AI in the Legal Industry

Christian Grant

Copenhagen, January 2026

AI in the Legal Industry: Use Cases, Benefits, and Challenges in the EU and US

Introduction

AI is increasingly transforming legal work in both the United States and Europe. After years of cautious experimentation, a new wave of **generative AI** tools (like GPT-based assistants) has accelerated adoption in law firms and corporate legal departments. Surveys indicate that while most lawyers were *not* using AI as of late 2024 (78% of US firms reported no AI use), usage is now rapidly rising – Thomson Reuters found the share of legal organizations integrating generative AI nearly doubled from 14% in 2024 to 26% in 2025 . Large law firms and in-house teams are leading the charge: 87% of large-firm lawyers had adopted some AI by 2025 , and global firms like Allen & Overy have rolled out AI assistants to **3,500+ lawyers** across dozens of offices . This report provides a comprehensive overview of how AI is used in the legal industry, focusing on key use cases (ranked by adoption and business value), real-world AI tools in each area, the technologies powering them, benefits they bring, and their limitations, risks, and emerging trends. Special attention is given to differences and developments in the EU and US contexts, including regulatory considerations.

Key AI Use Cases in Law (Ranked by Adoption & Value)

The table below summarizes the major use cases of AI in legal practice, ordered roughly by their current adoption and business impact. Each represents an area where AI delivers significant efficiency or insights for lawyers and legal staff:

AI Use Case	Adoption (Prevalence)	Business Value	Example AI Tools
Legal Research & Knowledge	Growing rapidly; a top use (61% of lawyers use AI for case law summarization)	Very high – saves substantial research time and improves thoroughness of case analysis	*Westlaw Precision* (Ask Westlaw), *Lexis+ AI*, *Casetext CoCounsel*
Contract Review & Analysis	High in contracts-heavy practices (44% use AI for contract review)	High – accelerates contract review, flags risks, and reduces human error in due diligence	*Kira Systems*, *Luminance*, *LawGeex*, *Spellbook*
Litigation Support (e-Discovery & Analytics)	Established in large-scale litigation (45% use AI for document review)	Very high – cuts thousands of hours in document review, uncovers key evidence, aids case strategy	*Relativity AI*, *Everlaw*, *Brainspace*, *Lex Machina*
Compliance & Regulatory Monitoring	Moderate adoption (primarily in highly regulated sectors and in-house teams)	High – monitors regulatory changes and ensures adherence, preventing costly compliance failures	*Compliance.ai*, *Regology*, *IBM Watson* (Regulatory Compliance)
Decision Support & Legal Analytics	Emerging/limited use (niche predictive analytics tools)	Moderate – offers strategic insights (case outcome predictions, risk scoring), but reliability concerns limit use	*Lex Machina* (outcome analytics), *Blue J Legal*, *Premonition*, *COMPAS* (court risk scores)

Each of these use cases is explored in detail below, with real-world examples, the AI technologies behind them, and their benefits and limitations.

AI in Legal Research and Knowledge Management

Legal research was one of the earliest and now fastest-growing areas for AI in law. AI-powered research tools help lawyers sift through vast databases of **case law, statutes, and legal commentary** much faster and more intelligently than traditional

keyword searches. In fact, **61% of lawyers using AI report trusting it for summarizing case law** – making it the top AI use case in one 2025 survey . These tools can parse a natural-language query (e.g. *"What are the latest EU cases on data privacy fines?"*) and return relevant authorities or even generate a summary memo.

Modern legal research platforms are increasingly integrating **large language models (LLMs)** and **natural language processing** to enhance their capabilities. For example, Thomson Reuters' new **"Ask Westlaw"** generative AI feature (built on technology from their 2023 acquisition of Casetext) allows attorneys to ask research questions in plain English and get draft answers with cited sources . This system was trained on Westlaw's extensive content, enabling it to retrieve and summarize authoritative legal materials in context. Similarly, LexisNexis has introduced **Lexis+ AI**, and startups like *Casetext* (with its **CoCounsel** assistant) pioneered GPT-4-based legal research assistants in 2023 . These LLM-driven tools can draft legal research memos, identify relevant cases based on a fact pattern, and even review documents uploaded by the user to find pertinent information .

Example – Casetext CoCounsel: Casetext's CoCounsel (now part of Thomson Reuters) can *"draft legal research memos based on natural language prompts," "summarize transcripts,"* and *"review documents"* among other skills . It was initially powered by OpenAI's GPT-4 and launched as a paid service in 2023 . One appellate lawyer's review found CoCounsel "impressive" at tasks like generating deposition questions and summarizing trial transcripts in minutes, although still requiring attorney oversight for accuracy . As CoCounsel's technology has been folded into Westlaw, it illustrates how **generative AI with legal-trained models** is becoming mainstream in research.

Under the hood, these tools use techniques like **Retrieval-Augmented Generation (RAG)** – meaning the AI will **search and retrieve relevant documents** (cases, statutes, etc.) from a database and feed them into the LLM to ground its answers with real sources. This greatly reduces the risk of "hallucinations" (fabricated answers) that pure chatbots suffer from. As one legal technologist notes, *"It's not a search engine… [an LLM] is going to predict an answer, but providing specific questions and additional context materials is the solution"* to improve accuracy . In practice, that means a query to an AI research assistant will pull up, say, five pertinent cases and have the LLM summarize or analyze them, rather than relying on the model's memory alone. This blend of **extractive search** and **generative summary** leverages AI's speed while keeping results tied to verified legal texts.

