The Law Society

AI and lawtech

New Practice note

Using artificial intelligence in legal practice

Understand how existing professional duties and regulatory obligations apply when artificial intelligence (AI) is used in legal practice.

21 Oct 2026 • 1 hour 12 minutes read

On this page

1 Introduction
2 Key terms
3 AI in legal practice
4 Professional duties
5 AI governance
6 Regulation and compliance
7 Key points
8 Case studies
9 Training, resources and more information

Legal status
Terminology

This practice note is to support solicitors, and the teams they work with, to understand and consider the key issues when using artificial intelligence (AI) within legal practice.

It is intended to help you decide whether, when and how AI can be used safely, responsibly and consistently with your professional obligations.

It covers practical topics such as:

- how AI is being used in legal practice
- the risks associated with different types of AI and ways of accessing it
- how existing regulatory duties apply when AI is used
- approaches to AI governance within legal practice, including staff training and supervision
- AI's impacts on confidentiality, privilege, data protection, insurance and cyber security
- issues to consider before using AI systems or dealing with AI-generated materials

This practice note is for all solicitors and types of legal practice.

We will provide further guidance for segments of the profession that have distinct needs and challenges.

This practice note is the Law Society's view of good practice in this area and is not legal advice. For more information, see the legal status.

Who should use this practice note?

- Solicitors interested in using, or already using, AI in their practice

- Anyone looking to procure or develop AI for a law firm or in-house legal team, including solicitors, practice managers and people working in IT or innovation

1. Introduction

Use of artificial intelligence (AI) and generative AI (gen AI) is now common within legal practice.

Solicitors, including those within in-house legal departments and general counsel functions, and law firms are using AI to support tasks such as:

- legal and other research
- document review and disclosure
- drafting and client communications
- early case assessment
- comparing and analysing documents
- contract analysis
- litigation data analytics and intelligence
- due diligence
- compliance monitoring
- workflow management

Artificial intelligence offers potential efficiency gains and could help meet evolving client expectations.

However, AI systems can vary significantly in capability, security and intended purpose. This can cause uncertainty about how to responsibly use the technologies in legal practice.

The use of AI also raises important professional, legal and ethical considerations for solicitors.

1.1 Guidance in this practice note

This practice note clarifies AI terminology and aims to support informed decision making.

It explains how existing professional duties, regulatory obligations and legal requirements apply when AI systems are used in practice.

It does not seek to either promote or discourage the use of AI. It aims to help you to understand your responsibilities and exercise professional judgement.

This practice note cannot address every issue or scenario in depth.

Depending on your role, organisation and intended use of AI, you may need to supplement this guidance with:

- appropriate training
- organisational policies, procedures and governance arrangements
- specialist advice

2. Key terms

This note uses the following key terms:

2.1 Artificial intelligence

An umbrella term for the theory and development of computer systems that can perform tasks that usually require human intelligence, such as visual perception, speech recognition and analysis of information.

There are subcategories of AI, some of which are outlined below.

Agentic AI

A form of AI that operates with a high level of autonomy and reduced or intermittent human intervention. It can plan and execute multi-step tasks and coordinate across systems and multiple AI agents towards a set objective.

The main output of agentic AI is to carry out actions, rather than produce data.

AI agents

Specialised software programmes that use AI techniques and tools to perform tasks autonomously or semi-autonomously. It does so within a constrained, task-specific scope guided by prompts or human direction.

Chatbot

A digital tool or software application designed to simulate conversation with users. This is usually through text or synthesised speech.

Some chatbots use predefined responses. Advanced versions use AI, including agentic AI techniques, to provide more dynamic and relevant interactions.

Generative AI (gen AI)

A subcategory of AI that uses deep learning algorithms to generate new outputs in response to inputs.

Given the same prompt, generative AI systems may produce different outputs each time because they generate responses based on probability (drawing on statistical patterns), rather than following a fixed set of rules.

These outputs are produced based on large quantities of existing or synthetic (artificially created) training data. Outputs can include text, images, audio, video or computer code.

Generative AI systems are not search engines or databases.

For more information on generative AI, read our generative AI essentials guide.

Large language model (LLM)

A large language model (LLM) is an underlying technology that enables generative AI.

LLMs are AI systems trained on a large amount of data. There is no consensus on what constitutes ‘large’.

LLMs work with language. They use ‘deep learning’ (a type of machine learning) to predict how characters, words and sentences function together.

Sometimes the term ‘foundation model’ is used to refer to an LLM. However, foundation models can have broader functions than LLMs.

Some well-known LLM providers include Anthropic, Google DeepMind and OpenAI.

Machine learning (ML)

A subset of AI that enables computer algorithms to refine their performance by learning from data over time.

2.2 Automated decision making (ADM)

The process of making decisions through automated systems, sometimes through algorithms or AI, with little or no human involvement.

Under UK data protection law, the term ‘ADM’ applies only to decisions with a legal or similarly significant impact, made without any human involvement.

2.3 Cloud computing

Cloud computing means the delivery of computing services (including the provision of servers, storage, databases, networking, software, analytics and intelligence) over the internet.

Some well-known cloud computing service providers include Amazon Web Services (AWS), Microsoft Azure, Google Cloud and Oracle Cloud.

2.4 Data

Data is the information, documents and knowledge resources that AI systems rely on to generate outputs.

Input data

The raw information you give to an AI system or tool during use, which it uses to produce an output. For example, a prompt submitted to a generative AI system.

Output data

The information or results produced by the AI system based on input data and its training data.

Training data

The initial set of data fed into an AI system to teach it how to produce outputs, make predictions, recognise patterns or generate content.

Bias within training data may occur when the information contains systemic errors, gaps, human prejudices or is based on unrepresentative samples.

2.5 Hallucinations

AI hallucinations are incorrect or misleading results generated by AI systems.

These errors can be caused by a variety of factors, including:

- insufficient or low-quality training data
- the model's statistical nature
- incorrect assumptions made by the model
- ambiguous, misleading or insufficient user prompts
- biases in the data used to train the model

3. AI in legal practice

There is currently no AI-specific regulator for legal practice.

The use of AI by solicitors is governed through existing legal, regulatory and professional obligations.

This includes the Solicitors Regulation Authority's (SRA) Standards and Regulations and the SRA's implementation of the UK's AI Regulatory Principles within its remit.

Some regulators have oversight of the development, implementation and use of AI more broadly, such as the Information Commissioner's Office (ICO).

The use of AI by solicitors may also trigger other existing legal and regulatory frameworks, such as data protection, equality and consumer protection legislation.

3.1 Regulator AI resources

Legal Service Board

The Legal Service Board (LSB) published its 2026-27 plan for AI-driven innovation (PDF 1.2 MB) in June 2026 and AI strategic plan (PDF 296 KB) in April 2024.

Both resources are directed toward the role and responsibilities of the regulators within its remit, including the SRA.

The LSB also issued statutory guidance on innovation and technology in April 2024, which set out three outcomes for regulators to ensure that regulation:

1. enables the use of technology and innovation to support improved access to legal services and to address unmet need
2. balances the benefits and risks, and the opportunities and costs, of technology and innovation in the interests of the public and consumers
3. actively fosters a regulatory environment that is open to technology providers and innovators

Solicitors Regulation Authority

The SRA published a warning notice on the misuse of AI in August 2026.

Warning notices are intended to alert regulated law firms and solicitors to emerging risks. They also clarify professional obligations and the expected standards that must be met to avoid regulatory action.

The misuse of AI warning notice sets out two areas of concern:

- AI hallucinations which lead to false or incorrect information (including citations) within legal work and submissions to the court
- inputting of confidential client information and clients' personal data into AI systems that do not have the contractual and technical safeguards to maintain client confidentiality

The SRA also sets out three impacts of not paying due consideration to these issues on consumers and the profession:

1. harm to clients where poor or inaccurate information or advice is given, resulting in a poor outcome in their matter
2. negative impacts in the courts, where legal professionals use incorrect or fake case law/citations generated by AI and not checked or verified before submission
3. harm to the standing of the profession from breaches of professional obligations and regulations through AI misuse, with high-profile occurrences risking damage to public trust and confidence in the legal profession

The SRA regulates solicitors and authorised firms. It does not:

- regulate technology
- authorise, certify or approve AI systems or other software

For more information on the SRA's approval of 'AI law firms', see the section on professional duties.

The SRA does not approve any AI systems to provide regulated services.

The first SRA approved AI-driven law firm is authorised to deliver narrow and standardised areas of law.

Named regulated solicitors retain accountability for outputs and compliance with professional standards.

In November 2023, the SRA published its risk outlook report focused on the use of AI in the legal market.

The report clarifies that the SRA focuses on the outcomes that solicitors' and firms' actions produce rather than the systems they use.

Supporting innovation and technology adoption within the profession is also one of its strategic priorities.

3.2 Distinction between legal tech and lawtech

The terms 'lawtech' and 'legal tech' are often used interchangeably. However, we draw a distinction between the two terms.

'Legal tech' is typically used to refers to technology specifically made for solicitors and legal professionals.

'Lawtech' is a broader term which refers to everyday tools used by solicitors and legal professionals, which other professionals also use.

Lawtech also covers technology that enables access to justice and legal innovation. This includes client-facing tools such as chatbots and legal advice platforms.

Lawtech is our preferred term because it encompasses all the technologies used in a legal context.

For more information on lawtech, see our:

- introduction to lawtech guide
- practice note on using lawtech

3.3 Uses of AI in legal services

The choice and use of AI systems may differ according to the:

- task
- area of law
- type of client or matter involved
- size, structure and risk appetite of the firm or organisation
- level of available investment

AI can be used in legal services in a variety of ways. It is important to distinguish between:

- AI used within the delivery of legal services – the practice of law
- AI used within the operational, administrative and commercial aspects of delivering legal services to clients – the business of law

AI used within the practice of law could include support with legal research, drafting or document review.

AI used within the business of law could include support with billing, matter management, marketing and firm administration.

This distinction is important because the level of risk, oversight, client disclosure and governance required will depend on how the AI system is being used.

AI systems used in the practice of law engage your professional regulatory obligations more directly than AI systems used solely in the business of law.

The uses of AI in legal services may change as AI technology develops. Below are some non-exhaustive examples.

Business of law

The commercial and operational activities that support the delivery of legal services, such as billing, matter management, marketing and firm administration are not regulated legal activities.

AI systems may be used to assist with content creation, client intake and tracking legal and regulatory developments.

AI systems can also help with general business productivity or efficiency, such as by:

- assisting with scheduling and diary management
- searching for internal content
- reviewing the tone of communications
- supporting recruitment and people management
- facilitating meetings
- indexing material, including some client information

AI systems can be applied across both the practice of law and business of law together.

For example, where AI-driven insights into the profitability of a single legal matter or case (the business of law) inform how legal work is resourced and delivered (the practice of law).

Client engagement

AI systems can, with appropriate safeguards, support client engagement through tools such as chatbots, virtual assistants and automated communication platforms.

These AI systems can improve responsiveness and client experience while freeing up time for more complex work by:

- responding to routine queries
- assisting with onboarding
- providing updates

Contract review

AI enabled contract review tools analyse agreements to identify key clauses, deviations from standard terms and potential risks.

Such tools can assist in reviewing contracts more quickly and consistently by:

- flagging inconsistencies
- extracting relevant information
- comparing documents

Document summarisation and transcription

AI systems can summarise, at speed, large volumes of text such as case files, correspondence or legal documents.

