

GENERATIVE AI: PRACTICAL SUGGESTIONS FOR LEGAL TEAMS

There has been much discussion in the legal industry and more broadly of the risks associated with generative AI (with ChatGPT in particular hitting the headlines), but much less guidance on how legal teams might start to use it in the near and longer term. This short guide will look at some of the possibilities and set out some practical ideas that you can try today.

We will be focusing on 'LLMs' as a form of generative AI (language models trained on large volumes of text and capable of generating new text in response to prompts) such as OpenAI's ChatGPT, Microsoft's Bing, Google's Bard and Anthropic's Claude. This guide is divided into three parts:



Part 1: How can you use LLMs today?



Part 2: What does this mean for organisations?



Part 3: What's on the horizon?

Before proceeding, there are some caveats. First, never share any confidential or sensitive information with an LLM unless you are confident about how it will be used, who can access it and that these circumstances are within the ambit of relevant information security and confidentiality obligations. As it stands, many businesses are putting restrictions on usage of LLMs by their employees. Second, the scope of regulation remains unclear, and you should approach these tools with caution: our [Regulating AI](#) series explores the range of legal issues arising from the use of algorithms and AI, including across IP, data protection and employment. This is a fast-changing technology and we will be considering the practical and legal issues more as they develop. With that said, there is perhaps no better time to start becoming familiar with generative AI while it is still an emerging technology.

SUMMARY

- LLMs, in their current state of development, can be useful for a number of common tasks, like planning work, refining text, summarising ideas and helping learn new concepts. They are best used as collaborators (as their output will need to be checked) and you should spend some time exploring different 'prompt' techniques to get the most out of them and to understand their capabilities.
- Generative AI will increasingly be built into enterprise tools (like [Microsoft Copilot](#)) and organisations will need to spend time assessing the suitability of these tools when they are made available. For more complex use cases, you should understand the wider processes and data sets involved and consider secure access to LLMs to mitigate privacy issues.
- There's no sign of the rapid pace of change slowing down. Look out for increasing regulation, security fears and autonomous AI agents in the coming months.



PART 1: HOW CAN YOU USE LLMs TODAY?



What tasks could I use it for now?

To understand what LLMs could be used for, it is helpful to understand their limitations. Notoriously, LLMs have a tendency to make up facts (to “hallucinate”) or miss key bits of information. Experts have concerns around bias and IP infringement and some LLMs, like ChatGPT, have a knowledge ‘cut-off’ meaning that they do not have access to recent factual information (in ChatGPT’s case, nothing more recent than September 2021). In general, it is therefore best to use an LLM as an assistant or collaborator, something that produces work for you to review and develop to create the best possible output.

Teach - learn the basics about a specific topic or concept:

- “Pretend I have no financial background and explain the concept of net present value”
- “Give examples of how this legal issue would apply to a business in the consumer goods sector”
- “Give examples of how this clause might be considered ambiguous”


Transform - reframe and repurpose:

- “Re-write this summary to remove legalese”
- “Extract and summarise the key obligations from this piece of legislation”
- “Convert this internal briefing into a 5 minute presentation and speaker notes”
- “Write an excel formula to remove blank spaces and join this text together”

Summarise - condense key information into a concise format:

- “Create a table of action points from this webinar transcript”
- “Write a single sentence summary of this article”
- “Extract the information in this article relating to IP law and convert it into flashcards”
- “Based on this email, suggest a single sentence summary of each action”

LLMs like Bing and Bard also have internet search capabilities, so can start to be used for ‘retrieval’ type tasks (e.g. “find me all the information on my panel firm’s website on [x] legal issue”). The advantage of using an LLM for this type of task (over just a standard search) is the LLM’s ability to access multiple search results and summarise the outputs. This can be a quick way of starting some research but requires particular care when using (for example, to ensure you have captured all relevant sources).





PART 1: HOW CAN YOU USE LLMs TODAY?



How to write a great prompt

A key aspect of using LLMs is what to say to get a desired output, an emerging skill which is sometimes referred to as “prompt engineering”. In many ways, engaging with an LLM is more like dealing with a person, as compared to a traditional machine which will follow logical rules in a clear and coherent fashion.

