




Guidelines for the **Use of Generative Artificial Intelligence** by the Legal Profession in Ireland

December 2025

Introduction



The aim of this guidance note (**Guidance Note**) is to raise awareness of what Generative AI (**GenAI**) is, and to explain the current uses and limitations, along with the risks and opportunities associated with the use of GenAI in legal practice. This guide specifically aims to highlight the use of GenAI while adhering to core professional obligations as set out in the *Solicitors' Guide to Professional Conduct (4th Edition)* (the **Code of Conduct**).

The Law Society of Ireland is aware that the field of GenAI is moving and evolving fast. Therefore, the focus of this Guidance Note is on the fundamentals of GenAI and risks most relevant to everyday legal practice. Additional guidance and further updates will be issued by the Law Society of Ireland as the technology continues to evolve and as new issues or use cases arise in legal practice.

This Guidance Note does not explore or approve sector-specific or specialist 'Legal AI' tools designed for use in legal practice, nor does it cover other subsets of AI or areas of practice that are already significantly impacted by the adoption of AI such as copyright, cybersecurity, intellectual property, and data protection which require further guidance from subject-matter experts.

Further guidance and resources will be made available on the Law Society's webpage regarding practical advice and a list of potential providers to the legal profession.

The Guide is relevant for all solicitors, whether in private practice or in employment in the in-house and public sector.

What is AI?

The term ‘Artificial Intelligence’ (“**AI**”) is an umbrella term that covers a wide range of technologies and tools that can vary significantly in their capabilities and applications.

For this reason, it’s important to specify what type(s) of AI is being used, as the ethical, practical, and professional risk profiles may differ considerably depending on the specific AI being used.

What is GenAI?

GenAI is a subset of AI that generates ‘new’ content. This content can be in the form of text, audio, images and video that did not previously exist in precisely that form. This ability to generate new material is what sets GenAI apart from other forms of AI, which might only sort, search, or filter information without creating anything new.

The most common use of GenAI in legal practice involves text-based models, commonly known as large language models (“**LLMs**”). For the purposes of this Guidance Note, references to ‘AI system’ ‘*GenAI*’, ‘*text-based GenAI*’, and ‘*LLM*’ should be understood as referring to the same type of technology.

How does GenAI work?

GenAI works by combining two key elements:

- **The data it is trained on:** GenAI models are trained on vast quantities of data (including books, articles, images, and even some laws). The quality, size, and diversity of this dataset determine what the model “knows” and how it is able to respond to an instruction given by the user to the model (known as a ‘prompt’).
- **Advanced mathematics/statistics:** GenAI models do not “think” or “understand” like humans. The model uses statistical and probabilistic methods to calculate *the most likely* way to respond to a particular instruction or prompt. For LLMs, this means predicting one word at a time, which builds into sentences and paragraphs based on the instruction or prompt that was provided by the user.


Examples of publicly available LLMs include *Copilot* (Microsoft), *ChatGPT* (OpenAI), *Claude* (Anthropic), and *Gemini* (Google). These LLM models are defined as “*general purpose AI systems*¹” and by their nature can be used across a number of business sectors for general application.

¹ Article 3 of the Regulation (EU) 2024/1689.

Key Characteristics of text-based GenAI/LLMs

LLMs are good at language manipulation tasks, which can help reduce the manual labour involved in a task that would otherwise require a human to start the drafting process from scratch. Below is a table summarising some uses and limitations for LLMs:

Uses	Limitations
• Writing • Summarising • Translation • Brainstorming ideas, analogies, strategy planning	• Accuracy and understanding • Hallucinations • Bias and sycophancy • Confidentiality and security • Transparency • Limited context window • Outputs differ each time • Knowledge cut-off



Accuracy and understanding

GenAI systems are not built for accuracy, even if they have access to accurate material within their dataset. For example, an LLM can produce something that is technically correct or accurate, but it can organise that data in incorrect ways due to its probabilistic nature and knowledge blind generation of text. Without subject-matter expertise and a human in the loop to review and verify all outputs, material facts and subtle nuances could be missed in the drafting process.

