Contents

Executive summary

AI is reshaping the legal industry, transforming practices, and influencing behaviours of legal tech buyers and vendors alike. However, AI adoption introduces complex risks that require careful consideration. *The Risks of Artificial Intelligence in Law* is devoted exclusively to AI risks in the legal profession, drawing upon the candid insights of over 30 legal professionals and technologists across various legal jurisdictions and the legal ecosystem.

Chapter one first outlines the scope and methodology of this book. It details the market validation survey process the authors conducted on social media and the interview process, including the selection of the interviewees and the conducting of the interviews. It then defines key terms of legal AI that feature throughout this book, and in providing an overview of the different roles across various jurisdictions covered in the interviews, the authors examine how the key personnel interviewed view their responsibilities in regard to AI risks. The authors then summarise the substantive chapters of this book.

Chapter two examines the data privacy risks faced by law firms and in-house legal teams when contracting with external vendors that develop and/or deploy AI products and services used by legal teams – first by differentiating the newer privacy risks specific to AI from the privacy risks concerning earlier technology used by lawyers. It covers how the strict professional and ethical obligations (particularly those imposed by external governance bodies like bar councils) to safeguard confidential client information impact legal teams' decisions regarding the use of AI, bearing in mind that a fundamental tenet of the legal system and process is lawyer–client privilege.

This chapter highlights the lack of alignment in the interests of legal teams and external vendors, delves into how legal teams and their vendors allocate risk and ensure alignment in risk appetites, values, and objectives, and covers the information asymmetry between the legal team and the external vendor regarding the technology, such that an external vendor can, in theory, violate ethical rules and contractual agreements regarding

how it uses the data provided by legal teams – all without the legal team's knowledge. This is contrasted with an underexplored topic – how legal professionals engage with their internal procurement processes (which legal professionals are often not held accountable for in practice). The authors consider the adequacy of existing safeguards – both practical (including privacy-enhancing technologies such as anonymisation) and legal (such as GDPR-compliant privacy policies and data protection agreements) – specifically by examining the situation in practice, including whether non-approved use of confidential client information by external vendors is rampant. Finally, the authors examine the extent to which legal teams successfully enforce such agreements and policies in practice as part of their overall vendor risk management process.

With external vendors accessing and using legal teams' proprietary data to fine-tune and improve the AI model, there are issues concerning not only data governance and privacy but also intellectual property. Chapter three examines the risks of intellectual property infringement by external vendors. Copyright infringement might, oddly enough, be a bigger problem in the legal context, as compared to other contexts like the literary/artistic – it is significantly easier to identify the authors of novels that are "stolen" to train AI models, whereas datasets in law mostly lack obvious "owners", given the informative nature of the data (which is also a factor against a successful finding of copyright infringement under various legal jurisdictions' laws, as compared to literary/artistic works). Through interviews with legal professionals, the authors examine the effectiveness of license agreements that limit external vendors' access to and use of data, and consider if the legal team has put in place appropriate accountability measures in instances of intellectual property infringement.

By a significant margin, the biggest complaint by the lawyers surveyed by the authors involves how untrustworthy and unpredictable AI is. In chapter four, the authors acknowledge that the picture is more nuanced than all of AI use and output being unreliable – the trustworthiness of AI depends on the specific use case. AI's utility is acknowledged for more straightforward tasks like summarization and information extraction (such as references and rules) in relation to documents fed into a "closed universe" system (although it can still be inconsistent in handling documents). However, legal AI notably struggles for higher order questions regarding industry-specific concepts and bespoke legal provisions, or tasks involving synthesis or analysis.

Acknowledging the rate at which legal AI is developing, chapter four

focuses on the inherent traits of both legal work and AI that raise serious questions regarding their compatibility. For one, the costs of unreliable information are significant – beyond mere embarrassment in front of clients, filing briefs containing cases or quotes or references made up by ChatGPT may be tantamount to committing a fraud on the court. Many decisions to be made in legal work require context that cannot be captured in PDF documents to be read by AI models (despite their larger and larger context windows as previously mentioned) – the client's personality is one such "soft" factor part of the necessary context. Related to that is the reasoning process leading up to those decisions, involving creative interpretation, argumentation, and application to human circumstances.

In contrast to human lawyers, AI (or perhaps technology in general) may appear reluctant to acknowledge ignorance and uncertainty, but rather hallucinates, often in a confident-sounding manner that compounds AI's untrustworthiness in the eyes of legal professionals. The data dependency of AI makes substantive legal research likelier to lead to unreliable output in legal jurisdictions where there is no publicly available dataset (or even if it is publicly available, the relevant law prohibits its bulk downloading or use for AI training), such that AI output is derived from law firm website commentaries rather than case law. Even as AI model capabilities improve, they may also rely on more data from the internet (including misinformation and disinformation), such that their output would not be significantly improved. In focusing on the above tensions, the interviews will delve into how legal teams mitigate the risks of reliability by validating both the use of and output from AI. Related to that is how legal teams deal with the opaque "black box" nature of AI antithetical to the legal reasoning process, specifically whether their AI systems are designed with clear mechanisms for transparency and explainability as to their decision-making processes.