Give it a persona - influence responses with character traits:

- “You are a detail-oriented regulatory lawyer”

Provide context - explain the purpose of the response:

- “Suggest improvements to this email summarising the key points for a busy non-technical person”

Give it limits - set parameters for answers:

- “In no more than 3 sentences, using language a non-lawyer would understand”

Provide style guidance - shape the communication style:

- “Write in an active voice, following British grammar conventions and emulating the tone and style of the Economist”

Iterate beyond the first response - refine and adjust:

- “Use a less formal tone”
- “Does this answer meet the criteria?”
- “How can this be improved?”

Chain instructions together - break down the request into separate steps:

- **Prompt 1:** “Summarise the top 5 principles for creating engaging presentations”
- **Prompt 2:** “You are a legal writing expert assisting with the creation of a 10-minute presentation on effective writing for a law class, first, provide a brief outline of the presentation, including the main topics to be covered”
- **Prompt 3:** “Next, create the text for each slide, focusing on key points and examples, following best practice”
- **Prompt 4:** “Finally, recommend suitable charts and images for each topic, describing why they would be beneficial for the presentation”

Designing effective prompts is an emerging area of practice and, to some extent, highlights the maturity of the technology: the difference between an LLM being highly effective or not can turn on the form of words used to instruct it. We should expect future iterations of the technology to place less reliance on effective prompting and for tools to address common use cases directly.





PART 2: WHAT DOES THIS MEAN FOR ORGANISATIONS?



Thinking strategically about use cases

How can organisations, and in particular legal teams, capture the benefits of using LLMs whilst acting strategically? We think there are three elements to this.

- **Leverage the conversation:** an indirect way to benefit from the developments in generative AI is to take advantage of the current interest in the technology to accelerate existing digital projects. For example, if you are implementing a contract management or knowledge system, can you use the interest in LLMs to encourage engagement from stakeholders and think creatively about future capabilities?
- **Identify commoditised uses:** some of the capabilities discussed in Part 1, like outlining a paper or summarising an article, will likely be made available directly in existing enterprise tools from major suppliers. For example, **Microsoft Copilot** is expected to be able to generate slide outlines in PowerPoint and create charts in Excel from simple user inputs. For these use cases, organisations might choose to be patient and wait and review new generative AI features of enterprise solutions that can be accessed within existing technical boundaries. In the meantime, resources can be invested in enabling projects (such as upgrading to the latest versions of software).

- **Identify organisation and sector-specific uses:** other uses, such as extracting particular types of data from documents or summarising content for a particular audience are unlikely to be made available 'out of the box' in enterprise tools. For these use cases, organisations will need to consider specialised vendor tools or deploy LLMs directly, as discussed below. However, before embarking on a technology project (or creating too many new job titles), organisations should focus on understanding the use case and the processes it supports. This will help quantify the scale of the opportunity and what might need to be put in place to address it in a scalable way.

More complex use cases, such as using LLMs to assist with analysis and question answering, will also require organisations to consider the relevant data sets. As well as the approach to personal and confidential data (discussed below), most organisations will also need to invest resources into preparing and curating their data. For example, if you want employees to be able to ask questions about your company policies, are those policies up to date? Can they be easily read by a machine? Do they have the right metadata for them to be easily retrieved for the LLM?





PART 2: WHAT DOES THIS MEAN FOR ORGANISATIONS?



Deployment options

There are several ways to access and deploy LLMs.

- **Public access:** whilst LLMs have the potential to augment certain tasks, using publicly accessible models (like ChatGPT, Bing or Bard), risks exposing sensitive data to providers whose terms of use have not been reviewed or negotiated. Most providers of public models reserve very broad rights as to how they might use user data and, whilst providers will likely continue to **iterate their terms of use** to address concerns, many organisations will wish to restrict employee access to LLMs or set guidance on their use.
- **Private access:** an alternative is to access LLMs via a private tenancy. For example, Microsoft is making some OpenAI services available to customers in a way that ensures users' data remains encrypted and within an organisation's boundaries (and Amazon is providing similar services with other LLMs). This option may be particularly attractive for organisations which have existing relationships with these vendors and we expect it to be widely taken up once these services are released and testing has been undertaken to ensure they are as secure as is expected.

