



Generative AI for Lawyers

Working smarter and faster – while complying
with legal professional obligations

Contents

Foreword	3
1 Introduction	4
2 AI for the legal profession	5
3 Addressing common legal considerations	8
4 Responsible AI	10
5 Current regulatory landscape	12
6 What's next?	13
Appendix: Prompt engineering guide	14



Foreword

“Like no technology before it, these AI advances augment humanity’s ability to think, reason, learn and express ourselves. In effect, the industrial revolution is now coming to knowledge work. And knowledge work is fundamental to everything.”

— Brad Smith, President and Vice Chair, Microsoft

Artificial intelligence (AI) is poised to revolutionize the legal profession. The potential of generative AI, which processes language and generates new content, to transform the provision of legal services is enormous, by automating routine tasks and streamlining processes. AI tools are being used by our Corporate, External and Legal Affairs team – comprising over 1,800 employees around the world – as a copilot to work in a whole new way, allowing our team to spend less time on routine tasks and focus on higher-value work and client engagement.

At the same time, it is crucial to ensure that AI is deployed and used responsibly and, for lawyers, in compliance with legal professional obligations. Human oversight to review, edit and quality-assure AI-generated content will be essential. At Microsoft, we are committed to the responsible use of AI and have developed principles, tools and resources to guide our work, as well as made these publicly available to help others navigate this evolving landscape. We also believe ongoing and proactive dialogue and sharing lessons learned are essential to build trust in this technology and enable its potential to be realized.

In that spirit, we have prepared this paper to explore the considerations associated with the use of generative AI by lawyers, share some use cases relevant to the legal profession and highlight early industry leaders in generative AI innovation. We also examine the status of AI regulation and share the work Microsoft is doing to promote the responsible use of AI. Our goal is that this paper will be a helpful resource for you to use in exploring AI technologies and unlocking the benefits that they can provide to your legal practice.



Mike Yeh

Regional Vice President
Corporate External and
Legal Affairs, Microsoft Asia

1 Introduction

Recent developments in the field of AI are a leap forward in digital capability at least as significant as the arrival of internet browsers in the mid-90s or the widespread adoption of smartphones.

However, far more than the technologies preceding it, AI has the capacity to advance our understanding of the world and our ability to learn and express knowledge.


AI describes the use of computer technologies to perform tasks such as learning, reasoning and problem-solving that would ordinarily be thought of as requiring human intelligence. Recently, AI technology has attracted a surge of renewed interest as the legal profession seeks to understand the impact and potential uses of generative AI. Generative AI can be distinguished from other AI technologies (that primarily use pattern recognition to make decisions) by its ability to create new and unique content, such as text, images and code, using learnings from wider datasets. The accessibility of apps like ChatGPT has catapulted this powerful technology into the hands of the general public.

There is an unprecedented opportunity for lawyers to use AI to enhance capabilities and drive productivity gains, freeing up time to focus on higher-value work. Yet lawyers must be mindful of their legal professional obligations when using these tools.

Harnessing the power of AI in a responsible way is paramount. Microsoft has been committed to this since 2017, when we launched our Aether Committee with researchers, engineers and policy experts to focus on responsible AI issues and craft our AI principles that we adopted in 2018. Such efforts are grounded in our company's mission to empower every person and organization on the planet to achieve more.

The technology is still in its early days. Currently, most countries have no AI-specific regulation, although professional associations and regulators are starting to offer preliminary guidance. Legal use cases are still emerging. No matter what the future holds, advanced human capabilities, empathy and insight will continue to be needed in the legal profession.

But observing the AI revolution from the sidelines is not a viable option. Lawyers must learn to work with this technology in order to adapt to the rapidly evolving legal environment or risk being left behind. With appropriate generative AI system selection and governance, lawyers can safely and responsibly use these technologies in a way that is fully consistent with their legal professional obligations.



2 AI for the legal profession

Generative AI has enormous potential to be a copilot to augment and complement a lawyer's work, allowing legal work to move up the value chain – but not replacing the role of a lawyer.