Benefits: AI legal research tools can dramatically **speed up finding on-point authorities** and reduce the chance of missing a critical case. They can analyze wording and context (not just keywords), which helps when the law uses varied terminology. By generating summaries or even drafting a portion of a memo, they save attorneys time in writing and analysis. This efficiency is reflected in surveys: 72% of lawyers using AI said it **improved the speed** of their work, and 60% said it **improved the quality** of their work . AI can surface insights (like a rarely-cited precedent or an obscure regulation) that a human might overlook, especially under time pressure.

Limitations: Despite advancements, generative AI in research **cannot be fully trusted to provide correct or comprehensive answers on its own**. High-profile incidents have underscored this – e.g. in June 2023, a New York lawyer was sanctioned after his brief, drafted by ChatGPT, **cited six fictitious cases** that the AI had completely made up . The court fined the attorneys $5,000 for submitting the fake citations . This example shows that **"hallucinations"** remain a risk if AI output isn't verified against actual sources. Consequently, many AI research tools now *always* provide the source documents for any answer, and professional rules are emerging that demand lawyers supervise and validate AI research. In the US, a federal judge in Texas even issued a standing order requiring lawyers to **certify that no part of a filing was drafted by generative AI, or that any AI-generated content was checked against original sources** , explicitly warning that chatbots *"make stuff up… even quotes and citations."* Lawyers must therefore use these tools as an aid, not a replacement for traditional research skills, double-checking all AI-provided authorities.

Another concern is **access to up-to-date or jurisdiction-specific data**. Many LLM-based tools have knowledge cutoffs or lack access behind paywalled databases. Vendors are addressing this by fine-tuning models on proprietary legal data and updating them continuously, but users must ensure the AI's legal content is current and relevant to the jurisdiction (especially important in the EU, where laws vary widely by country). **Privacy and confidentiality** is also a factor – using client facts in a public AI (like the free ChatGPT) is usually forbidden due to confidentiality rules. Law firms are mitigating this by using **private or on-premise models** and ensuring any cloud AI vendor signs confidentiality agreements. (For instance, some large firms license **Harvey AI** – a GPT-4 based legal assistant – in a way that keeps all data in a secure, private environment .)

AI in Contract Analysis and Review

Reviewing contracts and other legal documents for issues is a time-consuming, detail-intensive task where AI has proven hugely valuable. **Contract analysis** AI can rapidly **extract key clauses, spot deviations from standards, flag risks, and even suggest edits or deal points** – tasks that otherwise require armies of junior lawyers in due diligence or contract management. This is one of the most mature legal AI use cases: products like *Kira Systems*, *Luminance*, *Seal Software*, and *LawGeex* have been used by law firms and in-house teams for years to accelerate contract review. By 2025, about **44% of lawyers using AI reported using it for contract review** tasks , and adoption is highest in corporate and transactional practice groups. According to one industry report, *"80% of the top 25 global M&A law firms"* have used Kira or similar AI in due diligence , underscoring how common this has become in high-value transactions.

How it works: Earlier contract AI systems relied on **machine learning (ML) and natural language processing** trained on thousands of example contracts. For example, Kira (now part of Litera) uses a **predictive ML model** to automatically recognize and extract over *1,400 types of clauses and data points* (e.g. termination clauses, indemnities, dates, parties) across 40+ document types . The system comes pre-trained on a vast range of agreements, achieving over 90% accuracy in identifying clause language with no manual setup . As users review and correct the AI's findings, it "learns" and improves (the ML model adapts to the style of the documents and the specifics of what the user cares about, essentially getting *"smarter with use"*).

Newer generations of contract review tools now blend in **generative AI** as well – for instance, Kira has introduced "generative smart fields" that let lawyers query contracts in natural language to find custom information, and even a GPT-based assistant to **summarize contracts or suggest revisions** . Other tools like *LawGeex* focus on **automated contract review against playbooks**: LawGeex's AI compares incoming contracts (like a third-party NDA or vendor agreement) to the company's standard or legal policy, and then *"flags any discrepancies or risks,"* highlighting clauses that don't meet preset standards . This produces an instant report of which clauses are acceptable and which need negotiation, saving legal teams considerable time. Similarly, *Luminance*, a UK-based AI platform, uses ML to detect anomalies or unusual clauses in contract sets – it's used in due diligence to spot risky provisions hidden in hundreds of files, a task nearly impossible to do manually with the same speed .

Real-world examples: Kira Systems (acquired by Litera) is a flagship contract AI used by many law firms. It *"automatically identifies and extracts key clauses from large sets of documents, significantly reducing the time it takes to review contracts,"* and its machine learning improves as it processes more documents . Big firms use Kira in M&A due diligence to quickly find clauses like change-of-control, assignment, or termination for convenience across thousands of contracts – tasks that used to take weeks can be done in days or hours. **Luminance** has been used by firms like Slaughter and May to analyze entire document collections during corporate transactions; it can cluster documents by similarity and flag outliers (e.g. one lease in a batch of 100 that lacks a standard liability clause) . **LawGeex** has been used by in-house teams to automate NDA review – the AI checks each NDA against a predefined checklist (e.g. confidentiality term, jurisdiction, non-compete language) and returns a marked-up version indicating which clauses deviate from policy, so lawyers only need to focus on those issues . More recently, startups like **Spellbook** (by Rally) offer an AI add-in for Microsoft Word that not only analyzes contracts but also suggests **alternate language** for negotiation and ensures contracts meet regulatory requirements. Spellbook's compliance-oriented AI will *"scan contracts for missing protections, risky clauses, and jurisdictional misalignments"* in seconds – for example, flagging if a data processing agreement lacks GDPR-required clauses. This shows how contract AI is expanding from pure review into **smart drafting and negotiation support**.

