AI compliance for law firms: key regulatory considerations

by *Shawn Curran* and Natalie Cooksey, *Travers Smith LLP*

Practice notes | **Maintained** | England, Wales

A note discussing the key factors and overarching regulatory and compliance issues that SRA regulated law firms should consider when using artificial intelligence (AI) (including generative AI) and legal technology tools, both within their business and in the provision of legal services.

Scope of this note

Artificial intelligence (AI) allows computers to perform tasks that would normally involve human intelligence. Machine learning is a type of AI through which computers identify and learn from patterns within data (for example, to answer questions posed by the user), as opposed to being programmed to produce a particular output (see *Practice note, AI jargon buster*).

AI can be viewed in two categories:

- Legacy AI includes basic entity extraction and classification, such as detecting certain clauses within a contract.
- New AI includes new, transformer-based large language models that can carry out more complex cognitive tasks, such as (for example) assessing the language a customer uses in emails to predict the likelihood of that customer terminating an agreement.

The development of new AI effectively moves the technology closer to a position where it can understand concepts and draw inferences in a more human-like manner.

Reflecting a greater willingness among clients and consumers to engage with AI systems and legal technology tools, many law firms are increasingly looking to harness the benefits of such technology, both in managing their businesses and in providing legal services.

However, use of these tools creates risk and law firm risk and compliance professionals must ensure that their firm's deployment of AI (particularly new AI) is safe for their business and clients and is in line with relevant legal and regulatory obligations.

This note examines the key overarching regulatory and compliance issues that firms should consider when using AI and legal technology in their business. It focuses on issues that would be of concern to the firm's risk and compliance professionals and does not cover wider legal considerations arising out of the use or development of AI tools (for example, employment law or product development issues).

The note is aimed at SRA regulated law firms practising in England and Wales but will also be of interest to other legal services businesses.

The Legal Services Board has issued guidance to regulators on how they should ensure that their regulatory approach to AI use in their regulated communities encourages innovation and aims to widen consumer access to justice (*LSB: Guidance on*

promoting technology and innovation to improve access to legal services. The SRA's approach may therefore develop over time. We will update this resource when further information becomes available.

For information about the firm-level, client and matter-level issues that law firms should consider in relation to their use of AI and legal technology, see Practice notes:

- *AI compliance for law firms: firm-level considerations*.
- *AI compliance for law firms: client and matter considerations*.

Law firm use cases

The use cases for AI in law firms are varied and span usage for the firm's own internal purposes (for example, as part of the firm's staff recruitment or appraisal processes) and usage in connection with client matters. The SRA's *Risk Outlook* report on the use of artificial intelligence in the legal market sets out some specific examples of how AI is being used in the legal sector.

Examples of how AI may be used in law firms include:

- A chat bot on the firm's website to answer simple queries from potential clients.
- Predicting case outcomes.
- Assisting in risk assessment decisions around onboarding new clients or suppliers.
- Generating smart or automated contracts.
- Coding documents within a document review platform.
- Conducting legal research.

AI can often be used to speed up processes and reduce administrative work where a matter or task involves large volumes of repetitive action. However, such efficiency gains must be balanced against the need to safeguard against unnecessary risk, by ensuring that the AI system has been designed and implemented in accordance with "compliance by design" principles.

Appropriate training should be given to staff on any restrictions around how AI systems may be used and internal governance procedures should be established to manage and monitor risk.

In its own work, Travers Smith LLP (Travers Smith) has organised potential AI use cases into two categories:

- Generative use cases, for example, the production of a first draft legal agreement by a large language model with the ability to search through vast amounts of data and use the information sourced from that data to produce a human-like answer to a question or "prompt" posed to the system.
- Extractive use cases, whereby specific search terms are applied to a pre-defined dataset in order to extract information from that dataset.

At present, and due to the risks inherent in generative use cases (for example, hallucination and potential copyright breaches), we anticipate that many law firms will wish to focus on maximising the benefits of extractive use cases.

