



Lexis[®] AI Insider

Moving towards a future-ready enterprise

Practical tips for AI implementation

Introduction

You can't scroll through LinkedIn without reading posts about the '5 things you need to know about generative AI'. Every second podcast in your feed now seems to relate to AI regulation and the risks and benefits of using generative AI. You have even prompted Stable Diffusion to create some outlandish images and asked ChatGPT to write your best man's speech. You read some pieces written by technology thought leaders and it becomes clear to you that generative AI is probably a game-changer of the same magnitude as the internet and mobile computing.¹

There is a lot of talk at work about what AI – both generative and extractive – can do to help accelerate growth, improve your customer or client service, generate content, and make your backend processes more efficient. Before long, you have found yourself tasked with or involved in implementing AI solutions in your business or firm. It seems like a huge task in an uncertain technology landscape with no clear path forward, and you find yourself wishing for some guidance.

The purpose of this paper is to highlight three key principles to bear in mind when implementing AI solutions in a business context:

- » Prioritise discovery of a problem space or business opportunity before considering implementation
- » Test and learn while skilfully managing risks
- » Build alignment with key stakeholders at each stage of the process



Authored by

Claire Linwood

Product Manager
LexisNexis® Pacific

What is generative and extractive AI?

It helps to have a working definition of AI when considering how best to make AI work for your organisation.


“Artificial intelligence’ is not a term of art. It is an umbrella concept that captures the idea of using machines to perform tasks previously considered exclusively the domain of human intelligence: things like reasoning, translating, decision-making, categorising, creating text, and viewing images. Data is key to training machines to perform these tasks.”

It is not necessary to have a deep understanding of data science or the ways in which AI applications work in the background in order to understand their basic functioning. Here, the computer science concept of ‘abstraction’ is helpful: abstraction, also known as ‘information hiding’, is the reduction of complexity via the removal of unnecessary details.² Grasping some basic ideas about AI and data, along with understanding clearly what sort of AI model output you expect for a particular use, will in many cases be sufficient to enable you to make key business decisions about AI implementation.

¹ Bill Gates, ‘The Age of AI has begun’, <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>.

² Cornell University, ‘Abstraction’, http://www.cs.cornell.edu/courses/cs211/2006sp/Lectures/L08-Abstraction/08_abstraction.html.

Perceived impact of generative AI on the legal practice



→ Generative AI and the future of the legal profession

LexisNexis UK recently surveyed over 1,000 UK lawyers and legal professionals on generative AI, with 95% of respondents believing that generative AI will have a noticeable impact on the law.

Generative, as opposed to extractive, AI models are capable of creating new text, code, or images. One new type of generative AI model that has particular importance for the legal industry is the large language model, or 'LLM'. LLMs are AI models that have been trained on a vast, unstructured corpus of text to produce human-like text in response to prompts. They are a type of 'foundation model'.

Foundation models are AI models that can be fine-tuned and adapted to perform a range of different tasks across a variety of domains, as opposed to models that have been pre-trained to perform a specific task.³

Foundation models like LLMs that have been trained on vast quantities of data tend to exhibit unpredictable 'emergent' qualities that are not evident in smaller models and do not increase in a linear fashion compared to the performance of smaller models. These emergent qualities include the ability to perform multi-step reasoning and identify the meaning of a word in context.⁴

Extractive AI models are trained to retrieve the most relevant information from a pre-defined dataset for a specific purpose.⁵ LexisNexis® New Zealand offers multiple products that contain extractive AI features, including Leading Cases, Legislation Compare and the Lexis Answers feature available as a search enhancement in Lexis+®. The search engine powering both the recently launched, latest-generation platform Lexis+ and Lexis Advance® has used extractive AI for over a decade.

Soon to come in 2025 will be Lexis+AI, a generative AI platform that searches, summarises and drafts using authoritative LexisNexis content.

³ IBM, 'What are foundation models?', <https://research.ibm.com/blog/what-are-foundation-models>.

⁴ Google Research, 'Characterizing emergent phenomena in large language models' <https://ai.googleblog.com/2022/11/characterizing-emergent-phenomena-in.html>.

