etc).
2. Enrich your content with metadata, making it more identifiable and permitting the relationships to be more confidently inferred.
3. Remove redundant and outdated content and \*at minimum\* adopt proper file-level version control rather than named versions.

This solves a big part of the headache, but interestingly, we are still left with some more esoteric challenges. When looking at the content, faced with 100 examples of previous tender responses, which ones come first? (Don't say all of them, as this has massive cost implications.)

This moves us to the next level of content preparation, which might include:

1. Identifying the "right" types and rules for content that is to be made available for Gen AI.
2. Checking for accuracy, completeness, consistency, and reliability.
3. Cleaning and preparing content repositories, accordingly, including checking permissions.

Considering that most organizations will hold hundreds of gigabytes of data per employee, it is not surprising that most CIOs, CISOs and CDOs will look at these challenges as a very big mountain to climb. Possibly too big a mountain.

The alternative is to focus on making the higher graded content, the gold standard content, available to AI. The organizations with the most foresight and ambition will implement a modern knowledge platform to better create, maintain and manage (including disposing of) this higher graded content effectively and dynamically.

Your high-priority use cases will be those that determine if \*all content\* should be in scope or if you can prepare and make smaller sets of content available for specific use cases and tasks.

We believe the latter will be the popular choice for most, and indeed, in Atlas, we developed an innovative concept to support this. Called "knowledge collections," this allows organizations to efficiently manage how their content is made available for specific tasks and use cases.

## How an Intelligent Knowledge Platform Can Accelerate the Journey

To make AI sing to your tune, a lot of groundwork is needed.

A Knowledge Platform with capabilities to prepare for, utilize, or deliver on use cases supported by Generative AI can significantly accelerate the journey. Such a platform, also known as an Intelligent Knowledge Platform, should be seen as an orchestration tool for leveraging AI capabilities while delivering on hygiene requirements and advanced features for driving knowledge-centric productivity, collaboration, and communication. **ILTA**

ATLAS - AN INTELLIGENT KNOWLEDGE PLATFORM - ACCELERATES THE JOURNEY IN SEVERAL WAYS, INCLUDING THOSE IN THE TABLE BELOW:

INTELLIGENT KNOWLEDGE PLATFORM CHARACTERISTICS	HOW IT ACCELERATES THE AI JOURNEY	OTHER BENEFITS
Auto-tagging of content across at least five categories of tags	Content that is richly tagged and described will provide better quality outcomes in Generative AI, allowing for better anchoring and grounding.	Comprehensive tagging is the foundation for successful search scenarios, including greater usability and contextual search results.
Consistent and scalable content governance controls	Consistency is fundamental for confidence in AI generated outcomes. Therefore, efficient management of content across 1000s of Sites, Libraries, Teams, Channels, etc. is a must.	Reduces the risk of data leakage; increases ability for auto-labelling; improves ability to set “just enough” permissions.
Automated application of security or sensitivity labels	Areas of content or individual content items that have labels can be excluded from various scenarios and views – including Microsoft 365 Copilot usage.	Labels can help automatically exclude content, set ethical walls, and also assist in content lifecycle management by allowing policies for retention, archiving or disposal to be applied.
Collaborative knowledge bases for subject matter experts, with ring-fenced management and content grading	Give knowledge managers and subject matter experts full control of the management and grading of their knowledge.	Helps drive a firm wide knowledge agenda by enabling delegated and decentralized knowledge creation and management, underpinned by global governance controls.
Creation of dynamically maintained “knowledge collections”	Enabling authorized knowledge owners to define dynamic collections of knowledge that can be consumed by AI empowers business users to create their own GPT data sets.	Reduces reliance on IT operations and reduces overall cost of creating, running and operating RAG vectorizations at scale.
AI Assistant User Interface	A safe environment, fully controlled and governed by the firm, for users to work with Generative AI while leveraging organizational data.	Visibility and control of usage costs.



**Gabriel Karawani** is a Co-Founder of ClearPeople, the company behind Atlas - The Intelligent Knowledge Platform. Gabriel brings a wealth of practical experience in helping businesses get the most out of Microsoft 365 to enhance their knowledge management and information architecture (IA). With a strong foundation in engineering and a focus on real-world applications of AI, Gabriel's expertise in Information Architecture for AI in the enterprise has made him a valuable asset in the industry. .