Law Society of Scotland

Guide to Generative AI

October 2024

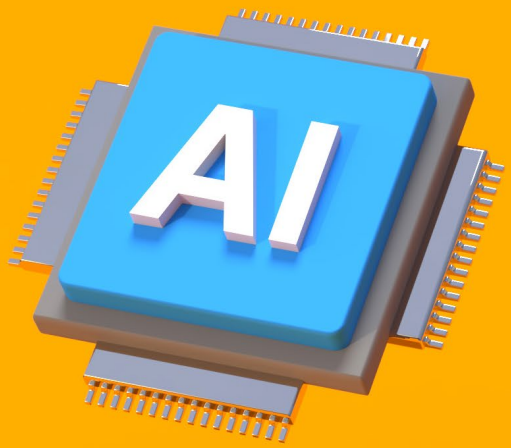

In association with

Wordsmith

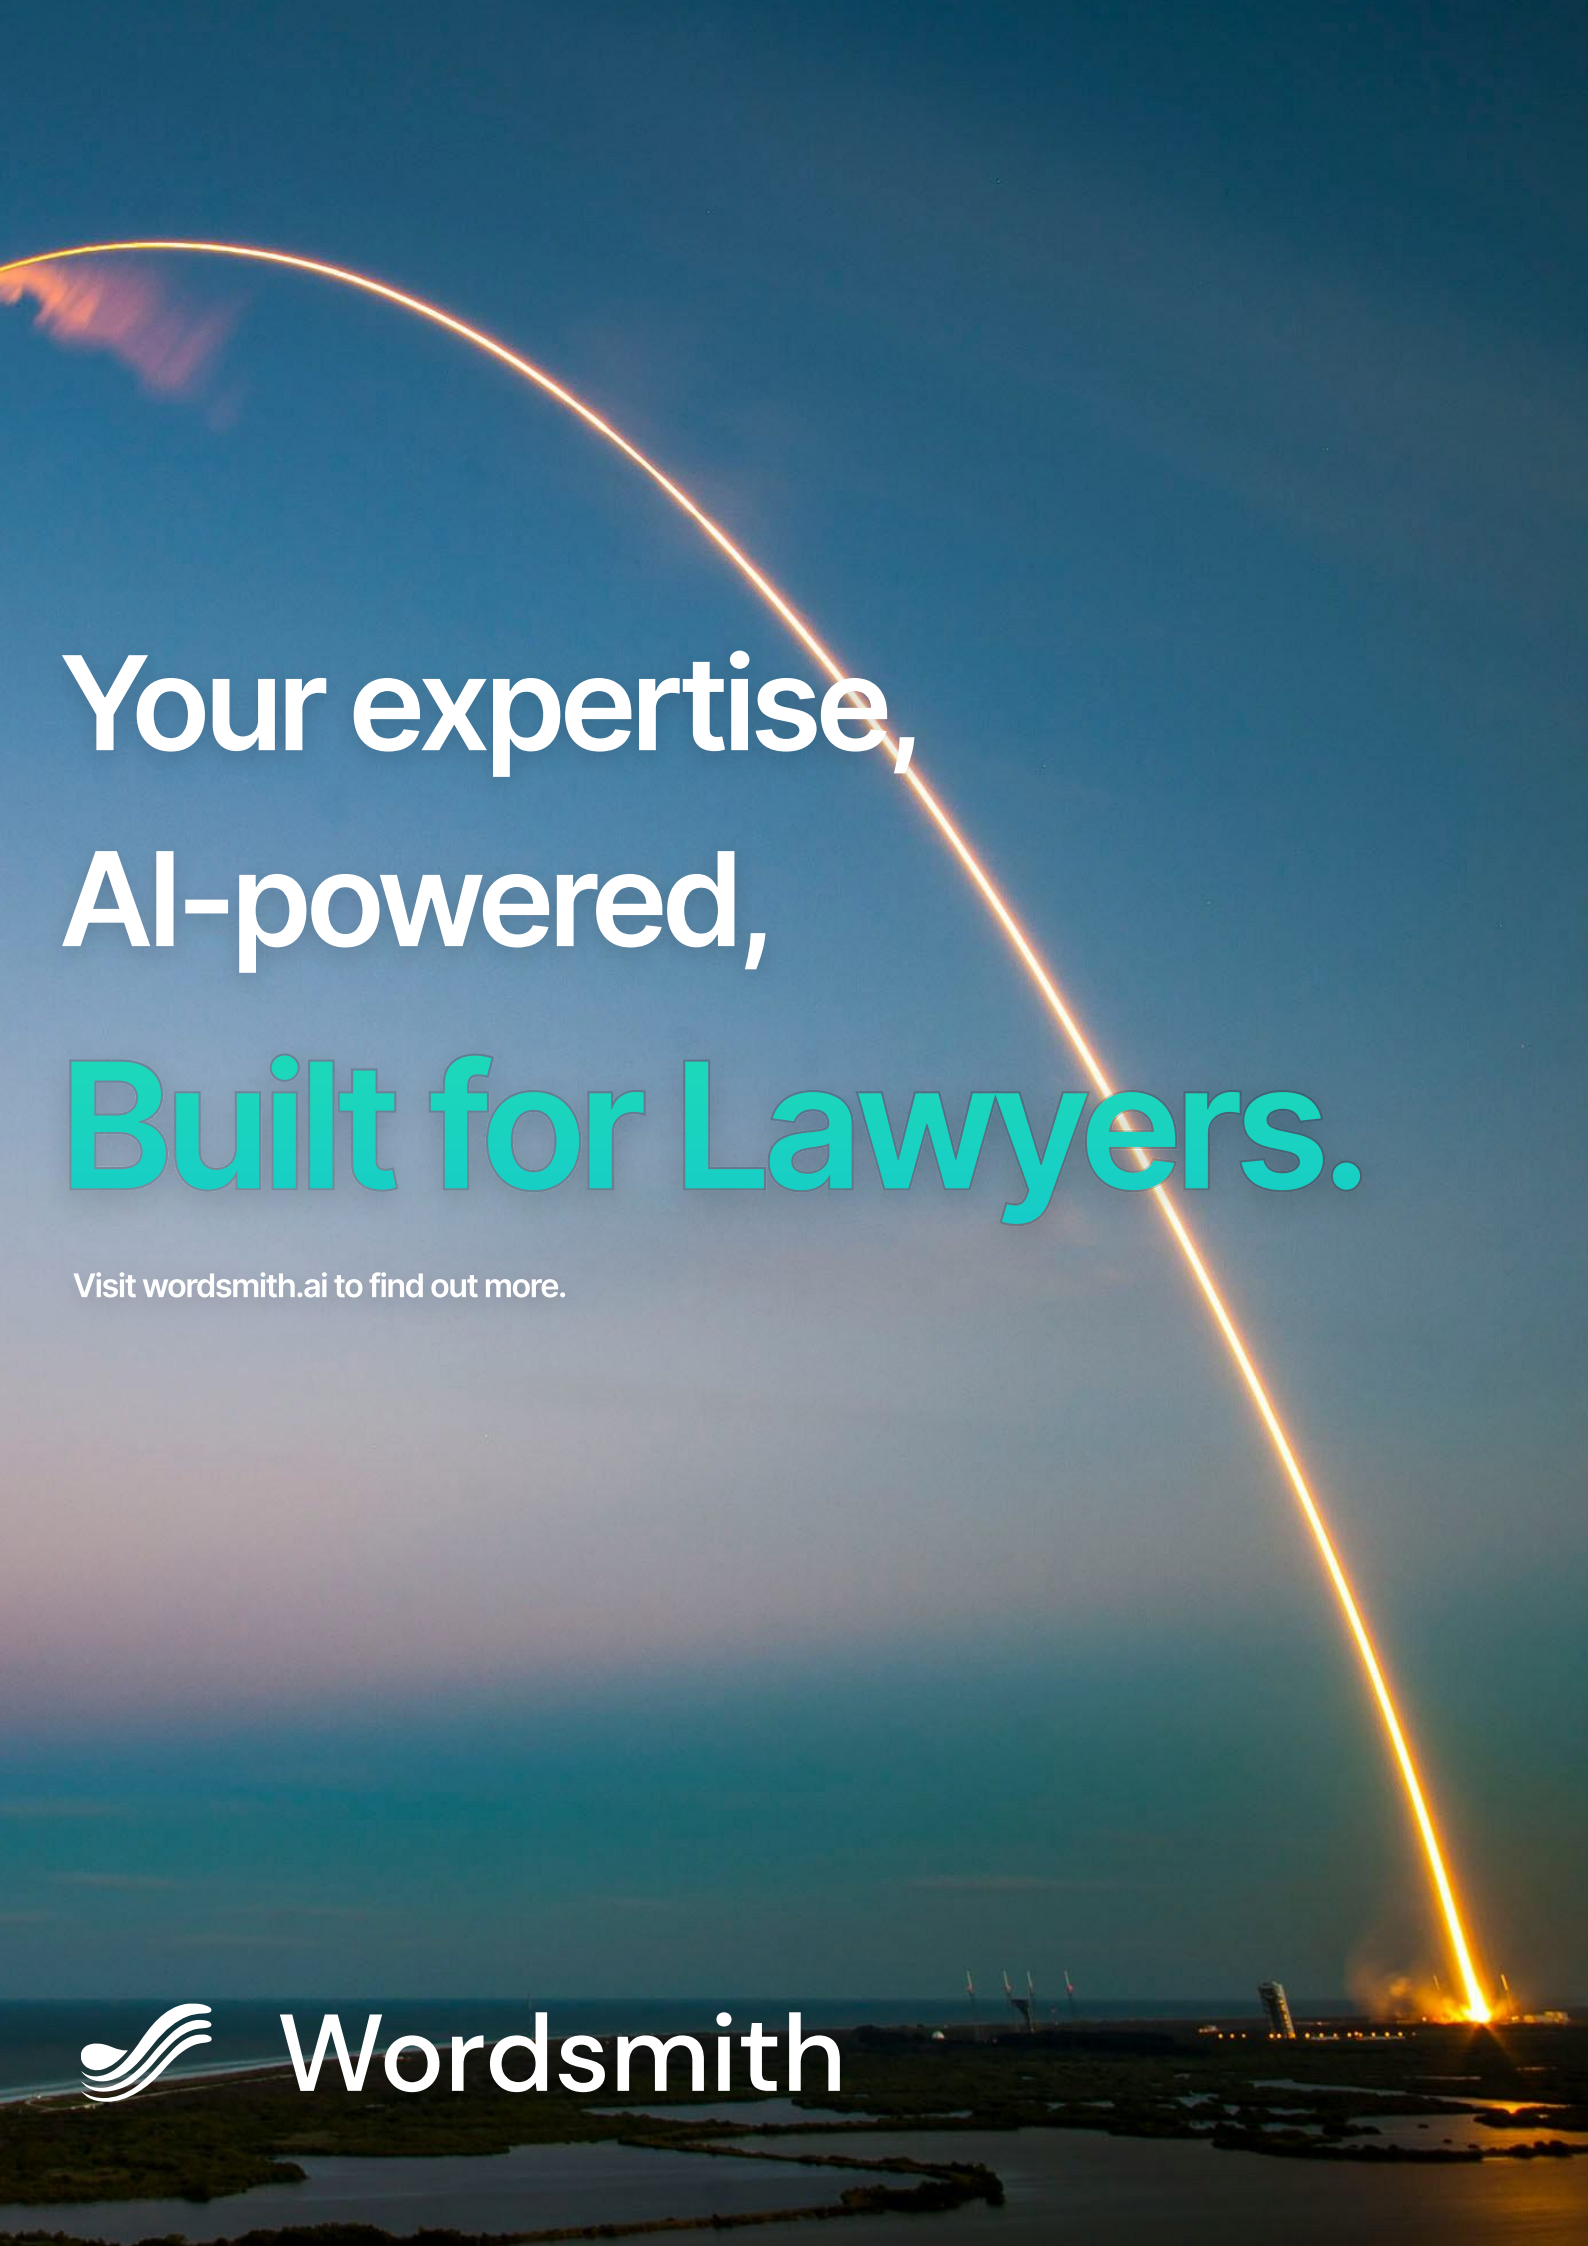
Your expertise,
AI-powered,
Built for Lawyers.
Visit wordsmith.ai to find out more.
Wordsmith

Contents

Overview and key issues 4

Q1: What is AI and how does it differ from other software/ systems? 5

Q2: How do generative AI systems like ChatGPT work? 6

Q3: What is prompt engineering in generative AI, and why does it matter? 7

Q4: What are the concerns when using generative AI systems? 8

Q5: Do I need client consent to use generative AI systems 12

Q6: A colleague used generative AI to review a document – should I be concerned?..... 13

Q7: Do I need to remove client information from documents before using generative AI? 14

Q8: Do client engagement letters need to be updated to deal with generative AI? 14

Q9: How does use of generative AI systems impact my insurance cover?..... 15

Q10: Which Practice Rules and Guidance are relevant to the use of generative AI? 16

Q11: What questions do I need to ask before purchasing /using generative AI technology?..... 19

Q12: How do I assess whether generative AI outputs are useful and appropriate?..... 20

Q13: What is automation bias and how do I guard against it? 21

Q14: What security issues might arise when using generative AI? 22

Overview and key issues

The use of AI, and in particular generative AI, has the potential to transform the way legal services are provided. The purpose of this guide is to provide answers to the key issues concerning the use of generative AI that are relevant to members of the profession.

