


AI Guide for Legal Professionals

A Foundational Overview



Objectives

This guide is intended as a starting point for lawyers and legal professionals who are considering incorporating artificial intelligence (AI) tools into their work.

It is a synopsis of a fast-moving landscape and designed to be broadly applicable — across practice areas, both in-house and private practice, and regardless of firm or department size.

How to Use This Document

This white paper is a primer, not a playbook, that aims to:

- Describe the current state of AI adoption in legal practice
- Explain the core AI platforms and types
- Identify the general circumstances in which AI tools are appropriate for legal work
- Examine professional responsibilities, including accuracy concerns and ethical considerations
- Explore what lies ahead for the profession's use of AI
- Provide a glossary of key terms and technical concepts (see Appendix A)

Current State & Background

Artificial intelligence has moved from the margins of legal practice to the center of the profession's technology conversation in a remarkably short period.

Adoption Is Uneven

AI adoption across the legal profession varies dramatically by firm size, practice area, client base, and individual practitioner. However, lawyers at every level — whether they use AI themselves or not — are now likely to at least be on the receiving end of AI-assisted work product from opposing counsel or clients.

Legal AI Tools Are Still Evolving

While some legal AI tools have accumulated a substantial body of industry protocols and case law, Generative AI (GenAI) tools capable of drafting documents, summarizing opinions, and answering legal research questions are far newer, and Agentic AI systems that can execute multi-step legal workflows with limited human intervention are newer still.

Consensus in the legal industry on what these tools can reliably do, boundaries for applicable use, and lawyers' professional responsibilities in using them is an ongoing process that is still in the early stages. Development of legal AI tools and their usage is evolving rapidly and continually changing.

Core AI Platforms & Types

It is important to differentiate AI platforms (categories of tools) from AI types (underlying technologies). These AI types can be overlapping: any single tool may incorporate multiple AI types.

Platforms	What It Is	Legal Usage Examples
General-Purpose AI Platforms <i>(e.g., ChatGPT, Claude, Gemini, Copilot, etc.)</i>	Widely available consumer-facing AI assistants built on large language models (LLMs). Capable of handling a broad range of tasks through a general purpose conversational interface. Unlike legal-specific or enterprise tools, they may lack controls for accuracy, confidentiality, data security, and workflow integration.	<i>Basic research on topics; initial orientation to unfamiliar legal concepts; summarizing a document or topic</i>
Legal-Specific or Enterprise AI Platforms	AI tools specifically designed for use by legal professionals and their organizations.	<i>Legal research with verified, jurisdiction-specific citation; contract drafting and review; document review; due diligence analysis; deposition and transcript analysis; drafting of briefs, memos, and pleadings</i>
AI Type	What It Is	Legal Usage Examples
Generative AI (GenAI)	A category of AI that produces new text, images, programming code, or other content in response to a prompt. Powered by LLMs.	<i>Drafting contracts or memos; summarizing opinions; generating legal research outlines</i>
Agentic AI	Executes sequences of autonomous actions across other tools and systems with minimal further step-by-step human direction.	<i>Due diligence pipelines; multi-database legal research agents; regulatory compliance monitoring; document management; time and billing; CRM; docket management</i>
Predictive AI	Analyzes historical data to produce patterns, classifications, scores, rankings, or outcome probabilities.	<i>Ediscovery; privilege review; litigation outcome scoring</i>

Applicability to Practice of Law

Research

Legal research is one of the highest-volume, most time-intensive tasks in legal practice — and one where AI can meaningfully reduce the time required to familiarize oneself with the relevant body of law, identify authority, and assess precedential weight. Some relevant examples:

- Generating an initial orientation to an unfamiliar area of law — identifying key concepts, major statutes, and doctrinal frameworks. Outputs need to be independently verified.
- Summarizing judicial opinions, regulatory guidance, or agency rulings for internal attorney consumption. Not for external transmission without review.
- Conducting citation analysis to evaluate precedential value, subsequent judicial treatment, and current validity, enabling researchers to identify controlling authority and assess whether cited cases remain good law.

Drafting Work Product

AI can be helpful for routine, high-volume, or templated documents where speed and consistency matter and the attorney's role is review and judgment, not initial composition. Examples include:

- Routine transactional documents where parameters are well-defined (e.g., NDAs, engagement letters, basic commercial agreements, board resolutions, standard form pleadings).
- Templated correspondence (e.g., demand letters, cease-and-desist letters, status updates).
- Administrative and workflow tasks (e.g., timekeeping narratives, billing summaries, due diligence checklists, closing checklists, meeting agendas).