Goldman Sachs estimates that 44 per cent of legal work stands to be automated by generative AI.¹ An analysis by Microsoft and the Tech Council of Australia has found that 10 per cent of legal tasks could be automated and 32 per cent augmented.²

Early applications of AI in the legal profession centered around data analytics, making predictions based on data or searching large volumes of information. AI technology automating tasks such as document discovery, due diligence and contract management has become widespread but, until now, many AI systems and tools have been rudimentary.

Generative AI offers unprecedented opportunity to advance innovation in the legal profession. By automating routine work, generative AI allows lawyers to focus on higher-value work – from offering strategic guidance to building trusted client relationships – where critical thinking, commercial acumen, empathy and experience, and a holistic understanding of the needs of clients come into play. Generative AI capabilities are becoming easily accessible. For example, Copilot for Microsoft 365 is integrated into apps commonly used by lawyers, including Microsoft Word, Outlook and Teams.

Generative AI systems and tools are only as good as the data they are trained on and use to perform tasks. If training or grounding data are inaccurate, output data may be too. Human oversight will be critical to review the quality and accuracy of AI-generated content.

Examples of AI use cases for lawyers

- **Legal drafting:** creating first drafts of contracts, legal advice and court pleadings;
- **Legal research:** researching and summarizing information drawn from a number of sources and extracting trends and insights;
- **Legal analysis:** evaluating large volumes of contracts to assess risks and compliance. Tools such as Kira, Imprima and Document Intelligence are gaining popularity as AI becomes a standard part of due diligence exercises;
- **Summarization:** generating synopses, summaries and fact sheets;
- **Negotiation support:** generating standardized responses; providing negotiation intelligence;
- **Knowledge management:** retrieving information and generating insights from knowledge database;
- **Communications:** preparing first drafts of client and other communications in the right "voice";
- **Meetings:** scheduling meetings; translation features for more inclusive meetings; transcribing and summarizing meetings and extracting insights;
- **Practice management:** generating insights into performance metrics and financial optimization;
- **Billing:** automating billing processes; estimating time to complete tasks based on past data; and
- **Marketing:** creating marketing content to showcase capabilities and achievements.

We have set out some further examples below.

¹ Goldman Sachs Economics Research, [The Potentially Large Effects of Artificial Intelligence on Economic Growth](#) (26 March 2023)

² Microsoft & Tech Council of Australia, [Australia's Generative AI Opportunity](#) (July 2023)



Practical implementation of AI use cases

Microsoft

Microsoft's Corporate, External and Legal Affairs (CELA) team has launched a department-wide initiative and crowdsourced over 250 ideas to harness the power of AI in three key areas – knowledge management and self-help capabilities; improving contracting; and complying with an increasingly complex regulatory landscape. Use cases being piloted include AI-powered Q&A apps trained on our knowledge database to answer high-volume and low-complexity queries; automating intake and triage processes; and generating draft contracts and identifying clause deviations from standard provisions.

King & Wood Mallesons

King & Wood Mallesons has established a working group to identify use cases to increase efficiencies using AI technology, such as Copilot for Microsoft 365. These include transforming routine processes such as producing client proposals, converting texts to diagrams to communicate complex concepts and utilizing pattern-matching to find previous advice on a topic. Patrick Gunning, Partner, says that "We should embrace generative AI's potential because it's coming, like it or not. The earlier you get your head around it, the more likely you are to benefit from it".

Singapore Ministry of Law (MinLaw)

MinLaw and the Infocomm Media Development Authority have launched a Legal Industry Digital Plan highlighting AI use cases for Singapore law practices with varying needs and levels of digital literacy. A joint release states, "Generative AI can help lawyers better perform knowledge work, such as conducting legal research by asking natural questions and generating first drafts of documents. Such efficiency gains will free up lawyers' time, so that they can focus on performing higher value work such as crafting case/negotiation strategies, identifying new business opportunities and building relationships with clients."³

Linklaters

Linklaters has built its own GenAI chatbot, Laila, using Microsoft's Azure OpenAI infrastructure. Laila, which can be used in multiple languages, streamlines tasks while keeping prompts and responses within the firm's controlled environment, maintaining confidentiality and security. Thousands of users generate queries daily through this handy tool, helping boost efficiency and creating more space for strategic thinking. Adrian Fisher, Partner and Asia Head of TMT, says that "GenAI has the potential to unlock our global institutional knowledge leading to faster turnarounds, enhanced responsiveness and, eventually, to an even more personalized client experience".

