

AI Governance: Ethical Considerations in the Transformative Use of AI in Your Law Practice

By: Kellye- Raé J. Fisher Campbell, Esq.

²⁵ Columbia University - Master of Science (MSc) in Information & Knowledge Strategy (Focus: Data Analytics & AI)

Master of Laws (LLM) in Intellectual Property - George Washington University

Attorney-at-Law (Bar Admissions - New York, USA & Jamaica, WI)

CIPP/E | CIPP/US | CIPM | Fellow of Information Privacy (FIP)

ABSTRACT

With the evolving and increasingly important role of AI in a modern and globalised world, not only are we on the cusp of a new transformative data-driven era, but we're also in the embryonic stages of understanding the harms, misuses and violations that these seemingly boundless technologies may also bring. This new reality demands constructive and evaluative discourse on the legal frameworks necessary to govern AI's integration into our increasingly interconnected world.

The potential integration of artificial intelligence (AI) into legal practice has ushered in a new era, with increased possibilities in efficiency and capability, but it also raises significant ethical considerations that lawyers must carefully navigate. As AI tools become increasingly sophisticated and ubiquitous, legal professionals are confronted with the challenge of harnessing these technologies while adhering to their ethical obligations. This paper looks at transformative use of AI in law practice and then examines key ethical consideration arises from such use cases.

INTRODUCTION

In the rapidly evolving landscape of legal technology, AI has emerged as a transformative force, challenging traditional paradigms and offering unprecedented opportunities for efficiency and innovation in legal practice. As we stand at the cusp of a new era in law, it is crucial to examine the ethical considerations that arise from the integration of AI into the legal profession and our day-to-day practice of law.

The concept of AI, first introduced by computer scientist John McCarthy in 1956 as "the science and engineering of making intelligent machines,"¹ has since evolved into a multifaceted field with far-reaching implications for the legal sector. Today, AI encompasses a wide range of applications, from automated legal research and document review to predictive analytics and intelligent decision support systems. For the purposes of this paper, we will adopt the definition of *AI systems* provided by the European Commission's High Level Expert Group on Artificial Intelligence Ethics Guidelines for Trustworthy AI:

¹ McCarthy J. What is artificial intelligence? Computer Science Department, Stanford University; accessible on <https://www-formal.stanford.edu/jmc/whatisai.pdf>

*"Artificial Intelligence (AI) systems are software (and possibly hardware) systems designed by humans, that given a complex goal, act in the physical and digital dimension by perceiving their environment through data acquisition, interpreting the collected structure of unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best actions(s) to take to achieve the given goal. AI Systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour, by analysing how the environment is affected by their previous action."*²

This comprehensive definition encapsulates the key aspects of AI relevant to legal practice, including its data-driven nature, task dependence, and capacity for continuous learning and adaptation. As we delve into the ethical considerations surrounding AI in law, it is essential to recognise that these systems are not merely tools but complex entities capable of influencing legal outcomes and shaping the future of the profession.

The integration of AI into legal practice raises profound questions about the nature of legal work, the role of human judgment, and the fundamental principles of justice and fairness. How do we ensure that AI-driven legal systems uphold the ethical standards and professional responsibilities that form the cornerstone of the legal profession? What safeguards must be put in place to prevent bias, maintain accountability, and protect client confidentiality in an AI-augmented legal environment?

This paper aims to explore these critical issues, examining the ethical implications of AI in law through the lens of professional conduct, client representation, and the administration of justice. We will analyse the potential benefits and risks of AI adoption in legal practice, discuss emerging regulatory frameworks, and propose guidelines for the responsible development and deployment of AI in the legal sector. As we navigate this complex terrain, it will be imperative that legal professionals, technologists, and policymakers collaborate to create a framework for AI governance that upholds the highest ethical standards while harnessing the transformative potential of this technology. Only through thoughtful consideration and proactive engagement can we ensure that AI serves as a force for positive change in the legal profession, enhancing access to justice and upholding the rule of law in the digital age.

² European Commission's High-Level Expert Group on Artificial Intelligence, 'Ethics Guidelines for Trustworthy AI' (2019) p. 36 EU Parliament, Artificial Intelligence Act: Deal on Comprehensive Rules for Trustworthy AI, Dec. 9, 2023, (Available at: <https://www.europarl.europa.eu/news/en/press-room/20231206IPR15699/>)

PART 1: AI IN LAW PRACTICE

A Brief History of LegalTech

The integration of technology into legal practice has evolved over the last several decades, transforming from basic administrative tools to highly sophisticated systems that increasingly rely on AI to optimise legal processes. The global history of legal technology, often referred to as "legaltech"³, is a testament to the legal profession's gradual adaptation to technological innovations, aimed at enhancing efficiency, reducing costs, and improving access to justice. This journey can be traced back to the 1960s when law firms first began to adopt mainframe computers for essential functions such as billing and accounting, laying the foundation for the modern era of legal technology.⁴

The 1970s marked a pivotal period in legaltech with the advent of electronic databases for legal research. LexisNexis and Westlaw, two of the earliest pioneers in this field, revolutionised legal research by making vast amounts of case law, statutes, and legal commentary searchable and accessible to legal professionals.⁵ Before these developments, I imagine that lawyers spent countless hours manually searching through physical legal libraries—a process that was both time-consuming and prone to error. With the introduction of electronic databases, legal research presumably became more efficient and systematic, significantly enhancing the speed and quality of legal work.

The 1980s, 1990s and early 2000s witnessed further technological advancements, including the proliferation of personal computers and the rise of the internet. These innovations transformed communication and information-sharing within the legal profession, allowing for faster exchange of documents and collaboration between legal teams, clients, and courts. During this time, document management systems and case management software became integral tools for law firms, providing a more structured and efficient way to handle case files, track deadlines, and manage legal workflows. The growing use of computers in law firms also prompted the development of specialised software tailored to specific legal needs, further streamlining administrative tasks and improving the organisation of legal practices.

³ LegalTech generally refers to the wide-ranging technology that is used to deliver legal services which can benefit the law practice and legal service industry by assisting in meeting the legal needs of consumers, improving access to legal services and the standards of service, and driving efficiencies and competition in the legal sector. (See *Practical Law Practice Compliance & Management* (n.d.). *Legal technology and AI toolkit: England and Wales*. Retrieved November 21, 2024, from <https://content.next.westlaw.com/practical-law/document/1dfaf633ff94311ed8921fbef1a541940/> - subscription needed)

⁴ Susskind, R. (2023). *Tomorrow's lawyers: An introduction to your future* (3rd ed. Chapter 2 – *Strategies for Success*). Oxford University Press.

As the 21st century approached, the widespread adoption of cloud computing and mobile technology radically altered the way legal professionals all over the globe have continued to work. These innovations have enabled lawyers to collaborate and access critical case information remotely, facilitating the rise of virtual law firms and real-time collaboration. As an example, cloud-based tools for file storage, document sharing, and communication have become standard in law practice, ensuring greater flexibility and efficiency for lawyers who must adapt to an increasingly globalised and fast-paced environment.

