Contents

Chapter 17: Giving clients what they want (as well as what they need)

By Joanne Brook, Lionshead Law

Chapter 2:
Navigating the ethics of AI and legal practice

By Harry Borovick, general counsel and AI governance officer, Luminance

Why humans matter

The nuances of human interaction remain central to the delivery of good legal services. Whether it is the ability to dig into questions or complex issues, the ability to understand context outside of the words spoken or the facts provided, or understanding the subtlety of human stress (typically in legal matters for all parties), AI systems are a long way off the capabilities of a competent or effective human.[1] Providing an experience equal to that of interacting with a human is unlikely by 2030, even if AI systems are rapidly approaching (and may exceed) the on-paper technical proficiency of many lawyers, particularly at scale. This will likely remain the case, even when digital avatars give end users or practitioners the *impression* of interacting with a human.

Applying AI in legal practice

The applications of artificial intelligence in law are varied and remain in the nascent stages of adoption. When considering the application of AI to legal practice:

> *"The reality, as most legal professionals know, is that the daily work of a litigator varies tremendously from that of a transactional lawyer. Even within similar areas of legal practice, the day-to-day can be diverse."*[2]

For example, within practice areas that can be broadly described as "transactional", the vast majority of work carried out by a commercial real estate lawyer involves distinct and specialized tasks compared to the work of a capital markets lawyer, an employment lawyer, or even a real estate finance lawyer. Within litigation, many of the tasks faced daily by a criminal defence lawyer in a small legal practice will be almost unrecognizable compared to those handled by a civil litigator in a large law firm, even if there are similarities regarding the high-level principles of the work.

That said, many tasks remain the same across legal practice areas. Lawyers still mostly work at desks, send emails, make notes, research case law and/or statutes, interpret guidance, analyse legal documents, and deliver an output of some kind to an end client. That end client may be a colleague within the same company or a third party.

Many responsibilities also remain the same. Therefore, for legal practitioners and law firms, when considering how to navigate the legal risks and ethical challenges of the future, it is worth starting with first principles and understanding that the primary duties of most lawyers remain as follows:[3]

- Duty to the courts;
- Duty to the client;
- Duty to the profession; and
- Duty to oneself and wider society.

For simplicity, this chapter will not extensively deliberate what it means to act with "good ethics" – this is far too subjective. Instead, we can use these four duties as the de facto ethical framework, and determine whether behaviour or actions are likely to be "ethical" based on how courts and regulators have previously spoken, published, acted or enforced on legal practitioners, as well as the providers and end-users of AI systems.[4]

Navigating ongoing ethical challenges

For legal practitioners developing, deploying, or offering AI systems to serve their clients, it may be hard to know where to begin in terms of practical guidance that can be applied across multiple legal practice areas. At the date of publication, this remains a major challenge – particularly where international elements affect either the location of the practitioners or their client base.

Looking to well-established regulatory guidance or internationally recognized risk management frameworks is a practical first step in setting up an AI risk management approach within a legal practice. Ethical issues should not be considered separately from AI risk management; rather, they should be at the core of balancing practical opportunities against tangible and intangible AI risks.

Beyond regulatory compliance, law firms should proactively consider ethical challenges posed by AI systems in legal practice, even in the absence of strict legal or regulatory requirements. A non-exhaustive list of key considerations includes:[5]

- *Transparency*. The potential need, or absence of need, for clients to understand when and how AI is being used in their cases or transactions.
- *Accessibility*. Whether clients are best served, and access to justice is best achieved, by removing human elements in provision of legal services (e.g. where the only affordable advice may be AI-generated) and how best to deliver that service safely.
- *Bias and fairness*. AI models trained on biased, insufficiently diverse, or excessively broad but unrefined legal data can drive, perpetuate, or cause disparities in legal outcomes.
- *Confidentiality and data security*. AI tools that interact with client data must adhere to appropriate security standards to ensure that privileged or sensitive information remains protected.
- *Human oversight*. While AI can enhance efficiency, final legal advice or decisions should always involve human judgment to maintain professional integrity, unless this materially conflicts with the other considerations and would cause detriment (notably if the human oversight pushes the cost of access to legal services out of reach).

Even recognizing that some ethical risks remain unknown within a legal practice is itself a positive step, as it demonstrates an understanding that the level of ethical risk is difficult to quantify. But, practically, legal practices need *somewhere* to begin – ideally, right now. So, what framework should be used?

The National Institute of Standards and Technology (NIST), part of the US Department of Commerce, produces one of the best-established frameworks for AI risk management and deployment. This framework[6] is designed to be relatively sector-agnostic and provides a strong foundation for legal professionals integrating AI into their practice. It is also well-suited, arguably better than any other internationally-respected framework, as a starting point for considering how to build, maintain, and enforce guardrails for ethical conduct by users and deployers of AI in legal practice. Between now and 2030, the AI Risk Management Framework is likely to be updated, and competing (or even, better) frameworks may become widely used. But for now, this chapter will lean on this framework to set out a four-part strategy for practically navigating the use and impact of AI on legal services.

Its four core components (Govern, Map, Measure, Manage) are heavily simplified below, but in each case it is critical to break these down into practical realities.