Shorter, structured outputs can help to quickly understand key issues, facts and arguments, particularly when dealing with extensive or complex materials.

AI systems can also help transcribe client meetings and interviews, which can be converted into an accurate summary of what was said (following human review).

Due diligence and risk assessment

AI systems can support due diligence by reviewing large datasets to find risks, anomalies or patterns that may not be immediately apparent.

In transactional and compliance contexts, this can include:

- analysing documents
- screening entities
- highlighting areas requiring further investigation

E-discovery and e-disclosure

Electronic discovery (e-discovery) and electronic disclosure (e-disclosure) are the processes of identifying, collecting and reviewing electronically stored information for legal work.

AI is widely used in e-discovery and e-disclosure processes. AI can help categorise, prioritise and analyse large volumes of documents to identify relevant material.

Legal document creation and automation

AI systems can assist with legal document creation by generating initial outputs based on templates, prompts or structured data.

AI systems can automate routine drafting tasks, such as producing standard agreements or correspondence, improving consistency and reducing time spent on repetitive work.

Legal research assistance

AI-powered legal research tools can complement traditional research methods by:

- identifying relevant case law, legislation and commentary using machine learning and large legal datasets
- surfacing key authorities and examining briefs by analysing large volumes of legal information
- providing conversational or query-based support, helping to locate relevant materials, navigate legal databases and synthesise findings

Predictive analytics

Predictive analytics tools use historical legal data to identify trends and estimate potential outcomes, such as:

- litigation risks
- potential financial liabilities
- case trajectories

These insights can support strategic decision-making. They should be used with caution and alongside professional judgement.

3.4 Uses of AI in legal practice and risk

Different uses of AI within both the practice of law and the business of law pose different risks.

The risk of each application of AI will depend on many factors, including:

- the factual circumstances of its use
- the governance framework established by your firm or organisation
- individual adherence to guidance and policies

Other firms and organisations you may work with will have their own policies and risk management processes, and different sectoral or jurisdictional requirements that will need to be considered.

Activities within the practice of law are subject to professional regulation and may fall within reserved legal activities, which can only be performed by authorised (human) persons.

You should apply an appropriate level of human involvement and safeguards to mitigate risk, for both business of law and practice of law use cases, based on the level of risk posed by the particular AI use.

Reliance on AI systems and risks

Reliance on third-party AI systems can create operational and commercial risks.

An AI provider may change its business model, alter its terms of service, discontinue products or features, be acquired by another organisation, become subject to legal or regulatory restrictions, or cease trading altogether.

These events may affect the availability, functionality, security or cost of AI systems that you or your organisation rely upon.

Dependence on a single AI provider, underlying model or cloud infrastructure may create concentration risk. A disruption affecting that provider could affect multiple systems, workflows and legal matters simultaneously.

Your organisation should consider whether it could continue to deliver legal services if an AI system became unavailable at short notice. This may include considering:

- whether alternative providers or solutions are available
- how data, records and work could be retrieved or transferred if the service ends
- whether contractual exit provisions are adequate

- whether staff retain the knowledge and capability needed to complete work without AI assistance
- how changes in ownership, control or investment of an AI provider could affect your organisation's risk profile

Overseas AI system providers

Reliance on AI systems developed, hosted or supported outside the UK may create risks.

Access to these systems could be disrupted by changes in external factors such as law, regulation, trade relationships, sanctions or export controls.

There may also be concentration risk where several AI products used by your organisation rely on the same overseas provider, model or infrastructure. A failure or withdrawal of that shared service could affect multiple systems and legal workflows at the same time.

Your organisation should assess whether reliance on an AI system could affect your operational resilience and ability to continue delivering legal services.

You and your organisation may wish to consider:

- where the provider, underlying model, infrastructure and support services are located
- how changes in other jurisdictions could affect access, support, data handling or contractual rights
- how data and records can be retrieved or transferred if the service ends
- whether staff retain the knowledge and capability needed to continue work without the AI system

Effective AI governance is one way to help mitigate operational risks. For more information, see the section on AI governance.

3.5 Liability for AI usage in the legal profession

The UK Jurisdiction Taskforce's (UKJT) legal statement on liability for AI harms under the private law of England and Wales provides clarity on whether and in what circumstances those who have not set out deliberately to cause harm may be liable for harms resulting from the use of AI.

The legal statement is not intended to be legal advice, nor is it binding. The UKJT seeks only to clarify the likely application of the law to novel scenarios.

We are providing only a factual summary of the UKJT legal statement within this practice note.

It sets out two legal bases on which English and Welsh law will impose liability on a person for loss caused by using AI where that person has not intended to harm.

- Firstly, through responsibilities that a person has voluntarily taken on. Contract is a common legal mechanism by which a person voluntarily takes on responsibility for a risk
- Secondly, through liability imposed by law. This type of liability may be attributed regardless of whether a person has chosen to take on that liability

Existing professional, regulatory and legal obligations continue to apply.

Liability cannot be transferred to an AI system as it does not have a separate legal personality.

The legal statement outlines three main ways that liability may be imposed by law.

1. Under the law of negligence: a person is liable where a duty of care is owed, where that person fails to meet the required standard of care, and where this failure caused foreseeable harm
2. Vicarious liability: an employer can be held vicariously liable for AI-related harm if that harm arose because a human employee acted wrongfully while using AI
3. Professionals, such as solicitors, have an obligation to exercise reasonable skill and care in carrying out their duties. This obligation applies to their use of AI just as it does to any other aspect of their work. 'Reasonable skill and care' is determined by the standards ordinarily exercised by competent members of that specific profession, informed by expert evidence and the guidance of professional bodies

The UKJT explains that a professional may be found negligent for using AI inappropriately, using an unsuitable model, failing to conduct proper due diligence or failing to test AI or validate its outputs effectively.

Potential liability for failing to use AI

The UKJT also states that a professional could be liable for failing to use AI in circumstances where a competent member of their profession would have done so.

The legal statement expresses only that there is a need to be "aware of the possibility that a professional could also be liable for failing to use AI…".

Any finding of negligence for failure to use AI will depend on whether professional bodies' regulations and/or guidance suggests that a comparable reasonable professional should have used AI in that context.

For legal professionals, this would also include the court's assessment of what a reasonable legal professional carrying out that task should have done.

There is no guidance from us or the SRA to declare instances in which solicitors should use AI, and none is planned at this time. We are not aware of any existing or planned positive duties for regulated solicitors or law firms to use AI in legal services.

If any guidance were to be developed on when you should use AI, it would need to take into consideration the range of attitudes on complex issues such as the:

- impact of AI on the climate and environment
- use of technology developed and controlled within other jurisdictions
- training of junior solicitors

Liability and false statements

The legal statement also addresses where liability may be allocated for false statements made by an AI chatbot.

In general, liability for negligent misstatement requires a statement to be made by or on behalf of a legal person, or for a legal person to 'adopt' the statement.

AI is not a legal person. There is no current UK legal authority on the question of whether AI can make a statement 'on behalf' of a legal person.

The question was addressed in the Canadian case of *Moffatt v. Air Canada*, 2024 BCCRT 149, where Air Canada suggested that a chatbot was a separate legal entity that was responsible for its own actions.

Whilst not binding in this jurisdiction, it is relevant to consider the tribunal's finding that "[i]t makes no difference whether the information comes from a static page or a chatbot" when considering liability for negligence.

The UKJT explains that, under the law of England and Wales, liability could be established if:

- a legal person had an AI chatbot communicating on their behalf, or
- if there is an express or implied representation that the chatbot's statements are correct

The core negligence may be the careless acts that permitted the output, such as human decisions behind the AI system's design, testing and deployment. However, this would require factual analysis of the specific scenario.

3.6 Client information requirements

Your organisation's client care letter, terms of business or retainer letter sets out the information about how you will provide your legal service to a client.

Your terms of business might impose additional requirements or restrictions, particularly in relation to:

- client confidentiality
- data sharing
- third-party suppliers
- supervision
- the client's expectations about how their matter will be handled

In your client care letter, consider being transparent about:

- the extent to which AI is used
- the potential impact of a client's own use of AI on the handling of their matter

This practice note covers suggested minimum requirements and suggested actions.

You should always consider:

- any additional obligations that may arise from your terms of engagement with specific clients
- what is in your clients' best interests

Professional duties always come first. These duties are explored further in the professional duties section.

3.7 Types of AI and risk

When assessing AI risk, be aware that risk levels are not fixed. They vary depending on the circumstances of your use of AI, including:

- the type of AI system you are using and how you access it
- the specific task you are using it for
- decisions made about AI adoption, deployment and governance within your organisation
- your adherence to training, governance processes and policies

Generative AI

In legal work, AI-enabled outputs from generative AI systems may include text, summaries, draft clauses, correspondence, research notes or other materials that appear authoritative.

Generative AI does not understand inputs and outputs in the same way a human does. It cannot reason, or consistently and reliably check the accuracy of its own results alone.

Generative AI is non-deterministic, which means:

- it can produce different outputs or take different paths even when given identical inputs on separate occasions
- outputs cannot always be reliably reproduced or independently verified by simply re-running the same query

Testing or validating a generative AI system once does not guarantee it will behave identically in future.

Any outputs from generative AI need appropriate human verification, oversight and professional judgement.

For more on the key risks of using generative AI in legal work, see our guide on generative AI: the essentials.

AI-powered chatbots

AI-powered chatbots within legal practice may:

- be embedded into websites, client portals, case management systems or legal technology products
- operate with varying levels of autonomy, from simple scripted responses to more advanced generative AI systems capable of producing unscripted responses

Chatbots may support client triage, onboarding, routine enquiries, information gathering, internal knowledge management or the delivery of legal information.

Chatbots may improve access to justice or legal support for some clients. But they can exclude those who are less digitally literate or confident.

The risks of using AI-powered chatbot tools depend on:

- what service they provide
- what data they process
- the human oversight of the chatbot and its outputs
- client understanding of the chatbot outputs and potential misunderstanding or misinterpretation

The use of a chatbot to deliver any part of legal services does not transfer or diminish your professional responsibilities.

You and your organisation remain responsible for legal advice provided, even if the advice was provided by a chatbot.

Where chatbots form part of a service that provides legal advice, you and your organisation are responsible for:

- the accuracy and appropriateness of information, guidance and advice delivered
- managing risks
- ensuring that clients are not misled

Legal-specific generative AI systems

Some generative AI systems are promoted as being specific to legal practice.

These systems are typically:

- trained on legal datasets
- designed to align with common legal workflows
- self-described as being well-suited for legal practice
- targeted at the enterprise level, meaning they are built with larger, complex organisations in mind
- protected by enterprise-grade security which can, if enabled correctly, safeguard the confidentiality of user-uploaded data

The SRA does not authorise or certify software. Self-labelling by AI or legal tech companies does not guarantee that the AI system is appropriate or safe to use in legal practice.

Many legal AI systems are layered on top of AI models from well-known generative AI providers, using LLMs.

Competing legal technology products can be using the same underlying AI model. If the underlying model fails, multiple legal technology products could be affected.

Organisations and individuals deploying AI systems should:

- understand whether multiple AI systems rely on the same third-party AI provider or model
- consider the operational implications of that dependency

AI agents and agentic AI

AI agents and agentic AI differ from other forms of AI because:

- the main objective is to carry out actions, rather than produce data
- they can operate with more autonomy, requiring reduced or intermittent human input

AI agents and agentic AI are often referred to interchangeably. However, they are not the same thing.