Travers Smith's view is that, on the extractive side, where the firm is able to supply a model with data and have the model conduct reasoning based on questions asked of that data, there can be more confidence in the model's accuracy and less risk of an IP infringement arising through the use of any output generated (given that the firm is able to control the input).

For further information about potential IP issues, see *Practice note, AI compliance for law firms: firm-level considerations: Intellectual property (IP) issues*.

Key regulatory and compliance obligations

However the firm decides to use AI within its business, it will need to ensure that its use complies with the firm's overall regulatory obligations.

The *SRA Standards and Regulations* set out the key regulatory and compliance obligations for SRA regulated law firms. In particular, firms (and the individuals working for them) must comply with the *SRA Principles*, which comprise the fundamental tenets of ethical behaviour that the SRA expects all those that it regulates to uphold.

The following Principles are particularly relevant to law firm AI use:

- **Principle 2:** you act in a way that upholds public trust and confidence in the solicitors' profession and in legal services provided by authorised persons.
- **Principle 3:** you act with independence.
- **Principle 6:** you act in a way that encourages equality, diversity and inclusion.
- **Principle 7:** you act in the best interests of each client.

In relation to Principle 7, the SRA says that client best interests must remain at the centre of law firm decisions about the use of technology. This means that firms should have appropriate governance, systems and controls to ensure they are using any such technology responsibly (see *SRA: SRA Innovate: Compliance tips for solicitors*).

Other regulatory obligations that are relevant to AI use are set out in:

- The *SRA Code of Conduct for Firms* under:
    - rule 2 (Compliance and business systems); and
    - rule 3 (Co-operation and accountability).
- The *SRA Code of Conduct for Individuals* under:
    - rule 1 (Maintaining trust and acting fairly);
    - rule 6 (Confidentiality and disclosure);
    - rule 7 (Co-operation and accountability); and
    - rule 8.6 to 8.11 (Client information and publicity).

The firm must also comply with data protection legislation, and individual solicitors must observe their fiduciary obligations under the common law.

Upholding public trust and confidence

Principle 2 of the *SRA Principles* requires firms and individuals to act in a way which upholds public trust and confidence in the legal profession.

Clients and consumers are often keen to engage with AI systems and legal technology tools where this may reduce legal costs (thereby enabling greater access to legal advice). To ensure compliance with Principle 2, the firm should be transparent about its use of AI tools and give consumers enough information to understand both the benefits arising from, and any limitations inherent in, their use.

Transparency with clients and others

In its 2023 consultation, the Legal Services Board stated that:

> "as adoption of technology increases, particularly technology that creates the perception of human interaction, the need for transparency about its deployment arguably becomes greater." (See *LSB: Consultation on draft guidance on promoting technology and innovation to improve access to legal services (10 July 2023)*.)

This concern for transparency, particularly with consumers, has been carried into the LSB guidance published following the consultation (*LSB: Guidance on promoting technology and innovation to improve access to legal services*).

This has a strong consumer focus and, in particular, suggests regulators consider how best to ensure that:

- Consumers are informed about the benefits of technology in legal services, costs, quality and routes of redress.
- Their regulated communities are transparent with clients about how technology may be used in their matter.

Examples of how the firm might achieve transparency with clients in practice include:

- Where the firm places a chat bot on its website, the firm should make clear to website visitors that the chat bot is software-driven and that any information produced:
    - is generated automatically;
    - does not involve consideration of the client's specific legal issue by a qualified adviser; and
    - is not a substitute for detailed legal advice.
- Where a client has instructed the firm and wants AI tools to be used on their matter to minimise legal spend, the firm should give a detailed explanation of the specific capabilities of the relevant technology and any risks that might arise, as well as providing information about how liability issues would be addressed (see *Practice note, AI compliance for law firms: client and matter considerations: Limiting liability*).