Hallucinations

The goal of an LLM is to continue the sequence of words in a way that looks ‘correct’ regardless of whether the output is factually correct or not. This process optimises for *fluency* (i.e. the answer sounds right and appears to be well drafted) and *relevance* to the prompt, but it does not account for accuracy. This limitation is known as a ‘hallucination’ which is why unverified AI-generated content must never be assumed correct or used without independent verification.

Hallucinations commonly occur in the following ways:

- 1. Using LLMs to assist with subject matter expertise:** where a party acts *pro se*² and uses an LLM to draft legal documents for legal proceedings, such documents could be well-drafted and persuasive but riddled with misleading or outdated references that a subject matter expert would notice.
- 2. Using LLMs to provide factual information:** this is common among users that use LLMs as a substitute for search engines or legal research that require authoritative sources (e.g. legislation, case law and academic journals). LLMs can only generate information within its training data, which if not reviewed and verified by a subject-matter expert, may generate misleading or fabricated sources that appear correct.

Bias and sycophancy

Bias in GenAI systems arises due to the vast quantities of data it is trained on. The quality of this data is unknown to users, and may embed historical or systemic biases. These biases can inadvertently reproduce or amplify existing societal prejudices, leading to unfair or inaccurate outputs.

² Pro se is latin for or “on one’s own behalf”

Confidentiality and Security

Solicitors must be able to clearly differentiate between free, paid, and enterprise-grade versions of GenAI systems offered by providers.

By default, free and paid consumer versions of GenAI systems are not suitable for securely handling personal data or client confidential data, and should not be considered compliant with data protection laws for professional use.

In most cases, the data entered into a GenAI system is processed by a third party, which can include personal information and confidential client details. AI providers may not be aware of the nature of the data provided, and that its nature may be personal data or client confidential data. Therefore, the AI provider may be unintentionally treating the data as ordinary data.

In addition, data entered into free or paid consumer GenAI systems may be stored and re-used by the provider for training and improving the GenAI system. This includes using the input and conversation history of its users and users' personal information. Data may also be accessed by the provider, its employees, or shared with third parties or other vendors for various other purposes.³

Enterprise-grade tools are specifically designed for professional use and can offer enhanced contractual and technical protections to support compliance with data protection laws and protect confidential data. Solicitors should *never* assume that using an enterprise-grade tool

guarantees compliance. It remains essential to carry out thorough due diligence, including a detailed review of all terms and conditions, on any provider before using their platform, and to ensure all professional and legal obligations continue to be met.

Transparency

At present, virtually all GenAI systems exhibit what is known as the 'black box' phenomenon, where their internal reasoning processes are opaque and difficult to interpret. For example, a user cannot ask an LLM why it provided a specific answer, as it does not 'reason' or 'understand' any of its outputs.

Limited context window

Context windows are the memory boundaries for LLMs which means they can only process a certain amount of text at one time. The size of the context window will affect how accurately an AI model can interpret complex or lengthy prompts, follow conversation threads, and maintain coherence across interactions. Any information outside the LLMs context window may be ignored or forgotten by the LLM.⁴

³ See for example [Docusign FAQs for AI](#) (accessed 16 October 2025), [OpenAI Privacy Policy](#) (accessed 16 October 2025), and [Anthropic Privacy Policy](#) (accessed 20 February 2025).

⁴ For example the context window for ChatGPT 5 is currently 128k tokens, which is roughly 150-170 pages of text. Source: <https://help.openai.com/en/articles/11909943-gpt-5-in-chatgpt> (accessed 17 October 2025)

Outputs differ each time

As LLMs optimise for fluency and language manipulation, users will not receive the same ‘word for word’ output in response to the same prompt each time. The below is an example where the user used the same prompt on two separate occasions:



Example 1 (<i>Copilot</i>)	Example 2 (<i>ChatGPT</i>)
Instruction/prompt: <i>“Summarise the EU AI Act in one sentence”</i> Output: <i>“The EU AI Act sets rules to ensure that AI systems used in the EU are safe and respect fundamental rights.”</i>	Instruction/prompt: <i>“Summarise the EU AI Act in one sentence”</i> Output: <i>“The EU AI Act regulates the development and use of AI in the EU to promote trust, safety, and accountability.”</i>

Knowledge cut-off

By default, LLMs have a knowledge cut-off date and do not have access to real-time information, subscription databases or paywalled information, including consolidated and up-to-date legislation.⁵

⁵ For example, Claude Sonnet 4.5 has a knowledge cut-off date of January 2025. See: <https://www.anthropic.com/transparency> (accessed 17 October 2025).