AI agents:

- are specialised software programmes that use AI
- can perform tasks autonomously or semi-autonomously
- operate within a defined scope and remain subject to human instructions or controls

Agentic AI:

- AI systems that can pursue objectives with a higher degree of autonomy
- operate with reduced or intermittent human intervention
- can plan and execute multi-step tasks, and adapt the approach in pursuit of a goal
- may involve the operation of multiple AI agents

For example, an agentic AI system could use multiple AI agents to review documents developed with generative AI. The agentic AI system could then generate recommendations to resolve any issues directly through other systems before notifying you. It could do so without needing human approval to progress through each stage.

As agentic capabilities are increasingly being adopted within legal practice, you and your organisation should consider:

- whether the level of autonomy granted to an agentic AI system is appropriate for the task being performed
- limiting the permissions available to the system to those strictly necessary for the task
- restricting access to client information, matter files and internal systems unless there is a clear need and appropriate safeguards are in place
- establishing clear human oversight and approval checkpoints before actions with legal, regulatory or client impacts are completed
- ensuring responsibility for outcomes remains clearly allocated to named individuals within the organisation

The use of AI agents and agentic AI does not transfer or diminish professional responsibility from you or your organisation.

Appropriate human oversight becomes increasingly important with the use of more autonomous AI systems such as AI agents and agentic AI.

You remain responsible for verifying final outputs and ensuring compliance with your professional, regulatory and legal obligations.

For more information, read our:

- research report on the future of agentic AI in legal practice
- explainer on agentic AI and how it could impact legal services

3.8 AI systems and data use

Different AI systems may appear similar, but they can operate under a range of contractual terms and data handling arrangements.

The level of risk when using AI within legal practice will depend on:

- how the AI system is accessed

- the terms of service agreed to with the service provider
- the data that the AI system can store, retain, reuse or expose
- where AI service providers operate
- where data being processed by AI is processed and transferred

You must be aware of how you are accessing AI and the respective implications for your duties regarding client information, confidential material, privileged communications or personal data.

You, as an individual solicitor, should:

- understand what AI systems do, their limitations and the risks they may present, at a level suitable to your role
- check the version of the system you are using

It is generally riskier to access AI through free, publicly available versions than through enterprise or paid-for licences.

You, as a decision-maker, and your organisation should:

- carefully review the AI system provider's terms of service at the time of adopting and implementing the technology
- regularly check for any modifications or updates to such terms, and the underlying product itself, once the technology is embedded within your practice
- monitor for AI or AI features being introduced into existing software without prior warning
- be aware of how terms of service, data handling practices and risk profiles can differ across each layer of an AI supply chain. A single vendor's assurances may not extend to every party involved in delivering the AI system

You must ensure that any use of client data with an AI system is consistent with your confidentiality obligations and responsibilities under UK data protection law.

For more guidance:

- review the section on data protection
- see our buying new technology guide

3.9 Types of access to AI

The broad categories for accessing AI are:

- free and general-purpose generative AI systems
- paid-for or enterprise-level AI systems
- internally developed or configurable AI systems

Free and general-purpose generative AI systems

Publicly accessible consumer-facing AI services, available at no cost and on a consumer-facing basis, are:

- typically able to generate content in response to user prompts (such as text, images, audio or code)
- often not designed specifically for legal practice
- used under standard terms of service rather than bespoke contractual arrangements

Standard terms of service for these systems may:

- allow providers to retain user prompts and review conversations
- process data in other jurisdictions (outside England and Wales)
- use inputs (including data uploaded by users) and outputs to improve their services, including training their models

You may have limited ability to negotiate how inputs and outputs are stored, retained, accessed or reused.

For example, the common providers of free and general-purpose generative AI systems often:

- store and use input data to train and improve models and personalise future responses, unless you specifically opt out through data control settings if that is an option
- may allow human evaluators to review input and output data within conversations to refine the quality and safety of the system
- retain conversation history for anything from one month to indefinitely
- make user content accessible only to authorised persons within the company, under strict access controls

Considerations when using free and general-purpose generative AI systems

The performance of free and general use AI systems may not always differ significantly from paid-for or enterprise versions. However, options to customise the tools' configuration, based on jurisdiction and data sources, will be limited.

Most free and general-purpose online generative AI systems will not have been designed for use in the provision of legal services or with organisational-level use in mind.

Their use may carry higher risks when used in legal practice than other types of AI systems. They may:

- lack enterprise-grade privacy controls
- store, retain or reuse input data for purposes incompatible with confidentiality and data protection obligations
- be more likely to produce inaccurate or misleading results
- lack jurisdiction-specific knowledge
- fail to protect confidentiality and lead to greater risk of losing any legal professional privilege that may otherwise have existed
- inadvertently share or index conversation history to publicly accessible webpages

Paid-for or enterprise-level AI systems

AI services accessed through a paid subscription, commercial licence or organisation-wide agreement with an AI service provider.

Paid-for or enterprise-level AI can:

- generate the same content as free tools in response to user prompts
- be offered with additional contractual, technical and administrative controls around data use, security, confidentiality, retention, access management and compliance, but these safeguards will vary by product, contract and specific configuration of AI systems
- be offered under bespoke or enterprise-level contracts, operated within the organisation's secure cloud environment
- contain additional product features and functionalities or integrations not available in free tools
- offer earlier or exclusive access to powerful and leading 'frontier' AI models, the most technologically advanced AI models at any given time
- provide users and organisations with access to customer service and support from the AI service provider

The data handling practices of paid-for or enterprise providers vary, but the common providers often:

- do not store your prompts or outputs to train models by default, or share them outside your organisation without permission
- allow enterprise administrators to control how long data is retained, with options for short or no retention
- encrypt all content and store chats only within your workspace
- retain conversation history, even if deleted by the user or organisation, for abuse monitoring and legal reasons

Some providers offer options for data residency and processing location. You should check where data is both stored and processed, as well as the relevant contractual and technical commitments.

Legal-specific paid-for and enterprise AI

AI systems designed specifically for legal practice are typically accessed in a similar way to paid-for and enterprise tools. These systems are typically offered at enterprise level.

Examples of this type of AI system may include standalone legal AI platforms, as well as AI integrations into existing legal research, document management or practice management systems.

These systems may offer:

- legal-specific workflows
- access to legal content or integration with an organisation's systems
- enhanced contractual, technical and security controls

Considerations when using paid-for or enterprise-level generative AI systems

Enterprise-level AI systems typically allow the supplier and organisation to enter into contractual commitments about how data is handled.

Using enterprise-level AI systems offers protections that are not available through free tools. However, you should not assume that an enterprise label alone will provide sufficient protections.

Prompts, documents and other information entered into the AI system can often be processed within the organisation's own secure cloud environment, rather than in the supplier's cloud environment or used to train the provider's public model.

Securing those protections can give you and/or your organisation greater control over:

- where data is stored
- who can access it
- how long it is retained
- whether it can be reused

These protections are not automatic, even within paid-for tools.

Your organisation should check the contract, privacy terms and technical settings to confirm that the AI system's configuration can meet your confidentiality, data protection and professional obligations.

If you wish to input client data into an enterprise AI system, it is important to be aware of how your inputs, as well as the data the AI system has access to, will be handled.

As a decision-maker in a firm or legal team, you can assess whether necessary protections are in place by:

- reviewing the terms of the contract with the AI service provider
- checking the technical configurations

For guidance on contract reviews, visit our buying new technology guide.

With AI in legal practice, special consideration should be given to:

- actively reviewing and understanding the AI system provider's contract terms
- configuring appropriate privacy settings
- restricting use of AI and data that AI systems can access to contexts appropriate to the level of risk involved

When using enterprise AI systems, the organisation is likely engaging the vendor as a data processor.

UK GDPR article 28 requires that your organisation (as data controller) only use processors providing sufficient guarantees of security and compliance, formalised in a written contract.

Internally developed or configurable AI systems

Internally developed AI systems are usually:

- intended for specific internal activities rather than being publicly available or under consumer-facing terms
- operate within the organisation's own technical environment and governance framework
- draw on internal data sources (such as precedents, knowledge materials, policies, templates or matter information) depending on their purpose

Developing AI internally allows more opportunities for organisations to build controls around the AI system. It can also give more control over where data is processed and stored than through a vendor's cloud-based service.

However, additional responsibility sits with the firm or organisation to have the hands-on and technical expertise required to design, test, document and supervise those controls.

The safety and reliability of an internally developed or configurable AI system depends on both the underlying model and how these surrounding arrangements are designed and governed.

It can also shift more responsibility onto your organisation for security, maintenance, updates, monitoring, performance testing and technical governance.

Internal development of an AI system does not necessarily mean that the whole system has been built from scratch.

Internally developed AI may still use third-party infrastructure, AI systems or cloud computing services.

Configurable AI

Configurable AI may also involve some internal development. This typically involves:

- connecting a firm or organisation's own systems or applications to a third-party LLM, for example, via an application programming interface (API)
- allowing generative AI capability to be built into internal workflows, client-facing tools or legal technology products

Configurable AI allows organisations to embed AI into their own systems, rather than relying on a standalone AI system. The interface may appear with internal branding, but:

- the underlying model will likely still be operated by a third-party provider
- data may still be processed outside the organisation's cloud computing environment

Increasingly, organisations are focusing on harness engineering rather than developing an underlying AI model.

Harness engineering means designing the wider environment around a model, including:

- the instructions and context it receives
- the data and systems it can access

- the actions it is permitted to take
- the controls used to test and verify its performance

Other examples of internally developed or configurable AI

You or your organisation may consider using on-premises AI. This typically means downloading, hosting and running an AI model within an organisation's own infrastructure and running it on its own hardware.

Considerations when using internally developed or configurable AI systems

Internal development of an AI system does not automatically mean it is secure.

The level of risk and protection required will depend on how the AI system is designed, configured and governed. This includes the data used, access controls, retention settings, auditability and arrangements with any third-party service providers.

Configurable AI should be treated as higher risk until appropriate, documented and verified controls and safeguards are in place.

When carefully implemented, internal systems may offer stronger protections such as:

- more controls over what data the AI system can access, including ensuring that client data is not retained or used for training by the external service provider
- use of approved legal information, including template documents, case law and legislation
- the ability to set organisation-specific rules, prompts and disclaimers
- better alignment with internal data policies and retention of data within the organisation's secure internal system
- application of own security measures, access controls and audit logging

AI-generated or assisted code

If coding with AI support or using AI coding tools to build bespoke systems or integrations, you and your organisation should not assume that working code is safe or reliable.

AI-generated code should be subject to appropriate technical review, testing, security assessment and documentation before it is deployed in any legal workflow.

4. Professional duties

The SRA takes an outcomes-based approach to regulating individual solicitors, law firms and licensed bodies.

The SRA sets the standards that you and your organisation must meet. It does not prescribe exactly how those standards should be met in different circumstances.

You and your organisation are free to use AI systems and adopt new ways of working if you ensure that they meet the SRA Standards and Regulations.

The SRA does not:

- authorise or approve AI systems
- certify or approve specific AI systems for legal practice use
- monitor the use of technology advancements for the purposes of regulatory action against solicitors

The SRA does monitor how legal services are being delivered, which includes where regulated persons and organisations may use AI systems inappropriately or without proper safeguards.