⁵ LexisNexis, 'Combining extractive and generative AI for new possibilities', <https://www.lexisnexis.com/community/insights/legal/b/thought-leadership/posts/combining-extractive-and-generative-ai-for-new-possibilities>.

PRINCIPLE #1

Prioritise the discovery of a problem space or business opportunity

A paradigm-shifting technology like generative AI tends to produce so much hype that it is tempting to apply it first and see what it can do later. In some instances, valuable use cases do emerge over time as real users come into contact with a new technology and experiment to discover what it can do for them.

However, technology in search of a problem to solve is not likely to produce optimal business outcomes in the real world on a consistent basis.

When considering the implementation of AI at the enterprise level, it is important not to reason from the technology towards a use case. Rather, consider the pressing business problems that your enterprise is facing and only then consider what technical solutions may address them. In this way, you can de-risk the AI implementation process by reducing the likelihood that your new solution will not address a real business problem or pain point. This will also assist in Principle #3 – ensuring a smooth change management process by getting key stakeholders on board.

There is a myriad of possible industry-wide applications of AI – and particularly of generative AI. The business value and economic impact of those applications are likely to be substantial. A June 2023 report by McKinsey Digital on the economic potential of generative AI estimated that current generative AI and other technologies could automate work activities that presently absorb 60 to 70 per cent of employees' time.⁶


⁶ McKinsey 'The economic potential of generative AI: The next productivity frontier', <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#business-value>.

The enormous natural language capacity of LLMs makes them particularly well suited to use in legal practice, augmenting the expertise of lawyers and streamlining time-consuming manual tasks. Here at LexisNexis, we anticipate that the most valuable use cases for our customers will, broadly, fall into the categories of search, summarisation and drafting.

Consider using a cross-functional (i.e. not just lawyers in the room!) design thinking process to brainstorm key business problems or opportunities that would be both feasible and valuable to solve. Start broad, taking into account both internal pain points (for example, difficulty accessing firm content and know-how when and where you need it) and external client needs and pain points (for example, the need to receive advice that is as simple as possible to understand and apply). Narrow down to a single problem or opportunity by voting, and ideate about how you might solve that problem.

As part of the process of selecting a viable use case, consider what a successful solution or output would look like. Would you like to reduce the amount of time that lawyers in the firm spend on research by 20%? Or would it be more valuable to the firm to generate more leads by increasing the production of thought leadership material from three articles per month to six articles per month? These are two of the many possible use cases for AI in a law firm.

Generative AI use cases with the most potential



→ Generative AI and the future of the legal profession

Respondents in LexisNexis UK's recent survey on Generative AI considered legal research to be the use case with the most potential.

PRINCIPLE #2

Test and learn while managing risk

You've gathered a cross-functional team to select a use case that is technically feasible and viable for the firm. You have a range of options for implementing it, including building a tool in-house, working with a third-party provider to develop the tool, or using a paid or free AI product. Once your business problem or opportunity is clearly defined, take an iterative approach to developing or selecting your new AI solution.

Much will depend at this stage on the organisation's resources and appetite for investment and risk. Very large organisations with significant in-house data science and software engineering capacity, as well as a large corpus of data, may choose to develop AI models to respond to specific internal needs or as a product offering for external customers. Bloomberg, for example, has developed BloombergGPT – a large language model trained specifically on financial data to perform a range of finance-related natural language processing (NLP) tasks, including sentiment analysis and named entity recognition.⁷

Other organisations may choose to fine-tune an existing LLM or other AI model for specific use cases. In this case, the key to implementation will be ensuring that the data used to train the model is adequate and free from unfair bias. Subject matter experts (SMEs) should be involved from an early stage to define the desired output, label training data (if necessary), and test that the model output conforms to expectations. Expect multiple rounds of testing and iteration in order to enhance model performance before the model is ready for deployment.

Most organisations, however, will be selecting a tool from a third-party provider – perhaps with some capacity for customisation. In this very early stage of technology adoption with significant uncertainty and many new platforms, partnering with a trusted third party is a safe and effective option.