Gen AI in legal practice: Opportunities and Risks

The use of GenAI in legal practice can be hugely beneficial to solicitors provided all users are sufficiently trained on its uses and limitations associated with its use. This in turn highlights opportunities for use by solicitors and staff, and ensures that ethical and professional standards are always maintained in the process.

LLMs can produce easy wins for solicitors and unlock efficiencies for simple administrative tasks that users would otherwise do manually. In some instances, LLMs can even assist the solicitor with their professional obligations, such as translating a piece of legal text into plain English for a client⁶ that is not legally trained. This can be beneficial and time saving, provided the solicitor retains full control of the output at all times.

Using LLMs in this manner is not the equivalent of delivering legal advice, which can only be provided by a practising solicitor issued with a current practising certificate from the Law Society of Ireland.⁷

Circumstances may arise where the LLM makes recommendations or provide outputs that the solicitor may use in the process of formulating his or her work product forming the basis for legal advice. The solicitor, in using their own skill, knowledge and professional judgement, must take full responsibility for its content and be able to fully explain their reasoning behind the final work product.

Below is a non-exhaustive list of tasks suitable for working with LLMs:

1. **Creating checklists:** Ask an LLM to prepare a checklist for key provisions in a contract review.
2. **Summarising materials:** after reviewing a document, ask the LLM to summarise the document in under 500 words.
3. **Adapting Clauses:** Provide a sample clause to an LLM and instruct it to produce a 'seller friendly' or 'purchaser friendly' version.
4. **Translating Legalese to Plain English:** Provide the LLM with a piece of legal text and ask it to provide a plain English version for a client that is not legally trained.
5. **Enhancing Persuasion:** Provide the LLM with a list of points and request suggestions for counterarguments or ways to make your position more persuasive.
6. **Rephrase tone:** Ask the LLM to rewrite a paragraph to be more formal, concise, or convey an assertive tone.
7. **Brainstorming:** Provide the LLM with a list of points you want to discuss at your next meeting and ask it to give you feedback or additional ideas.

⁶ Page 25, Code of Conduct.

⁷ See [the Solicitors Acts eCompendium](#) (accessed 17 October 2025).



Below is a non-exhaustive list of tasks not suitable for working with LLMs, including tasks that are more likely to generate hallucinated outputs:

- 1. Legal advice:** LLMs do not understand context and should never be used to form the basis for legal advice. Their outputs may miss jurisdiction specific nuances and generate inaccurate advice.
- 2. Citing specific legislation and case law:** LLMs do not have access to verified subscription databases. Further, it does not ‘understand’ what information is relevant when researching legal content. An LLM is more likely to produce a hallucinated output when used for this purpose.
- 3. Document review:** asking an LLM to perform a legal gap analysis, verify defined terms, or cross-check defined terms across documents is not a suitable task for LLMs.
- 4. Numerical calculations:** unlike dedicated spreadsheet tools (e.g. Excel), LLMs are not designed to perform mathematical computations or complex numerical analysis.
- 5. Real time information retrieval:** LLMs have a knowledge cut-off date and by default do not have real-time internet access. All outputs are generated from the LLMs underlying training data from a fixed period. Some LLMs offer a ‘search’ function to pull recent information from the web, this is a secondary feature layered on top of the core LLM model. Users are relying on the LLM decide what information is ‘relevant’, and therefore more likely to misinterpret or entirely missing relevant information.

Finally, it is worth noting that there is an emergence of ‘Legal AI’ tool providers that combine GenAI with other subsets of AI to ensure more accuracy in their outputs and tackle current limitations associated with LLMs. This topic is beyond the current scope of this Guidance Note. It is intended that further information regarding potential providers will be made available in due course.