Benefits: For law firms, the business value of contract review AI is very high: it **cuts review time by 20–60%** (Kira claims to *"cut review time in half"* while maintaining accuracy equivalent to senior associates). This efficiency means firms can handle due diligence or contract markup faster and more cost-effectively – often a competitive advantage when clients expect quicker deal closings. AI can handle *"mundane, rote work"* like scanning for boilerplate, freeing lawyers to focus on deal-specific issues. It also **reduces human error**: an AI is unlikely to *forget* to check a clause in the 50th contract after hours of drudgery, whereas humans can slip. By flagging risky language consistently, AI helps ensure no important issue is overlooked in a mountain of paperwork. In-house legal departments benefit by processing high volumes of routine contracts (NDAs, supplier agreements, etc.) with minimal lawyer intervention, saving legal staff for higher-risk negotiations. Additionally, AI tools can create a **searchable repository of contract data** (key terms extracted from all contracts), which is valuable for business insights and compliance (e.g. quickly finding all contracts affected by a new law or all vendor contracts with a certain liability cap).

Limitations: Despite strong accuracy for many standard clause types, contract AI isn't perfect. **Unique or highly complex agreements** may confuse the model – as Gartner

analysts note, *"unique and complex agreements may require additional training or custom AI models"* for optimal performance . If a contract is in an uncommon format or a novel legal issue is buried in text, the AI might miss it. Thus, lawyers cannot assume the AI catches *everything*; critical deals still demand human review of key documents. Another limitation is that AI tools might not fully grasp **context or intent** – for instance, a clause that is acceptable in one commercial context might be problematic in another, which the AI wouldn't know without being told. Therefore, legal judgment is needed to evaluate AI-flagged risks (some could be false positives or not a real concern given the business context). There are also **integration and workflow challenges** – making AI outputs fit into the format of a deal memo or contract management system requires some work. Many firms find they need to invest in training attorneys to effectively use the tools and interpret results, which is a change-management issue as much as a tech issue .

Confidentiality and data privacy concerns also arise: feeding sensitive contracts into a cloud AI platform triggers worries about data security. Reputable vendors address this with encryption and by operating as GDPR-compliant processors in the EU, but some firms still opt for on-premise deployments. Finally, **AI doesn't negotiate on its own** – while some solutions now suggest edits or even *"automated negotiation"* features, human oversight is crucial for complex agreements . The AI can draft a redline, but a seasoned lawyer must decide if it's appropriate and aligns with business goals.

AI in Litigation Support and E-Discovery

Litigation often involves massive amounts of data – evidence documents, emails, deposition transcripts, case law, and more. AI has become indispensable in managing and extracting value from this trove, primarily through **e-discovery document review** and **litigation analytics**.

E-Discovery & Document Review: In large disputes or investigations, lawyers might have to review *millions* of documents to find those relevant to the case (a process known as discovery). **AI-driven e-discovery tools** use ML algorithms to streamline this. This was one of the earliest successful applications of AI in law, sometimes called **Technology-Assisted Review (TAR)** or predictive coding. Essentially, lawyers label a subset of documents as "relevant" or "not relevant" to a matter, and the software then **learns patterns** to rank or classify the remaining documents. Systems like *Relativity*

and *Brainspace* can *"sort, classify, and prioritize documents, making it easier for legal teams to locate critical information quickly."* Instead of manually reading every email, attorneys can focus on the top few thousand that the AI deems most likely to be relevant, drastically reducing labor. Courts in the US and UK have approved this approach for over a decade now, given evidence that a well-trained algorithm can recall as many (or more) relevant documents as human reviewers, in a fraction of the time.

Modern e-discovery platforms incorporate continuous active learning – the AI refines its model as reviewers code more documents – and clustering techniques that group similar documents together. For example, *Everlaw* (a cloud e-discovery tool) uses AI to cluster emails or files by topic and can visually show communication networks, which helps lawyers uncover patterns (like a particular custodian who is a hub of relevant communications) without reading every message. AI can also **detect anomalies** – e.g. highlighting if a custodian suddenly stopped emailing the company (perhaps indicating deleted emails) or if there are near-duplicate documents with subtle differences (which could signal tampering). These capabilities extend human ability to spot issues in huge datasets.

AI is also used to **summarize and transcribe** content in litigation. Transcription AI (like AWS Transcribe or Otter.ai) is now often used to transcribe depositions or hearings. Then LLM-based tools can summarize those transcripts. One lawyer testing Casetext's CoCounsel had it summarize a trial transcript – the AI produced a *"credible summary of what happened"* in about 8 minutes . While he wouldn't rely on it without reading the actual transcript, it was extremely useful for quickly getting the gist and identifying important witness testimony . This kind of AI-driven transcript summarization can save litigators many hours and allow quicker strategy pivots after a long deposition or court day.