3 MinLaw and the Infocomm Media Development Authority, [New Legal Industry Digital Plan to further enhance Singapore Law Practices' digital readiness](#) (accessed 22 February 2024)



Rajah & Tann

Rajah & Tann has rolled out generative AI tools, including Copilot for Microsoft 365, to help employees streamline tasks such as meeting minutes; action items lists; contract clause drafting; “horizon scanning” of regulatory frameworks; and generating client reports on legal developments. “These tools are advancing at lightning speed and we cannot adopt a waiting game,” says Rajesh Sreenivasan, Partner. “Besides intensive initial evaluation by IT specialists, leveraging and refining these tools is now a firm-wide exercise to allow our lawyers to direct their attention and learning towards more value-added aspects of their jobs and offer greater service value to clients.”

Attorney General’s Office of São Paulo (AGU)

The AGU currently reviews 20 million lawsuits, averaging 10,000 summonses a day. By adopting GPT-4 technology, embedded in the Azure OpenAI Service solution, the agency aims to enhance lawyers’ work and generate more efficiency and accuracy in judicial processes. “The assistant has the role of generating summaries and analysis of processes, which can often be long and have several pages, and assist in the preparation of appeals and decisions. AI is being used to become a copilot of our teams,” says Eduardo Lang, director of AGU’s Legal Intelligence and Innovation Department.⁴

JLL

JLL Technologies (JLLT), the technology division of JLL, has developed a bespoke generative AI model – the first of its kind for the commercial real estate industry – which is used by JLL’s 103,000+ workforce around the world, including JLL’s in-house counsel who are regularly using it to support their day-to-day work. “I am so excited to have JLL at the forefront of this AI sector. While still in the early stages, our teams identify and use generative AI to increase our efficiency and productivity as lawyers. We should embrace and evolve (with controls) rather than block. This will soon be the norm,” says Christopher Y. Chan, APAC General Counsel.

Jardines

Jardines’ legal team recently partnered with the University of Hong Kong’s Law, Innovation, Technology and Entrepreneurship Lab to explore productivity gains from the use of generative AI technology, streamlining routine tasks like summarizing contracts, translation and drafting. It has also partnered Jardines’ innovation and technology team to create hands-on exercises using Microsoft Azure OpenAI Service for its 70+ lawyers to explore potential use cases. “The enthusiastic feedback highlighted untapped potential across industries and jurisdictions. While AI augments work, it doesn’t replace our lawyers’ rich experience and deep-seated corporate knowledge”, says John Koshy, Head of Group Legal Affairs, Jardine Matheson.⁵

4 [With the adoption of AI, AGU seeks to improve efficiency in legal proceedings](#) (June 2023)

5 [Asian Legal Business, *Innovation Unleashed*](#) (January-February 2024 Asia Edition)

3

Addressing common legal considerations

We explore below some common issues and concerns associated with the use of AI by lawyers:



Confidentiality

Lawyers are subject to a duty to maintain the confidentiality of client information. Legal professional privilege depends on the confidentiality of privileged communications being maintained. Lawyers need to consider the data protection commitments and controls associated with an AI solution before inputting confidential data (either as part of a prompt or as part of grounding data used to provide additional context).



Accuracy

Lawyers are subject to a duty to act in the best interests of the client. There is a need for human oversight and implementation of accuracy techniques (such as retrieval-augmented generation) to deliver services competently and ensure the AI technology is not “hallucinating” – that is, delivering false or misleading information. Generative AI should be used as a copilot, not an autopilot.




Data privacy

Where personal data is used in an AI system – whether to develop, train or test an AI system or to provide prompts or input – developers and users will need to ensure that its use is fair and lawful under applicable data protection laws. Lawyers should carefully review the data protection and privacy commitments of their AI service providers.