Document Review

AI can significantly reduce review time and improve consistency across large document sets:

- Generating summaries of individual documents for attorney review in high-volume transactions or productions.
- Pattern recognition across large document sets — identifying all documents containing a particular clause, obligation, or factual issue; or those lacking a required provision.

Ediscovery & Fact Development

GenAI-assisted workflows are evolving from document classification into areas such as substantive fact development:

- Prioritizing and classifying documents for relevance, privilege, or responsiveness.
- Document summarization, issue spotting, and timeline construction.
- Continuous Active Learning (CAL) workflows, where the model is refined as reviewers code documents, improving accuracy over the course of a review.

Legal Administrative & Workflow Tasks

Many repetitive workflows can achieve greater efficiency through use of AI by legal operations professionals and other staff, as well as attorneys:

- Time entry and billing narrative drafting.
- Client intake, matter opening, and routing workflows.
- Document management.
- CRM, deadline tracking, and docket management.
- Internal memos, meeting agendas, and status reports.

Professional Responsibility: Accuracy & Ethical Considerations

AI use in legal practice implicates a number of professional responsibility obligations, which are evolving along with the technology.

The accuracy limitations of GenAI tools are significant and well-documented. Large language models can produce legal citations that are confident, well-formatted, and yet entirely fabricated — a phenomenon commonly referred to as hallucinations. Courts have imposed sanctions on attorneys who filed AI-generated briefs containing fictitious case citations. Accuracy concerns also extend to legal analysis, jurisdiction-specific conclusions, and factual summaries.

Responsibility for use of AI rests with the supervising attorney. No AI output should be submitted to a court, transmitted to a client, or otherwise relied upon without independent verification.

- [ABA Formal Opinion 512](#) (2024) — Use of Generative Artificial Intelligence Tools
- [ABA Model Rules of Professional Conduct](#) — Rules [1.1](#), [1.6](#), [3.3](#), [5.1](#), [5.3](#) govern supervising AI usage, verifying outputs, and protecting client data.
- State bar ethics opinions on AI are available through jurisdictions' state bar websites; a compilation is maintained by the [ABA Center for Professional Responsibility](#).
- [AI Policy Consortium for Law and Courts](#) (Thomson Reuters Institute and the National Center for State Courts)

Appendices B & C provide additional information regarding professional responsibility concerns.

Additional Resources:

AI Fundamentals & Legal Technology

- [Stanford Center for Legal Informatics](#) (CodeX) — academic research on AI in law, including working papers on LLM accuracy and legal benchmarking.
- [AI Risk Management Framework](#) (National Institute of Standards and Technology) — the federal framework for AI risk assessment; increasingly referenced in legal and compliance contexts.
- [2026 AI in Professional Services Report](#) (Thomson Reuters Institute) — a broad view into the current usage and planning, sentiment towards, and business impact of AI for legal, tax & accounting, corporate functions, and government agencies.
- [LegalTech Hub](#) — vendor-neutral directory of legal technology tools and categories.

Appendices B & C provide additional information regarding professional responsibility concerns.

What Lies Ahead

The legal profession is at an early and consequential stage in its use of AI. The tools available today are genuinely useful — capable of enabling lawyers and legal professionals to work more thoroughly and efficiently. They also present real risks, including accuracy, confidentiality, and the temptation to substitute AI output for professional judgment.

The profession is still determining what appropriate AI use looks like across different practice areas, matter types, and client contexts. The frameworks provided here reflect current best thinking, but the field is moving quickly. Courts are issuing new standing orders on AI disclosure. Bar associations are issuing new ethics opinions. Vendors are continually releasing new tools that bring both new capabilities and new risks. Staying current is an increasingly critical component of the professional obligation of competence in AI.

Appendix A — Definitions & Core Technical Concepts

This is not a definitive glossary but provides key terms that are essential for understanding legal use of AI and functioning effectively in AI-enabled legal environments.

Understanding them at a working level — not a technical level — can help you ask the right questions, understand how AI models produce their outputs, and evaluate those outputs more critically.

Terminology is evolving rapidly, and different vendors, platforms, and professional bodies use the same terms in different ways.