The use of AI does not diminish or transfer your professional responsibilities.

The existing Standards and Regulations and Codes of Conduct for Solicitors and Firms apply. You remain accountable for all work and advice delivered to clients, the court, third parties and your regulators whether or not AI has been used in its preparation.

Appropriate human oversight, informed professional judgement, and a proportionate, risk-based approach are essential to ensure compliance with your regulatory and legal obligations.

You need to engage with the SRA, through its usual authorisation processes, if you or your organisation intends to use AI systems to deliver legal services:

- that include reserved legal activities
- in a way that changes your business model, structure or authorisation status

4.1 Duties to the client

Client care

When using AI systems to interact with or provide information or advice to clients, you should make sure that clients understand:

- where AI systems will be used in relation to their matter
- that they should not place undue reliance on AI outputs which have not been verified by a solicitor

Client use of AI

You may receive AI-generated content from clients or potential clients. They might use AI systems to:

- develop questions to ask you
- generate documents, summaries or legal arguments before seeking legal advice
- draft communications to you
- summarise communications from you
- draft complaints about the service they have received

You may encourage clients to disclose any use of AI systems with you.

There are risks that you may wish to explain to clients about the use of AI, including that:

- AI-generated materials may contain inaccuracies, omissions or references to authorities from other jurisdictions or lack proper legal reasoning or evidential support
- inputting paid-for legal advice into AI systems without adequate data handling safeguards may undermine the protection of confidentiality and privilege offered by obtaining legal advice from a regulated professional
- potential additional costs may be incurred if the use of AI results in a high quantity of AI-generated material that will take longer than you expect to review

You may wish to cover these risks to clients within your terms of engagement or client care letter. This might include an explanation of the risks and recommended actions to mitigate or avoid them.

When working with client materials that you know or suspect to have been AI-generated, you should:

- speak with the client to check the suitability and accuracy of any AI-generated materials
- clarify any inaccuracies in the material
- review and validate the legal and factual content before relying on the material

You should always ensure that any advice you provide is based on your independent assessment. Professional duties apply regardless of the source of the information or the ways clients develop their communications with you.

You may also set out your organisation’s preferences regarding the use of AI within communications with clients.

For example, you may include a request that if a client uses AI to help draft communications:

- they disclose their use of AI
- the correspondence remains reasonable and proportionate in length and content

For more information, see the SRA’s guidance on transparency in price and service.

Client consent

Specific disclosure or client consent may not be required where AI is used for internal administrative purposes, such as searching publicly available information or summarising internal materials.

The use of AI within your organisation’s own secure systems, where data is not used to train an AI system, is not likely, of itself, to require informed consent.

The appropriate approach in those circumstances is likely to be transparency around use.

If you use identifiable client data in AI systems, you and your organisation must ensure that the use is lawful and consistent with your legal and regulatory obligations, including SRA regulations and wider data protection and other information-related obligations.

The use of client information that could breach confidentiality, including to train AI, is likely to require specific client consent.

A client care letter or terms of engagement may be an appropriate place to address this.

You and your organisation may wish to take independent advice to ensure that any approach taken is lawful, safe and responsible.

Vulnerable clients and digital accessibility

Digital and AI-enabled client-facing services should be designed and delivered in a way that is accessible to clients.

You should consider whether your clients can understand, access and use the service effectively, and whether reasonable adjustments may be needed under the Equality Act 2010.

You should not assume that all clients will be able or willing to use digital tools, including lawtech or AI-enabled services.

Where a client cannot reasonably use a digital service, or chooses not to do so, you should consider offering a suitable non-digital alternative so that the client can still access legal services and make informed decisions about their matter.

For more information, review:

- our guide on meeting the needs of vulnerable clients
- the SRA's compliance tips on the use of AI and technology
- the SRA's guidance on accepting instructions from vulnerable clients

Charging for AI-assisted work

Your organisation's billings for legal services should accurately reflect the charging basis agreed with the client.

Where fees are time-based, only time actually and properly spent may be recorded and billed. If AI is used in billing arrangements of this kind, you and your organisation should not bill for the time that a task would have taken had AI not been used.

Time spent appropriately operating the AI system, reviewing and verifying its output, making revisions and applying professional judgement may be charged where consistent with the client care letter.

Any separate AI or technology charge should be clearly explained and agreed in advance with the client. It should not duplicate an amount already recovered through time-based rates or another fee.

For fixed, capped or value-based arrangements, you and your organisation should explain the scope of work and charging basis clearly to the client, ensuring that the overall fee remains justified.

The impact of AI on charging for legal work is an evolving area. You and your organisation may want to consider:

- how to adapt appropriate commercial models and billing arrangements that reflect the changing ways of delivering legal services
- how AI-enabled efficiencies are reflected in the value delivered to the client and ways to communicate this

For more information, see the SRA's guidance on transparency in price and service.

Confidentiality and risks

You are required to protect client confidentiality and handle client data appropriately whether you are using AI systems or not.

You should continue to adhere to your existing obligations under the SRA's guidance on the confidentiality of client information.

If you enter client data into AI systems that do not have the appropriate safeguards, your duty to protect confidentiality could be breached.

Risks to confidentiality may vary depending on whether the AI system is free and general-purpose, paid-for or enterprise-level, or an internally deployed secure system.

AI use and client confidentiality

There is yet to be a substantive judgment on AI use and client confidentiality within England and Wales.

However, the Upper Tribunal in *UK and R (on the application of Munir) v Secretary of State for the Home Department* [2026] UKUT 81 (IAC) observed that uploading certain client materials to a publicly available AI system was "...to place this information on the internet in the public domain, and thus to breach client confidentiality and waive legal privilege".

Whereas "[c]losed source AI tools which do not place information in the public domain... are available … without these risks".

Regardless of your method of accessing AI, in most cases, it will not be appropriate to use client data, without their consent, to train AI systems or where it might be reused to train AI systems. To do so presents significant risks both to confidentiality and data protection.

You should not focus only on whether client information will be used to train an AI system. Risk may also come from:

- retention
- logging
- monitoring
- troubleshooting

- abuse detection
- product improvement
- support requests
- other provider-side processing activities

If you are considering how to effectively protect client confidentiality through AI use, you should understand:

- where information entered into AI systems will be processed
- whether it may be transferred across other jurisdictions and the potential implications
- who within the service provider or its sub-processors can access your inputs
- what contractual safeguards apply before allowing client information to be entered into an AI system
- which operational safeguards you have in place to limit the risk of ‘shadow AI’ (including the use of systems which staff may not consider to be AI systems, such as publicly available AI-powered translation tools)

It is also important to understand whether contractual terms, technical settings and service provider’s due diligence are sufficient to protect confidential information and are consistent with your data protection obligations.

For more information on how different types of access to AI systems may change the data handling processes, see the section on types of access to AI.

Shadow AI use and confidentiality

Use of AI on any device without the formal approval, visibility or oversight of an organisation’s IT and security departments can be referred to as ‘shadow AI’.

Consider the risks of staff using personal or ‘bring your own devices’ to access unauthorised AI systems, or systems that are blocked at an organisational level.

Appropriate measures may help mitigate the risks of client or commercial data being inputted into unauthorised AI systems. For example:

- operational controls
- staff training
- prohibiting the copying or downloading of data from approved accounts or systems

Accidental inputting of confidential information

Inputting confidential client information into an AI system with insufficient data controls may result in a breach of confidentiality.

This may result from using a public AI system, or an enterprise system without the appropriate safeguards.

If this happens, you and your organisation should treat this as a potential breach of confidentiality. You should consider both regulatory and data protection implications.

If you consider the incident to amount to a serious breach of SRA regulatory arrangements, you and your organisation must report it to the SRA promptly.

The SRA’s reporting guidance requires serious breaches to be notified without undue delay. For more, see our guidance on reporting a breach of personal data.

Legal professional privilege

Use of AI systems can have implications for legal professional privilege, particularly where client information is shared with third-party systems.

AI systems may process or store data outside your or your organisation’s direct control. This can affect whether communications remain confidential and therefore privileged.

You should understand how AI use might impact your ability to uphold privilege. This includes risks of disclosure or waiver where information may be entered into systems without appropriate safeguards.

The decision of *UK v Secretary of State for the Home Department* [2026] UKUT 81 (IAC) addressed issues relating to confidentiality and privilege.

The judge held that:

> “[u]ploading confidential documents into an open-source AI tool, such as ChatGPT, is to place this information on the internet in the public domain, and thus to breach client confidentiality and waive legal privilege, and any such conduct might itself warrant referral to the regulatory body and should, in any event, be referred to the Information Commissioner’s Office.”

The case mostly dealt with the hallucinations present within the materials presented to the court. There is not yet clear case law that deals directly with legal professional privilege and the use of AI in legal practice.

Questions remain about whether:

- prompts, draft outputs, audit logs, internal AI-assisted analysis or translation records are treated as lawyers' working papers and attract privilege
- the treatment of materials depends on the type of AI system used and the safeguards in place

Until the position is clarified, you should take a cautious approach.

Consider whether using an AI system could create additional records, disclosures or third-party access that may affect privilege. Ensure that any privileged material is handled only through systems and processes that provide appropriate safeguards.

For more information on how different AI systems handle data, see the overview of types of AI section.

Client identity checks

AI can increase the risk of clients using false or manipulated information to bypass identity checks.

This may include the use of AI-generated identity documents, synthetic images or deepfakes to deceive remote onboarding, video calls or liveness detection processes.

A deepfake is an AI generated or manipulated video, image or sound recording that realistically replaces a person's face or voice with that of someone else's.

You should consider whether your client due diligence and electronic verification processes are resilient to these risks, particularly where you do not meet clients in person.

This may include reviewing:

- how identity documents are checked
- how video or biometric verification is used
- whether additional checks are needed for higher-risk matters
- whether staff are trained to identify and escalate suspected AI-enabled fraud

For more information, see the SRA update on clients using AI to bypass identity checks.

Transparency of AI use

It is important that clients are given appropriate information about how their matter will be handled, including in relation to the use of AI.

Providing appropriate information includes not misrepresenting the use of AI to clients, the court or third parties.

Routine and administrative use of AI may not require specific disclosure or consent. However, AI use may need to be disclosed as part of privacy information to be provided under UK data protection law if client personal data is processed by AI systems used by your organisation.

You may wish to take an upfront transparent approach about how AI is used in your organisation within your terms of engagement or client letter. This might include:

- that your organisation uses AI
- which systems are used
- how the AI is used, including the tasks and activities it is applied to
- how you protect client information

If a client asks how AI is being used in their matter, you should respond clearly and accurately, at a level of detail appropriate to the client and the circumstances.

This may include:

- explaining what type of AI system is being used
- what task it supports
- whether client information is entered into the AI system
- whether the output is reviewed by a solicitor
- how responsibility for the work remains with you or your firm

4.2 Duties to the court

Use of AI in the preparation of court documents has come under increasing scrutiny.

The High Court judgment in *Ayinde v London Borough of Haringey and Hamad Al-Haroun v Qatar National Bank* [2025] EWHC 1383 (Admin) identified the need for the legal profession to address the misuse of generative AI in legal proceedings.

The president of the King's Bench Division highlighted the dangers of AI-generated inaccuracies, including:

- the citation of non-existent cases
- the importance of lawyers complying with their duties to the court

Your professional duties and obligations apply in the same way whether AI or other technologies were used:

- to assist with that work
- by the solicitor personally or by anyone under that solicitor's supervision

You must make sure all submissions, documents and representations are accurate and based on genuine and verifiable sources.