- **Build your own:** organisations might also invest in developing their own LLMs, giving them more control over the data that is used to train the model as well as the way the model is deployed and used. Today, training an LLM from scratch is an exceptionally time and resource intensive process (with Bloomberg's '**BloombergGPT**', trained on an array of financial data and designed to assist with finance related tasks, being one of the few examples). However, the barriers to building models will likely continue to fall and it may become more feasible for organisations to train their own LLMs. In any case, that is likely to be an organisational level decision, rather than something specifically for the legal team, where the risks of building bespoke (e.g. the costs of keeping the model updated and secure) will need to be weighed against the benefits of a tailored solution.

It is also important for organisations working with vendors of tools which use LLMs to understand which of these deployment options the vendor is using, as this will impact the data privacy analysis.





PART 3: WHAT'S ON THE HORIZON?



The pace of change in this area and the rapid exploration of LLMs has taken the technology world by surprise. We should expect to see an array of new digital tools, and people putting LLMs to new uses. Organisations will need to adopt a flexible approach to the evolving landscape.


Tools building on LLMs

Unsurprisingly, a great many start-ups have emerged which integrate LLMs. A key area of focus is finding ways to apply LLMs to the large volumes of data often associated with legal work. For example, can LLMs be used to diligence the contents of a data room or analyse documents during disclosure? Can LLMs interact with legislation and other primary sources to support legal research? Or perhaps we can give an LLM access to a set of internal policies, guidance notes and precedents to help users answer questions?

These use cases are among the most exciting in the legal context and a range of tools are starting to show promise in tackling these tasks. However, these use cases are far from being 'solved': responding to questions that require complex reasoning or information from many different sources remains challenging. Organisations procuring these capabilities will likely need to start with less complex tasks and invest time in working with vendors to develop their products as they mature.

Other interesting areas to keep an eye on in the coming months include:

- **Prompt injection and prompt stealing:** an area of risk for LLMs that currently remains relatively unexplored is the ability to attack them (or tools using them) using false, untrusted or misleading text as part of an input, a technique known as "prompt injection", resulting in a false instruction. This method can also be used to effectively "steal" the user's original prompt (for example, by simply asking the tool to ignore previous instructions and provide the original prompt), which may be valuable business data.
- **Agents:** LLMs are currently being experimented with as semi-autonomous agents, such as [AutoGPT](#) and [BabyAGI](#). These agents can communicate with themselves, create multi-step plans, interact online, search the internet, and more. Although they are currently of limited use without human intervention and often struggle to stay on-task, a great amount of energy is being invested in improving them.
- **Regulation:** new laws are on the horizon but as calls for AI-specific regulation mount, it will be interesting to see how quickly legislators and regulators roll out new rules (and, indeed, some regulators have taken immediate action, such as [Italy's Garante](#) and the [UK's CMA](#)). We are also yet to see whether the UK Government's sector-specific approach, and overarching cross-sectoral principles, will be successful in practice. Please see our [Regulating AI series](#) for more information.





Our **Client Innovation Network** offers a forum for members of in-house legal teams at our client organisations to connect and share ideas and experiences on innovation topics, including generative AI. If you would be interested to join the Network, please register using [this link](#) or speak to our **Innovation Team** or your usual Slaughter and May contact.

CONTACTS




Sally Wokes

Partner

+44(0) 20 7090 5312

sally.wokes@slaughterandmay.com



James Cook

Partner

+44(0) 20 7090 4216

james.cook@slaughterandmay.com




Nick Johnston

Associate

+44(0) 20 7090 3271

nick.johnston@slaughterandmay.com



Jezah Khamisa

K&I Manager

+44(0) 20 7090 3201

jezah.khamisa@slaughterandmay.com



Richard Batstone

Senior K&I Manager

+44(0) 20 7090 3669

richard.batstone@slaughterandmay.com



Jack Eames

Trainee Solicitor

+44(0) 20 7090 5435

jack.eames@slaughterandmay.com

© Slaughter and May

This material is for general information only and is not intended to provide legal advice. For further information, please speak to your usual Slaughter and May contact.