The firm should also consider the need to be transparent with opponents, tribunals and other external stakeholders about the use of AI, in line with the following provisions of the *SRA Code of Conduct for Individuals*:

- Rule 1.2, which says practitioners must not abuse their position by taking unfair advantage of clients or others.
- Rule 1.4, which says practitioners must not mislead or attempt to mislead their clients, the court or others.

Transparency and data protection

Transparency is furthermore a key principle under data protection legislation and the firm should assess whether its privacy policy should include reference to the use of AI and legal technology tools insofar as personal data is uploaded to those tools (see *Confidentiality*).

If the firm uses legal technology to make automated decisions, it should maintain as complete a record as possible of how each decision was arrived at (for example, a list of the relevant factors taken into account), and should regularly check the outputs for any evidence of unintended discrimination. This will help the firm justify its decisions and actions as required by rule 2.2 of the *SRA Code of Conduct for Firms* and rule 7.2 of the *SRA Code of Conduct for Individuals*.

The firm must also comply with any relevant obligations pertaining to an individual's right to opt-out of automated decision making and their right to request a human review. The ICO has published guidance on this (see *ICO: Automated decision-making and profiling*). For further information, see *Data protection*.

Service and competence

Regulatory obligations

The *SRA Code of Conduct for Firms* requires firms to:

- Keep up to date with, and follow, the law and regulation governing the way the firm works (*rule 3.1*).
- Ensure that the service the firm provides to clients is competent and delivered in a timely manner, and takes account of client needs, attributes and circumstances (*rule 4.2*).
- Ensure that the firm's managers and employees are competent to carry out their role, and keep their professional knowledge and skills, as well as their understanding of their legal, ethical and regulatory obligations, up to date (*rule 4.3*).
- Have an effective system for supervising client matters (*rule 4.4*).

These requirements are reflected in rule 3 of the *SRA Code of Conduct for Individuals*, which also contains a requirement that where a practitioner supervises or manages others providing legal services, they:

- Remain accountable for the work carried out through them (*rule 3.5(a)*).
- Must effectively supervise work being done for clients (*rule 3.5(b)*).

Implicit in these rules is a requirement that firms and practitioners should keep up to date with technological advances that may be used for the benefit of clients.

Additionally, firms should ensure that they have the expertise to appropriately configure any legal technology systems they use, as they will remain accountable for any work generated by such systems.

Human supervision of output

A firm's use of an AI system does not in any way reduce the supervisory burden, or the obligation on the firm or on individual practitioners to provide a competent service. In most cases, AI output should therefore be curated and monitored by a human to ensure the accuracy and completeness of the information provided. Any constraints should be fully outlined to the client. For example, where an extractive AI model (that is, one that is designed and trained to recognise patterns to extract relevant results) is used to review a suite of hundreds of documents, it may not be feasible for the supervising individual to go through and manually check the entirety of the output. In these circumstances, the firm should give the client a clear explanation at the outset of how the AI tool will work, so that they have a comprehensive feel for the reliability of the results of the review. This explanation should include:

- Information about how the system has been trained.
- The details of any search terms to be used.
- The scope of any second level review to be undertaken.

Fallibility and appropriate use

While AI tools can offer significant time savings, they are not infallible. This is particularly evident in the field of legal research. Generative AI systems (GenAI) are known to "hallucinate" results, creating outputs that are false or non-sensical, often as a result of the AI system having incorrectly perceived a pattern within the data on which the system was trained (for a further discussion of this topic, see: *Hallucination is the last thing you need*). For these reasons, the firm should not use AI tools for legal research in isolation and should ensure any results generated are reviewed by humans and validated against other research sources in the usual way.

Practitioners may see generative AI systems as a way to extend their practice into areas in which they have less experience. However, practitioners should not be tempted to stray outside of their area of expertise in reliance on these tools, as human expertise will always be needed to competently monitor and assess the accuracy and quality of any AI outputs. The use of any form of AI does not reduce practitioners' obligation to keep the knowledge and skills required for their role up to date, under rule 3.3 of the *SRA Code of Conduct for Individuals*.