Intellectual property

A key legal consideration is whether or not the particular solution utilizes data that is protected by copyright. Lawyers should seek contractual protections against potential third party copyright infringement claims from their AI service providers.





Microsoft supports in addressing these common issues and concerns as follows:



Data security

Microsoft's AI solutions leverage the entire complement of data security and privacy safeguards available in Azure and/or Office 365. Azure OpenAI Service foundation models do not store customer data. Copilot for Microsoft 365 presents only data that a user can already access using the same underlying controls for data access used in all Microsoft 365 services.



Provenance and accuracy

Microsoft is a co-founder of Project Origin and the Coalition for Content Provenance and Authenticity (C2PA) standards body to address media provenance. Our watermarking and metadata techniques enable Microsoft Copilots to include links in output content to source materials so that users can verify these. We also share information about ways to reduce the likelihood of "hallucinations", such as our [guidance on retrieval-augmented generation \(RAG\)](#). RAG grounds the model vectors – relationships between words – with data that is more relevant to the user, such as data specific to a person's role, company or industry.



Privacy safeguards

Azure OpenAI and Copilot for Microsoft 365 are core online services protected by the highest levels of Microsoft's data protection and privacy commitments, including European Union General Data Protection Regulation (GDPR) and data boundary commitments. Customer data is not used to train foundation models. Data that the customer uploads to create a customized model is not shared with any other customer or third party, nor is it used in the processing of any other customer model.



Customer copyright commitment

The large language models that Microsoft makes available have been trained in compliance with copyright laws. If a third party sues a commercial customer for copyright infringement arising from the output content, Microsoft contractually commits to defend the customer and pay the amount of any adverse judgments or settlements, as long as the customer has used the required guardrails and safety measures. Microsoft has taken steps in the design process to mitigate the risk that the output content will violate copyright laws, and publishes guidance for how customers developing their own models can protect against such risks.

4 Responsible AI

Every organization, including law firms, that creates or uses AI systems will need to develop and implement its own governance processes to ensure AI tools are used and developed responsibly.

Microsoft has been on a responsible AI journey since 2017 when we established our research-led AI, Ethics, and Effects in Engineering and Research (AETHER) committee in order to explore AI technologies and the standards that should apply to their development.

In 2019, Microsoft established the Office of Responsible AI in order to foster a comprehensive approach to Responsible AI. This office has dedicated significant resources to understanding the principles that form the basis for a responsible approach to AI and the standards that should inform development and implementation of AI systems.

Microsoft's six AI principles are as follows:



Based on these principles, Microsoft has launched its [Responsible AI Standard](#). This standard consists of goals and requirements for each of the six principles and is intended to function as a checklist or scorecard for companies developing and implementing AI systems.

To further help customers on their responsible AI journey, Microsoft recently launched a set of [AI Customer Commitments](#). These commitments focus on three areas:

- Sharing and providing AI resources to help customers deploy AI responsibly;
- Creation of an AI assurance program; and
- Developing responsible AI partner programs. For example, Microsoft recently launched a partnership with PwC and EY to leverage responsible AI expertise to help mutual customers deploy their own responsible AI systems.

Learn more about AI principles and explore responsible AI resources at: aka.ms/ResponsibleAIResources

“The point is to recognise that our growing use of generative AI tools will come with various ethical risks, and resolutely commit to their appropriate and ethical use, so as to preserve our status as an honourable profession.”⁶