In particular, this guide will help members make informed decisions about how to safely incorporate the use of generative AI products into their legal practice.

The guide also aims to allow members to understand the key issues surrounding the use of generative AI, namely:

- The need for good governance
- The need for care around the use of any confidential, personal or sensitive information
- Why it is necessary to carry out a risk assessment around data protection and information security requirements in respect of any generative AI system
- Why proper review and oversight of any output of generative AI systems is important, particularly given the risks of bias, inaccuracy or unfairness in the output
- When it is appropriate to inform a client or seek a client's consent
- What other risks exist when using generative AI

Question 1:

What is AI and how does it differ from other software/systems?

AI itself is an umbrella term which refers to a suite of technologies capable of performing tasks or operations that would otherwise require human involvement.

Common examples of AI-enabled technologies include:

- Predictive text and translation tools
- E-commerce platforms which analyse customer preferences and behaviour and make recommendations
- Chat-bots used on websites to facilitate communication with customers
- Certain document management and information retrieval systems; and
- Fraud detection systems used in the financial services and insurance industries

In the legal profession, there is a wide range of tools and systems that use AI, such as tools for contract review and analysis, legal research, e-discovery, compliance monitoring and handling initial client queries.

Interest in AI is predominantly due to the large number of AI-enabled products that can be easily accessed by a wide range of people. This can produce results that are relevant to all manner of business and personal uses, including use by the legal profession.

The AI systems that gain the most attention use machine learning models and techniques that enable computers to learn from and make predictions or decisions based on data, without being explicitly programmed for a specific task.

Machine learning models "learn" from input data during training of the model. The performance of generative systems built on these models can be improved after model training e.g. by model re-training, fine-tuning of the model on domain-specific data, through prompt-tuning or prompt engineering. It is important to note that AI is not "conscious" and does not have "agency" as these terms are generally understood. While the output may seem like it has been produced by a human, it is achieved by machine learning models operating on data. At their heart, AI systems are probability based. As a result, the nature of the output it produces always has some possibility of error.

Question 2:

How do generative AI systems like ChatGPT work?

Generative AI usually refers to a particular form of AI that can be used to generate new content, such as text, images, videos and code, based on the input of users. This input is in the form of "prompts". These systems work by utilising the power of Large Language Models (machine learning models trained on very large amounts of text) to give a very realistic approximation of content that ordinarily could only be produced by humans. This guide will focus on generative AI. There are many different generative AI systems, some of which are publicly available (public generative AI), and others which are part of a closed system, either residing on a firm's own IT system or part of a cloud-based system that is ring-fenced to a particular firm (private generative AI). It is important to be aware that different systems may present different legal issues, depending on how they work and the contractual terms which govern their use. Additionally, generative AI systems can show abilities beyond the ambit of their original training or produce unexpected results. When a generative AI system produces an unexpected and incorrect result it is referred to as a 'hallucination'.

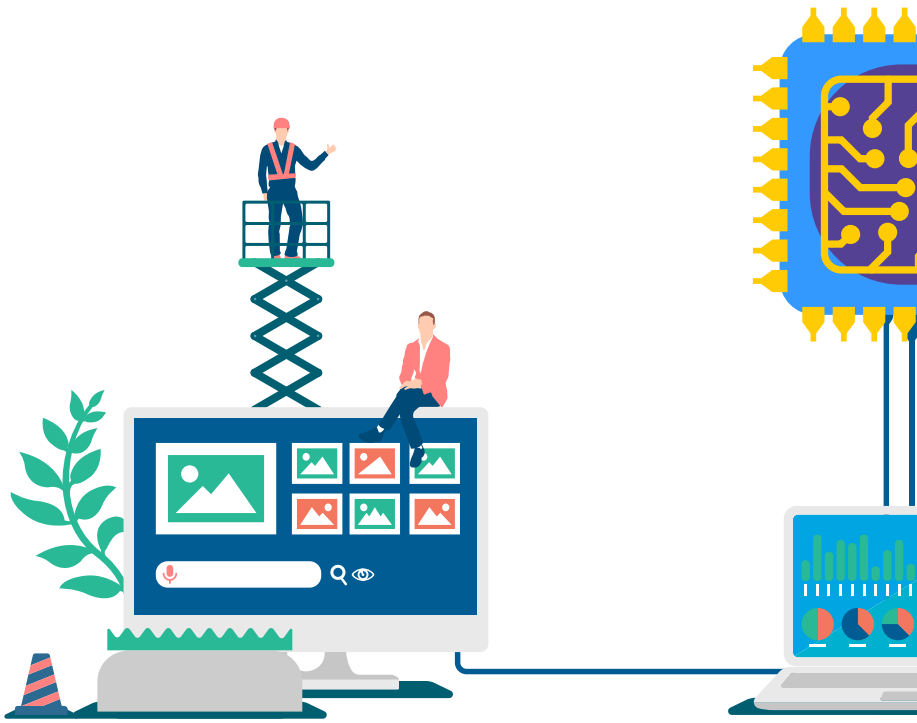

Question 3:

What is prompt engineering in generative AI, and why does it matter?

Users interact with generative AI by entering questions or instructions into a user interface. These inputs are known as 'prompts'.

The wording of the prompts matters; the system can only respond to the question or instruction that is input, not the question or instruction a user meant to input, or had in mind but did not clearly express. 'Prompt engineering' refers to the art of creating effective prompts – those which clearly communicate a user's intentions so that the system, in turn, generates the most useful responses. The output of generative AI systems can be heavily influenced by the content of prompts, which is why prompt engineering is now an area of focus in its own right.