The current phase in the global evolution of legaltech seems to be heavily impacted by the integration of AI and machine learning, which are beginning to revolutionise numerous aspects of legal practice. AI-powered applications are now capable of automating complex tasks such as legal research, contract review, and document analysis, tasks that once required hours of labour-intensive work by legal professionals. Machine learning algorithms can analyse vast quantities of data, identify patterns, and even predict outcomes, empowering lawyers to make more informed decisions and enhancing the overall effectiveness of legal services. This transformative phase of legaltech is not merely an incremental improvement of previous technologies but represents a fundamental shift in the way legal services are delivered.

Types of AI applications in the Practice of Law

The application of AI in law holds immense promise, but also raises important questions about the future of legal practice. As the profession continues to embrace these advancements, it becomes critical to examine the key types of AI applications currently in use, their implications for the profession, and the ethical challenges they present. Below, I briefly explore the specific types of AI applications that are currently impacting and reshaping legal practice globally and will later delve into some of the benefits of these use cases, and then some potential challenges generally associated with the use of AI. The discussion does not seek to endorse any solution shared, however by examining how AI has begun to be utilised globally in different ways within the industry, (i.e. legal research, document review, predictive analytics, contract management, etc.) and explore additional use cases so that we can better understand the broader implications for the future of the legal profession and the potential transformative uses of AI in law; and thereafter discuss the ethical considerations.

1. Legal Research

Legal research is perhaps the most prolific and significant application of AI in legal practice, enabling lawyers to access, analyse, and synthesise vast amounts of legal data more efficiently than ever before.

Traditionally, legal research required lawyers to manually search through extensive libraries of case law, statutes, and legal commentary, whether digitally or hard copy. With AI, this process is becoming increasingly automated and has the potential to become even more precise.

AI-driven legal research tools leverage Natural Language Processing (NLP)⁶ and machine learning⁷ to understand the nuances of legal language and context. These tools can parse large volumes of text, recognise relevant cases, statutes, and legal doctrines, and provide sophisticated, context-aware suggestions. Unlike traditional search engines, which rely on directly matching key words, AI research platforms possess the ability to analyse the meaning and intent behind the search query to identify the most relevant legal materials.

For instance, Thomson Reuters' Westlaw Edge⁸ claims to leverage the power of artificial intelligence to enhance legal research and litigation strategies. By combining advanced AI with what it describes as over a century of editorial expertise, the platform asserts that it provides lawyers with fast, accurate answers and valuable insights. According to Westlaw, its WestSearch Plus feature allows users to quickly address questions across thousands of legal topics, offering predictive research suggestions as users begin typing in the search bar. The Quick Check feature is said to enable lawyers to upload their briefs and identify highly relevant legal authority that may be overlooked in traditional research methods.⁹

Additionally, Westlaw Edge promotes its Litigation Analytics tool, which offers data-driven insights into judges, courts, lawyers, law firms, and case types, helping lawyers refine their strategies and better manage client expectations. The KeyCite Overruling Risk feature claims to ensure that legal professionals can trust their citations by alerting them when a point of law in a case has been implicitly undermined. The Statutes Compare and Regulations Compare features are presented as tools for quickly comparing any two statutes or regulations, allowing users to track and understand changes over time. Through these AI-powered tools, Westlaw Edge asserts it is transforming legal research, offering more precise, efficient, and insightful resources for legal practitioners.¹⁰

⁶ A Glossary has been included at the end of this paper to assist in understanding key AI terms generally. Some of these terms have been used within this article. See Glossary for NLP definition.

⁷ See *Glossary*

⁸ Thomson Reuters. (n.d.). Legal research with Westlaw Edge. Retrieved November 21, 2024, from <https://legal.thomsonreuters.com/en/c/legal-research-westlaw-edge>

⁹ *Ibid*

¹⁰ *Ibid*

Similarly, LexisNexis claims that Ravel View¹¹, a legal research tool, also uses data visualisation to map case law, offering lawyers the ability to quickly identify precedents and key legal trends. According to LexisNexis, Ravel allows lawyers to find key cases more efficiently and uncover those that may have been overlooked in traditional research. Additionally, the platform is said to provide insights into how cases have been treated—whether positively or negatively—by other courts (*see also Predictive Analytics below*). This integrated view, as claimed by LexisNexis enhances the research process with a comprehensive and intuitive approach to legal analysis.¹²

These tools claim not only improve the speed of legal research but also enhance its quality by incorporating advanced analytics that provide deeper insights into legal arguments.

2. Document Review

There have also been advances and new solutions document review, electronic disclosures¹³ and due diligence which can traditionally be time-consuming and costly, particularly in large-scale litigation or compliance audits. AI has emerged as a powerful solution to automate and accelerate some of these processes, particularly in identifying, collecting, and reviewing data and electronic documents for legal proceedings¹⁴.

Thus far, AI systems specifically designed for legal document review rely heavily on machine learning algorithms learn from human inputs and progressively improve their ability to classify documents accurately. These systems are capable of distinguishing between relevant and irrelevant documents, identifying privileged content, and flagging potential compliance issues.

3. Contract Analysis and Management

AI tools equipped with machine learning and NLP can automate the review and analysis of contracts, assisting lawyers with managing legal risks, ensuring compliance, and identifying key dates, deliverables and clauses across numerous documents.

¹¹ LexisNexis. (n.d.). Ravel view: How-to literature. Retrieved November 21, 2024, from <https://www.lexisnexis.com/pdf/lexis-advance/Ravel-View-How-To-Literature.pdf>

¹² *Ibid*

¹³ For more on AI for document review and e-discovery, see Dillen, B. (n.d.). How AI transforms document review in eDiscovery. KPMG Switzerland. Retrieved November 21, 2024, from <https://kpmg.com/ch/en/insights/cybersecurity-risk/e-discovery.html>

¹⁴ Platforms like Relativity (<https://www.relativity.com/>) and Logikcull (<https://www.logikcull.com/>) offer AI-enhanced e-discovery tools that automate the classification of electronic documents.

AI-powered contract analysis platforms, such as Kira Systems¹⁵ and Luminance, boast algorithms that automatically extract key provisions from contracts, including termination clauses, indemnity provisions, and confidentiality agreements. These companies say that their tools can assess whether a contract complies with regulatory requirements and flag potential risks, such as ambiguous language or unfavourable terms. They can also be used to streamline contract negotiations by suggesting language or edits that align with the client's goals.

In addition to these platforms, tools like ContractPod AI¹⁶ and DocuSign AI¹⁷ also market enhanced contract management. Both solutions claim to provide end-to-end contract lifecycle management, offering solutions for drafting, reviewing, and managing contracts with AI-driven insights that optimise the contract creation and approval processes. DocuSign AI also assists with automating document workflow and ensuring compliance through smart contract management features, such as automatic tracking of signatures and deadlines.

These AI-driven contract management platforms are not only designed to streamline the immediate contract review process but also provide ongoing oversight. This level of automation significantly reduces the administrative burden on legal teams, allowing them to focus on more strategic legal tasks.