Litigation Analytics & Decision Support: Beyond document review, AI is applied to **legal analytics** – parsing historical case data to find insights that inform litigation strategy or predict case outcomes. Thomson Reuters' *Lex Machina* (and similar tools from Bloomberg and others) exemplify this: it mines thousands of court dockets and decisions to deliver statistics on things like how often a particular judge grants motions to dismiss, the average time to trial in a patent case, or win rates for certain types of lawsuits. Lawyers use this data to **guide their decisions** (e.g., whether to settle a case, which venue is more favorable) and to manage client expectations. For instance, Lex Machina can show that a company suing for trademark infringement in a given district historically wins, say, 70% of the time, and that Judge X has never awarded more than

$1 million in damages in such cases – extremely useful info for deciding on litigation or settlement posture. As an example, *"Lex Machina uses data from past cases to help lawyers understand trends, predict outcomes, and develop stronger strategies, giving insights into opposing parties and judges."* This kind of **predictive analytics** is still emerging, but it's gaining traction, especially in the US. Another example is *Blue J Legal*, which uses ML to predict case outcomes in tax and employment law by modeling the factors that courts have considered important. Blue J's tax law prediction tool, for instance, asks the user a series of questions about the fact scenario (mimicking the factors in past cases) and then predicts how a court would rule (with a stated confidence level), also showing similar past cases. While not 100% accurate, it provides a researched starting point – and often points lawyers to analogous cases they might not have known.

In the **EU**, litigation analytics is less developed on a broad scale (due to fragmented jurisdictions and less public data), but some national initiatives exist. For example, France has *Predictice* and *Jurisprudence Analytics* for predicting outcomes in certain areas of law. However, a notable difference is regulatory: France in 2019 **banned the publication of judge-specific analytics** – making it a crime to publish statistics on individual judges' decision patterns . This law was a reaction to legal tech companies potentially "profiling" judges (e.g., predicting how Judge X rules) which raised judicial privacy and impartiality concerns. Violators face up to five years in prison . As a result, analytics tools in France can provide trends in the aggregate but **cannot disclose a particular judge's tendencies** . This contrasts with the US and UK where such judge-focused analytics, like Westlaw's Profiler or Lexis's Judicial Brief Analysis, are allowed and used. The EU's forthcoming AI Act may also classify some litigation prediction tools as high-risk (discussed under Trends and Regulation below).

Benefits: In e-discovery, AI's value is well-proven – it can **save millions of dollars** in large cases by cutting down on manual review. One study found that a well-executed predictive coding process could review the same dataset with *several orders of magnitude* fewer human hours than a linear review, with similar recall of relevant documents. For corporate clients, this means lower billables or lower costs passed on. It also speeds up the discovery phase significantly, which can lead to faster case resolutions. For litigation strategy, analytics provide **data-driven insights** that were not readily available before. Lawyers can make informed decisions rather than relying solely on gut or anecdotal experience. For example, knowing that *"Judge Smith grants summary judgment only 5% of the time in employment discrimination cases"* is concrete data a lawyer can use to advise a client whether to settle or proceed. AI can also improve **consistency** – ensuring that similar documents or issues are handled in similar ways, which is important in large multi-district litigations where hundreds of

lawyers might be working. In summary, AI in litigation support brings **efficiency, deeper insight, and better management of legal risk**.

Limitations: One limitation of predictive coding is that it requires a reasonably large and rich dataset of documents and some supervision. If the collection of documents is small or if there are only a few truly relevant documents (needles in a haystack), the algorithms may struggle or require a lot of upfront training by humans. Additionally, parties and courts must agree on protocols for AI-driven discovery – while now widely accepted, there can be disputes if, say, one side thinks the other's model missed important documents. It's essential to have transparency in the process (often sharing recall statistics or samples of unreviewed docs) to satisfy legal discovery obligations. For analytics and outcome prediction, **correlation is not certainty**: just because similar cases had certain outcomes doesn't guarantee the same result in a new case. Lawyers must be careful not to over-rely on a predicted win probability, especially if key facts differ. Moreover, any predictive model is only as good as the data – biases in past cases will be reflected. For instance, if historical data shows a bias (conscious or not) against a certain group in a type of case, the AI will predict outcomes reflecting that bias, which could be problematic or even **unethical to rely on without context**. The COMPAS example is instructive: COMPAS is an AI risk assessment used in some U.S. courts to predict reoffending and assist in decisions like bail or sentencing, but it was found to have significant racial bias – **mislabeling Black defendants as high-risk at nearly twice the rate of white defendants** . This kind of result urges caution: if an AI tool's recommendations might be biased or unfair, lawyers and judges must use independent judgment and not treat the algorithm as infallible. Indeed, awareness of AI biases is leading some jurisdictions to require **explainability** for any AI used in justice – a black-box prediction that a defendant is "high risk" is not acceptable without a transparent rationale.

AI in Compliance and Regulatory Monitoring

Legal compliance – ensuring that organizations follow laws and regulations – is a critical, ongoing challenge for companies, especially in heavily regulated industries (finance, healthcare, data privacy, etc.). **AI is increasingly deployed to help track regulatory changes, assess compliance obligations, and even check contracts and processes for compliance gaps.** This domain overlaps with the burgeoning "RegTech" (regulatory technology) sector.

Regulatory Change Monitoring: One of the most labor-intensive aspects of compliance is keeping up with new and changing regulations. For example, a bank might need to monitor hundreds of regulatory bodies and legislatures for any rule changes that could affect its operations. AI can assist by **automatically monitoring and summarizing regulatory updates**. Tools like *Compliance.ai* apply **machine learning models to scour thousands of sources – government websites, regulatory bulletins, financial authority notices – and identify changes relevant to the user's business** . The software then *maps those changes to the company's internal policies, procedures, and controls* . For instance, if the SEC issues a new rule on cybersecurity disclosures, the AI can flag it and even highlight which internal policy documents might need updating to reflect the new rule. These platforms often provide dashboards and alerts, so compliance officers are instantly notified of pertinent developments rather than manually reading issue after issue of the Federal Register. This is highly valuable in the EU as well, where companies face a mosaic of EU regulations and member-state laws – AI can monitor multiple jurisdictions in parallel. Some providers (like *Regology*) offer a **global AI-driven regulatory knowledge base**, where an AI agent can answer queries like "What are the key obligations for data retention in Germany vs France?" by pulling from the respective laws.