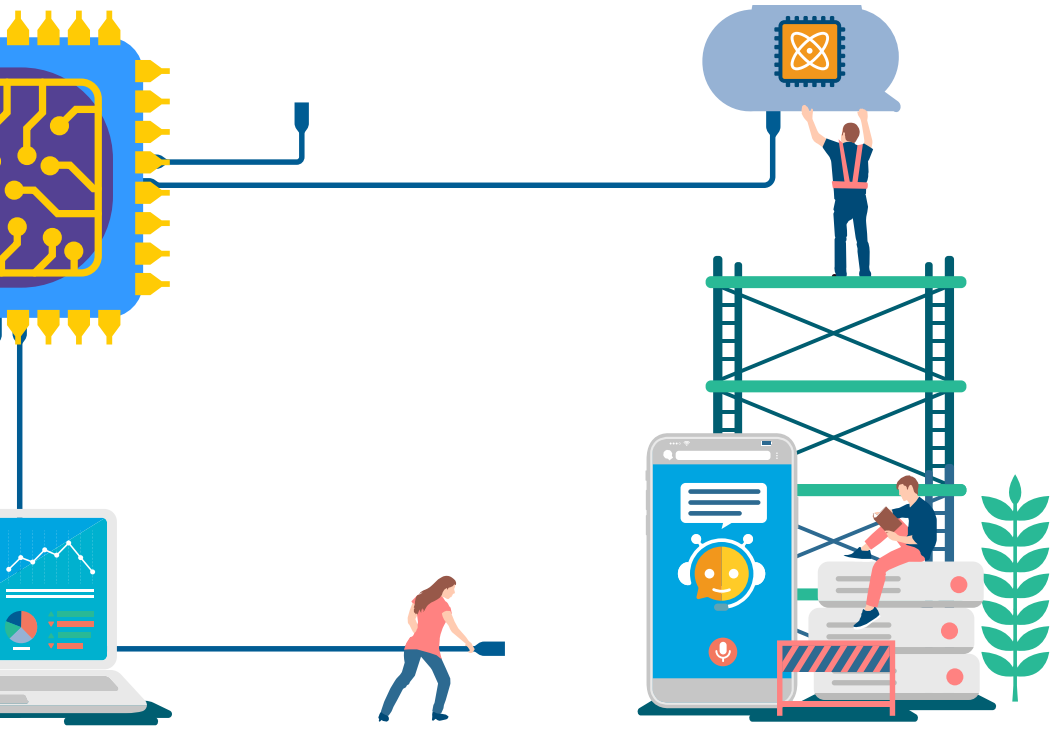

Question 4:

What are the concerns when using generative AI systems?

Among the key concerns and areas to be understood are:

- The terms of use of the generative AI systems
- Issues around quality, accuracy and freedom from bias regarding the output of the system
- Accountability for the use of generative AI systems
- What issues generative AI poses for client confidentiality, information security and data protection
- The use and ownership of intellectual property used and generated by the generative AI systems
- How to explain the workings of a generative AI system to parties who could be affected by its use
- Issues around sustainability and how those fit with a firm's own policies

Terms of Use of generative AI systems

Each generative AI system will be governed by its own Terms of Use. Some have been designed specifically with the legal market in mind, but many have not. It is important that any member of the profession using a generative AI system reads the terms and understands their implications. For example, certain generative AI systems may reserve the right to use any prompts you provide, and the answers generated, to further train the system. As a result, you should carefully consider what information you include in the prompts, taking account of confidentiality, data protection, intellectual property and other rights. In addition, some generative AI systems disclaim all and any warranties for the output, so if you were exposed to a claim because the output was inaccurate, breached intellectual property, data protection or other rights, your rights of redress against the system provider may be limited or wholly excluded.

Accuracy, bias and hallucination

One of the major challenges with the use of generative AI is that there is no certainty that the output is correct. There have already been reports of individuals relying on generative AI output which has "hallucinated" responses, such as making up case citations or omitting material information in the output.

While the output of generative AI systems may offer a useful starting point, care, skill and knowledge is required to ensure that any advice incorporating generative AI output is accurate, does not omit anything material, and applies to the jurisdiction relative to the work undertaken and advice provided for your client. Generative AI systems will base their output on all information available to them, and in relation to legal questions, it is not always clear whether the output provided by a generative AI system is appropriate to the relevant jurisdiction. Asking the generative AI systems if the output is applicable to your jurisdiction still runs into the same issue, which is there is no way of being sure it is accurate without further checking. While it may be safer to ask generative AI to perform certain tasks indirectly associated with the legal advice,

Concerns (cont.):

such as summarising information, using generative AI for the core legal advice is like relying on an unqualified member of staff – the output may be very good, or it may not be. Ultimately, it requires to be checked by someone legally qualified.
It should also be noted that the output of generative AI is only as good as its training data and the prompts that are used. Generative AI is trained on the basis of the information which is provided to it and if a generative AI systems learning is based on any inaccurate or outdated information, the answers provided by it are also likely to be inaccurate or outdated. This feature can also lead to issues of bias. There have already been well publicised examples of errors or bias, such as facial recognition systems which have a much higher error rate for dark-skinned women than light-skinned men, or AI powered systems which exhibit gender bias in supporting hiring decisions. If generative AI has been trained using data which contains hidden biases or to reflect certain principles or policies, those will influence the generative AI output and may give rise to output which is not impartial or objective. Generative AI systems have developed advanced countermeasures for various forms of bias in outputs, but this remains a risk to be aware of and to guard against.

While most generative AI systems will have been trained on vast amounts of data, many are only trained on data up to a specific date. As a result, the output of the system may not take account of the most up to date position in law. You should establish whether this is applicable to the model you are using. This issue can be mitigated to an extent by using generative AI systems which have "plugin" access to sources of information other than their original training data.
Taking all this into account, in using generative AI the question of how you will check the output is crucial. At present, the nature of the errors that generative AI could make are likely to be different from those that (say) a solicitor would make in researching or reviewing a bundle of documents. Therefore the methodology you use to check the work might need to be different from that which you would usually use – and the skills required may be even greater – given the lack of experience and training in properly reviewing generative AI output. If there is no way that you can be sure without doing the work yourself, then you need to factor that into any evaluation of what benefits come from the use of generative AI in the first place.