When selecting an AI tool developed by a third party, we recommend following a framework along the lines of Gartner's step-by-step guide to enterprise tech buying:⁸


- Plan, including creating an evaluation team, defining objectives, and identifying a budget.
- Research the vendor landscape and technology options.
- Specify detailed use case requirements and rank by importance.
- Evaluate specific vendors according to your requirements.
- Finalise contract parameters and pricing.

Ensure, if possible, that a cross-section of your stakeholders is present at any product demos in order to mitigate the risk that there is an unseen roadblock to implementing the solution and help to encourage adoption. A trial period including intended end users will also help you to decide whether the solution is right for your selected use case.

Some questions that can assist law firms and legal teams in the process of procuring an AI solution from an external vendor are:

- How accurate is the output? Where is the tool getting its information?
- How private are my interactions? How is my data treated?
- How customised is the tool? Does it understand legal terms? Does it understand different practice areas?
- Do I get citations so that I can check the output?⁹

Biggest barriers to adopting legal technology



→ To escape the legal labyrinth, in-house teams need to be constantly innovating

LexisNexis UK's recent survey of in-house legal professionals revealed that lack of budget was considered a barrier to adopting legal tech. Choosing the right use case for the budget that you do have is essential.

⁷ Bloomberg, 'Introducing BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance', <https://www.bloomberg.com/company/press/bloomberggpt-50-billion-parameter-llm-tuned-finance/>

⁸ Gartner, 'Enterprise tech buying: a step-by-step guide', <https://www.gartner.com/en/topics/successful-tech-buying-process>.

⁹ LexisNexis, 'AI for growth: harnessing the power of large language models in the legal industry', <https://www.lexisnexis.com.au/en/insights-and-analysis/practice-intelligence/2023/ai-for-growth-harnessing-the-power-of-large-language-models-in-the-legal-industry>.

PRINCIPLE #3

Collaborate and bring everyone along for the ride

Even the most valuable and well-conceived applications of AI can fail to gain traction in an organisation if key stakeholders and their teams are not on board.

This is where effective change management and stakeholder engagement can make enterprise AI implementation a success.

What are some of the key things to bear in mind when planning change management for AI implementation?

First, bring stakeholders from different functions and levels of seniority in as early as possible. Ideally, representatives from across impacted functions will be present and involved in selecting the pressing business problem or opportunity that you will use as a test case for AI implementation. This has the double impact of both ensuring that the use case selected will have real business value and be feasible and creating the conditions for shared ownership of the implementation project.

As part of the ideation and planning process, it helps to align on a success metric that you can use to track progress towards the target state. Your success metric will represent the resolution of the identified business problem or opportunity. This metric might be made visible and transparent as the project progresses. Similarly, track adoption and usage once the solution has been rolled out to a larger group – sustained usage of the solution is a great indication that it is meeting a real need. If usage is low, investigate why that might be. Is the initial onboarding too complex? Does the output not meet user expectations? Or are users simply not aware that the solution exists because it has not been publicised or is not discoverable?

Employees will not use AI-enabled tools if they are sceptical of the value of their output or concerned about their roles. For lawyers, trust in AI tool output is essential – as the advice delivered to clients is ultimately a reflection of their brand. Behaviour change and adoption will not occur if trust is not established at the outset. Encouraging upskilling in the use of AI tools may assuage some scepticism about AI tools in general. Prompt engineering for LLMs is a particularly accessible skill for legal professionals, as it leverages verbal reasoning ability.

It is also crucially important to involve functions including Legal, Finance, Compliance, Data Privacy and HR (where relevant) at as early a stage as possible. Besides the fact that internal AI tools may be used by or intended for those teams, they are best placed to advise on any risks involved in implementation.

Conclusion

Now is a great time to prepare yourself and your organisation to embrace emerging technologies, using them to solve internal pain points and provide the best possible service to your clients or customers. The newest generation of AI models represents a technological step change. With their capacity to reason in complex ways and generate human-like language, LLMs have the potential to transform the nature of legal practice and legal operations. Effective implementation of AI solutions is key to obtaining a competitive advantage and freeing up employees' time to focus on those tasks that require creativity, strategic thinking, and relationship building. While caution and careful implementation are necessary, doing nothing is likely not an option that will bear fruit in the long term.