Risk Identification and Compliance Checks: AI also helps **identify compliance risks** within an organization's own documents and data. For example, in data privacy compliance (GDPR, CCPA, etc.), AI tools can scan company records to find personal data or detect if certain data processing activities lack the required consent records. An AI might flag, say, that an HR spreadsheet contains national ID numbers stored in plain text, which is a security concern, or that a marketing list has contacts from EU countries without documented consent – issues a compliance officer should address. Natural language processing can parse policy documents or legislation to extract obligations, which can then be turned into checklists for auditors.

A concrete example is **contract compliance**: We discussed Spellbook above in the contract context – it not only flags negotiation issues, but explicitly *"helps ensure contracts meet regulatory requirements"* . Spellbook's AI will check if a contract includes all clauses needed for compliance with laws like GDPR or industry standards, and if not, it flags the omission. For instance, under GDPR, a data processing agreement needs certain clauses (like subprocessor obligations, audit rights, etc.); an AI compliance tool can instantly review an agreement and highlight if any of those mandatory provisions are missing or weak. This use of AI saves lawyers from manually cross-referencing each contract against a regulatory checklist.

Case Study – Financial Services: Banks have been leveraging AI for compliance in areas like anti-money laundering (AML) and fraud detection for years (using ML to detect suspicious transactions). Extending that, banks now use AI to parse complex regulations. *IBM Watson* was used by some banks for **regulatory compliance** – Watson could ingest regulatory text (e.g. Basel banking rules or OFAC sanctions lists) and answer questions about what the bank needs to do, or classify which parts of a regulation apply to which internal policies. While early expectations were high, banks found that AI needed significant training to truly understand legal nuance. Nonetheless, gradually, AI is assisting compliance teams by taking on the initial heavy lift of reading and interpreting regulatory text, letting humans focus on high-level risk decisions.

Benefits: AI in compliance offers the benefit of **breadth and speed** – it can monitor an enormous range of information sources in real time, something no team of humans could do as effectively. This means organizations can react faster to legal changes (reducing the risk of non-compliance due to being unaware of a new rule). It also provides **consistency**: AI will apply the same standards in checking documents or processes, whereas human reviewers might have varying approaches or fatigue. By automating routine checks (e.g. scanning every new vendor contract for data privacy clauses), AI **frees compliance officers and lawyers to focus on complex, judgment-intensive issues**. Many companies also find that AI helps with **audit and documentation** – e.g. a compliance AI platform can log every change in law and how the company responded, creating a clear trail that regulators appreciate during inspections. Another benefit is **risk reduction**: early identification of compliance gaps (like a policy that's out of date, or an employee practice that violates a rule) can prevent costly violations and fines. Given EU's stringent penalties (GDPR fines can be up to 4% of global turnover), having AI "watchdogs" to catch issues is very valuable.

For example, a 2025 report noted that legal AI tools are helping address human issues like burnout – lawyers felt AI was taking away some of the grunt work of compliance and document review. In the Ironclad survey, **76% of respondents said AI decreased feelings of burnout**, and a significant number pointed to automating routine tasks (like contract compliance checks) as a key relief . This highlights a softer benefit: AI can improve job satisfaction by letting legal professionals concentrate on more meaningful work rather than endless check-the-box compliance reviews.

Limitations and Risks: Compliance is a sensitive area – getting it wrong can mean legal penalties. Thus, AI's limitations here must be managed. One issue is **accuracy and context**: an AI might flag too many false positives (e.g. labeling something as a risk which upon human review is actually fine under an exemption in the law) or false

negatives (missing an issue). Over-reliance on AI without human oversight could give a false sense of security. Therefore, best practice is to use AI outputs as a starting point, then have compliance experts review critical findings. Another challenge is that **regulations are often ambiguous or require interpretation**. AI can struggle with the subtleties – for instance, determining whether a new regulation "applies" to the company might require a legal interpretation that an AI can't definitively make. In cross-border contexts, AI needs careful tuning to understand different languages and legal structures; while language models have multilingual capabilities, legal terminology doesn't always translate cleanly.

Data privacy laws also impact the use of AI in compliance: if the AI is cloud-based, feeding it confidential compliance data could itself raise compliance issues (a bit of a Catch-22). Companies have to ensure that using the AI doesn't violate any privacy laws – for example, if uploading personal data to an AI service, GDPR might consider that a data transfer, requiring proper safeguards. Vendors are addressing this by offering EU-hosted services and by allowing on-prem installations for sensitive data.

Finally, regulators are beginning to scrutinize the **use of AI by companies in compliance efforts**. There's an emerging expectation that while AI can help, it **cannot replace responsible officers**. For instance, financial regulators might ask: who validated the outputs of your AI compliance tool? If a bank were to blame an "AI glitch" for a compliance failure, that likely would not be a defense – regulators will say the company should have had proper human oversight and testing of the AI. So, organizations must implement governance around these tools, including regular audits of the AI's decisions and outcomes (as Spellbook's guide suggests: *"Regular audits and transparent documentation are essential – AI systems should be reviewed frequently to avoid outdated outputs and ensure compliance with changing standards"*).