Oversight and accountability

If a firm were to rely on the output of a generative AI system, they would be exposed to potential claims for the advice given or services provided based on the output of the generative AI system. Therefore, extreme caution is required before any such use. The main point is that while such generative AI systems may be useful as a starting point, they cannot replace or minimise the need for proper oversight by qualified solicitors.

Firms that incorporate generative AI into their services must carefully consider the implications of accountability and liability. The ultimate responsibility for providing accurate legal advice rests with the solicitors. Irrespective of the systems used, solicitors should still exercise oversight and verify the accuracy and suitability of the information provided by generative AI systems.

When using generative AI systems provided by a vendor with a services contract in place, firms may have more information about the system and its data sources and may have appropriate liability terms in place. While this may give a firm more confidence to use such generative AI systems as opposed to the public generative AI systems like ChatGPT, whichever system is used, solicitors must recognise that they are ultimately accountable for the advice given, and exercise proper diligence and oversight in using generative AI to assist in their services.

Where firms intend to allow clients to directly use generative AI systems (e.g. via the firm's website) where there is no

Concerns (cont.):

oversight, it would be prudent for firms to explain, prominently, that this is not the provision of legal advice and that the firm takes no responsibility for the output.

Client confidentiality, security and data protection

Firms must ensure that they comply with duties of client confidentiality and do not disclose information to any third parties, including generative AI providers, which would breach those duties.
For this reason, confidential or client sensitive information should not be shared with public generative AI systems. Any documentation shared should have such information removed.

If a solicitor is considering sharing confidential information with private generative AI systems at a minimum they should undertake appropriate checks to satisfy themselves that (a) appropriate terms are in place with the vendor so that the information inputted will not be accessible by the vendor or used for any other purposes (b) the security arrangements meet appropriate information security standards and (c) that the use is compliant with the firm's own terms of business with clients.

In contrast to the terms typically offered on free to use generative AI, paid enterprise options will tend to have terms that address these points to some extent. Using technical means to "mask" the confidential information that is being shared with vendors may also be necessary as a way of addressing many of these concerns.

When using generative AI to perform tasks, solicitors must also comply with their data protection obligations. Any personal data which is provided to generative AI systems will be processed to provide an answer. Depending on the system in question, there is a risk that some personal data may also be stored and re-used by the generative AI system as training material to enable it to continue to learn and provide solutions for other users outside your firm. There is also a risk (known as "model inversion") that the generative AI system could be attacked in an attempt to reverse engineer or infer sensitive information of individuals if their data was used to train the model.

While it may be possible to use anonymised data, any time where an individual can be identified from any data, and particularly when generative AI is being used to review documents containing personal data, it is inevitable that the generative AI systems will be processing personal data.

Otherwise, the same considerations apply as with use of any personal data, which requires appropriate data protection terms to be in place with any vendor and the same checks and restrictions that would apply to any disclosure of personal data to a third party. Information must also be given to any clients whose personal data is being processed in this way to ensure that this type of processing is transparent and allows data subjects to exercise their rights in relation to automated processing of personal data.

Intellectual property

Generative AI systems will not discriminate between information which is protected by intellectual property ("IP") rights (and therefore cannot be copied or re-used without permission) and information which can be re-used. The reason this is an issue is that data which is subject to IP rights may be stored, copied and re-used by generative AI systems. Firms should therefore be aware that providing data which contains IP to generative AI systems could result in infringement of IP rights. Firms should ascertain how any IP provided to generative AI systems will be used, who has access to it and take steps to put in place agreements with generative AI providers to protect that IP.

Separately, generative AI output may also be based on the

Concerns (cont.):

unauthorised use of content which has been used to train generative AI. It is therefore possible that content which has been generated by a generative AI system includes content protected by IP rights. There are already multiple legal cases where rights-owners are suing generative AI producers for misuse of their intellectual property in training the systems. Firms should take steps to clarify what data the generative AI system has been trained on and have appropriate agreements in place with generative AI providers in relation to the protection of IP rights.

Explainability

Understanding and being able to explain how a generative AI system arrived at an output and what data it used is seen as an important safeguard for those who may be impacted by the outputs of such systems. While granular details of how precisely, for example, Large Language Models, operate are unlikely to be required, a basic understanding of the operations of the system should be obtained. If firms are procuring generative AI systems from third parties, this information will need to come from these third parties. Be aware that clients may be interested in receiving this information depending on the use you are making of the generative AI system.

Impact on the rule of law

The independence of solicitors is essential to the maintenance of the rule of law. Solicitors should be careful to ensure that their independence is not compromised by overreliance on the output of generative AI systems. They should also be aware of the systemic implications of use of these systems: wide-scale use of these systems in the domain of law necessarily increases the power of providers, including multinationals that already dominate the legal publishing market.

Sustainability

Generative AI uses a large amount of computing power, even though this is not necessarily obvious to the user who simply sees a response to a question they have asked pop up on their screen with seemingly minimal effort.

Significant computing power is required both to train the models on which generative AI is based, and to allow these systems to carry out their core function of generating content in response to prompts.

Any firm which is thinking of using a generative AI system may wish to consider the sustainability issues raised. Those can include the 'green credentials' of any proposed provider and consideration should be given to the issues raised in the General ethical and sustainability considerations section of the IT Procurement Guide.

Question 5:

Do I need client consent to use generative AI systems?

The first consideration is whether or not personal data, client confidential or sensitive information relating to the client would be inputted into any generative AI system.

If not, using generative AI to assist in the provision of legal services, whether this is through automating workflow or supporting the provision of advice, such as carrying out research, summarising documents or producing draft materials does not, of itself, require you to inform the client or obtain consent, **unless required by your terms of engagement with the client in question.**

While consent is not required in these circumstances, you may still wish to ask yourself whether the client would expect to know that generative AI has been used, and consider informing your client accordingly.

Where, however, you do intend to use personal data, confidential or sensitive information relating to the client with generative AI systems, further considerations apply.

First and most importantly, personal data, client confidential or sensitive information should not be inputted into any public generative AI system. If you are unsure about the nature of the generative AI system, you should assume it is a public generative AI system.