Bias: Systematic skewing of AI outputs based on patterns in training data. In legal contexts, examples of biased outputs include: skewed litigation outcome predictions, biased document review prioritization, or disparate output quality across jurisdictions, practice areas, or demographic groups.

Black Box: An AI system whose internal decision-making process is not transparent or explainable to users. Many predictive AI models operate as black boxes, which can complicate validation and challenge discoverability in litigation.

Context Window: The maximum amount of text an AI system can “see” and process in a single session, including prompts, uploaded documents, and chat history. A short context window means the system may not be able to process a lengthy contract, deposition transcript, or case file in a single pass, potentially causing it to miss information or lose coherence across a long document.

Data Sovereignty: The legal and contractual principle governing where data is stored, processed, and controlled. Lawyers must understand data sovereignty implications before inputting client information into any AI platform operating across jurisdictions.

Hallucination: The tendency of Generative AI to produce confident but factually incorrect output — including fabricated case citations, invented statutes, and inaccurate legal holdings. This is among the most significant accuracy risks in GenAI for legal workflows.

Human-in-the-Loop (HITL): A workflow design where a human reviews and approves AI-generated outputs before they are acted upon. The ABA and most state bars effectively require HITL for consequential legal tasks.

Large Language Model (LLM): The underlying technology in most Generative AI tools, including legal drafting and research. LLMs are trained on enormous text datasets. When given a prompt, the LLMs use learned patterns to generate responses. They are predictive language tools; they are not thinking systems. Unlike search engines, LLMs do not retrieve verified information, so their output must always be independently checked.

Machine Learning (ML): The branch of AI in which systems learn from data rather than from explicit programming. Predictive coding and technology-assisted review are examples of machine learning applications.

Natural Language Processing (NLP): The capability that allows AI systems to understand, interpret, and interact with human language. For example, it is what allows a user to query a research database in plain English rather than Boolean search strings.

Predictive Coding: Use of machine learning models trained on human-reviewed examples to classify, rank, or prioritize large volumes of documents based on relevance, privilege, or responsiveness, including ediscovery, contract review, and legal research.

Prompt: The instruction or query input a user provides to a Generative AI system. Prompt construction is critical — the quality and specificity of the prompt significantly affects the quality of the output — a discipline sometimes called prompt engineering.

Reasoning Model: Advanced models that work through a problem step-by-step — considering multiple angles, statutory elements, and testing possibilities — before producing its output, similar in structure to how a lawyer might think through an issue before writing a conclusion. Also called Extended Thinking.

Retrieval-Augmented Generation (RAG): A technique where an AI system searches a specific set of documents (e.g., a client’s contract library, a firm’s prior work product, or a defined body of case law) before generating its response, grounding the output in authentic source materials rather than relying solely on what the model learned during training. This allows you to “ask questions” of your own documents. May reduce risk of hallucinations.

Temperature: A parameter that controls how creative (or variable) a model’s outputs are. Higher temperature = more varied responses; Lower temperature = more consistent, predictable outputs.

Token: The unit of text (roughly a word or word fragment) that LLMs process. Token limits constrain how much text an LLM can consider at once, which affects its performance on long documents.

Training Data: The dataset on which an AI model was trained (e.g., case law, contracts, legal briefs). The data quality and recency directly affect the model’s capabilities and biases. LLMs trained on data through a given date may lack knowledge of subsequent legal developments. Data use policies and default settings determine whether information that the user inputs, including privileged or confidential information, may be retained by a third-party vendor or used to train future model versions, potentially exposing it to other users.

Appendix B — Examples of Professional Responsibility Risk

Certain tasks or contexts may present elevated professional responsibility risks in using AI in legal practice.

Here are some relevant examples (not a comprehensive list):

Task or Content	Examples of potential concerns	Examples of Risk Mitigation
Court filings citing AI-generated authority	Sanctions; Rule 3.3 candor obligation	Independent citation verification; disclosure where court rules require it
Inputting client data into public AI tools	Confidentiality (Rule 1.6)	Use of enterprise tools with data isolation and signed data processing agreements
AI-generated jurisdiction-specific analysis	Hallucination risk in unfamiliar courts	Supplementation with jurisdiction-specific human expertise
Agentic systems acting externally on client matters	Unauthorized commitments; supervision failure	Attorney review and approval prior to any external-facing action
AI-generated outcome predictions as sole basis for advice	Competence; client harm	Treatment as one input among several; human judgment governs final advice

Appendix C — Court Cases on Use of Legal AI

Court cases and orders involving use of legal AI.