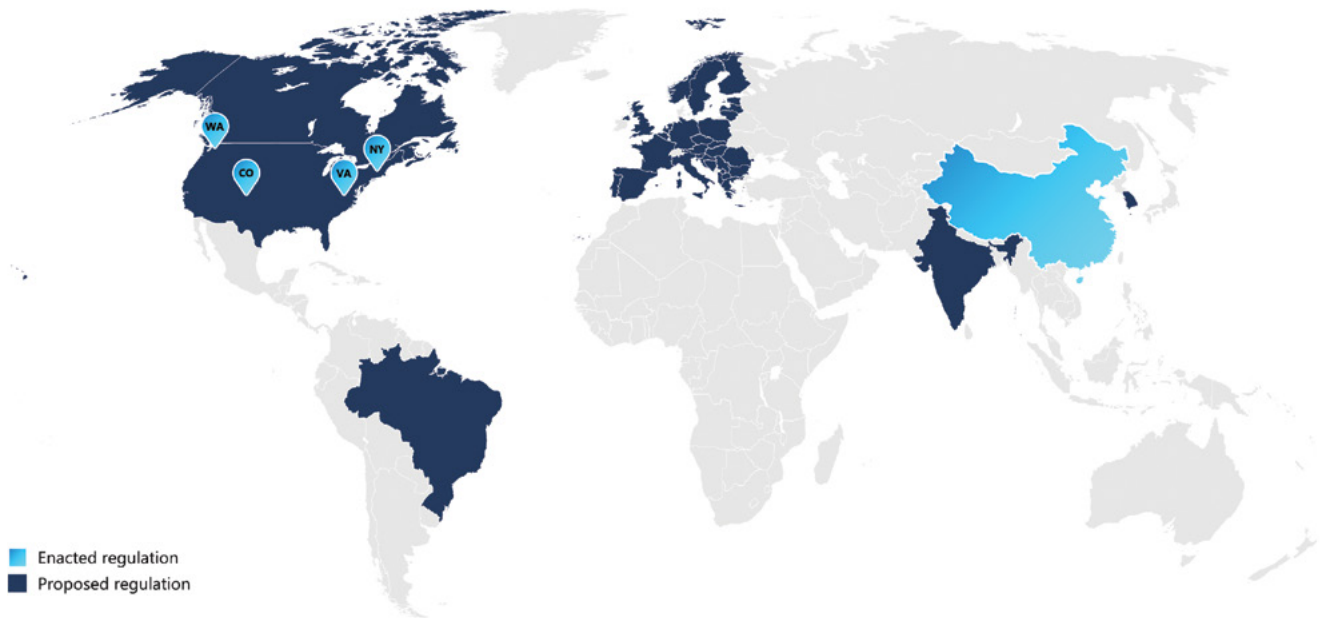
– The Honourable the Chief Justice Sundaresh Menon, Supreme Court of Singapore

6 Sundaresh Menon CJ, “Answering the Call in the Age of Artificial Intelligence”, Mass Call Address 2023 (21 August 2023), para 23



5 Current regulatory landscape

Governments are adopting a range of approaches to consider how best to govern the use of AI technologies.



AI regulation in most countries is at an early stage, albeit evolving rapidly. Generally, regulators in Asia have focused on “soft” guidance in the form of ethics principles, guidelines or codes of conduct for the use and development of AI technology. **Hong Kong** and **Singapore** are examples of countries taking a pro-innovation, principles-based approach. A notable exception is **China**, which has taken a prescriptive approach and introduced three AI-specific regulations to regulate specific AI technologies (including, most recently, generative AI). Although other countries in Asia have started discussing whether AI-specific laws should be introduced (including **India**, **Japan**, **South Korea** and **Taiwan**), these discussions have not yet crystallized into binding laws or regulations.

Some indication of what future regulation could look like comes from the **European Union’s** proposed AI Act, one of the first comprehensive laws of this nature in the world.⁷ The proposed AI Act follows a risk-based approach. AI systems that contravene European Union

values, such as violating fundamental rights, will be deemed an “unacceptable risk” and banned. Those that pose significant potential harm to health, safety, fundamental rights, environment, democracy and the rule of law will be classified as “high-risk” and subject to additional compliance requirements. General purpose AI models will have to adhere to transparency requirements, and additional requirements will apply to general purpose AI models with systemic risk.

While AI policy is still in its early stages of development for many countries, some key common principles are emerging, which align closely with Microsoft’s AI principles described in Section 4 above. Microsoft’s [blueprint for AI governance](#) presents our proposals for AI regulatory frameworks, which we believe should be risk-based and outcomes-focused, focusing safeguards on the highest risk applications and utilizing impact assessments and system testing to identify and address risks and ensure systems are performing appropriately.

⁷ European Parliament, [Artificial Intelligence Act: deal on comprehensive rules for trustworthy AI](#) (9 December 2023)

6 What's next?

Generative AI promises to place exciting new capabilities in the hands of every lawyer and organization.