However, AI can be extremely helpful for ensuring compliance with straightforward tasks, such as electronic court disclosure and filing obligations.

Monitoring and controls

Regulatory obligations

The *SRA Code of Conduct for Firms* provides that firms must:

- Have effective governance structures, arrangements, systems and controls in place that ensure they comply with all the SRA's regulatory arrangements, as well as with other applicable regulatory and legislative requirements (*rule 2.1(a)*).
- Keep and maintain records to demonstrate compliance with the firm's obligations under the SRA's regulatory arrangements (*rule 2.2*).
- Identify, monitor and manage all material risks to the business (*rule 2.5*).

When developing and implementing any new AI or legal technology system, the firm should build in measures to ensure compliance with these obligations. The firm should document the steps that have been taken to build in compliance and to implement risk management procedures and governance structures when designing a new AI or legal technology system.

Assessing the risks and opportunities

There are various resources available that can assist the firm to record how it has assessed the legal issues and regulatory framework that apply to its business in respect of a particular AI project. Examples include:

- Microsoft's responsible AI impact assessment template, which sets out some of the key questions that the firm may wish to consider (see *Microsoft: Responsible AI Impact Assessment Template (June 2022)*). This includes a section for capturing any potential system benefits and harms, as well as sections to record the details of any human oversight and control and how the system will be used to inform decision making.
- The NIST AI Risk Management Framework, which links to a playbook setting out suggested actions for achieving outcomes in relation to governance, mapping, measurement and management of AI functionality (see *NIST: AI Risk Management Framework* and *NIST: AI RMF playbook*).
- In relation to data protection considerations, the CNIL's self-assessment guide for AI systems, which covers topics such as: "Asking the right questions before using an artificial intelligence system", "Collecting and qualifying training data", and "Achieving compliance" (see *CNIL: Self-assessment guide for artificial intelligence (AI) systems*). For further information, see *Data protection*.
- Small and medium-sized firms may also wish to consult The Law Society's recent guidance entitled *Generative AI – the essentials*.

Ensuring compliance

The firm should:

- Maintain records covering the source of the data on which the system has been trained, whether that data has been cleansed and how the firm has tested the system before making it available for general use by the business.
- Implement a structured quality assurance programme, both before and after roll-out, so that the technology can be refined as needed. A key issue to look out for will be whether any output generated by the system is influenced by any implicit biases present in the data.
- Ensure that individuals using these systems receive an appropriate warning about the risk of biased outputs impacting on the fairness or appropriateness of any decisions made in reliance on the outputs.

A benefit of developing AI solutions in-house (as opposed to procuring a third-party system) is that firms with their own in-house system may be able to exercise more control over the testing process and the way in which any issues are remedied. They will also be better placed to explain to clients and regulators how the system operates and the logic deployed in producing outputs.

On this final point, "agent-based" AI technologies (that is, technologies that utilise independent entities or "agents" that can operate autonomously and interact with each other in complex ways to solve problems) enable law firms to establish and define the reasoning applicable to their AI workflows by breaking those workflows down into various steps, with each step being more easily explained individually (using the agent-based AI) than when looking at the model as a whole. When using such technologies, there should be a human in the loop to validate and verify the reasoning, and firms should ideally focus on lower-level tasks, so as to avoid limiting their ability to explain the AI's reasoning (which may be the case when setting the AI a higher-level agent-based goal).

When designing appropriate governance structures, firms may wish to consider the following:

- Appointing a senior individual to have oversight of the use of the system. The SRA has said that it expects *compliance officers for legal practice* "to be responsible for regulatory compliance when new technology is introduced" (see *SRA: SRA Innovate: Compliance tips for solicitors*).
- Setting up a committee with responsibility for AI usage at the firm, with membership comprised of senior stakeholders and technical system experts, and ensuring that all staff are aware of how and when AI-related concerns should be escalated to the committee.
- Carrying out regular audits of the firm's deployment of AI systems to assess functionality and effectiveness and to identify any potentially worrying outcomes or risks. Where a specific risk is identified, it should be recorded in the firm wide risk assessment and the risk register.