For private generative AI systems, personal data, client confidential or sensitive information may be used, provided that use of generative AI systems is not prevented or restricted by any client specific engagement terms, but only where the following conditions are met:

- You have confidentiality terms in place with the provider of the generative AI system, so that the data you input and the outputs will only be accessible by you and used for your purpose
- You have satisfied yourself as to the security controls in place, location of data and data transfer arrangements if relevant
- Any processing of confidential or personal data by these generative AI systems is consistent with the terms of your own privacy policy and/or data protection policy, complies with your obligation to maintain confidentiality and the applicable data protection regime.
- None of the data you input will be used to train other generative AI models other than the one which is private to you. Bear in mind that even the suppliers of private generative AI systems might wish to use some of the prompts or responses to train other models

As stated above, you may also wish to consider whether you should inform clients that you are using their confidential information with generative AI, to reassure them that appropriate safeguards are in place. In such case, you should understand the basic operations of such systems and be ready to provide answers to clients on the way their information will be used and secured.

In addition, you should check whether the terms that govern the use of a private generative AI system can be changed by the provider, and if so, whether this will impact your ability to use it or change the nature of client consent you need.

Question 6:

A colleague used generative AI to review a document – should I be concerned?

You should refer to the answers to What are the concerns when using generative AI systems? and Do I need client consent to use generative AI systems?

In addition to the information contained within those answers, key issues to be considered include:

- Was the system used in accordance with any firm policies and procedures on generative AI, and for the purpose it was intended?
- What system was used and what, if any, terms and conditions are in place with the provider of the generative AI system? In particular, what is the provider's position with regard to using any data which is input, or the output of the system, to train the model and can this be shared with other parties?
- Was any confidential information and/or personal data provided to the generative AI system and how will that be used by the system?
- What does the firm letter of engagement say about the use of generative AI?
- Does the client know or would the client expect generative AI to be used? If not, should the client be informed?
- Does such use involve a breach of confidentiality, rights in personal data, intellectual property or other rights of a counterparty to the contract?
- Has the generative AI output been reviewed by a solicitor?
- What reliance will be placed on the generative AI output by the firm and/or the client?

Clarity for solicitors and staff on the firm's approach to the use of generative AI is likely to be helpful in internally regulating use of generative AI, ensuring that standards are consistent and use of generative AI by the firm and its employees complies with the firm's duties.

Consideration should therefore be given to adopting firm-wide policies on the use of generative AI and training of staff to ensure that they know if and how generative AI systems can be used.

These policies would typically cover (a) an overview of generative AI and the risks of use (b) the purpose for which generative AI may and may not be used in the firm (c) guidance on how to use generative AI safely and appropriately, including care regarding the use of confidential information and the need to ensure human oversight of the output (d) the need to obtain client consent to the use of generative AI where appropriate (e) the need to have completed training and (f) consequences of non-compliance.

Question 7:

Do I need to remove client information from documents before using generative AI?

As stated in the answer to Do I need client consent to use generative AI systems? client specific information should not be inputted into public generative AI systems.

However, provided the conditions as set out in the answer to the above referenced question are met, you can use client specific information in private generative AI systems.

Question 8:

Do client engagement letters need to be updated to deal with generative AI?

Standard letters of engagement at present, frequently state that the firm may make use of third-party IT vendors and service providers. While strictly speaking this may suffice to also cover generative AI systems, firms are advised to consider updating these to make clear that this also includes the use of generative AI systems.

If the use of generative AI significantly alters the basis on which any work was originally scoped, you should consider updating the basis of your engagement. See Q10 Which Practice Rules and Guidance are relevant to the use of Generative AI? for further commentary on client communication generally.

If you start to offer services providing your client with direct access to specific generative AI systems, you should consider whether new standard terms are needed. For example, is it clear who is taking responsibility for the output and what the expectations are regarding use of the data which the system collects and produces?

In any negotiated (or client produced) letter of engagement, the specific terms should be also checked – they may impose their own restrictions on the use of generative AI systems without consent/notification, or such a restriction may be implied by virtue of the protocols that are to be followed when advice is produced - for example, the need to notify the client if any third party is involved in the delivery of the advice.

Question 9:

How does use of generative AI systems impact my insurance cover?

As regard's firms own insurance, you should carry out a review of terms to ensure coverage is adequate and whether any additional provisions are required in your engagement terms or your day-to-day practices when using generative AI systems, to ensure coverage is unaffected.

While generative AI systems are a step-change given the ease of use and wide-ranging impact (and hence the reason for this guide), the Master Policy insurers, RSA have confirmed:

Although Artificial Intelligence (AI) systems are beginning to emerge as an acceptable method of delivering legal services, their impact on the legal sector is still unknown and unpredictable. However, as with all other forms of technical tools and resources, where AI systems are used by solicitors to carry out work that is ordinarily undertaken by the firm, provided such usage complies with Law Society rules, then the work will continue to be covered under the existing terms of the Master Policy, again, subject to the terms and conditions of the policy.

It should be noted that the Master policy will not extend to cover any third-party supplier of AI systems or services and law firms will want to assess the liability and insurance cover related to the use of any tool and undertake appropriate due diligence on suppliers.

If you have any concerns regarding the insurance position, you should contact Lockton to raise those with them. Lockton's contact details can be found at Master Policy - Scottish Solicitors.

Regarding other insurance policies held by your firm (i.e. non Master Policy certificates), you should carry out a review of terms to ensure that coverage is adequate and whether any additional provisions are required in your engagement terms or your day-to-day practices when using generative AI systems, to ensure coverage is unaffected.

Question 10:

Which Practice Rules and Guidance are relevant to the use of generative AI?

B1.3 Independence

Summary:

You must give independent advice and have a duty not to allow your independence to be impaired.

Guidance:

This means that when you give advice you must not defer to the outputs or recommendations of a generative AI system. You must not allow the system to dictate your advice, the process by which you reach the conclusions relevant to your advice or the course of action you have taken on behalf of your client. Generative AI outputs should generally be checked and amended as appropriate by qualified persons, and any areas of greater reliance on machines specifically identified.

Examples:

You use a generative AI system to produce the first draft of a contract. You check the draft against your client's instructions, your firm's bank of styles or contract playbook, and consider whether recent changes in the law mean that certain clauses should be changed to achieve the desired legal effect and/or allocation of risk. You use or amend the draft in line with your assessment. In this situation your advice is independent of the outputs or recommendations of the system.