4. Predictive Analytics

This use case generally involves using historical data and statistical models to forecast the likely outcomes of legal proceedings, offering valuable insights that assist in legal strategy and decision-making. AI tools that apply predictive analytics look at past case data, judicial behaviour, and other relevant factors to predict how a particular case might unfold.

AI predictive tools like Premonition¹⁸ and Lex Machina¹⁹ offer lawyers the ability to examine how similar cases have been adjudicated, including the success rates of particular lawyers or law firms, the tendencies of specific judges, and the strategies that are most likely to succeed. By using data-driven insights, lawyers can develop more effective litigation strategies and improve client outcomes. Moreover, these tools can identify trends in the law, offering lawyers a competitive edge by allowing them to anticipate shifts in legal interpretations and precedents.

¹⁵ Kira Systems. (n.d.). *How Kira works*. Retrieved November 21, 2024, from <https://kirasystems.com/how-kira-works/>

¹⁶ ContractPodAI. (n.d.). Contract management solutions to empower legal teams. Retrieved November 21, 2024, from <https://contractpodai.com/>

¹⁷ DocuSign. (n.d.). DocuSign AI: Agreement process AI. Retrieved November 21, 2024, from <https://www.docusign.com/products/platform/ai>

¹⁸ Premonition. (n.d.). Legal analytics: Unfair advantage in litigation. Retrieved November 21, 2024, from <https://premonition.ai/>

¹⁹ Lex Machina uses vast datasets to create insights about specific legal domains, helping law firms predict outcomes based on the strategies of opposing parties or the characteristics of a particular judge. Lex Machina. (n.d.). Legal analytics by Lex Machina. Retrieved November 21, 2024, from <https://lexmachina.com/>

5. Chatbots and Virtual Legal Assistants

Chatbots and virtual legal assistants powered by AI are being used to provide 24/7 client support, offering immediate answers to client queries and assisting in the completion of legal forms or the processing of simple legal procedures.²⁰ These tools may be particularly valuable in increasing access to justice by providing affordable legal information and services to individuals who may not have the resources to engage a lawyer for routine matters.

AI-based chatbots can also help law firms automate client intake processes, answer frequently asked questions, and guide clients through legal procedures. By offering legal services through automated channels, chatbots and virtual assistants help democratise access to legal information, particularly for those who face barriers to traditional legal counsel.

6. Compliance Monitoring

Leveraging AI system to assist in compliance monitoring have become indispensable in industries like finance and healthcare, where organisations are required to adhere to complex regulatory frameworks. AI tools can automatically monitor compliance obligations, assess whether an organisation is meeting regulatory requirements, and identify potential compliance risks before they become problematic.

For law firms, AI-powered compliance monitoring tools may be used to track changes in regulations, assessing the potential impact of these changes on client operations, and ensuring that clients remain compliant with evolving laws. AI tools can also automate the process of checking contracts, policies, and other documents against evolving regulatory standards, reducing the need for human oversight; providing real-time analysis of compliance risks, enabling law firms to act swiftly and prevent costly regulatory violations.

7. Sentiment Analysis and Social Media Monitoring

By leveraging NLP and machine learning, AI has the capability to vast datasets of social media posts, news articles, and other online content to gauge public sentiment about a particular legal case, company, or issue.

²⁰ Juro. (2024, May 13). The 7 best legal AI chatbots for 2024. Retrieved November 21, 2024, from <https://juro.com/learn/legal-ai-chatbot#exit> . Popular platforms like LegalZoom and Rocket Lawyer also use AI to automate document preparation and provide basic legal consultations . This often comes with a warning and a caveat that these tools cannot and should not be used to replace complex legal counsel. As such independent legal advice is often recommended. Notwithstanding this, tools such as these can provide access to legal services for individuals dealing with less complex matters such as the creation of wills, incorporating a business or registering a trademark. Always with condition fine print of *caveat emptor*.

Sentiment analysis can be particularly valuable in litigation, where understanding public opinion might influence legal strategies. For instance, a law firm representing a client in a high-profile client or case may use sentiment analysis tools to monitor how media coverage or social media discussions could affect the public's perception of the case.²¹

8. Risk Assessment and Fraud Detection

AI is also being used to detect fraud and assess risks which in turn assists and complements legal investigatory work, risk management and compliance within legal practice and strategy. This is particularly prevalent in areas like financial law, banking, and insurance. AI tools can analyse transactions, legal documents, and communications to identify fraudulent behaviour or potential risks. These systems can recognise patterns indicative of fraudulent activity and flag them for further investigation.²²

9. AI in Intellectual Property Law

The use of AI systems in the practice of intellectual property law also present new opportunities, where patent searches, trademark analysis, and patent management are becoming increasingly automated. AI-systems can quickly analyse vast patent databases, compare inventions, and assess whether a new invention is likely to be granted a patent based on existing prior knowledge²³.

10. Automation of Routine Legal Tasks

AI is also streamlining routine, repetitive tasks within law firms, automating mundane activities that previously consumed a significant amount of time. These include tasks like document drafting, contract creation, timekeeping, minute-taking, proofreading, and client communication. By automating such tasks, AI frees up lawyers to focus on higher-value work such as strategy development, client counselling, and courtroom advocacy. This law firms to optimise their operations, cut down on overhead costs, and improve the overall efficiency of their practice.

²¹ Platforms like LexisNexis and Thomson Reuters offer tools that provide sentiment analysis and monitoring services so that law firms can track public opinion, understand media coverage, and predict how public sentiment may impact ongoing legal matters.

²² Many payment gateways, online terminals and processors utilise AI systems to detect fraudulent activity. Stripe's Radar is a good example of this (<https://stripe.com/radar>). Companies like Darktrace and Palantir have also developed AI tools for entities to monitor financial transactions and legal documents for signs of fraud. In legal practices dealing with financial institutions, AI-driven risk assessment tools can proactively detect suspicious activity and mitigate financial risks before they escalate into legal issues.

²³ Corsearch's acquired platform TrademarkNow, for example, uses AI to analyse trademarks and assess their chances of being successfully registered in various jurisdictions.

Benefits of AI in Law Practice

The integration of AI systems in law practice offers numerous benefits that can significantly transform the delivery of legal services, enhancing both efficiency and effectiveness. The main advantages can be summarised in four key areas: efficiency, improved service quality, cost-effectiveness, and increased accessibility.

AI's ability to automate routine, time-consuming tasks is perhaps the most impactful benefit. By processing vast amounts of data quickly and accurately, AI may outpace traditional methods, reducing administrative workloads and improving the thoroughness of legal work. These systems also drive cost savings by streamlining workflows and minimising the need for manual labour. With AI's predictive capabilities, lawyers can better manage client expectations and refine strategies with greater confidence. Additionally, AI helps ensure compliance with evolving regulations by providing real-time updates and flagging potential issues before they escalate. Overall, AI boosts efficiency, reduces costs, and enhances service quality in legal practice.