The firm should also ensure that any governance structures can operate agilely, so that the firm can respond quickly to an ever-evolving landscape of regulatory and competitive requirements and challenges.

Some firms have taken a "bottom up" approach to deploying AI tools, providing access to such tools to everyone within the firm and allowing them to decide how best to use it, However, from a governance perspective, a better approach may be to permit access on a controlled basis to specified user groups that are assigned to investigate specific issues. This will allow feedback to be easily collated, giving a centralised overview of where the AI performs well (and can be deployed in ways that generate cost savings and benefits for clients), as well as the use cases where it performs less well, enabling senior management to focus development efforts appropriately.

For more information, see *Article, AI governance, risk and compliance: shaping an unknown future: Tips for effective AI governance*.

Confidentiality

Maintaining client confidentiality is one of the key requirements that firms and practitioners must be able to satisfy under rule 6.3 of both the *SRA Code of Conduct for Firms* and the *SRA Code of Conduct for Individuals*.

Over time, as reliance by firms on third-party technology systems has increased, compliance teams have had to balance the risks to client confidentiality that arise from sharing data with third parties against the efficiencies and cost-savings generated by legal technology solutions.

For information about some of the key data protection issues to consider when procuring an AI or legal technology system from a third-party supplier, see *Data protection*.

As far as client information is concerned, Rule 6.3 of the SRA Codes of Conduct imposes an obligation to "keep the affairs of current and former clients confidential unless disclosure is required or permitted by law or the client consents".

Training an in-house AI system

Where AI capability is being developed in-house using an enterprise version of an AI model (instead of a consumer version), the firm should consider:

- Whether the model the firm is using already has sufficient capability to perform a specific task, such that no fine tuning is required (because the dataset on which the base model has been trained was sufficiently broad or detailed, as appropriate).
- Where fine tuning is required in order to further train the model, whether any of the data the firm intends to upload to the AI system may be confidential.

For example, if the firm is seeking to teach the system how to identify specific provisions typically found in a share purchase agreement, it would upload multiple examples of that type of agreement into the platform. To comply with the firm's duties of confidentiality, the firm should cleanse the dataset by redacting any client-specific information from the agreements before uploading them to the system. This may also negate any potential need to obtain client consent to use the documents in this way.

If a particular precedent that the firm intends to upload contains the intellectual property (IP) of another law firm, the firm should consider whether uploading the material (even with redactions) could potentially infringe that firm's IP rights.

In all cases, the output generated by the system should always be checked by a human to ensure that it does not contain any confidential information relevant to other clients or any third-party IP rights.

Using external AI systems

Where the firm is using an external or consumer version of an AI system, the risks to client confidentiality are potentially greater. There has been a lot of concern in the legal industry that large language models such as those deployed by ChatGPT may use data uploaded to the tool by users for the purpose of re-training the model and may also make that data available as an output to all users across their customer base.

For example, if a solicitor subscribes to a consumer instance of ChatGPT and uploads a client contract to the platform, that contract will be ingested into the ChatGPT system. In the worst-case scenario, this may lead to extracts of the contract being made available to other users across the globe in their own outputs from the system.

For this reason, many firms have established hard parameters around staff use of such technology. Examples of measures firms have taken include:

- Blocking the use of external AI tools, by ensuring that websites through which the software is available are not accessible to users on firm-provided devices. This has the advantage that users will not easily be able to upload documents from the firm's systems into such sites. However, the speed at which new offerings are being launched can made it difficult to keep track of which sites need to be blocked. In addition, there is always the risk that users may access the sites using other devices. (although the fact that most firms have data loss prevention tools in place to ensure

that information is not improperly extracted means that the risk of client information being uploaded to such sites is smaller).