Gen AI in legal practice: Professional Obligations

The use of GenAI engages several core values of the profession. The purpose of this section is to highlight the most important areas where the use of GenAI could potentially breach these core values under the existing Code of Conduct. This Guidance Note does not amend or update the Code of Conduct, nor does it create new rules or impose new obligations on solicitors.

Duty of Confidentiality and Privilege

Client confidentiality represents a fundamental professional obligation that remains paramount when using AI tools, applying to all communications passing between a solicitor and their client or former client, and to the existence of the relationship.⁸ Confidentiality also extends to the staff of the solicitor.⁹

Solicitors are already required to use reasonable endeavours to prevent a breach of security and confidentiality, adhere to data protection requirements and ensure a level of security appropriate to risks within their practice.¹⁰

It is recognised that, in the course of practice, firms will inevitably have to give limited access to client data to their professional advisers, IT maintenance, contractors, and others.¹¹

Regarding the use of GenAI systems, solicitors should be aware that:

- users are responsible for the information they provide to any GenAI system, which may be in breach of their professional obligations and/or relevant data protection laws at the point of inputting the data into a system.
- users should refrain from entering any personal, confidential or other data relating to a client into a GenAI system (prompts), unless there are appropriate safeguards in place with the provider of that GenAI system.

Appropriate safeguards can include but are not limited to, for example:

- contractual obligations for the provider to treat data as confidential and apply zero data retention periods;
- a data protection agreement with the provider whereby the data entered will only be used for the purposes of the firm;
- setting up relevant technical safeguards and if available, apply internal system settings to avoid or limit data sharing; and/or
- setting up AI systems to run locally or within a secured environment controlled by the firm.

⁸ Page 52, Code of Conduct.

⁹ Page 57-58, Code of Conduct.

¹⁰ Page 112, Code of Conduct.

¹¹ Page 57, Code of Conduct.

Regarding privileged communications (written or oral) that pass between a solicitor and their client or a prospective client, the privilege is that of the client, and the solicitor cannot be compelled to disclose those communications unless ordered to do so by a court.¹² Intentionally providing privileged communications to a free or paid GenAI model without appropriate safeguards in place may lose the benefit of privilege due to the intentional release of a privileged document to a third party outside of the firm.¹³

GenAI tools have increasingly been incorporated into many everyday tools which solicitors use, such as PDF readers, text editors, and transcription or recording features (e.g. Microsoft Teams, and Zoom). It may be that users do not realise the extent to which AI is embedded in these applications. In particular, a solicitor should not record a conversation via a video platform without the express consent of the other party.¹⁴ If consent is obtained, then the same care should be taken with confidential information and protecting privileged communications when using these tools.

Professional Competence¹⁵

Solicitors are responsible for their work, and in doing so are required to keep their knowledge and skill up to date on a continuing basis during the whole of their professional career¹⁶, which includes keeping abreast with the technological developments affecting their legal practice.

Having sufficient knowledge and training on the use of GenAI tools will help solicitors to assess and mitigate the risks related to the use of GenAI tools. To support solicitors, the Law Society of Ireland will continue to run professional training courses on the use of AI tools in legal practice.

Independence

The independence of a solicitor is a core value of the profession, and it is a duty of the solicitor to ensure that their independence is not compromised.¹⁷ Solicitors who are owners of firms also have a responsibility for the good management of their firm.¹⁸ Overreliance on GenAI tools by solicitors and staff has the potential to undermine independence, particularly if there is insufficient review and verification of the output by a subject-matter expert.

¹² Page 49, Code of Conduct

¹³ Page 51, Code of Conduct

¹⁴ Page 86, Code of Conduct

¹⁵ See *McKechnie J. in Law Society v Carroll* [2016] 1 IR 676 and Page ii, Code of Conduct.