AI's impact on legal research extends beyond efficiency and accuracy, potentially levelling the playing field between large and small firms by providing access to sophisticated analytical tools previously reserved for well-funded practices. Furthermore, AI has the potential to democratise legal services, offering more affordable options for individuals and small businesses that might otherwise be unable to afford traditional legal counsel. AI systems also improve client relationship management, enabling lawyers to focus on strategic interaction while automating administrative tasks.

While AI is revolutionising the legal industry, improving efficiency, reducing costs, and enhancing access to legal services, its adoption brings with it inherent risks and challenges that must also be carefully considered.

Potential challenges and risks associated with AI

Modern AI systems, which rely primarily on large language models (LLMs) and machine learning algorithms, process vast datasets through statistical analysis. Though seemingly intelligent these systems are limited, task dependent systems that lack human-like capabilities such as common sense reasoning, background knowledge integration, abstract thinking, and causal inference. Foundation models, often trained on large corpora of text, serve as the foundation for many AI applications.

Understanding the nature and scope of the training data used to develop these models is crucial for assessing the associated risks and recognising the limitations in their application.

Most current AI applications are examples of "narrow AI," meaning they are designed to perform specific tasks rather than exhibit general intelligence across diverse domains. The development of artificial general intelligence (AGI), capable of human-like cognition in a wide array of contexts, remains a distant goal. This limitation makes AI less reliable for complex legal scenarios requiring nuanced judgment, context interpretation, and emotional intelligence.

A key concern is the potential for bias within AI systems. As the old adage goes, "garbage in, garbage out," meaning that biases present in the training data can be perpetuated or even amplified by AI algorithms, leading to biased or discriminatory outcomes in legal decision-making.

Another challenge is the lack of transparency in AI decision-making processes. Some AI systems, particularly those based on complex machine learning models, operate as what is often called a "black box"²⁴ making it difficult to understand or explain how conclusions are reached. This opacity raises significant concerns regarding accountability, particularly in legal contexts where decisions need to be justified and verifiable. The inability to challenge or scrutinise AI-generated recommendations can undermine the trustworthiness of legal decisions.

The integration of AI into legal workflows also poses significant data privacy and security risks (*discussed further in Part II*). Legal professionals handle personal and sometimes sensitive client information, and the use of AI systems in this context may expose this data to potential cybersecurity threats or unauthorised access. Ensuring robust cybersecurity measures and compliance with data protection regulations is critical to safeguarding client confidentiality and maintaining the integrity of legal practice.

A Note on Gen-AI & Hallucinations in Court Cases

Hallucinations refer to AI-generated content that seems highly plausible but is, in fact, inaccurate or entirely fictitious. In the legal field, this presents a grave concern, as the use of AI to generate pleadings,

²⁴ The lack of visibility or understanding into how some deep learning systems make their decisions is known as the "black box problem." This is sometimes falsely associated with most or all AI systems. It's important to understand that not all AI systems are "black boxes", as there are many models which are explainable. Due diligence in investigating and selecting AI tools and platforms is the responsibility of the user.

case law, or legal arguments could lead to incorrect submissions being accepted by courts as legitimate, undermining legal precedent.

In the proliferation of generative AI platforms in 2023, this has illustrated by a few cases, including the UK's *Felicity Harber v The Commissioners for HMRC*²⁵, and the US *Roberto Mata v Avianca Inc*²⁶. In these cases, AI-generated fictitious legal precedents which were in turn submitted counsel in legal proceedings. The UK court cited in agreement the earlier decision of the US Court of Appeals finding generally that:

"Many harms flow from the submission of fake opinions. The opposing party wastes time and money in exposing the deception. The Court's time is taken from other important endeavours. The client may be deprived of arguments based on authentic judicial precedents. There is potential harm to the reputation of judges and courts whose names are falsely invoked as authors of the bogus opinions and to the reputation of a party attributed with fictional conduct. It promotes cynicism about the legal profession and the...judicial system. And a future litigant may be tempted to defy a judicial ruling by disingenuously claiming doubt about its authenticity."

Even though AI tools like ChatGPT can produce content that appears authoritative, they may "hallucinate" cases, legal citations, or facts that don't exist, leading to the submission of unreliable and erroneous pleadings or motions. While legaltech AI tools are bespoke and specifically programmed to provide sound and detailed legally *trained* output, these bespoke AI tools still have hallucinations. A recent study by Stanford researchers²⁷ showed that a recognised legaltech system hallucinated up to 34% of the time. The study demonstrated that whilst there were better results when compared to general purpose GPT models, the high incidence of hallucinations and inaccuracies highlighted the important need for continued human oversight.

A study conducted by legaltech owner Thomson Reuters in 2023²⁸ found that law firms are increasingly exploring the potential of generative AI like ChatGPT to revolutionise legal practice, with many recognising its capacity to streamline tasks such as legal research, drafting, and contract review. However, this enthusiasm is tempered by concerns around data privacy, accuracy, and the ethical implications of AI use in client-facing matters (*See Part II: Ethical Consideration in AI Governance*).

²⁵[2023] UKFTT 1007 (TC)

²⁶ 22-cv-1461 (PKC), US DC Southern District of New York

²⁷ Magesh, V., Surani, F., Dahl, M., Suzgun, M., Manning, C. D., & Ho, D. E. (2024). *Hallucination-free? Assessing the reliability of leading AI legal research tools*. Retrieved from https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf

²⁸ Thomson Reuters. (2023, April 17). New report on ChatGPT & generative AI in law firms shows opportunities abound, even as concerns persist. Retrieved November 21, 2024, from <https://www.thomsonreuters.com/en-us/posts/technology/chatgpt-generative-ai-law-firms-2023/>

Firms are understandably cautious, often adopting a phased approach to implementation, balancing innovation with risk management. In contrast, another study²⁹ revealed that corporate legal departments, appear more proactive in embracing generative AI. They, somewhat understandably focus on using these tools to cut costs and drive efficiency, reflecting a somewhat more pragmatic and less risk-averse attitude compared to law firms.³⁰

PART II: ETHICAL CONSIDERATIONS IN AI GOVERNANCE

Bias and Fairness

Bias refers to systematic errors in decision-making that lead to unfair outcomes, while fairness in AI denotes the absence of such bias or discrimination. In AI systems, bias can stem from multiple sources, including the data used for training, algorithmic design, and human interpretation. Machine learning models, a common form of AI, can learn and perpetuate biases present in training data, resulting in discriminatory decisions. While bias is often unintentional, fairness is a deliberate objective that requires conscious efforts to prevent discrimination against any individual or group.