1. Govern

What it means
Building a culture within an organization that prioritizes risk management, with clear processes, policies, guidance, and reporting across the entire lifecycle of a company's AI usage, whether for internal or external purposes.

In practice
Committees: Establish AI ethics committees within firms to oversee AI applications.

Training and competence: Encouraging new or existing staff to take on roles of AI governance. Practically, they should do so when appropriately educated so firms paying/supporting staff to engage in academic or professional training courses and certifications is sensible. Many such courses already exist, and a whole new market of AI governance training is springing up. Reputable courses such as those provided by the IAPP,[7] various universities, or even tech companies can also be attractive to employees as a career-enhancer.

Policies: Develop AI governance policies that align with professional legal ethics. Many companies and legal practices have developed AI governance policies, or are considering doing so. It is critical to not create rigid policies in a vacuum that neglects to consider existing regional and international professional standards (e.g. SRA guidance in the UK, CIaRB guidance for arbitrators, ABA guidance in the US, etc.).

Transparency: Actively encourage client-facing and internal transparency in AI decision-making, particularly when AI tools assist in legal analysis or case prediction. Whether it is appropriate, or perhaps even legally required, to make transparent disclosures to third parties, such as the court or the counterparties, will be a matter of law and regulatory guidance that varies between jurisdictions (but definitely something to bear in mind).

2. Map

What it means
Conducting an initial assessment of AI usage, opportunities, and risks to

effectively govern AI deployment. This is essential to build the initial understanding of how AI usage within a legal practice aligns with ethical and legal standards.

In practice

Legislation, regulation, and guidelines: Identifying applicable or relevant legal and ethical frameworks, e.g. human rights legislation such as the European Convention on Human Rights, bar/legal practice codes of conduct such as the SRA Code of Conduct, data privacy legislation such as the GDPR, customer-specific legal frameworks such as DORA (in terms of hard law), or guidance by regulators such as the FCA and SEC.

Appropriate applications: Identify legal tasks where AI can improve efficiency while managing ethical compliance.

- This should be done systematically (overlapping with the "manage" stage of the NIST AI Risk Management Framework below).
- Common applications of AI in legal practice include:[8]
  - Document review and management;
  - Contract negotiation;
  - Evidence analysis;
  - Case evaluation;
  - Drafting of legal documents (e.g. legal letters, filings, and opinions);
  - Legal research;
  - AML/KYC management;
  - Case management; and
  - Billing.
- In future, AI in legal practice may more commonly include many more complex applications such as:
  - Case success probability analysis;
  - Auto-negotiation of negotiations, settlements and disputes; and
  - Low-value/low-risk judicial operations and judgments.

Bias: Recognising the potential biases within AI systems and implementing measures to mitigate them. Quite simply, AI systems are built on data. Where that data is reliant on particular subsets of data, it may be more prone to bias. Bias can be good. It is important to recognise that bias isn't always bad (a bias can be trained into an AI system to give a desired or advantageous outcome, such as the way a contractual clause's risk should be considered).

The negative impacts of bias in AI systems is materially aggravated where the output creates a greater risk to individual rights and freedoms or systematically discriminates against individuals or groups of individuals.[9]

As an example – with respect to access to justice:

- A legal practice may be entirely meeting its duties (1-4 set out earlier in this chapter) if it uses an AI system to assess the probability of success in a case, and therefore whether it should take on that case, if it can be relatively confident that the AI system is not providing biased evaluation based on the individual characteristics of the potential client (such as ethnicity, religion, or other personal characteristics of this kind). The probability that the legal practice is acting within its duties is increased if it has properly "mapped" the risks of bias within the AI system's data or at least recognised the potential for bias. Practically, this can be as simple as having a clear understanding of where the training data comes from (e.g., is it generalist data or a specialist data set to minimise risk). Without mapping such risks, they cannot be addressed or mitigated.
- If the legal practice decides not to provide legal representation for an individual and that decision was based solely or materially on the AI system's determination, without any human mapping of, or mitigation of, biases by the system, it is unlikely that it is acting justly and within its duties.

Oversight: Assessing the extent of human oversight required to maintain legal integrity in AI-assisted tasks is critical to success in deploying AI systems in a manner that complies with overriding legal duties.[10] Once potential legal frameworks, use cases, and bias risks have been mapped, it is then appropriate to determine and plan for how much human oversight is productive to usage of AI in an ethical manner. For example, it is relatively easy to predict that full AI automation (rather than augmentation/acceleration and then verification by a human) of legal research and the provision of a legal opinion or court filing on that basis is not acceptable in the eyes of the courts, clients, the wider profession, or society as a whole.

There have been numerous cases of lawyers reprimanded for "use" of unverified legal cases that do not exist. While that kind of incident is likely to become less common due to increasing sophistication and accuracy of the underlying technology, it remains a strong illustration of what is expected

of a good legal practice. Courts, clients, and the wider profession expect that when legal practitioners act for clients they are doing so while applying genuine professional expertise and appropriate human oversight. Therefore, mapping where that oversight is best applied (and how/when), is critical for successful usage of AI in legal practice – both practically and ethically.

This extract from the chapter 'Navigating the ethics of AI and legal practice' by Harry Borovick is from the title The Law Firm of 2030 – How the future law firm might look, *published by Globe Law and Business.*

www.globelawandbusiness.com/books/the-law-firm-of-2030