Conduct your own review of AI-generated materials and check primary sources where required.

In *Ayinde*, the court pointed out in the strongest terms that those using AI to conduct legal research have a professional duty to "check the accuracy of such research by reference to authoritative sources, before using it in the course of their professional work (to advise clients or before a court, for example)".

You must not allow AI use to influence the substance of evidence, for example, by:

- generating false evidence
- presenting unsubstantiated material as evidence
- citing non-existent or inappropriate case law in support of a point or argument

Failure to use AI appropriately when producing documents for use in court will be regarded by the courts as a serious matter. It is very likely to lead to a referral to the SRA for investigation and the court may impose its own sanctions.

AI and witness statements

It is important that witness statements are an accurate reflection of the actual words of the witness. Using AI to help prepare witness statements may have an impact on this.

If you use AI in any aspect of the preparation of witness statements, you must consider how AI summaries or AI-suggested edits to witness statements have impacted the original evidence provided.

Even subtle changes to wording can change its meaning and affect the integrity of the evidence being given.

Coaching a witness risks their evidence no longer being classified as their own (*R v Momodou* [2005] EWCA Crim 177).

In *R v FGD* [2026] EWCA Crim 918, the Court of Appeal considered whether the use of AI to prepare a witness for cross-examination may be classified as impermissible "coaching". The court was firm in:

- its explanation that coaching witnesses on the substance of their evidence, including through the use of AI, is prohibited
- discouraging witnesses from using AI to prepare for giving evidence at trial as doing so "could run the risk of unfairness up to the point, in an extreme case, where the court may have to consider excluding evidence"

Expert witnesses have their own independent responsibilities regarding AI use in expert reports.

Checking AI outputs and AI-assisted court documents

It is important that work that is relied on by, or provided to, a court does not mislead whether through inaccuracy, omission or the use of fabricated material.

You should take steps to satisfy yourself that AI-assisted work or AI outputs are accurate, reliable and fit for the intended purpose before they are presented to the court, relied on, or shared externally.

What amounts to a reasonable and proportionate check will depend on the task and the risk involved. You may want to:

- review final versions to make sure that all citations to reference material (such as law reports or quotations from books, journals or articles) are verified
- ensure that legal arguments are properly reasoned and supported by authoritative and verified sources
- review any AI-generated information against reliable and authoritative primary sources
- exercise independent professional judgement on the arguments posed and conclusions reached

Consequences of a breach of professional standards

The court has made clear that any breach of professional standards resulting from the misuse of AI to carry out legal research or draft court documents is extremely serious.

You should expect the court to consider the entire range of sanctions available to it in exercising both its inherent jurisdiction over solicitors as officers of the court and the court's *Hamid* jurisdiction.

The *Hamid* jurisdiction, which stems from the case *Hamid, R (on the application of) v Secretary of State for the Home Department* [2012] EWHC 3070 (Admin), refers to the courts' jurisdiction to govern their own procedures. This includes ensuring that lawyers act properly and in accordance with their codes of conduct.

The court's sanction will depend on the facts and circumstances of each case, and could include:

- public censure
- public admonishment
- striking out of proceedings
- referral to the relevant regulator
- contempt of court or referral to the police

To date, courts have mostly considered the misuse of AI in this context to be so serious that public admonishment alone is insufficient.

The court has made it clear that it will not tolerate the misuse of AI in legal proceedings. It has indicated that practitioners should expect such conduct to be subject to the more serious sanctions at the court's disposal, including:

- initiation of proceedings for contempt of court
- referral to the appropriate regulator and, where appropriate, the police for a criminal investigation

Inaccurate information being presented through misuse of any AI system will breach the standards required by the SRA Code of Conduct for Solicitors. The code requires solicitors to:

- act in a way that upholds the rule of law and the proper administration of justice
- uphold public trust and confidence in the legal profession
- act with honesty, independence and integrity

AI hallucinations and inaccurate outputs

AI, particularly generative AI systems, can produce outputs that appear plausible, persuasive or convincing, but are inaccurate, incomplete or fabricated. This is often referred to as an AI hallucination.

In its misuse of AI warning notice, the SRA expresses a particular concern for AI hallucinations in legal work and submissions to the court.

The SRA has received reports of potential breaches of the code of conduct from senior members of the judiciary.

There have also been several instances of solicitors self-reporting where they have relied on inaccurate or misleading content generated by AI systems.

An AI system cannot reliably validate the accuracy of its results on its own. Relying on AI to verify AI outputs will not satisfy your professional obligations if there are inaccuracies.

You remain responsible for ensuring that work is accurate, defensible and fit for purpose.

Submitting inaccurate or misleading information to the court, clients or third parties may breach:

- principles to act honestly and with integrity
- duties not to mislead the court, clients or others

Where errors are identified, you should take appropriate steps to correct the position as quickly as possible.

For more information on hallucinations and the risks for the legal sector, read our article on how AI systems hallucinate.

Disclosure of AI use to the court

There is currently no specific duty to disclose to the court that AI has been used in the preparation of court documents. Existing duties, principles and codes of conduct apply.

You remain professionally responsible for a document's content where your name, as a regulated professional, appears. This position stands whether or not AI has been used to assist with research, drafting, review or document production.

The Civil Justice Council (CJC) consulted in 2026 on whether additional rules are needed to increase transparency where AI is used to prepare court documents.

It proposed that:

- different categories of court document, including statements of case, advocacy documents, disclosure statements, witness statements and expert reports may require different rules regarding disclosure of AI use
- existing professional responsibility may be sufficient for some document types, provided the document bears the name of the legal representative taking responsibility for it
- there may be a need for specific declarations of AI use, where it may affect the presentation of evidence
- there may in the future be a need to consider disclosure of the prompts inputted into AI

Our position is that professional responsibility remains central but that it is not sufficient on its own in all circumstances.

We responded to this consultation saying that we:

- favour an incremental and proportionate approach, rather than immediate wide-ranging changes to the Civil Procedure Rules
- consider that greater transparency is needed where AI use may materially affect the substance of court documents, evidence presented or fairness of the litigation process
- believe disclosure requirements should be clear, practical and targeted and should not capture low-risk uses of ordinary digital tools (such as those to assist with spelling, grammar and formatting, or provide administrative assistance)
- support further measures to improve transparency around AI use in court documents, alongside further AI-specific guidance from the SRA and HM Courts and Tribunals Service (HMCTS)

However, we cautioned against relying primarily on disclosure obligations. Equal focus should be on the governance, training and supervision needed to make obligations meaningful.

We believe you should be supported by court guidance to understand:

- when disclosure may be required
- how AI outputs should be checked
- what organisational safeguards are needed to ensure responsible use

In June 2026, the CJC published an update on the consultation findings. The Use of AI in Preparing Court Documents Working Group is now:

- consolidating its position on AI use in witness statements
- preparing a draft report setting out its recommendations to the CJC
- refining a final report which will be presented to the CJC, with publication anticipated later in 2026

Litigants in person

You may increasingly receive AI-generated or AI-assisted correspondence, complaints, letters of claim or court documents from litigants in person.

These may be lengthy, inaccurate or non-compliant with procedural rules.

If you are engaging with a litigant in person that you consider may be using AI to generate their court materials, you should:

- assess the substance of the material carefully
- identify any procedural defects
- avoid engaging unnecessarily with irrelevant or repetitive content
- ensure that genuine issues are addressed

Where materials submitted by a litigant in person appear to include fabricated authorities, false factual assertions or misleading content, you should consider whether and how this should be raised with the court. You should take into account:

- the litigant in person’s duties to the court
- your duties to the court
- your client’s interests, and
- procedural fairness

In our response to the CJC consultation on the use of AI for preparing court documents we recommended that HMCTS issue short, principles-based guidance on the use of AI in court documents.

We proposed that this guidance apply to litigants in person, as well as legal representatives, experts and other third parties involved in court proceedings.

The guidance should explain, in accessible terms, the types of AI, potential uses and risks. It should also outline a principles-based approach to AI and the use of AI in court. This could include principles such as:

- transparency
- explainability
- equality
- confidentiality
- human review
- accuracy
- security and standardisation

This guidance by HMCTS could provide a ‘ready guide’ to signpost litigants in person to help advise them of their duties and responsibilities.

Barristers’ duties

Barristers are subject to their own professional duties under the Bar Standards Board (BSB) handbook and codes.

The duties of a barristers are similar to those of solicitors. They include:

- duties to the court
- to act in the best interests of each client
- to act with honesty and integrity
- to maintain independence
- to preserve client confidentiality
- to provide a competent standard of work and service
- to maintain public trust and confidence in the profession

The use of AI does not change these duties.

The BSB has published guidance on the use of AI and other technologies by barristers.

That guidance explains how existing duties in the BSB handbook apply when barristers use AI. This includes support to understand professional considerations when solicitors and barristers work together on AI-assisted litigation or advisory work.

The instruction of a barrister does not transfer responsibility for your own professional obligations.

You should be aware of potential intellectual property and ownership considerations when instructing barristers. For example:

- whether the barrister’s chambers or any AI service provider could assert rights over AI-assisted drafting used in the opinion, and
- how this interacts with your and your client’s expectations about ownership of, and reliance on, the final product

4.3 Professional skills, oversight and supervision

When AI is used in legal practice, competence, oversight and supervision work together to help ensure that legal services are delivered to the required professional standard.

You must maintain your competence and keep your professional knowledge and skills up to date. Your organisation must ensure that managers and employees are competent to carry out their role. The use of AI should not erode this.

What constitutes appropriate oversight of AI-assisted work and the supervision of those using AI depends on the task and level of risk.

Even where precautions have been taken to ensure safeguards around AI use, you must maintain appropriate ongoing oversight of AI-assisted work and supervision of those for whom you have responsibility in providing legal services.

Competence

You have an ongoing duty to maintain your competence and keep your professional knowledge and skills up to date.

This may include maintaining an appropriate level of AI literacy and understanding of the AI systems you use, their capabilities, limitations and associated risks.

The level of understanding will depend on your role and the extent to which AI is used in your work or the work of those you supervise.

Your organisation has a duty to ensure that managers and employees:

- are competent to carry out their role
- keep their professional knowledge and skills up to date

- keep their understanding of their legal, ethical and regulatory obligations up to date

This may include providing appropriate training on AI literacy, the organisation's approved AI systems, relevant policies, and the professional, legal and regulatory implications of AI use.

Effective AI training is not a one-off exercise. As AI capabilities and regulatory expectations evolve, your organisation should consider how ongoing learning, guidance and engagement can support responsible use of AI.

The use of AI to complete a task more efficiently should not prevent you or those you supervise from retaining or developing the competence needed to perform and scrutinise that work without AI support.

Your organisation may wish to consider:

- when a task should be completed without AI for training or assessment purposes
- whether staff understand and can explain work produced with the assistance of AI
- how supervisors can identify and address gaps in knowledge or skills
- how training and work allocation can preserve the skills needed to have effective oversight of AI-assisted work
- whether staff retain sufficient capability to continue work if an AI system is unavailable

Professional skills and development

AI should support rather than shortcut your professional growth and learning as a solicitor.

Overuse or overreliance of AI systems within legal practice may contribute to a loss of legal, analytical or practical skills. This risk applies to you regardless of the level of your career.