Bias and fairness are critical ethical considerations in the use of AI in legal practice, particularly as these systems become more involved in legal decision-making. Addressing biases and ensuring fairness are essential for upholding the integrity of legal outcomes. When AI systems in law are biased, they can lead to discriminatory results that violate the core principles of justice, equality, and human rights. The potential for AI to reinforce or magnify existing biases further amplifies concerns about equal treatment under the law and the equitable access to justice. AI systems can inherit biases from multiple sources:

- a. Biased training data: AI models trained on historical legal data can inadvertently reflect and perpetuate past societal biases and discriminatory practices. For instance, if AI tools used in legal research or predictive analytics are trained on case law or judgments that disproportionately reflect biases—such as racial, gender, or socio-economic biases—the AI system may unknowingly replicate these prejudices. This is particularly problematic when legal datasets are drawn from

²⁹ Thomson Reuters. (2023, May 22). Corporate legal departments see use cases for generative AI & ChatGPT, new report finds. Retrieved November 21, 2024, from <https://www.thomsonreuters.com/en-us/posts/technology/chatgpt-generative-ai-corporate-legal-departments-2023/>

³⁰ Ferrara, E. (2023). Fairness and Bias in Artificial Intelligence: A Brief Survey of Sources, Impacts, and Mitigation Strategies. *Sci*, 6(1), NA. <http://dx.doi.org.ezproxy.cul.columbia.edu/10.3390/sci6010003>

biased sources or when they contain errors, lack critical information, or are incomplete. In such cases, AI systems may unintentionally reinforce existing disparities in the legal system.³¹

- b. Algorithmic bias: Even with neutral training data, the design and structure of AI algorithms can introduce bias. This occurs when the assumptions and choices made during the development of the algorithm—such as how data is weighted or which features are prioritised—lead to unintended biases. Certain model architectures or optimisation techniques may inadvertently favour or disadvantage specific groups, perpetuating disparities in the outcomes generated by the AI system.³²
- c. Deployment context: The way an AI system is implemented and used in practice can lead to biased outcomes, even if the underlying model itself is relatively unbiased. This is often referred to as user bias, which occurs when the individuals using AI systems introduce their own biases—whether consciously or unconsciously—into the system. Such bias can arise when users provide biased training data or interact with the system in ways that reflect their own prejudices. For example, predictive policing tools may disproportionately target certain communities based on how they are deployed or influenced by biased input data.³³

To ensure AI systems are fair, equitable, and serve the needs of all users, it is essential to identify and mitigate bias at every stage of their development and deployment. This responsibility lies primarily with:

- i) Developers, who must ensure that AI systems are designed, trained, and deployed in a fair, transparent, and unbiased manner, taking steps to eliminate or reduce biases from the outset.
- ii) Legislators and policymakers, who play a crucial role in creating ethical guidelines and regulatory frameworks that promote fairness, transparency, and accountability in the development and use of AI systems.

While these concerns may not directly affect legal practitioners in their use of AI within the practice of law, there are several approaches that developers can integrate into their development lifecycle, or that

³¹ Ferrara, E. (2023)., supra note 30

³² Sreelatha, A., & Choudhary, G. (2023, September). Exploring the use of AI in legal decision-making: Benefits and ethical implications. Woxsen University. Retrieved November 21, 2024, from <https://woxsen.edu.in/research/white-papers/exploring-the-use-of-ai-in-legal-decision-making-benefits-and-ethical-implications/>

³³ Sreelatha, A., & Choudhary, G. (2023, September). Supra note 32

regulators can mandate, to help systematically address and mitigate bias in AI systems. In the absence of clear legislative mandates, these practices become crucial for legal professionals when conducting due diligence on AI tools and legal tech suppliers. Legal practitioners enquiring into whether these steps are followed can help verify that the AI systems utilised within their practice are fair, reliable, and ethically sound:

- i) **Diverse and representative datasets:** One of the most effective ways to reduce bias in AI systems is by ensuring the data used to train these systems is diverse and representative of all demographics. It is essential that training datasets encompass a wide range of case types and outcomes to avoid reinforcing systemic biases. This may involve supplementing historical data with more recent, diverse samples to ensure that marginalised groups, such as racial minorities or lower-income individuals, are not disproportionately affected by AI decisions.
- ii) **Fair algorithm design:** AI models should be developed with fairness as a primary consideration from the outset. This includes incorporating fairness constraints and bias mitigation techniques directly into the algorithm design. Developers can use methods such as algorithmic auditing to ensure transparency in decision-making processes, ensuring they are free from bias. Additionally, adjusting model architectures or loss functions to prioritise equitable outcomes across different groups can further promote fairness. Regular testing is essential to identify and correct any biases, ensuring that algorithms do not unfairly disadvantage certain groups.
- iii) **Regular audits and monitoring:** Ongoing audits of AI systems are crucial for identifying and addressing potential biases. Implementing continuous bias detection and mitigation processes, including both automated testing and human review of AI outputs, helps ensure that any emerging issues are promptly addressed. This ongoing monitoring is essential for maintaining the fairness and accountability of AI systems over time.
- iv) **Multidisciplinary development teams:** Involving diverse perspectives in the development and implementation of AI systems is key to ensuring fairness. This includes engaging legal experts, ethicists, and representatives from communities that may be impacted by AI decisions. A multidisciplinary approach helps identify potential biases and ensures that the design of AI systems accounts for a wide range of social, cultural, and legal considerations.

In the UK, the Information Commissioner's Office (ICO) has provided guidelines on data protection and fairness in AI use, promoting transparency and accountability³⁴. Similarly, in the Caribbean, local regulatory bodies and data protection authorities could play a pivotal role in providing similar oversight, ensuring that AI tools are regularly assessed for fairness, compliance, and ethical standards.

Privacy and Data Protection

Legal practice inherently involves the handling of personal, highly sensitive, and confidential information, making privacy and data security a critical concern in the use of AI within the legal profession. The deployment of AI in law raises significant questions about data protection, client confidentiality, and compliance with evolving regulatory frameworks.

AI systems in legal practice rely on vast datasets for training and operation, while simultaneously processing sensitive client information during use. This dual nature of data interaction amplifies the importance of robust privacy safeguards. The integration of AI systems into legal workflows raises significant ethical and practical questions regarding both the data used to train AI algorithms and the client data entered into these systems during legal practice. On one hand, AI systems rely on large datasets to function effectively, and the quality and security of these datasets are crucial for ensuring that AI-driven decisions are fair and reliable. On the other hand, when AI tools are used by legal practitioners, there is the risk of exposing confidential client information if data is improperly collected, stored, or shared.

Legal professionals and AI developers share a responsibility to ensure that AI systems and their use comply with data protection regulations at every stage. This includes safeguarding personal data used to train AI models, ensuring that client data entered into the system is adequately protected, and implementing rigorous safeguards to prevent unauthorised access. Both practitioners and developers must work to ensure that privacy is maintained and that the ethical standards of confidentiality and data protection are upheld throughout the AI deployment process. The ethical use of AI in law requires a delicate balance between leveraging technology for improved efficiency and maintaining the highest standards of data protection and client trust.