Morgan v. V2X, Inc., No. 25-cv-01991 (D. Colo. Mar. 30, 2026)

United States v. Heppner, 2026 WL 436479 (S.D.N.Y. Feb. 17, 2026)

Warner v. Gilbarco, Inc., No. 2:24-cv-12333 (E.D. Mich. Feb. 10, 2026)

Park v. Kim, 91 F.4th 610 (2d Cir. 2024)

Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023)

Acknowledgments

Contributors

Co-Authors:

- **Andrew Shulman**
AI for Justice Innovation Lead & Program Manager, Thomson Reuters
- **Brendan Miller**
Former AmLaw 100 Partner, Legal Innovation Leader
- **Caitlin Moon**
Professor of the Practice of Law; Co Director, Program on Law & Innovation; Founding Co Director, Vanderbilt AI Law Lab
- **Dan Surowiec**
Global Chief Information Officer, Bryan Cave Leighton Paisner LLP
- **Doug McGhee**
Advanced GenAI Solutions Engineering, Thomson Reuters
- **Frank Schilder**
Senior Director, Thomson Reuters Labs
- **Gillian Glass**
Director of Legal Services, Farella Braun + Martel
- **Kassi Burns**
Senior Attorney, Trial and Global Disputes, King & Spalding
- **Kim Wolfe**
Chief Administrative Officer of Legal, State Street
- **Sean Easter**
Senior Manager, Applied Research Technology, Thomson Reuters Labs
- **Zafar Khan**
Global Enterprise Architect, Thomson Reuters

ILTA:

- **Jannise Vinson**
Vice President of Content, ILTA
- **Brian Balistreri**
Senior Education Programs Manager, ILTA
- **Kelly Thomas**
Editor of Content & Publications, ILTA

Thomson Reuters:

- **Stephen Seemer**
Senior Director, Thought Leadership Content and Strategic Relations, Thomson Reuters
- **Bryce Engelland**
Enterprise Content Lead for Innovation & Technology, Thomson Reuters
- **Joseph Kubes**
Director, Strategic Alliances, Thomson Reuters Legal

About ILTA

Technology solutions for law organizations and legal departments gets more complex every day. Connecting with your peers to exchange ideas with those who have “been there done that” has never been more valuable. For over three decades, the International Legal Technology Association (ILTA) has led the way in sharing knowledge and experience for those faced with challenges in their organizations and legal departments.

ILTA members include organizations of all sizes and all areas of practice, all sharing a common need to have access to the latest information about products and support services that impact the legal profession. Learn more at iltanet.org.

For more information, or suggestions or comments about this publication, contact education@iltanet.org.

ILTA Statement of Purpose

ILTA provides the premier environment for peer connections, education and collective intelligence to leverage the strategic advantages of technology in the legal profession.

Thomson Reuters

Thomson Reuters (TSX/Nasdaq: TRI) informs the way forward by bringing together the trusted content and technology that people and organizations need to make the right decisions. The company serves professionals across legal, tax, audit, accounting, compliance, government, and media. Its products combine highly specialized software and insights to empower professionals with the data, intelligence, and solutions needed to make informed decisions, and to help institutions in their pursuit of justice, truth, and transparency. Reuters, part of Thomson Reuters, is a world-leading provider of trusted journalism and news.

For more information on Thomson Reuters, visit tr.com and for the latest world news, reuters.com.

Thomson Reuters Institute

The Thomson Reuters Institute brings together people from across the legal, corporate, tax & accounting and government communities to ignite conversation and debate, make sense of the latest events and trends and provide essential guidance on the opportunities and challenges facing their world today. As the dedicated thought leadership arm of Thomson Reuters, our content spans blog commentaries, industry leading data sets, informed analyses, interviews with industry leaders, videos, podcasts and world-class events that deliver keen insight into a dynamic business landscape.

Visit thomsonreuters.com/institute for more details.

DISCLAIMER: This report is designed for use as a general guide and is not intended to serve as a recommendation, constitute legal advice, or to replace the advice of experienced professionals. If expert assistance is desired, the services of a competent professional should be sought. Neither ILTA nor any author or contributor shall have liability for any person's reliance on the content of or any errors or omissions in this publication.