AI in Legal Decision Support and Judicial Systems

A more nascent application of AI in law is in **decision support** – helping lawyers and even courts make decisions by predicting outcomes or suggesting resolutions. We touched on litigation outcome analytics in the Litigation section (which is one form of decision support for lawyers). Here we consider a broader scope: AI advising on legal strategy or even making quasi-judicial decisions.

Lawyer Decision Support: Some AI tools attempt to guide legal strategy. For example, an AI might analyze the merits of a case and predict a probability of success, helping an attorney advise a client on whether to settle or proceed. Another area is **sentencing analytics** for criminal defense – an AI could analyze past sentencing by a judge for similar crimes and suggest what sentence is likely, informing plea negotiations. There are also experimental AI that recommend **legal arguments**: given a fact scenario, the AI could suggest potentially winning arguments or relevant precedents (this overlaps with research assistance). Startups like *CaseCrunch* in the UK claimed their AI could predict outcomes of simple legal disputes (like insurance mis-selling complaints) with high accuracy, even beating junior lawyers in a controlled test. These are still relatively limited and often work only in narrow domains.

One tangible tool is *Blue J Legal* (mentioned earlier), particularly its employment law and tax modules. For instance, Blue J's "Employment Classification" tool asks about a worker's conditions (hours, control, etc.) and predicts if they'd legally be an employee or contractor, based on prior case law. This helps lawyers and in-house counsel make decisions in advisory contexts with more confidence, backed by data. While not a replacement for analysis, it's a **second opinion from the machine** that can highlight factors maybe the lawyer didn't consider, drawn from dozens of prior cases.

Judicial Decision-Making: Perhaps the most controversial aspect of AI in law is its use *by* courts or administrative agencies to make or inform decisions. In the US, as noted, some courts employ **algorithmic risk assessment tools** like **COMPAS** during sentencing or bail hearings. COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) uses a questionnaire and historical data to score a defendant's risk of reoffending. Judges in some jurisdictions see these risk scores when deciding things like pretrial detention or probation conditions. The promise is to bring data-driven consistency to decisions that were often subjective. However, COMPAS became notorious when a 2016 investigation by ProPublica found it to be biased against black defendants – it was *wrongly labeling black defendants as "high risk" at nearly twice the rate of white defendants*, while white defendants who did reoffend were more often misclassified as "low risk" . This highlighted the **serious risks of bias** in AI decision support used in the justice system. In Europe, such uses are viewed with great caution: the EU AI Act explicitly classifies AI systems intended to assist judicial authorities in legal decision-making as **"high-risk"** – meaning they will face strict requirements for transparency, oversight, and accuracy if used at all . For example, **any AI that helps a judge interpret facts or apply the law in a case will be high-risk under EU law** , reflecting concerns about due process and fairness.

There have been pilot projects in this vein: *Estonia*, famously, floated the idea of a "robot judge" for small claims court to clear backlogs . The concept was that for minor disputes (maybe under €7,000), an AI could render a decision which could then be appealed to a human judge. As of 2023, that project had not fully materialized and faced skepticism – the Estonian Ministry of Justice clarified it was not moving to replace human judges, though some semi-automated online dispute resolution was being tried . China reportedly has implemented AI in courts to a degree – with some courts using AI to recommend judgments or check consistency of new rulings with prior ones (China also uses AI and analytics in prosecution decisions, raising human rights concerns). These developments show that technically, AI can be used as a judicial aide, but culturally and ethically, most jurisdictions (especially democracies) are not ready to hand over core judicial decision-making to a machine.

Benefits: In decision support roles, AI's benefit is providing a **data-informed perspective** that can counteract human biases or blind spots. For lawyers, having an AI evaluate a case is like doing massive case law research and statistical analysis in minutes – it might surface, for example, that *80% of similar cases in this district settled before trial*, or that *your assigned judge has ruled for plaintiffs only 2 out of 10 times in analogous cases*. This kind of insight can be invaluable for strategy (it might steer a plaintiff to settle if odds look unfavorable, or encourage a defendant to fight a case where the AI suggests the plaintiff's position is weak historically). AI can also **increase consistency** in decision-making: if all judges had access to a tool that, say, warns them "this sentence you're about to give is much harsher than the norm for similar cases," it could prompt a reevaluation and lead to more uniform justice.

For simpler high-volume decisions (like small claims or administrative rulings on benefits), AI decision-makers could theoretically resolve cases faster and at lower cost, freeing up human judges for complex disputes. They could also operate without emotion or fatigue. A much-cited benefit is backlog reduction – e.g., if an AI could handle thousands of minor cases per year, it might clear dockets. However, these potential benefits are largely speculative at this stage, as fully automated decision-making in law is exceedingly rare in practice.

Limitations and Concerns: The limitations here are profound. Legal decisions, especially in court, are not just about patterns – they involve **normative judgment, morality, and individualized considerations** that AI is ill-equipped to handle. An AI might predict outcomes based on past data, but that can entrench the past's injustices

(like perpetuating racial biases in sentencing). Also, when an AI is wrong in a decision support scenario, the consequences can be grave (wrongly jailing someone, or advising a client poorly). Hence the tolerance for error is extremely low. AI systems lack **explainability** in many cases – a judge or lawyer needs reasons, and if the AI can't articulate a reason a human finds acceptable, it undermines trust in the outcome.