³⁴ Information Commissioner's Office. (2023, March 15). How do we ensure fairness in AI? Retrieved November 21, 2024, from <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/how-do-we-ensure-fairness-in-ai/>

Data Privacy Concerns in AI System Development & Use

When engaging with AI vendors, legal practitioners must conduct thorough due diligence to ensure that the development of AI systems used within their legal practice adheres to critical data privacy standards. Some of these considerations include:

- i) **Compliance with Applicable Data Protection Laws:** Legal professionals must utilise commercially reasonable efforts to ensure that AI systems they use are compliant with applicable data protection laws and that they have implemented certain minimum standards such as data encryption, access controls, and data masking. AI systems must be designed to safeguard data integrity and confidentiality, ensuring that only authorised personnel have access to sensitive client information.
- ii) **Consent:** Legal professionals must ensure that clients are adequately informed and made aware of what data is being collected, how it will be used, how it will be stored and what systems may be utilised to assist with the delivery of their services, particularly where AI systems are involved. Consent must be explicitly obtained, especially when sensitive data, such as health or financial information, is involved. Beyond being a requirement under data protection laws, the ethical principle of informed consent requires transparency regarding data usage, with clients fully understanding how their data will be handled.
- iii) **Data Collection:** AI systems in legal practice rely on diverse and representative datasets to function effectively. Collecting such data often involves aggregating information from multiple sources, which can include sensitive personal data like client details, case records, and court decisions. It is essential that vendors provide assurances that their data collection processes respect privacy rights and obtain the necessary consent where applicable. Data must be gathered lawfully and ethically, ensuring that only relevant information is included and that personal data is minimised to avoid discriminatory outcomes.
- iv) **Data Anonymisation:** To mitigate privacy risks, AI systems should have in place robust anonymisation or pseudonymisation techniques for the personal data used in training datasets. This reduces the likelihood of identifying individuals and ensures that information remains protected while still allowing the AI to learn effectively from the data.

- v) **Data Storage:** Secure storage solutions are critical to safeguarding sensitive data used in AI development. Legal practitioners must be vigilant in ensuring that AI systems have implemented strong security measures to protect stored data from unauthorised access or breaches. Special attention must be paid to the risks associated with cloud-based storage, which may involve data being stored across multiple jurisdictions, each with different levels of data protection. Practitioners should verify that AI systems they use have secure data storage protocols in place and are compliant with relevant data protection laws and standards.

- vi) **Data Sharing:** Collaboration between legal entities and AI developers often requires the sharing of data between various stakeholders, such as law firms, AI service providers, and third-party vendors. Strict data protection protocols must be followed to maintain confidentiality and comply with applicable data protection laws. It is essential that all parties involved adhere to the highest standards of data protection to prevent the unauthorised sharing or misuse of sensitive client information. In our reality, where lawyer-client privilege and confidentiality are paramount, ensuring the secure transfer and handling of data is crucial to avoid compromising legal integrity.

- vii) **Retention and Deletion:** The use of AI systems in legal practice also raises concerns about data retention and deletion. Legal professionals must establish clear understanding for how long client data will be stored in AI systems and how it will be securely deleted when no longer needed. The legal obligation to avoid excessive data retention and to protect the right to erasure is also crucial to maintaining client trust. It's important that any such exercise adheres to the practice's overall retention and deletion policies.

The use of AI systems in legal practice—whether for document review, predictive analytics, or legal research—raises additional data privacy concerns, particularly in how client data is inputted into these systems and how it is processed and protected. While the actual platform or system must have safeguards in place, it is also imperative that practitioners validate the adequacy of such measures and also put in place their own internal technical and organisational measures to protect client data.

Professional Responsibility

As technology continues to advance and becomes more deeply integrated into the legal profession, it is essential for lawyers to adapt while maintaining their commitment to professional ethics and

responsibilities. This ensures that new technologies, including AI are utilised ethically, effectively, and in compliance with relevant laws, codes and ethical standards. Below is a brief discussion on the key professional responsibility considerations that lawyers must bear in mind when incorporating AI into their practice

Duty of Competence

While lawyers are not expected to possess exhaustive knowledge of every aspect of the law, they do have a fundamental duty of competence in representing their clients. This professional obligation applies universally, regardless of jurisdiction. For example, Canon IV of Jamaica's Legal Profession Act (Legal Profession - Canons of Professional Ethics) Rules states that *"An attorney shall act in the best interest of the client and represent them honestly, competently, and zealously within the bounds of the law. He shall preserve the confidence of his client and avoid conflicts of interest."* Arguably, this duty of competence extends beyond traditional legal knowledge and includes technological competence, particularly in the digital age. In Canada, the Federation of Law Societies of Canada (FLSC)'s Model Code of Professional Conduct³⁵ has emphasised the importance of technological competence, requiring lawyers to stay abreast of emerging technologies like AI and to use them responsibly in their practice:

"To maintain the required level of competence, a lawyer should develop an understanding of, and ability to use, technology relevant to the nature and area of the lawyer's practice and responsibilities. A lawyer should understand the benefits and risks associated with relevant technology, recognising the lawyer's duty to protect confidential information..."

Failing to maintain competence, particularly in the use of AI tools, not only jeopardises the quality of client representation but also exposes lawyers to potential legal or disciplinary consequences. For instance, a British Columbian lawyer³⁶ was ordered to personally pay costs after citing AI-generated, non-existent case law in a ruling by the BC Supreme Court. The court emphasised the importance of "competence in the selection and use of any technology tools, including those powered by AI," reinforcing that the integrity of the justice system requires no less. This underscores the critical responsibility that lawyers have in understanding and appropriately applying AI and other technological tools in legal practice.

³⁵ Federation of Law Societies of Canada. (2024). Model code of professional conduct. Retrieved from <https://flsc.ca/wp-content/uploads/2024/11/2024-Model-Code-of-Professional-Conduct.pdf>

³⁶ Zhang v Chen. (2024). BCSC 285. CanLII.

Lawyers must be aware of the risks associated with AI use, including potential inaccuracies, biases, and "hallucinations" in AI-generated content. They are not expected to be IT experts but should understand the benefits and risks of relevant technology. They also ought to regularly update their knowledge about AI technologies and their implications for legal practice.

Duty of Confidentiality

The duty of confidentiality is a cornerstone of legal ethics, requiring lawyers to protect the information shared by clients during the course of their professional relationship. Canon IV (t) of Jamaica's Legal Profession Act (Legal Profession - Canons of Professional Ethics) Rules, provides that:

"An attorney shall not knowingly:

- i. Reveal a confidence or secret of his client; or*
- ii. Use a confidence or secret of his client to:*
 - 1. the client's disadvantage,*
 - 2. his own advantage, or*
 - 3. to the advantage of any person,*

unless it is done with the consent of the client after full disclosure."

This ethical mandate also requires that lawyers exercise utmost caution when utilising AI systems to process client information. This is particularly important when involving third-party vendors or cloud storage solutions. This overlaps with the lawyers duty of diligence, as lawyers must ensure that any technology used does not compromise the security of client data or violate confidentiality requirements. This includes thoroughly validating data security protocols and making reasonable enquiries into the technological structure of the AI systems in use. Ensuring compliance with relevant data protection laws is critical and forms part of the lawyer's broader ethical obligation to safeguard client confidentiality.

When employing AI systems, lawyers must ensure that appropriate safeguards are in place to prevent unauthorised access or data breaches. Furthermore, lawyers must review the terms and conditions of any AI service provider, especially those that may retain or share client data, to ensure that these platforms offer sufficient protection against unintended disclosures of confidential information. Adhering to these principles is essential in maintaining trust, integrity, and compliance with legal and ethical standards when incorporating AI into legal practice.