Chapter five examines the risks associated with the cost effectiveness of AI on operations in legal teams. On the one hand, the opportunity cost of not using AI effectively and missing out on its efficiency gains is a risk in itself. On the other hand, if learning and using AI tools results in a low or even negative return on one's time, that presents another risk affecting the bottom line of the organization. As the previous chapter on reliability risks has demonstrated, the time taken for legal teams to check the accuracy of AI-generated output or even redo parts of it can be substantial. This is less of an issue for specific tasks (such as a search for a limited number of legal provisions) and more of an issue for others (such as due diligence requiring

review of the entire document for mistakes and/or omissions). Other time costs include the time taken to iteratively refine AI prompts to ensure accurate output, and the time taken to learn new AI features that may or may not work better than established workflows that legal professionals are already used to.

To address the competing tensions outlined above, this chapter examines how operating models should be transformed to get the most out of AI adoption – both in terms of efficiency gains and the ability to continue investing in AI adoption. This chapter covers various settings, including law firms that price work according to the billable hour model. In that specific context, such law firms that decide to use AI may need to revisit their pricing arrangements – taking into account both the time costs discussed above, the time savings (if any), and the possibility of the time metric becoming redundant through AI use. This chapter also considers the additional layer of complexity (and associated operational risks) whereby lawyers of institutional clients that use AI bear the liability, responsibility, and time costs of checking the accuracy of their client's AI-generated output.

Chapter six examines the cultural risks of AI adoption in legal teams. The cultural risk on which the interviews primarily focus is the impact of AI on the growth and learning of lawyers, with a focus on junior lawyers specifically – a concern expressed by both junior and senior lawyers surveyed. The question is whether there is a risk of AI/algorithmic/automation bias – an over-dependence on AI that adversely impacts lawyers' skills and knowledge. Traditionally, lawyers have paid their dues in the earlier stages of their careers by undertaking the tedious but necessary research and drafting – and learning from the process (which includes making a whole range of mistakes). With the introduction of AI tools, there is some concern that client pressure to reduce fees would lead to law firms using AI for work that would otherwise be assigned to junior lawyers, preventing them from developing the skills and knowledge necessary for a lawyer to recognize when an AI system hallucinates and generates inaccurate output.

Based on the premise that AI is replacing some of what junior lawyers do (hence negatively impacting their growth and learning), the interviews in this chapter explore whether these lawyers can still learn as much from prompting and reviewing AI's performance of tasks, instead of completing those tasks themselves. The authors also examine other cultural risks of AI adoption, such as the threats legal professionals perceive to both their jobs and their professional identity as lawyers. Unlike chapter five, which focuses

on process efficiency and "hard" operational metrics (such as costs and pricing), this chapter on culture focuses on human development and "soft" cultural impact (such as mentorship and identity in the AI era) – while explaining the link between both categories of risks in a logical and comprehensible way.

Chapter seven wraps up the discussion by providing holistic observations of our findings across the various risks and themes covered in this book. Brief mention is made of several other AI risks that are worth bearing in mind in the legal context, before summarizing the findings from the interviews regarding the range of AI risks covered in this book, and synthesizing the "best practices" adopted by interviewees, which readers can turn to next to better address the risks of using AI in legal work.

About the authors

Matthew Seet is the author of the complete three-book trilogy on the ISO/IEC 27001/27701/42001 trifecta by Springer Nature (Apress), including *ISO 42001 and Legal Compliance: A Principled Implementation of the AI Management System, ISO 27001 in the Age of AI and IoT: Principles, Regulations, and the Information Security Management System*, and *ISO 27701 in Practice: Align Privacy Standards, Principles, and Legal Requirements Worldwide*. He is also the founder of LAIRisk, a legal AI risk management edtech company guiding legal professionals and legal technologists in managing AI risks by using the ISO 27001/27701/42001 framework.

Matthew was formerly an international law lecturer at the National University of Singapore where he taught law for seven years, with human rights writings published in international journals like the *Cambridge Law Journal*, *Journal of International Criminal Justice* and *International Journal of Refugee Law*, cited in the *Financial Times*, and awarded the 2017 Foundation for the Development of International Law in Asia Prize for Young Scholars. He obtained his Master's in International Law from the Graduate Institute of International and Development Studies in Geneva on a Swiss Government Scholarship, and is certified as an AI Governance Professional (IAPP) and ISO/IEC 27001/27701/42001 Lead Auditor and ISO 31000 Risk Manager.

An-Ru Stevens is a specialist in legal operations and transformation, providing advisory services to law firms and in-house legal teams across the APAC region. Her expertise centres on addressing both practical and cultural risks associated with AI implementation within legal services, focusing on issues such as reliability, capability erosion, and ineffective adoption strategies.

Previously, An-Ru held senior roles in legal operations and change management at prominent international law firms Herbert Smith Freehills and Eversheds Sutherland. During this period, she led regional digital transformation efforts, embedded legal operations capabilities, and designed training programmes for technology adoption across multiple jurisdictions.

At the time of publication, AnRu works at the intersection of legal operations and AI adoption, advising legal leaders on how to implement AI without undermining quality, capability, or trust. Her current work spans leading AI adoption within a major APAC renewable energy company, advising law firms on process redesign, and speaking regularly on AI transformation in legal services. She brings a riskaware, operationsled perspective grounded in an LLB (Hons) from the National University of Singapore and formal training in PRINCE2, Lean Six Sigma, and change management. Her work focuses on ensuring AI enhances, rather than erodes, the foundations of effective legal practice.