However, it may be particularly significant for junior and aspiring solicitors who are still developing the professional skills, knowledge and judgement needed for competent legal practice.

Your organisation, and you if you supervise trainees, should consider how the use of AI may affect the training and development of those entering the profession. This includes considering any differences for those pursuing traditional training contracts and those undertaking qualifying work experience as part of the Solicitors Qualifying Examination (SQE).

Those responsible for junior colleagues should ensure that AI does not replace opportunities to:

- develop legal research, drafting and analytical skills
- understand how legal conclusions are reached
- exercise independent professional judgement
- identify errors, omissions and unsupported propositions
- learn through appropriate responsibility, feedback and reflection
- develop the skills needed to have effective oversight of AI-assisted work in future

Oversight and supervision

For the purposes of this practice note, we are making a distinction between the acts you may carry out of checking the work of those you are managing or responsible for (supervision) and the checking of AI-assisted work before relying on it (oversight).

Supervision is about the people: the arrangements for managing those carrying out the work.

Oversight is about the work: the checks applied to AI-assisted work before it is relied on.

Supervision

If you have supervision responsibilities, your firm should ensure that you have sufficient understanding of how AI is being used within the firm and the implications for your supervision responsibilities.

This is best achieved through formal and mandated training on:

- AI literacy
- the use of the specific AI systems authorised or available at your firm or organisation (and which AI systems may not be used)
- how the supervision of more junior staff interacts with the need for oversight of the use of AI
- the impact of AI use, supervision and oversight on work within the specific workflows in your practice

Supervision should be active and embedded into the relevant workflow.

It will likely not be sufficient to instruct managers to supervise AI-assisted work without providing the time, training, tools and escalation routes needed to do so effectively.

If you are a decision-maker in an organisation, you should consider building proportionate review checkpoints into existing processes, consistent with your supervision obligations.

This is particularly relevant where AI is used in the process of substantive legal tasks, client communications, court documents or work carried out by junior colleagues.

These checkpoints should make clear:

- to what extent AI may be used in the relevant task
- to what extent AI has been used in the relevant task, such as which AI system has been used for which purpose
- who is responsible for reviewing the AI-assisted work
- what must be checked, verified and how to do so
- what sources should be used to verify the output produced by an AI system and included in the work, and
- when, how and to whom concerns relating to AI-assisted work should be escalated

For more information, see:

- the SRA's effective supervision guidance
- our practice note on supervision

Oversight

The effectiveness of oversight of the use of AI will depend on individual circumstances and conditions.

Oversight is needed for work that is communicated externally or that shapes the advice given, such as written advice to a client, a document provided to a third party, or a court document. It does not have to extend to every individual interaction with an AI tool.

Required levels of oversight are fact sensitive. They will depend on the:

- type of AI system being used
- specific legal task you are carrying out
- safeguards and governance mechanisms put in place by your firm or organisation, including which AI systems are authorised for use

Your organisation should ensure that you understand:

- how AI systems authorised for use in your organisation work
- what data they have access to and what data they should not have access to
- the limitations and risks of those AI systems
- which AI systems are not authorised or approved for use

Having effective oversight means taking reasonable and proportionate steps to confirm that AI outputs are accurate, appropriate and fit for purpose. Within the practice of law, this should contain, but is not limited to:

- checking the accuracy of citations, sources, legal reasoning and statements of facts and law by reference to primary sources
- testing if outputs are appropriate for the client's circumstances and instructions provided
- identifying any omissions, unsupported propositions, bias or over-confident conclusions
- applying independent professional judgement when reviewing AI outputs rather than accepting AI outputs as definitive

Automation bias

Be aware of the risk posed by automation bias.

Automation bias refers to the tendency of individuals to over-rely on, or defer to, the outputs of automated systems, including AI. This may occur even where those outputs are inaccurate or where the individual has the knowledge or means to identify an error.

In legal practice, automation bias is a significant risk because professional liability and regulatory accountability for legal work remain with you, regardless of whether AI was used to help produce it.

You should be alert to the risk of automation bias in your own work and in that of those you may supervise.

It may be helpful to treat AI outputs with the same level of professional scepticism and independent judgement that would be applied to work produced by a trainee or junior colleague.

It is best practice to have oversight of AI-supported outputs. How a human carries out oversight and the level of oversight required might depend on the tool, task and associated risks.

You may not need the same level of human oversight of lower-risk tasks, if the tasks do not affect the substance of:

- legal advice provided

- client outcomes
- materials submitted to the courts, clients or third parties

Examples of lower-risk tasks might include formatting, checking grammar or internal administrative or operational support.

Check how others are using AI systems and whether you need to adjust your level of review and scrutiny to suit the risks posed.

If adopting AI, your organisation should put in place the necessary training, support and guidance to allow you to understand what meaningful human oversight means in the context of your work and the specific AI systems you have access to.

5. AI governance

AI governance refers to the systems, policies, processes and oversight arrangements that manage how AI is used within an organisation.

This means putting in place technical and human controls to make sure AI systems are used lawfully, safely and consistently with your professional obligations.

Your organisation should communicate what is, and is not, acceptable AI use. This is needed even where your organisation has not formally adopted AI, because staff may use free public AI systems independently (shadow AI).

Effective AI governance can help you and your organisation to understand:

- what AI systems are being used
- what data AI systems can access and process
- how AI use that processes personal data is communicated to data subjects
- who is responsible for supervising the use of AI systems
- how risks are identified and managed
- how staff are supported to use AI appropriately

This is especially relevant for solicitors with supervision responsibilities.

Find more guidance on risk assessments for new technology in our buying new technology guide.

5.1 Approaches to AI to AI governance

The right approach to AI governance will depend on the size and structure of your organisation, the types of AI used and the tasks it is being used for.

Having some form of governance arrangements for AI use is fundamental and should include:

- appropriate safeguards to manage risks
- ongoing review of governance
- identification of any training needs

This can form part of your organisation's broader approach to risk management and supervision.

Accountability for your organisation's approach to AI should sit with your organisation's leadership, who should:

- set the overall risk appetite for AI adoption
- ensure adequate resource is allocated to AI governance, training and oversight
- work with data protection officers and AI risk officers (where relevant)

Everyone will have a responsibility to ensure effective compliance with the agreed arrangements.

Developing an AI governance framework can help you and your organisation remain compliant with professional and legal obligations.

It may also assist with demonstrating to insurers that the use of AI in your organisation is supervised, documented and subject to proportionate risk assessment and mitigation mechanisms. For more information, review the below section on insurance.

Engaging staff in AI governance processes

Your organisation may consider engaging with staff and staff representatives on the development of an AI governance framework.

Consultation with staff can help identify practical issues such as training requirements, and oversight and supervision challenges that may not be apparent at senior leadership level alone.

Depending on the circumstances, organisations may also need to consider any applicable obligations relating to workforce consultation or engagement when introducing AI systems that affect staff or working practices.

It is effective to engage staff on AI use and governance on an ongoing basis, rather than as a one-off exercise at the point of adoption.

It may also help you to provide assurance to staff, clients and insurers.

AI governance frameworks

There is no single universal framework or standard for AI governance.

You may consider several recognised frameworks to help create a governance approach suited to your organisation.

For example:

The USA's National Institute of Standards and Technology's (NIST) AI Risk Management Framework

- a structured, sector-agnostic approach focused on the pillars of govern, map, measure and manage
- treats AI risk management as a lifecycle activity, where AI governance must be considered and applied across the full lifecycle of an AI system

ISO/IEC 42001:2023

- treats AI governance as a management system: a set of policies, objectives and processes for the responsible development, provision or use of AI systems
- includes maintaining and continually improving that system

Considerations for AI governance

AI governance arrangements typically include:

- an AI policy
- data handling rules
- escalation procedures for concerns or instances of misuse of AI systems
- clear lines of responsibility for oversight and decision-making concerning use of AI systems
- regular review of quality of AI outputs and AI system performance
- assessing and managing risks related to data protection, accuracy and client impact resulting from the use of AI systems
- business continuity and resilience planning for any AI-dependent processes

They may also include staff training:

- to improve AI literacy and ensure a good understanding of the authorised AI systems used in your organisation
- for supervisors and managers on how to meaningfully review AI-supported outputs

AI in existing software

Where AI functionality is embedded within existing software, your organisation may not always have clear visibility of:

- where AI systems are operating
- what data the systems can access and process
- how underlying AI models are changing over time

This can make governance, inventory management and risk assessment more challenging.

You may be able to mitigate this by requesting regular updates from your existing technology suppliers regarding any AI functionalities.

For more guidance on managing supplier relationships, see our buying new technology guide.

AI registers

Your organisation may wish to develop an AI register to support with AI governance, reporting and insurance.

This may be a register of:

- approved AI systems
- approved use cases
- the owner of each AI system
- dates of approval and future reviews

AI policies

Your organisation is not required to have a standalone internal AI use policy.

However, the pace of change of AI capabilities means that AI poses unique risks. Its adoption and use may justify a separate policy.

Your AI policy should not be limited to generative AI only but capture all use of AI systems within your organisation.

An effective AI policy might include:

- which AI systems are approved for organisational use
- how to access the approved AI systems
- limitations on use of those AI systems including clear provisions regarding the use of client data in AI systems
- any prohibited uses

A standalone AI policy allows for more agile updates as the technology and regulatory landscape evolve, without requiring you to revisit your entire technology framework. It also provides a clear reference point for staff training and accountability.

It is important to make new starters aware of your technology and/or AI policies as part of any formal induction process. You should direct existing staff to review any new or amended policies.

If you're updating or creating a new policy that will cover the use of AI, you could consider:

- how to verify the accuracy and reliability of AI outputs
- how confidential or client sensitive information will be handled and who is responsible for checking this
- safeguards to prevent inappropriate use of AI or use of unapproved tools or systems
- how to maintain appropriate human oversight and what this means for your practice and organisation
- any wider legal or regulatory requirements, such as data protection obligations
- how you will assess and manage AI-related risks, including legal, regulatory and ethical considerations
- whether there is any need for additional training for staff using or overseeing the use of AI
- whether any AI system requires mandatory training before access can be granted

Audit and explainability

You don't need to understand exactly how the AI system works but you should understand where and how it is being used, and the associated risks.

Effective audit processes can help your organisation to demonstrate that AI systems are being used consistently with professional obligations, and regulatory requirements.

Audit processes may also assist in responding to enquiries from clients, insurers, regulators, courts and other stakeholders regarding your organisation's use of AI.

To support effective audit, your firm or organisation should periodically review:

- which AI systems are being used
- who is using them
- what information is uploaded
- performance of AI systems
- recurring errors
- changes to risks
- effectiveness of controls
- introduction of new functionalities
- compliance incidents
- client complaints

If your organisation is exploring the use of AI to process personal data, you may want to obtain support with a voluntary audit (PDF 164 KB) from the ICO's audit team.

As part of these audit processes, your organisation should consider the explainability of AI-supported work.

Explainability refers to your ability to understand, account for, and communicate or document how an AI system has produced its output.

Limited explainability can undermine your ability to have effective oversight of the work.

To support both audit and explainability, your organisation should keep adequate records to demonstrate that you know:

- where and how AI is being used
- the associated risks, and
- how you mitigate them

Where AI systems are used to undertake ADM, your organisation will need to, in addition to the above, be able to explain the 'logic' of a system to the people who are impacted by it.