Duty to Act with Integrity and Transparency

Lawyers have an ethical duty to act with integrity and transparency in all aspects of their professional conduct, particularly when incorporating technology like AI into their practice. This obligation extends to ensuring that clients are fully informed about the tools being used in their cases, as well as the potential implications for their legal matters. Lawyers ought to communicate clearly and effectively with clients about the use of AI, providing a thorough explanation of how it will be applied, its benefits, and the potential risks involved. This includes discussing any limitations of AI systems and the possible impact on the quality of legal services provided.

Clients must be made aware of the role AI will play in their case, especially when it may influence key decisions or aspects of their matter such as legal research, predictive analytics, or contract analysis. Lawyers are ethically bound to ensure that their clients understand the extent to which AI is involved, and to obtain their informed consent before proceeding, particularly when there is a legal obligation to do so under applicable data protection laws. This ensures that clients make well-informed decisions about the technologies used in their legal matters.

In some jurisdictions, lawyers are also prohibited from discriminating against or harassing colleagues, employees, clients, or others. The duty to act with integrity may extend to guarding against societal harms such as discrimination and bias, particularly when using AI systems that could inadvertently perpetuate these issues. Lawyers must provide meaningful oversight when using AI systems, ensuring that the results generated by these systems are carefully screened for any evidence of bias. They should not rely solely on AI-generated content, especially when there are significant risks to certain groups or when the AI tool may inherently produce biased outcomes. This responsibility is crucial in upholding ethical standards and ensuring that AI is used in a way that serves justice and equity.

Duty to Maintain Dignity of Legal Profession & Courts

Under Canon I of Jamaica's Legal Profession Act (Legal Profession - Canons of Professional Ethics) Rules, lawyers are required to "*assist in maintaining the dignity of the legal profession*" and avoid any behaviour that could discredit the profession or give the appearance of professional impropriety. Additionally, Canon V emphasises that lawyers must "*assist in maintaining the dignity of the courts and the administration of justice.*" This includes treating tribunals and courts with respect and ensuring the integrity of materials submitted to them.

A key aspect of this duty is the responsibility to present only accurate and truthful information to the court. Lawyers must ensure that the authorities and materials they provide, including case law and legal opinions, are legitimate and are reasonably free from errors, omissions, or misleading content. This duty of candour requires that lawyers rigorously review all documents submitted to the court to confirm their accuracy. Lawyers are also obligated to familiarise themselves with the practices and procedures of the relevant tribunal and adhere to them in order to demonstrate respect and professionalism.

The *Roberto Mata v. Avianca Inc.*³⁷ case illustrates the serious consequences of failing to conduct proper due diligence when preparing for legal matters. Submitting false or inaccurate opinions may lead to significant harm, including wasted time and resources for all parties involved. Such actions not only damage the reputation of the court and judges but also undermine public trust in lawyers and the judicial system. When lawyers submit fabricated or misleading materials, they risk discrediting the legal profession and eroding confidence in the administration of justice.

Ultimately, failing to properly prepare cases, submitting inaccurate materials, and wasting the court's time constitutes a serious ethical breach. Lawyers must ensure that their actions uphold the dignity of the profession, demonstrate respect for the court, and maintain the public's confidence in the legal system.

Part III: AI Governance Frameworks and Best Practices

Effective AI governance is crucial for organisations to manage the known and unknown risks associated with AI technologies. The primary objective of AI governance should be to establish the trustworthiness, reliability, safety, security, resilience, privacy, and fairness of each AI system. While implementing an AI policy and training staff offers an initial layer of protection, law practice entities should also negotiate robust contractual protections with their AI service providers to ensure these standards are upheld.

Existing frameworks and standards, such as the NIST AI Framework and OECD AI Principles, provide valuable guidance for developing a governance structure. Engaging a consultant to tailor these frameworks to the specific needs of your legal practice can help create a customised approach. Below are some best practices and tips for lawyers to consider when integrating AI into their practice:

³⁷ 22-cv-1461 (PKC), US DC Southern District of New York

- i. **Be Competent:** AI systems are not infallible and may generate false or misleading information. Lawyers have a professional duty to verify and validate any information before advising clients or presenting it to a tribunal. This duty to fact-check is fundamental to maintaining the integrity of legal practice.
- ii. **Be Aware:** AI is a tool, not a solution. Understanding the capabilities and limitations of AI systems is essential to using them effectively and ensuring they align with your professional obligations. Familiarity with your AI tools enhances their utility and minimises risks of misuse or errors.
- iii. **Comply with Confidentiality and Data Protection Obligations:** Lawyers must stay vigilant regarding how the use of AI systems impacts compliance with confidentiality and data protection laws. When in doubt, consulting with an expert on privacy and data protection is prudent to ensure adherence to legal and ethical standards.
- iv. **Policies, Training, and Awareness:** Compliance is not only the responsibility of individual practitioners but also extends to staff involved in the delivery of legal services. Implementing clear policies and guidelines on the appropriate use of AI systems, along with conducting regular training, ensures that all staff understand and comply with the ethical and legal implications of using AI.
- v. **Negotiate Key Contractual Terms with AI Vendors:** Ensure that contracts with AI service providers address critical issues such as warranties, indemnities, and limitations of liability. These contractual provisions should protect the organisation and its clients in case the AI systems fail or cause harm.
- vi. **Conduct Due Diligence and Security Screening:** Before adopting an AI system, conduct thorough due diligence to assess its security and data protection practices. It's important to understand how the AI vendor builds and trains the system, how data is managed, stored, and retained, and what safeguards are in place to protect client confidentiality.
- vii. **Audit and Monitor AI Tool Usage:** Regular audits and monitoring are essential to ensure that AI systems are being used responsibly and in compliance with organisational policies. Continuous oversight helps identify any potential misuse or risks associated with AI tools in practice.

To conclude, the Federation of Law Societies of Canada (FLSC)'s Model Code of Professional Conduct offers a pertinent perspective:

"The practice of law continues to evolve. Advances in technology, changes in the culture of those accessing legal services, and the economics of practicing law will continue to present

challenges to lawyers. The ethical guidance provided to lawyers by their regulators should be responsive to this evolution. Rules of conduct should assist, not hinder, lawyers in providing legal services to the public in a way that ensures the public interest is protected. This calls for a framework based on ethical principles that, at the highest level, are immutable, and a profession that dedicates itself to practice according to the standards of competence, honesty, and loyalty."

This emphasises the need for ongoing ethical reflection as technology advances, ensuring that lawyers can adapt to these changes while safeguarding the integrity of the legal profession.