You use an AI system to carry out large-scale document review for your client. You provide advice about the issues flagged by the system but do not carry out human review of the entire set of documents. In this situation your advice concerning the risks that are flagged may be independent. However, to the extent that you offer advice on the risks associated with the suite of contracts as a whole, your advice is not independent of the system. You do not know what the system may have missed. In this situation you should inform your client about the use of the system and make clear the scope of the work that you are agreeing to carry out. You should also consider whether you need to obtain your client's consent for the use of the system and whether you have outsourced the work or an aspect of the work to the provider of the system (see also Section E, Division B: The Management of Files, Papers and Information (Outsourcing) below).

B1.6 Confidentiality

Summary:

You must satisfy yourself that your use of the generative AI system will not compromise client confidentiality.

Guidance:

You should:

- Make sure you have a written contract with the provider of the generative AI system
- Check that the contract provides clear assurances that inputted information will be treated as confidential and not used or disclosed to third parties
- In many cases the providers will contract on standard terms. You must check the standard terms including what they say about security, confidentiality, ownership of and access to

Rules and guidance (cont.):

data and satisfy yourself as to their suitability for your intended use case

- If the data input into the system includes personal data for which you are the data controller, you will be required to comply with data protection legislation. Where you are the data controller and the provider of the generative AI system is the data processor, you will be required to enter into a written data processing agreement with the provider of the system. As to the suggested requirements of data processor agreements, see the guidance contained in the Cloud Computing Guide. You should also consider whether data is being transferred to third countries (i.e. those outside the UK and/or European Economic Area (EEA))

Example:

The use of 'free to use' variants of public Large Language Models such as ChatGPT poses particular risks as regards security, confidentiality and data protection. For example, OpenAI informs users that information input into the ChatGPT interface are, by default, used as training data, may be reviewed by OpenAI and disclosed to 'affiliates, vendors and service providers, law enforcement, and parties involved in Transactions.' Different terms apply to (paid for) enterprise terms for access to these systems.

B1.10 Competence, diligence and appropriate skills

Summary:

You must only act in those matters where you are competent to do so. You must only accept instructions where the matter can be carried out adequately and completely within a reasonable time. You must exercise the level of skill appropriate to the matter.

Guidance:

These obligations cannot be delegated to a generative AI system. To the extent that you rely on a generative AI system for the purposes of advice to a client, the requirements for competence and exercise of appropriate skills means that you must be capable of independently assessing whether the outputs of the system are useful and appropriate.

Example:

You use a generative AI system which claims to predict the outcome of cases. It is impossible to know why the system makes the prediction it offers, though the system might be able to indicate e.g. which of several clusters of words used to train the system are most relevant to its predictions, or which words or series of words are strongly associated with the predicted outcome. If, in giving advice to the client you defer to the output of the system, you are not exercising the level of skill appropriate to the matter. If you report the output of the system to the client, then in the exercise of the appropriate level of skill you must be able to offer an adequate explanation of how the system works, its capabilities and limitations, including e.g. its reliance on training data, on statistical inferences, its inability to engage in legal reasoning or argumentation, that it does not possess understanding including understanding of text, that it is incapable of anticipating how law might develop, that the accuracy metrics reported by the system provider may be unrelated to the task of prediction of outcomes of cases that are still to be heard. Any such report must be accompanied by your own independent assessment of the prospects of success in a case.

B4 Client Communication Generally

Summary:

Your terms of business letter should specify what work is being taken on and who is primarily responsible for doing the work.

Guidance:

A distinction may be drawn

Rules and guidance (cont.):

between using a generative AI system to carry out aspects of the work and using generative AI merely to augment your own analysis. You should consider whether you need to obtain your client's consent for the use of the system and whether you are outsourcing the work to the provider of the generative AI system. If you use generative AI only to augment your own analysis and work you should consider informing your client about use of the system. As a matter of fairness, you should inform clients about the use of a generative AI system in the context of delivery of legal services when such use materially impacts on the cost of your provision of services to the client. Your fee arrangements must be fair and reasonable (see B1.11: Professional Fees) and should be transparent (see Section E Division A: Standards of Service).

Example:

You use a generative AI system to carry out large-scale document review for your client. You provide advice about the issues flagged by the system but do not carry out human review of the entire set of documents. You do not know what the system may have missed. In this situation you should inform your client about the use of the system. Your terms of business must make clear the scope of the work you are agreeing to carry out.

Section E Division B: The Management of Files, Papers and Information (Outsourcing)

Summary:

It is important for members to carefully consider the option to outsource any part of their operational functions or service provision.
At this point, it may be helpful to define two classifications of 'outsourcing' - business processes and legal processes. For example, 'business process' outsourcing includes human resources, cashroom and accounting, IT services/support and 'legal process' outsourcing includes document production and review, legal research, drafting motions and briefs, commercial contracts, legal due diligence and litigation support.

While outsourcing can provide additional resources and expertise, it has to be remembered that this does not remove the regulatory obligations imposed by the Society on its members. In addition, the management of outsourced providers requires specific skills and resources, including sourcing appropriate providers, contractual arrangements and project management and reporting. It is also important to consider its implications on service quality provision, client contractual arrangements and professional indemnity insurance requirements.

Guidance:

If you are outsourcing aspects of legal process work to the provider of a generative AI system you should:

- Make sure you have a written contract with the provider of the generative AI system. The contract should set out the full extent of the obligations and responsibilities of both parties
- Carry out due diligence on the generative AI system provider, including their financial viability and professional indemnity cover
- Where practicable, advise the generative AI system provider of your specific regulatory and compliance requirements
- Check the terms of your professional indemnity insurance and advise your insurers
- Consider informing your clients in your terms of business letter

Question 11:

What questions do I need to ask before purchasing/using generative AI technology?

As well as the general questions contained in the Guide to IT Procurement, the following specific questions should also be asked:

- How does the generative AI system work in practice? How is the data I input and the outputs of the system used?
- What supporting materials can the provider supply to help me explain its workings to clients, if required?
- What terms and conditions govern the use of the generative AI system? Particularly consider confidentiality, availability, Intellectual Property rights implications and data protection
- How does the provider manage/check the quality of the output from its generative AI systems?
- Where are the computers that run the generative AI system physically located (remembering that this may be separate to the location of other servers used by the same provider), and will I need to consider data transfer arrangements?

Question 12:

How do I assess whether generative AI outputs are useful and appropriate?