5.2 Investment and funding

When deciding whether and how to adopt AI, your organisation should also consider the source and purpose of the investment behind the AI system.

Your firm or organisation may fund AI adoption through retained profits or through external financing, strategic partnerships or investment.

Funding arrangements may shape:

- organisational priorities
- procurement decisions
- risk appetite
- expectations about speed of deployment, return on investment or productivity gains

There is growing interest of private equity in legal service providers. Organisations should remain mindful of their obligations (and those of the solicitors working in them) regarding independence and governance responsibilities when considering such arrangements.

6. Regulation and compliance

6.1 Data protection

The use of AI in the workplace increases the risk of breach of UK data protection laws by introducing new data flows and processes.

It is important to consider how AI use might impact the seven key principles of UK data protection law:

- lawfulness, fairness and transparency
- purpose limitation
- data minimisation
- accuracy
- storage limitation
- integrity and confidentiality (security)
- accountability

For more information, see the ICO's guide to the data protection principles.

Organisations should also be clear about whether client data will be retained, reused, or incorporated into future versions of an AI system.

AI service providers may be able to see your input and output data or use it to train their AI systems.

The introduction of an enterprise system is an organisational risk to UK data protection law compliance that must be managed by the board and/or senior partners.

Review the section on types of access to AI for more detail on how this risk can vary.

Personal and confidential data

AI systems are trained using large volumes of data. It may be possible that sensitive information that is confidential or subject to legal professional privilege is exposed.

Metadata might be shared with the providers of underlying technology or models that support your AI system, not only with the specific service provider you are using.

Metadata is 'data about data'. For example, the creator of a piece of data and the date and time of creation. Metadata can qualify as personal or confidential data.

Where client data includes their own personal data or that of a third party, your obligations under UK data protection law will be triggered.

A breach of requirements under UK data protection law may also give rise to professional misconduct issues.

You, or your organisation, may wish to consider if and how you want to put personal and confidential data into an AI system, subject to the terms of use and the operational processes you establish with the provider.

For more information on how to handle personal data lawfully, see our guide on GDPR for solicitors.

International data processing

Many AI service providers are located outside of the UK or use servers outside of the UK. Data, including personal or confidential client data, may be transferred outside UK borders and may be classified as international data processing.

This could raise data protection questions about:

- what personal data was processed
- whether such personal data may be present in the outputs

Anonymisation and pseudonymisation

An attempt to anonymise client data does not reduce all risk or remove the data from the scope of data protection law. Data can only be considered anonymised when individuals are no longer identifiable. If this threshold is not met, data will be pseudonymised only.

Pseudonymisation refers to techniques that replace, remove or transform information that identifies people, and keep that information separate.

Pseudonymised data remains personal data and is therefore still subject to data protection law.

Any client data used to train an AI system may still be capable of reidentification even after pseudonymisation. For more information, see the ICO's guidance on anonymisation.

AI and automated decision-making (ADM)

With the introduction of the Data (Use and Access) Act 2025, the general prohibition of automated decision-making (ADM) was liberalised. ADM is now mostly permitted, with mandatory safeguards required.

The safeguards include enabling the data subject to:

- access information about the decision
- make representations about such decisions
- obtain human intervention
- contest such decisions

For more information, see the ICO's guidance on the ADM safeguards.

Individuals have the right not to be subject to decisions based solely on automated (including AI-driven) processing where those decisions are based entirely or partly on special category data, such as health records.

Risks may arise where you or your firm rely heavily on AI systems to make or recommend decisions about clients, matters or case strategy without sufficient human involvement.

It is important to consider the risks of automation bias in the context of ADM.

For more information on automated decision-making, explore the ICO's guidance on ADM, including profiling.

ICO data protection resources

For further guidance on data protection and AI use, see the ICO's resources:

- AI and data protection
- AI and data protection risk toolkit
- automated decision-making, including profiling
- ICO response to the consultation series on generative AI

6.2 Equality, discrimination and human rights

AI systems can produce discriminatory or biased outputs. This may be because they:

- were trained on unrepresentative data
- can reflect existing societal biases
- are used without appropriate safeguards

Your organisation should consider its obligations under the Equality Act 2010 when procuring, developing and using AI systems. The duty not to unlawfully discriminate covers all protected characteristics.

The risk of discrimination, and how to mitigate it, may be particularly relevant where AI is used in recruitment, performance management, client onboarding and the direct provision of legal services.

The use of AI may also engage human rights considerations including in relation to:

- privacy and data usage
- discrimination and bias
- existing remedies for violations of human rights

You and your organisation should consider whether appropriate testing, oversight and governance arrangements are in place to identify and mitigate these risks.

For more information, see:

- our equality, diversity and inclusion (EDI) hub
- our submission to the Joint Committee on Human Rights' (JCHR) Human Rights and the Regulation of AI inquiry (PDF 124 KB)
- the JCHR inquiry report on Human Rights and the Regulation of AI (PDF 1 MB)

6.3 Intellectual property

A key concern when using AI systems is determining the ownership of training and output data.

AI systems may, in some circumstances, reproduce third-party copyright material as outputs. Human review of outputs is important to mitigate against intellectual property infringements.

Issues may also arise regarding who owns AI-assisted work between your organisation, client and the service provider of your AI system. This can include whether any of it attracts copyright or other intellectual property protections.

You and your organisation may wish to use contractual arrangements to allocate ownership rights where copyright status is unclear.

Data scraping

In addition to the data you control and to which you have rights, data scraping is often performed in the training of AI systems and can cause concerns.

Data scraping is the process of extracting data from various sources, such as databases, files or websites (web scraping), often using automated tools.

This is a common technique for data gathering and is not exclusive to AI. Its use in training AI systems can lead to:

- intellectual property challenges
- copyright infringement (if input data containing copyright works results in recognisable outputs)

The World Intellectual Property Organisation has published a factsheet on generative AI and how the protection of intellectual property should be navigated (PDF 1.5 MB).

Mitigating the risk of data scraping in AI systems you use

To minimise the risk of intellectual property infringements, your organisation may want to consider:

- seeking clear information regarding the sources of training data
- if there are providers willing to offer indemnities for third-party training-related intellectual property infringement claims
- evaluating the extent and suitability of any indemnity agreements with AI provider(s)
- implementing staff policies and training to minimise the risk of producing infringing outputs

Mitigating the risks of your own data being scraped

You or your organisation may also be impacted by data scraping for AI training.

To mitigate the risk of potential data scraping and infringements of your intellectual property, you can:

- check the settings on generative AI tools to minimise the risk that your provider stores or trains using your prompts
- speak to your service provider about rate limiting website requests from scrapers (rate limiting is a strategy for limiting network traffic and can be used to cap how often someone, including bots, can repeat an action within a certain timeframe)
- deter crawlers (automated software tools that navigate the internet to gather information) from scraping your website by limiting access through editing the robots.txt file where possible

Ownership clauses in terms of services

AI service agreements may contain provisions or agreements allowing the service provider to:

- reuse input data to refine its AI system
- retain ownership of output data

You should ensure that such clauses do not conflict with your professional standards concerning the ownership and control of legal advice or risk compromising confidential data.

For more information, see the overview of types of AI section.

Procurement considerations

If you are considering purchasing a license for an AI system, it is essential to obtain clarity on who will own the intellectual property rights over input and output data.

You should review the terms of service and privacy policy of each tool to check:

- the usage rights
- ownership of outputs
- any disclaimers

For more guidance, visit our buying new technology guide.

6.4 Insurance

The use of AI may have implications for professional indemnity and cyber insurance.

Professional indemnity insurance (PII) and cyber insurance serve different purposes. Your organisation should understand the extent to which AI-related risks are covered by such policies.

PII cover must comply with the SRA minimum terms and conditions. For more information, see the SRA's guidance on adequate and appropriate indemnity insurance.

Professional indemnity insurance (PII)

AI use may affect your organisation's professional indemnity risk profile, particularly where it is used to support client work, legal research, drafting, document review, client communication or other tasks that may affect the substance of legal services.

You and your organisation should consider whether your use of AI changes your risk profile.

If there is any uncertainty about this, you should discuss it with your broker or insurer, particularly when arranging, renewing or updating your professional indemnity insurance.

Before using AI, you may wish to check whether your existing policy, proposal form or renewal documentation asks about:

- the use of AI
- automation
- legal technology
- outsourcing
- cyber security
- data processing

Information to provide to insurers

Under the Insurance Act 2015, your organisation is required to disclose not only risks you are aware of but also those that they ought to be aware of, or to have discovered after making reasonable enquiries.

The onus is on your organisation to properly investigate and understand all the risks in the business, including those posed by AI adoption or use.

For more information, see the SRA's guidance on the Insurance Act 2015.

You should answer any questions accurately and provide sufficient information to enable insurers to understand how AI is being used and what safeguards are in place.

Insurers may be interested in your AI governance processes. Being able to demonstrate that AI use is supervised, documented and subject to proportionate risk assessment and mitigation mechanisms will be helpful when engaging with insurers.

Your organisation may find that insurers seek information about the extent of human oversight of AI-assisted work. Clear supervision arrangements, documented processes and evidence of review may assist in demonstrating sound risk management.

You should also consider whether any AI-related activity could be seen to fall outside the considerations on which your insurance was arranged.

Notification requirements

If a problem arises within your legal practice from AI use, you should consider any notification requirements under your policy and take advice where appropriate. This might also include the need to notify the ICO.

If you believe that cover is available under one or more of your insurance policies, you should notify the insurer(s) as soon as possible, even if you have not been able to quantify the loss completely.

Cyber insurance

Your organisation may wish to have cyber insurance as part of your cyber security protection and prevention guards against damaging cyber losses.

Cyber insurance is increasingly relevant where your organisation relies on AI systems, third-party AI service providers and cloud computing service providers.

If your use of AI depends on external infrastructure and providers, a cyber incident affecting one of these service providers could:

- disrupt access to your services
- expose data
- affect multiple matters simultaneously

Cyber insurance policies may include cover for cyber-attacks on third parties that result in operational losses to the insured. Check your policies to see whether they include this type of cover.

Cyber insurance can provide important protections where events fall outside the scope of your minimum terms and conditions PII.

For more information, see our cyber insurance guidance for solicitors and law firms.

6.5 Cyber security

Using AI systems can introduce cyber security risks, including the potential for malicious actors to exploit vulnerabilities.

Your organisation should ensure staff are aware that simply reviewing a document through an AI system could trigger hidden instructions within that document.

Other cyber security risks include data source corruption that may compromise the generative AI outputs.

Generative AI systems can also be weaponised by adversaries to produce more sophisticated phishing and cyber security attacks in greater volume.

The National Cyber Security Centre has a hub for more information about AI and cyber security.

Visit our cyber security hub for further guidance specific to solicitors and law firms.

6.6 Consumer protection

Consumer protection law may be relevant where:

- AI-enabled services are supplied to consumers
- your organisation buys digital products and services from technology providers

As a decision-maker within your organisation, you should consider your role in the relevant transaction – for example, whether the user of the tool, client or purchaser is treated as a consumer for the purposes of consumer law.

If your organisation offers AI-enabled services directly to individual clients, you should ensure that your marketing, onboarding, pricing and cancellation processes are clear and not misleading.

This may be particularly important where AI is used in client-facing tools, automated triage, chatbots or online legal information services.