ANNEX

Glossary

*** Adopted from Coursera AI Glossary page³⁸ with a few additions from other cited sources. ***

- **AGI (artificial general intelligence):** Artificial general intelligence (AGI) is a field of theoretical AI research that attempts to create software with human-like intelligence and the ability to self-teach. The aim is for the software to be able to perform tasks that it is not necessarily trained or developed for. Current artificial intelligence (AI) technologies all function within a set of pre-determined parameters.³⁹
- **AI (artificial intelligence):** AI stands for artificial intelligence, which is the simulation of human intelligence processes by machines or computer systems. AI can mimic human capabilities such as communication, learning, and decision-making.
- **AI ethics:** AI ethics refers to the issues that AI stakeholders such as engineers and government officials must consider to ensure that the technology is developed and used responsibly. This means adopting and implementing systems that support a safe, secure, unbiased, and environmentally friendly approach to artificial intelligence.
- **Algorithm:** An algorithm is a sequence of rules given to an AI machine to perform a task or solve a problem. Common algorithms include classification, regression, and clustering.
- **Application programming interface (API):** An API, or application programming interface, is a set of protocols that determine how two software applications will interact with each other.
- **Big data:** Big data refers to the large data sets that can be studied to reveal patterns and trends to support business decisions. It's called "big" data because organisations can now gather massive amounts of complex data using data collection tools and systems. Big data can be collected very quickly and stored in a variety of formats.
- **Chatbot:** A chatbot is a software application that is designed to imitate human conversation through text or voice commands.
- **Chat-GPT:** is a generative artificial intelligence chatbot based on the multimodal generative pre-trained transformer developed by OpenAI. It uses natural language processing to create humanlike conversational dialogue. The language model can respond to questions and compose various written content, including articles, social media posts, essays, code and emails.
- **Cognitive computing:** Cognitive computing is essentially the same as AI. It's a computerised model that focuses on mimicking human thought processes such as pattern recognition and learning. Marketing teams sometimes use this term to eliminate the sci-fi mystique of AI.
- **Computer vision:** Computer vision is an interdisciplinary field of science and technology that focuses on how computers can gain understanding from images and videos. For AI engineers, computer vision allows them to automate activities that the human visual system typically performs.
- **Data mining:** Data mining is the process of sorting through large data sets to identify patterns that can improve models or solve problems.
- **Data science:** Data science is an interdisciplinary field of technology that uses algorithms and processes to gather and analyse large amounts of data to uncover patterns and insights that inform business decisions.
- **Deep learning:** Deep learning is a function of AI that imitates the human brain by learning from how it structures and processes information to make decisions. Instead of relying on an algorithm

³⁸ Coursera. (2024, November 14). Artificial intelligence terms: A to Z glossary. Retrieved November 21, 2024, from <https://www.coursera.org/resources/ai-terms>

³⁹ Amazon Web Services. (n.d.). What is AGI (Artificial General Intelligence)? Retrieved November 22, 2024, from <https://aws.amazon.com/what-is/artificial-general-intelligence/>

that can only perform one specific task, this subset of machine learning can learn from unstructured data without supervision.

- **Emergent behaviour:** Emergent behaviour, also called emergence, is when an AI system shows unpredictable or unintended capabilities.
- **Generative AI:** Generative AI is a type of technology that uses AI to create content, including text, video, code and images. A generative AI system is trained using large amounts of data, so that it can find patterns for generating new content.
- **GPT:** Generative Pre-trained Transformer
- **Guardrails:** Guardrails refers to restrictions and rules placed on AI systems to make sure that they handle data appropriately and don't generate unethical content.
- **Hallucination:** Hallucination refers to an incorrect response from an AI system, or false information in an output that is presented as factual information.
- **Hyperparameter:** A hyperparameter is a parameter, or value, that affects the way an AI model learns. It is usually set manually outside of the model.
- **Image recognition:** Image recognition is the process of identifying an object, person, place, or text in an image or video.
- **Large language model:** A large language model (LLM) is an AI model that has been trained on large amounts of text so that it can understand language and generate human-like text.
- **Limited memory:** Limited memory is a type of AI system that receives knowledge from real-time events and stores it in the database to make better predictions.
- **Machine learning:** Machine learning is a subset of AI that incorporates aspects of computer science, mathematics, and coding. Machine learning focuses on developing algorithms and models that help machines learn from data and predict trends and behaviours, without human assistance.
- **Narrow AI⁴⁰:** Narrow AI, also known as Weak AI or Artificial Narrow Intelligence, refers to AI systems that are designed to perform a specific task or a set of closely related tasks.
- **Natural language processing:** Natural language processing (NLP) is a type of AI that enables computers to understand spoken and written human language. NLP enables features like text and speech recognition on devices.
- **Neural network:** A neural network is a deep learning technique designed to resemble the human brain's structure. Neural networks require large data sets to perform calculations and create outputs, which enables features like speech and vision recognition.
- **Overfitting:** Overfitting occurs in machine learning training when the algorithm can only work on specific examples within the training data. A typical functioning AI model should be able to generalise patterns in the data to tackle new tasks.
- **Pattern recognition:** Pattern recognition is the method of using computer algorithms to analyse, detect, and label regularities in data. This informs how the data gets classified into different categories.
- **Predictive analytics:** Predictive analytics is a type of analytics that uses technology to predict what will happen in a specific time frame based on historical data and patterns.
- **Prescriptive analytics:** Prescriptive analytics is a type of analytics that uses technology to analyse data for factors such as possible situations and scenarios, past and present performance, and other resources to help organisations make better strategic decisions.

⁴⁰ DeepAI. (n.d.). Narrow AI definition. Retrieved November 22, 2024, from <https://deepai.org/machine-learning-glossary-and-terms/narrow-ai>

- **Prompt:** A prompt is an input that a user feeds to an AI system in order to get a desired result or output.
- **Quantum computing:** Quantum computing is the process of using quantum-mechanical phenomena such as entanglement and superposition to perform calculations. Quantum machine learning uses these algorithms on quantum computers to expedite work because it performs much faster than a classic machine learning program and computer.
- **Reinforcement learning:** Reinforcement learning is a type of machine learning in which an algorithm learns by interacting with its environment and then is either rewarded or penalised based on its actions.
- **Sentiment analysis:** Also known as opinion mining, sentiment analysis is the process of using AI to analyse the tone and opinion of a given text.
- **Structured data:** Structured data is data that is defined and searchable. This includes data like phone numbers, dates, and product SKUs.
- **Supervised learning:** Supervised learning is a type of machine learning in which classified output data is used to train the machine and produce the correct algorithms. It is much more common than unsupervised learning.
- **Token:** A token is a basic unit of text that an LLM uses to understand and generate language. A token may be an entire word or parts of a word.
- **Training data:** Training data is the information or examples given to an AI system to enable it to learn, find patterns, and create new content.
- **Transfer learning:** Transfer learning is a machine learning system that takes existing, previously learned data and applies it to new tasks and activities.
- **Turing test:** The Turing test was created by computer scientist Alan Turing to evaluate a machine's ability to exhibit intelligence equal to humans, especially in language and behaviour. When facilitating the test, a human evaluator judges conversations between a human and machine. If the evaluator cannot distinguish between responses, then the machine passes the Turing test.
- **Unstructured data:** Unstructured data is data that is undefined and difficult to search. This includes audio, photo, and video content. Most of the data in the world is unstructured.
- **Unsupervised learning:** Unsupervised learning is a type of machine learning in which an algorithm is trained with unclassified and unlabelled data so that it acts without supervision.