Firms and in-house legal operations of the future may be filled with knowledge engineers, data analysts, technologists, design thinkers and transformation experts alongside lawyers. In light of the evolving definition of what it means to be a legal professional, education and training models will need a refresh.

Some organizations will be happy to lead while others wait until there are established industry-wide frameworks. However, the rapid advancement of generative AI indicates this is an area organizations must prioritize to maintain a competitive advantage. The practice of law is changing fast.

*"We're at a point in time where we're shifting from the abacus to the voice-activated calculator and it's a beautiful thing. The question is whether we have the vision, aptitude for adoption, innovation readiness, and change management strategy to get there thoughtfully, responsibly, and efficiently from where we are today."*⁸

– Ilona Logvinova, Associate General Counsel and Head of Innovation for McKinsey Legal

⁸ McKinsey, *Legal innovation and generative AI: Lawyers emerging as 'pilots,' content creators, and legal designers* (11 May 2023)

Checklist for adoption of generative AI

Safe use

- Implement internal policies to provide guidance to employees on permitted uses of generative AI.
- Provide training to employees on responsible use of AI, including the need to verify outputs.
- Factor the use of AI into governance and risk management frameworks.
- Identify high-risk cases and escalation paths.

Data issues

- Ensure compliance with privilege/confidentiality obligations and data protection regimes, including confirming data classification and labelling.

Ethical challenges

- Consider applicable responsible AI guidance and standards.
- Be transparent to third parties about the use of AI in work products.
- Ensure there is human review of outputs.

IP issues

- Consider IP rights in relation to outputs.
- Review contractual protections from providers to cover potential copyright infringement risks.

Regulation and liability

- Ensure compliance with applicable professional standards and regulatory requirements with respect to the use of AI.
- Consider providers' compliance with applicable responsible AI standards and regulatory requirements.

Appendix: Prompt engineering guide

By knowing the basics of prompt engineering – the practice of formulating instructions to obtain desired results from the generative AI tool being used – lawyers will be better able to write effective prompts and generate more accurate and relevant results.

A good prompt will provide a clear description of the task, explain the role the AI tool needs to play, describe the audience, provide guidance on the tone, style and length of the expected output, and any additional context to be taken into account. An iterative process will help to refine the result.



Basic concepts:

- A user enters a **prompt (or "input")**, the question or instruction for the model.
- The model uses this information to generate a **response (or "output")**, in the model's format (e.g., natural language; images).

Other terms you may have heard of:

- **Grounding**: Data provided with a prompt to provide it context for responding to that prompt.
- **Fine-tuning**: A process of using data to further train a model so that it provides more specific responses in a given domain.
- **Metaprompt**: A master instruction to the model to apply in all prompts/responses.



We have set out below some examples of AI prompts that can be used in Copilot for Microsoft 365 – which incorporates generative AI capabilities into Microsoft apps that are commonly used by lawyers, such as Word, Outlook and Teams – to demonstrate how these can be used to support day-to-day work.



Copilot in Word

- “Please draft a news bulletin on the joint position paper published by the Information Commissioner’s Office and the Competition and Markets Authority on harmful online designs, with a similar structure to [attach reference file].”
- “Please write an email to [insert name] explaining what this agreement is and a bulleted list of the buyer conditions and also a bulleted list of the seller conditions, and please suggest a meeting next week to discuss next steps and ask them to prioritize it as there is time sensitivity.”
- “Generate a client proposal from [insert reference file], aligning it with my template [insert reference file].”



Copilot in Teams

- “Can you recap this meeting so far?”
- “Create a table of the options discussed with pros and cons.”
- “Can you draft an email to [insert name] summarizing action items from the call, and that I’m looking forward to following up on them during our discussion on Wednesday?”



Copilot in PowerPoint

- “Create presentation from file [attach reference file].”
- “Add one slide summarizing the proposed EU AI Act.”
- “Organize this presentation into sections.”



Copilot in Outlook

- “Summarize this thread.”
- “What’s the latest from [insert name]?”
- “Please write a short email in a formal tone to congratulate the team on the deal closing.”