There is also the issue of **accountability**: who is responsible if an AI-influenced decision goes wrong? With a human judge, we know who made the call. If a judge followed an AI recommendation that turned out to be biased or erroneous, is it the judge's fault or the AI developer's? This tangle is one reason the **ABA and other bar associations stress that AI is a tool, not an decision-maker** – a lawyer must not delegate professional judgment to a machine. The ABA's Formal Opinion 512 (2023) underlined that **generative AI should be treated as a non-lawyer assistant – the lawyer must supervise its work and take responsibility for the results** . The **Model Rules of Professional Conduct** in the US now make clear that using AI doesn't remove the lawyer's duty of competence and candor; lawyers must understand the AI's limits and **review its outputs just as they would the work of a junior colleague** .

In the EU, the forthcoming AI Act (agreed in late 2023, entering force likely 2025-2026) will impose specific obligations on high-risk AI, including those used in justice. Providers of such AI will have to implement risk management, testing, transparency (possibly even **disclosure to the subject that AI is involved in a decision**). For instance, if an AI were used in sentencing, a defendant might have to be informed and possibly given the right to a human review of the AI-augmented decision – similar to GDPR's provisions on automated decision-making. These regulatory safeguards aim to ensure that **AI supports human decision-making but does not override human control and fundamental rights**. It's telling that in the EU AI Act's risk classification, **"AI systems intended to be used by a judicial authority… to assist in researching and interpreting facts and the law"** are explicitly listed as high-risk and thus subject to strict oversight .

Given all these concerns, the prevailing view is that **AI can assist but not replace human legal decision-makers**. As one Gartner analyst put it, *"this isn't a technology that can simply replace people,"* but those who harness it effectively will have an edge . The key is maintaining human final judgment. We see that in practice: Judge Starr's order in Texas (mentioned above) emphasizes that while AI can be used in other parts of legal work, **for the core task of legal reasoning in briefs, the humans must verify everything** . Likewise, Judge Castel in New York (in the ChatGPT sanctions case)

made clear that the **duty of candor and competence lies with the attorney**, and that oath cannot be fulfilled by an AI which *"is not bound by any duty, honor, or justice"* . In short, AI can augment decision-making with data and speed, but its limitations – especially around bias, explanation, and accountability – mean that for the foreseeable future, **legal decisions must remain either fully or primarily in human hands, with AI as an advisor**.

Underlying AI Technologies in Legal Tools

The AI applications above are powered by a mix of advanced technologies. Understanding these **underlying AI technologies** helps clarify what the tools can (and cannot) do:

- **Large Language Models (LLMs):** These are at the heart of the generative AI revolution in law. LLMs (like OpenAI's GPT-4, Google's PaLM, etc.) are neural networks trained on massive text corpora, enabling them to generate human-like text and analyze language. Legal AI tools now frequently incorporate LLMs to draft documents, answer natural-language questions, and summarize texts. For example, Casetext CoCounsel's GPT-4 model was fine-tuned with legal data to better handle legal jargon and writing. Thomson Reuters' new CoCounsel (formerly Casetext) was *"trained on… Westlaw's content"* to ensure its generative answers have a legal grounding . LLMs excel at **language tasks**: drafting a research memo, rewriting a clause in simpler language, extracting a summary of a deposition, or holding a conversational Q&A (like a chatbot for legal FAQs). Their strength is the ability to **understand context and produce fluent text**, which is why they're used in almost all new legal AI products for tasks like document drafting, summarization, and chat-style interfaces to databases. However, as discussed, they have a tendency to **hallucinate** facts or sources when not properly constrained, so legal implementations usually pair them with retrieval mechanisms or limit their output to avoid fabrication.

- **Retrieval-Augmented Generation (RAG):** This is a specific design pattern that combines **search engines with LLMs**. In RAG, when the AI is asked a question, it first uses traditional information retrieval (or a vector similarity search) to fetch relevant documents from a knowledge base (e.g., case law database, statutory code, internal document repository). Those documents (or key excerpts) are then provided to the LLM as additional context, and the LLM is prompted to formulate an answer using *only* that provided context (often with instructions to cite the sources). The result is an answer that is backed by actual retrieved text, reducing the chance of error and enabling source citations. Many legal AI

systems adopt this: for instance, Harvey AI and others answer user queries by retrieving relevant law or guidance and then letting the LLM shape a response around it. The **"extractive and generative blend"** is considered a best practice for legal AI – you get the benefit of the model's language eloquence and ability to synthesize, but anchored in the *correct* knowledge. RAG relies on technologies like **vector embeddings** (which convert text into numerical vectors so that semantic similarity can be computed) to find text that is topically related to the query. This technique addresses the major limitation of standalone LLMs by injecting up-to-date, specific legal knowledge into the generative process.

- **Knowledge Graphs and Symbolic AI:** Some legal AI solutions, particularly older ones, use structured knowledge bases or rule-based systems. For example, compliance AI might use a **knowledge graph** of regulations – a network of legal rules and concepts – to reason about obligations. Early "expert systems" in law (popular in the 1980s-90s and seeing a resurgence in niche areas) require human experts to encode legal logic as rules (e.g., an immigration visa eligibility tool that asks questions and, based on a decision tree of rules, gives an answer). These are not machine learning but rather **symbolic AI**. They're highly precise in narrow domains but inflexible. Modern approaches may combine them with ML – e.g., a system might use a knowledge graph to represent key legal definitions, while using ML to classify which definitions apply to a scenario.