- Implementing user policies, setting out what appropriate usage of AI tools looks like. For example, this might involve an express restriction on uploading client information or personal data to any tool that has large language model capabilities. However, having a user policy on its own (without any associated technical controls such as blocking consumer instances of generative AI websites) is likely to be insufficient in controlling how AI tools are used, given that it is not possible to track which data is inputted into these websites.

  For more information about user policies, see *Practice note, AI compliance for law firms: firm-level considerations: Policies, values and appropriate use*.

- Negotiating with AI providers for terms that offer better protection than is offered under their standard consumer licence. For example, Microsoft offers an enterprise licence that makes the functionality of the ChatGPT models available to users, without any data that is uploaded into the system being used to train those models.

Depending on the approach taken, it may be necessary to implement a system for checking what is being inputted into such systems, in addition to checking the output. One way of doing this is to introduce a technology system that applies governance rules at a technical level, by controlling what users can put into the AI system, and what the AI system is permitted to generate as output. One example of this is YCNBot, an open source web application developed by Travers Smith, which sits above users' access to AI tools and acts as a gateway that is capable of enforcing an organisation's policies around the use of AI (for example, the system can be set up to prevent users from submitting prompts that very obviously include personal details, such as a person's name, and to prevent users from copying case law summaries generated by the tool, so as to avoid any hallucinations being unintentionally replicated). These types of systems also allow greater auditability in terms of the content of, and output generated by, specific prompts.

Privilege

When using AI and legal technology systems in the context of any client matters where litigation is reasonably in prospect, the firm should consider the importance of preserving clients' potential litigation privilege. For a client to successfully claim litigation privilege in relation to information the firm holds for them, the confidentiality of the relevant information will need to remain intact. It is therefore necessary to ensure that the sharing of such information, and any third-party access to it, is appropriately limited. It would therefore be sensible to restrict any privileged information from being uploaded into any large language model used by the firm

Data protection

Data protection issues often go hand in hand with confidentiality considerations for law firms. Conducting appropriate due diligence on any proposed third-party supplier involved in the provision of an AI service is an important task.

Data protection principles

The UK General Data Protection Regulation sets out seven key data protection principles:

- **Lawfulness, fairness and transparency:** you must ensure that you have an appropriate legal basis for processing data, and must process personal data fairly and in a transparent manner in relation to individuals.

- **Purpose limitation:** you must collect personal data only for specified, explicit and legitimate purposes and must not further process it in a manner that is incompatible with those purposes.
- **Data minimisation:** you must limit your processing of personal data to that which is adequate, relevant and limited to what is necessary in relation to the purposes for which the data is processed.
- **Accuracy:** you must ensure that personal data held by you is accurate and, where necessary, kept up to date.
- **Storage limitation:** you must keep personal data in a form which permits identification of data subjects (that is, the individuals to whom any personal data that you are processing belongs) for no longer than is necessary for the purposes for which the personal data is processed.
- **Integrity and confidentiality (security):** you must process personal data in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures.
- **Accountability:** you must be able to demonstrate your compliance with the obligations set out above.

(*Article 5, UK GDPR*.)

The ICO has published detailed guidance on *Artificial intelligence (AI) and data protection*, which provides a detailed overview of how to apply the principles of the UK GDPR to the use of information in AI systems.

Using anonymised data

If data used to train an AI model has been fully anonymised, it falls outside the scope of the data protection legislation. If the firm is training an in-house model, it should therefore consider whether anonymisation is possible before uploading any training data.

It may also be possible to build the model in such a way that it recognises and blocks any personal data from being transmitted over the system, either through a "human in the loop" monitoring of the prompts used and resulting outputs, or by using software that can automatically detect where inputs or outputs contain, for example, an individual or business name, as with YCNBot, or other recognisable information such as a home address or details of a person's occupation (see, *Using external AI systems*).