¹⁶ See, the Solicitors (Continuing Professional Development) Regulations 2017 (S.I. No. 529 of 2017), as amended by the Solicitors (Continuing Professional Development) (Amendment) Regulations 2023 (S.I. No. 419 of 2023)

¹⁷ Page 16, Code of Conduct

¹⁸ Page 16, Code of Conduct

Discrimination

A solicitor should not discriminate based on gender, civil status, family status, sexual orientation, religion, age, disability, race, or membership of the Traveller community.¹⁹ To avoid undetected bias being incorporated into their work, a solicitor should be particularly mindful of how they use GenAI tools to assist their work, and avoid using GenAI tools to assist them with any tasks that demand scrupulous fairness and areas where impartiality is essential.

Duty to Supervise

Solicitors are professionally responsible for all legal work carried out within their office, whether performed by qualified assistants, trainees, or support staff. A solicitor will not escape responsibility for work carried out in their office by delegating the relevant matter to the staff employed by them, even though they are qualified to do the work.²⁰

As with any new technology, it is prudent for solicitors to foster open and transparent conversations within their teams about the use of GenAI tools and ensure staff can make informed decisions about how GenAI tools work at a basic level, and what their limitations and risks may be (such as hallucinations or bias). Maintaining a culture of open discussion and basic

AI literacy across the firm ensures that GenAI outputs are always subject to appropriate review and verification.

Duty to not mislead the court and assist with the administration of justice

Particular caution is required before any GenAI assisted work is submitted externally to clients and the courts. The use of GenAI in legal submissions has been documented across various jurisdictions²¹, including the Irish courts²², as producing inaccurate or misleading results.

As already outlined in this Guidance Note, GenAI tools such as LLMs are not built for accuracy. Therefore, the use of GenAI and its potential benefits should never be at the expense of accuracy or truthfulness. The duty to ensure accuracy and reliability remains firmly with the solicitor.

For further information, see the Law Society's Guidance Note '*Mitigating AI Hallucinations*'²³ which provides tips to reduce the risks of hallucinations in case-law searches and outlines the differences between AI-driven legal research tools and general-purpose AI models such as LLMs listed in this Guidance Note.

¹⁹ Page 84, Code of Conduct

²⁰ Page 107, Code of Conduct

²¹ Jurisdictions include: the United Kingdom (R (Ayinde) v London Borough Council [2025] EWHC 1040), Australia (Valu v Minister for Immigration and Multicultural Affairs (No 2) [2025] FedCFamC2G 95) and the United States (Mata v Avianca 678 F. Supp. 3d 443)


²² As of the date of this Guidance Note, the following cases imply the use of Generative AI in legal submissions: Reddan & An Bord Pleanála v. Trustees of Nenagh Golf Club [2025] IEHC 172, John Coulsto et al. v Elliott [2024] IEHC 69, and Malone & McEvoy v Laois County Council, ABP & Booth [2025] IEHC 345.

²³ See Law Society Gazette article 'Reducing 'hallucination' risk in AI case-law search' Link: <https://www.lawsociety.ie/gazette/top-stories/2025/july/reducing-hallucination-risk-in-ai-case-law-search/> (date accessed: 16 October 2025)

Communications with clients and other solicitors²⁴

A solicitor should be honest in all their dealings with their clients and other solicitors.²⁵ A solicitor should be transparent as to the reasons for using GenAI tools to assist in their work and should consider referring to the use of GenAI tools in their terms and conditions letter to clients.

Notwithstanding the above, the duty of honesty does not create a positive obligation to disclose GenAI use in the same way that a solicitor is not required to disclose their service providers or IT maintenance arrangements to clients or other solicitors unless specifically asked. This does not apply to situations where the solicitor is required to disclose GenAI use as required by transparency obligations under the AI Act or otherwise, which is beyond the scope of this Guidance Note.²⁶ Solicitors are responsible for ensuring that they remain up to date with their obligations under the AI Act.



²⁴ Page 89, Code of Conduct

²⁵ Page 89, Code of Conduct

²⁶ See Article 50 of the AI Act for further information regarding transparency obligations for deployers.

The EU AI Act: What does it mean for the legal profession?