You should also consider:

- whether clients understand what service is being provided
- the limits of any AI-enabled tool
- when a solicitor will review or take responsibility for the work

Subscription contracts

The Digital Markets, Competition and Consumers Act 2024 will introduce a new regime for certain subscription contracts in spring 2027.

If your organisation develops or offers AI-enabled legal services on a subscription basis, such as access to an AI chatbot, client portal or recurring digital advice service, you should consider whether the incoming subscription contract rules will apply.

These rules may affect:

- pre-contract information
- reminder notices
- cooling-off rights
- cancellation processes
- renewal arrangements

You should take advice where needed and design any subscription model so that clients can understand what they are buying, how charges will apply and how they can cancel.

6.7 Environmental impact of AI

The rapid adoption of AI systems carries an environmental impact that solicitors and their clients may increasingly be expected to understand. This impact may arise from:

- the data centres that train and run AI systems
- the water consumed in cooling systems
- the embodied carbon in the underlying hardware

They could sit in tension with the net zero commitments your firms, organisation or clients may have made.

Our guidance on the impact of climate change on solicitors sets out climate-related risks that fall broadly into physical, transition and liability risks. These give rise to climate legal risks that may be relevant across nearly all areas of practice.

The environmental impact of AI adoption sits within this landscape, affecting how organisations manage their own transitions to net zero and the advice you may give to clients about their use of AI systems.

Your organisation and clients should be aware that use of AI may affect an organisation's emissions, which in turn can carry legal and regulatory implications.

For example, organisations and clients that have committed to reducing their indirect ('scope 3') emissions should consider that the energy cost associated with AI procurement:

- may count towards those emissions, and
- could bear on commitments they have made or reporting obligations they are subject to

If you are advising on AI adoption, procurement and service provider contracts, you may wish to take these considerations into account, consistent with the general duty to exercise reasonable care and skill.

The environmental impacts of using AI raise fast-moving technical, regulatory and ethical questions. We plan to produce more detailed guidance on this in the future.

6.8 International and cross-border considerations

AI law and regulation is developing differently across jurisdictions.

Some jurisdictions have introduced AI-specific laws or sector-specific requirements, while others rely on a patchwork of existing data protection, consumer protection and professional regulation.

You should consider whether any additional legal or regulatory requirements apply, and any additional risks may be involved, if you or your organisation:

- carries out cross-border work
- provides services to clients in other jurisdictions
- uses AI systems that are developed, hosted or supported outside the UK

Where relevant, you should take proportionate steps to identify and manage these risks.

For more information on risks, see the section on reliance on overseas AI systems providers.

7. Key points

This section summarises the main points from across the practice note.

You can use it as a quick reference to help identify the issues to consider when using or implementing AI in legal practice.

7.1 AI in legal practice: key points

- AI terminology is fast-moving. You should keep up to date and understand what terms mean to make informed decisions.
- Use of AI by solicitors is governed through existing legal, regulatory and professional obligations. There is guidance available from us, the SRA and the LSB.
- The SRA does not approve AI or any other technology systems. Named regulated solicitors retain accountability for outputs and compliance with professional standards, even within an AI-powered law firm.
- There is a difference between AI systems used in the business of law and in the practice of law. AI systems used in the practice of law engage your professional regulatory obligations more directly than AI systems used solely in the business of law.
- Generative AI is prone to hallucinations – incorrect or misleading results generated by AI systems.
- Your existing legal, regulatory and professional obligations continue to apply when using AI in your practice. Liability cannot be transferred to an AI system as it does not have legal personality.
- The UKJT explains that a professional may be found negligent for using AI inappropriately, using an unsuitable model, failing to conduct proper due diligence or failing to test AI or validate its outputs effectively.
- The UKJT also foresees a future where a professional could be liable for failing to use AI in circumstances where a competent member of their profession would have done so. However, there is no guidance from us or the SRA to declare instances in which solicitors should use AI, and none is planned at this time. We are not aware of any existing or planned positive duties for regulated solicitors or law firms to use AI in legal services.

Review the section on AI in legal practice.

7.2 Types of AI, access to AI and risk key points

- Risks posed by AI are not fixed. They depend on the type of system, how it is accessed, the task, the data involved, the level of autonomy and the organisation's governance and controls.
- Chatbots, legal-specific AI, AI agents and agentic AI present different risks. Greater autonomy requires tighter permissions, access controls, oversight and approval checkpoints.
- You and your organisation should understand how an AI system and its supply chain handle, store, retain, reuse and transfer data, and whether the arrangements protect confidentiality, privilege and personal data.
- Free, paid-for or enterprise, and internally developed or configurable AI systems offer different contractual and technical protections.
- Check contracts, privacy terms and technical settings before adoption and throughout use. Limit access to systems and data according to the risk involved.

Review the section on types of AI and ways to access AI.

7.3 Professional duties key points

- Using AI does not diminish or transfer your professional responsibilities. Existing standards, regulations and codes of conduct apply.
- Be transparent with clients about your use of AI and its impacts.
- You may encourage clients to disclose to you any use of AI systems and explain the risks of their use of AI. Specific disclosure or client consent may not be required where AI is used for internal administrative purposes.
- You and your organisation are required to protect client confidentiality and handle data appropriately regardless of the tools you use to do your work.
- AI systems may process or store data outside of your or your organisation's direct control, which can affect whether communications remain confidential and therefore privileged.
- Inputting confidential client information into an AI system with insufficient data controls may result in a breach of confidentiality. If this happens, you should treat this as a potential breach of confidentiality and escalate as appropriate.
- Your duty to ensure the accuracy of documents submitted to the court remains, regardless of whether AI was used to prepare them.
- Inaccurate information being presented to the court, client or third parties through misuse of any AI system will breach the standards required by the SRA Code of Conduct. It may also expose you to severe penalties initiated by the court.
- You must maintain appropriate ongoing oversight of AI-assisted work and supervision of those carrying out the work, even when precautions have been taken.

- Competence, oversight and supervision work together to make sure legal services are delivered to the required professional standard.
- Keep in mind that AI, particularly generative AI, can produce convincing but incorrect outputs, known as hallucinations.
- You should understand how an AI system has been used, what tasks it supported, and the risks associated with any outputs.
- Be alert to automation bias in your own work and in the work of those you supervise. It is a significant risk that may negate the safeguard of human oversight.
- Although AI systems can enhance your provision of legal services, continuing to develop your own legal skills is essential. AI should support not shortcut your growth and learning as a solicitor.

Review the section on professional duties.

7.4 AI governance key points

- AI governance refers to the systems, policies, processes and oversight arrangements that manage how AI is used within an organisation.
- Your organisation should communicate what is, and is not, acceptable AI use.
- This may include putting in place controls to make sure AI systems are used safely and consistently with your professional obligations.
- Having some form of governance arrangements for AI use is fundamental – the right approach will depend on the size and structure of your organisation, and the tasks AI will be used for.

Review the section on AI governance.

7.5 Regulation and compliance key points

- The use of AI in the workplace increases the risk of breach of data protection laws by introducing new data flows and processes. Where AI systems access and/or process data that includes personal data, your obligations under UK data protection law will be triggered.
- AI systems are trained using large volumes of data. It may be possible that confidential or sensitive information is exposed.
- Be aware that AI systems may produce discriminatory or biased outputs, which may impact obligations under the Equality Act 2010 and human rights considerations.
- AI and the ownership of data and materials is a complex area and there can be challenges around intellectual property, copyright and data scraping.
- Use of AI may have implications for professional indemnity and cyber insurance.
- Consumer protection law may be relevant where AI-enabled services are supplied to consumers, particularly by subscription and where your organisation buys digital products and services.
- AI carries an environmental impact. This can lead to tension between its use and existing environmental commitments made by you, your organisation or your clients.
- AI law and regulation is developing differently across jurisdictions – consider whether any additional legal or regulatory requirements apply, especially for cross-border work.

Review the section on regulation and compliance.

8. Case studies

8.1 How solicitors are using AI tools for contract review

What are the real AI use cases that solicitors are embedding – and what pitfalls should they avoid?

Understand how AI can be used in legal contract review and management.

8.2 Building your legal team's AI skill set

How best can leaders encourage their teams to adopt AI in an effective and safe way?

Explore how legal leaders can support their teams to use AI.

8.3 Lawtech, AI and innovation: how to stay ready for change

It's easy to feel overwhelmed by the pace of technological change in legal services.

Find out how you can increase technological awareness and build an innovation mindset.

9. Training, resources and more information

In 2024, we developed an AI strategy with our members to guide our AI work and technology is mission one of our 2025-2028 corporate strategy. Our first strategic goal under this mission is to support our members in technology adoption.

Read our:

- lawtech practice note
- generative AI: the essentials guide
- buying new technology guide
- in-house solicitors in an age of change: risk, AI, ethics and influence feature
- how AI tools hallucinate – and why it matters in law article

Find more insights and resources on our AI and lawtech pages.

Sign up for our training:

- Essential skills for practice: conducting legal research in the age of AI
- AI in litigation: key issues and considerations
- Risks and challenges of artificial intelligence in legal services

Explore:

- SRA Innovate – the SRA's latest research, resources and guidance
- LawtechUK – a Ministry of Justice-backed initiative for driving lawtech innovation
- CCBE – guides and recommendations from the Council of Bars and Law Societies of Europe's IT Law Committee
- ICO and AI – guidance and practical resources from the Information Commissioner's Office on AI

Expand all

Legal status

Practice notes represent the Law Society's view of good practice in a particular area. They are not intended to be the only standard of good practice that solicitors can follow. You are not required to follow them but doing so will make it easier to account to oversight bodies for your actions.

Practice notes are not legal advice, and do not necessarily provide a defence to complaints of misconduct or poor service. While we have taken care to ensure that they are accurate, up to date and useful, we will not accept any legal liability in relation to them.

For queries or comments on this practice note, contact our Practice Advice Service.

SRA Principles

There are seven mandatory principles in the SRA Standards and Regulations which apply to all aspects of practice.

The principles apply to all authorised individuals (solicitors, registered European lawyers, registered foreign lawyers and registered Swiss lawyers), authorised firms and their managers and employees, and to the delivery of regulated services within licensed bodies.

Terminology

Must – a requirement in legislation or a requirement of a principle, rule, regulation or other mandatory provision in the SRA Standards and Regulations. You must comply, unless there are specific exemptions or defences provided for in relevant legislation or regulations.

Should – outside of a regulatory context, good practice, in our view, for most situations. In the case of the SRA Standards and Regulations, a non-mandatory provision, such as may be set out in notes or guidance.

These may not be the only means of complying with legislative or regulatory requirements and there may be situations where the suggested route is not the best route to meet the needs of a particular client. However, if you do not follow the suggested route, you should be able to justify to oversight bodies why your alternative approach is appropriate, either for your practice, or in the

particular retainer.

May – an option for meeting your obligations or running your practice. Other options may be available and which option you choose is determined by the nature of the individual practice, client or retainer. You may be required to justify why this was an appropriate option to oversight bodies.

AI and lawtech | Business management | Client care | Climate change | Cybersecurity | Data protection

Cookies
Privacy policy
Accessibility
Legal notice

© The Law Society 2026. All rights reserved. - 2026.10.02.6

This page was printed on 23/10/2026 and the up-to-date version can be found online at https://www.lawsociety.org.uk/topics/ai-and-lawtech/using-artificial-intelligence-in-legal-practice