Using personal data

The use of personal data may be an integral part of some AI or legal technology systems (for example, AI tools used as part of staff recruitment or appraisal processes). Where the firm plans to use AI for a purpose which requires personal data in this way, it should carry out a thorough data protection impact assessment before implementing the relevant system (see *Data protection impact assessment*).

Data protection impact assessment

The firm should complete a data protection impact assessment before implementing any new AI or legal tech system. This is a legal requirement in certain circumstances under Article 35(3)(a) of the UK GDPR if the firm's use of AI involves:

- Systematic and extensive evaluation of personal aspects based on automated processing, including profiling, on which decisions are made that produce legal or similarly significant effects.
- Large-scale processing of special categories of personal data.

- Systematic monitoring of publicly-accessible areas on a large scale.

The impact assessment should:

- Identify the purpose of the data collection.
- Identify the legal basis for processing.
- Document any key risks that may arise.
- Record how those risks will be addressed, including any information security measures that may need to be implemented.

Completing the impact assessment process enables the firm to identify and record where data will be held and who will have access to it. It will also help the firm ensure that the system it intends to use has appropriate technical and organisational measures in place to be able to effectively implement the data protection principles (see *Data protection principles*) and safeguard individual rights, in line with the relevant ICO guidance (see *ICO: Guidance on Data Protection by Design and Default*).

For further information on data protection impact assessments, see *ICO: Guidance on Data Protection Impact Assessments* and *Practice note, Data protection impact assessments (DPIA) (UK)*.

Particular data issues to consider

Particular data protection issues that firms should consider in relation to their use of AI include:

- **Jurisdiction.** If any personal data uploaded to the system will be held in a jurisdiction other than the UK, or an overseas jurisdiction that has been designated as providing an adequate level of protection for personal data, the firm will need to ensure that an appropriate transfer mechanism is in place and consider whether the personal data is likely to be accessed by the courts or law enforcement in that jurisdiction. For further information, see *Toolkit, Data transfers toolkit (UK and EU)*. The ICO has published guidance on this (see *ICO: Guide to International Transfers*).
- **Retention and deletion.** The firm should consider the length of time for which data will be retained in the system and how personal data may be deleted from the AI system if requested by a data subject. These are difficult issues to address, particularly where personal data may have been used in training a large language model, since it will likely be extremely difficult to track the way in which personal data may have been incorporated into the model. For further information, see *Practice note, Data retention policies (UK)*.
- **Input parameters.** The firm may wish to introduce some red lines around the types of information that can be inputted into the system. These red lines should be set out in a policy governing use of the system (for example, each time a person accesses the system, it may be helpful to ask them to agree a set of terms outlining what appropriate usage looks like) (see *Practice note, AI compliance for law firms: firm-level considerations: Appropriate staff usage*). For a template usage policy, see *Standard document, Generative artificial intelligence in the workplace policy*.
- **Client transparency.** The firm should consider how it will inform data subjects about exactly how their data will be handled and how any decisions about them involving the use of AI will be made (see: ICO: Explaining decisions made with AI). While firms ordinarily achieve this via a privacy policy on their website or in their terms of business, the firm should also consider "point of use" notifications that provide transparency information to clients immediately before they upload any of their personal data to the system. The SRA has warned that consumer trust will be eroded if data is used in ways that clients do not anticipate or which do not benefit them, even if that use complies with data regulations (see *SRA: SRA Innovate: Compliance tips for solicitors*). This can also damage the firm's reputation. Transparency is

therefore key to maintaining public trust and confidence in the profession's use of AI. For a template privacy policy, see *Standard document, Law firm privacy policy*.

For further information about data protection considerations in supply chain compliance, see Checklists:

- *Contracting for AI: anatomy of an AI project*.
- *What to look out for in a LawTech implementation contract*.

For further guidance on AI and data protection, see *ICO: AI and data protection risk toolkit*.

Further reading

For further Practical Law resources on the use of AI and particular issues for law firms, see:

- *Legal technology and artificial intelligence toolkit*.
- *AI toolkit (UK)*.

END OF DOCUMENT