The EU AI Act (**AI Act**) is the first comprehensive framework of its kind to regulate all actors involved in the AI value chain, which include deployers²⁷, providers²⁸, importers²⁹, distributors³⁰, product manufacturers³¹ and authorised representatives.³²

AI systems are classified into four risk categories: unacceptable risk, high risk, limited risk, and minimal risk. The LLMs referred to in this Guidance Note are separately categorised as general-purpose AI systems, which have their own set of obligations under the AI Act.

In simple terms, the level of ‘risk’ attributed to an AI system is assessed by reference to the potential impact the system could have on fundamental human rights, health, and safety of individuals.³³ The level of risk associated with a general-purpose AI system depends on the level of ‘systematic risk’ it poses.

The scope of this Guidance Note is limited only to ‘deployers’ of general-purpose AI systems used in a limited risk manner, which is in line with the examples provided throughout this Guidance Note as they relate to professional obligations under the Code of Conduct.

For example, a solicitor that uses an LLM in their daily work should assume they are a ‘deployer’ of a general-purpose AI system under the AI Act, except where it is used in the course of a personal non-professional activity.

As of 2 February 2025, deployers have a mandatory obligation to ensure a sufficient level of ‘AI literacy’ among staff and other persons dealing with AI systems on their behalf. This includes ‘other persons’ broadly under the organisational remit that are dealing with the operation and use of AI systems on behalf of the deployer. These persons do not have to be directly employed by the firm and can include, for example, a service provider, client, consultant, or contractor.³⁴

²⁷ See Article 3(4), a “deployer” means “a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity”.

²⁸ See Article 3(3) of the AI Act. A “provider” means any organisation “that develops an AI system or a general-purpose AI model, or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge”

²⁹ See Article 3(6), an “importer” means “any organisation located or established in the EEA that places on the market an AI system that bears the name or trademark of any organisation outside the EEA”.

³⁰ See Article 3(7), a “distributor” means “any organisation in the supply chain, other than the provider or the importer, that makes an AI system available on the EEA market”.

³¹ The concept of a “product manufacturer” is not explicitly defined in the EU AI Act (instead, it is defined in the EU harmonisation legislation listed in Annex I to the AI Act – see Recital 87). Product manufacturers are within the scope of the EU AI Act when they place an AI system on the EEA market together with their own products and under their own name or trademark.

³² See Article 3(5), an “authorised representative” means “any organisation located or established in the EEA who has received and accepted a written mandate from a provider of an AI system or a GPAI model to, respectively, perform and carry out on its behalf the obligations and procedures established by the EU AI Act”.

³³ See Recital 28 of the AI Act.

³⁴ See ‘AI Literacy – Questions and Answers’ <https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers> (accessed 17 October 2025)

The minimum requirements for AI literacy require staff and persons to have the relevant “*skills, knowledge, and understanding that allow...deployers, and affected persons to make an informed deployment of AI systems, and to gain awareness about the opportunities and risks of AI and possible harm it can cause*”.³⁵

There is no ‘one-size-fits-all’ for AI literacy training, however recent EU guidance confirms that at minimum, deployers of AI systems should ensure a general understanding of AI within their organisation, taking into account their role; the level of risk attributed to that AI system; the technical knowledge, experience, education, and training of staff, and build AI literacy actions based on those factors.

It is worth noting that the requirement for ‘AI literacy’ would similarly arise from the duty of the solicitor to be competent in using GenAI tools, requiring them to have the necessary skills, knowledge and understanding to use those tools in a way maintains professional standards.

Solicitors should be aware that there are specific transparency obligations under the EU AI Act that apply to deployers of LLMs referred to in this Guidance Note, which are outside the scope of this Guidance Note.

Conclusion

GenAI is a tool that may be of assistance to solicitors. However, as with any tool a solicitor should conduct a risk assessment to ensure the tool is used correctly and in accordance with office policy and procedures. Your office policy and procedures should set out permitted uses, accountability, safeguards and guidelines to ensure compliance with GDPR, client confidentiality, privilege, etc. The onus is on firms and solicitors to ensure that ethical and professional standards are always maintained in the process.

³⁵ See Recital 20 and Article 3(56) of the AI Act.



Law Society
of Ireland