Unless the client has otherwise agreed, you are responsible for assessing whether the outputs of generative AI-enabled legal technologies are useful and appropriate. You should not take the output of a generative AI system (whether a recommendation, draft, summary, review or any other output) at face value. Generative AI systems are constrained by their training data, may output information that is incomplete, inaccurate, out-of-date or misleading, produce biased or harmful outputs and can be vulnerable to adversarial attacks.

How you go about assessing the outputs of a generative AI system will depend on the use case. For example, if you use a generative AI system for research you might check its outputs by comparing them with what you know from your own experience and from other resources (books, journals, case law) or systems. It may be more difficult to assess the outputs of a generative AI system used for large scale document review or contract analytics, particularly with regards what the system may have missed. In that case you should inform the client about your use of the system and explain its capabilities and limitations.

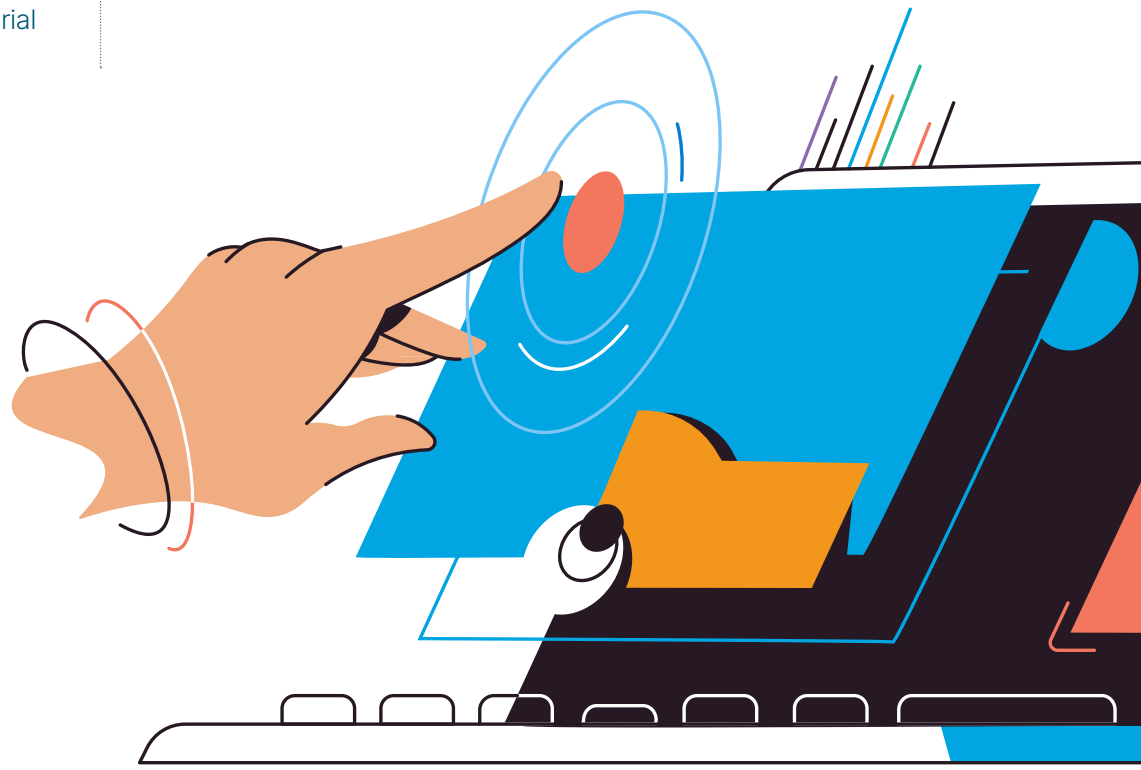

Question 13:

What is automation bias and how do I guard against it?

Automation bias refers to the tendency of humans to place too much reliance on the outputs or decisions of automated systems. This problem is not confined to systems that use generative AI – it is equally relevant for other software systems. However, the risk of overreliance is greater where, as in generative AI systems, the output of the system tends to have an air of plausibility.

Educate yourself

One way of guarding against automation bias and overreliance on the outputs of a generative AI system is to educate yourself about how the system works. Ask the provider for information about the system, make use of tutorials about how the system works, its capabilities and its limitations, attend relevant CPD courses. Continually evaluate the performance of the system.

Protocols

Another may be to put in place explicit protocols about checking the outputs of such systems, especially if these outputs will inform advice to clients, form part of the content of legal documents (e.g. contracts, pleadings) or communications with other solicitors, Counsel, or the courts.

System design

Some systems may be designed to reduce the risk of automation bias, for example, by warning users that the output is generated by a Large Language Model, may be inaccurate and does not constitute legal advice, or by prompting users to carry out independent checks to assess whether the output is useful or meaningful.

Read the contract

Do the providers of the system offer warranties about the reliability, accuracy and completeness of the system's outputs? It is very unlikely that they do because it is impossible for providers to ensure that generative AI systems behave in that way. Familiarity with the contract terms might go some way to discouraging overreliance.

Question 14:

What security issues might arise when using generative AI?

As stated in the answer to What is prompt engineering in generative AI, and why does it matter?, ensuring that the generative AI system complies with your information security policy is paramount.

At a basic level, generative AI systems create another potential route by which attacks may be made to a firm's IT systems. Firms need to ensure they have carried out their due diligence on the provider and the system, and that their own IT systems are connecting to the generative AI systems via validated secure protocols. In effect, all of the same security protocols that firms would deploy when connecting to any other cloud vendor should be followed.

As well as the general information security risks, specific issues that apply to generative AI systems include the concept of "injection attack", which can arise when outputs from generative AI systems are automatically fed into the other IT systems of the firm. Firms should have protocols in place to ensure care is taken in this regard.

Firms should also be aware that generative AI systems used by others can make 'social engineering' attacks easier, with very convincing content generated which may then be used to influence you or members of your firm. Some organisations have seen instances of attacks in the form of text messages, social media posts and convincing voicemail recordings created using generative AI purporting to be from senior personnel asking more junior personnel to reset passwords or send emails.

Authors

This guide was written by the Generative AI Working Group of the Law Society of Scotland's Technology Law and Practice Committee.

The working group will continue to review and update the guide as necessary and users of the guide should be aware that amendments may be made in the future.

Please send any feedback to: antonymcfadyen@lawscot.org.uk

Law Society of Scotland

The Law Society of Scotland
Atria One
144 Morrison Street
Edinburgh
EH3 8EX
T:+44(0) 131 226 7411

www.lawscot.org.uk