- **Supervised Machine Learning (Traditional ML):** Many legal AI tasks reduce to classification or prediction problems that traditional ML can handle. For instance, identifying whether a clause is a change-of-control clause can be done with **supervised learning**: feed the algorithm thousands of examples of clauses labeled "change-of-control" vs "not," and it learns to classify. Earlier contract analysis products used algorithms like support vector machines or random forests with hand-crafted text features. Now, neural network models (like BERT-based classifiers) are more common – these are *pretrained language models* (not as large as GPT, but similar idea on a smaller scale) fine-tuned for tasks like **document classification** (e.g., is this email privileged or not), **named entity recognition** (e.g., find all party names and dates in a contract), and **clustering** (grouping similar documents). In e-discovery, logistic regression and other classifiers have long been used to predict relevance based on word frequencies or embeddings. In litigation analytics, regression models or decision trees might predict an outcome (win/lose) from features like case type, judge, claims, etc. These traditional ML methods are **powerful for well-defined tasks** and often easier to explain than deep LLMs – an important consideration if output needs justification. They do, however, require quality training data (labeled examples), which in law can be scarce or expensive to produce.

- **Deep Learning (Neural Networks beyond LLMs):** Outside of language, if we consider other data types – e.g., analyzing **images or voice** in evidence – AI uses deep learning models (CNNs for images, for example) to do things like recognize text in scanned documents (OCR), identify logos or objects in images (for IP or e-discovery purposes), or transcribe spoken words in audio recordings (using speech recognition models). While not unique to legal, these are enabling technologies. For instance, an AI might perform OCR on a trove of scanned PDFs to turn them into text, then feed that text into an NLP analysis pipeline.

- **Hybrid Human-in-the-Loop Systems:** Worth noting is that many "AI" solutions in law are actually **hybrid systems** where human expertise is tightly looped in. The term **"Expert-in-the-Loop"** is used by some compliance platforms . This means the AI does initial heavy lifting (e.g., finds 100 potentially relevant documents), then a human (lawyer or trained analyst) reviews and gives feedback, and the AI refines its output. This iterative approach leverages the best of both worlds: AI's speed and scale, plus human judgment for quality control. The continuous feedback also improves the model over time (active learning). Most successful deployments of AI in law embrace this approach – fully automatic outputs without any human review are rare in mission-critical tasks.

In summary, the legal AI tools we see are powered by a **convergence of NLP innovations and classic ML**, often orchestrated together. Generative LLMs provide fluent language understanding and generation; retrieval systems ensure factual grounding; predictive models handle classification and analytics; and human oversight frameworks keep the AI aligned with legal standards. The result, when done right, are systems that can parse complex legal texts, reason (to an extent) about them, and communicate insights in a useful way to lawyers – effectively functioning like **supercharged legal assistants** that never tire.

Benefits of AI for Legal Practice

When implemented properly, AI solutions deliver a range of **tangible benefits** to law firms and in-house legal teams:

- **Increased Efficiency and Speed:** Perhaps the clearest benefit is doing work faster. Tasks that once took dozens or hundreds of hours (like reviewing documents, researching case law, or drafting a base contract) can be

completed in a fraction of the time. In surveys, **72% of lawyers reported AI has improved the speed of their work** . This efficiency means lawyers can meet tight deadlines more easily – for example, responding to a regulatory inquiry by sifting through emails with e-discovery AI in hours instead of weeks. It also allows firms to handle *more* matters without increasing headcount, improving scalability. One global law firm partner noted they see AI as saving *"a couple hours a week (per lawyer) on client work"* even now, with the potential for more . Those hours add up across a firm.

- **Cost Savings and Better Resource Allocation:** By automating routine, labor-intensive tasks, AI reduces the need for large teams doing low-value work, thereby cutting costs or freeing those resources for other tasks. Corporate legal departments under budget pressure value this highly – if an AI tool can review contracts at a fraction of the cost of outside counsel, that's direct savings. Law firms can also pass savings to clients (or take on alternative fee arrangements profitably). Moreover, **AI allows smaller firms or solo practitioners to punch above their weight** by handling large volumes of data or complex research that previously only big firms with armies of associates could manage. This is a business advantage and potentially increases access to justice (if routine legal services become cheaper thanks to AI).

- **Improved Work Quality and Consistency:** Well-designed AI tools can **reduce human error**. They don't get tired or distracted, so they are less likely to miss a critical clause on page 97 of a contract or a key phrase in a case. 60% of lawyers using AI said it **improved the quality of their work** . AI's ability to cross-verify and analyze systematically leads to more consistent outputs – for example, consistency in how contracts are reviewed against a checklist, or in how legal research memos are formatted with all relevant points. By catching overlooked issues (like an AI flagging a clause that everyone thought was benign but actually poses a risk), AI can raise the quality bar. Additionally, AI can surface **data-driven insights** that improve strategy quality – e.g. litigation analytics might reveal a trend that informs a winning strategy that a lawyer relying only on anecdotal experience might miss.

- **Enhanced Legal Research and Knowledge Management:** AI's speed and breadth in research means lawyers are better informed. They can consider more authorities and perspectives in the same amount of time. This leads to **more robust legal arguments and advice**. AI summarization also helps digest lengthy documents or transcripts quickly, so lawyers grasp the key facts faster. Furthermore, AI can help capture and reuse a firm's internal knowledge. For instance, some law firms are deploying GPT-based tools on their brief banks or prior work product, so that any lawyer can query "Have we ever dealt with issue X before?" and instantly find relevant prior work. This democratizes knowledge within the firm, making every lawyer as good as the firm's collective experience.

- **Reduction of Drudgery & Improved Job Satisfaction:** A less quantifiable but important benefit is taking away the most soul-crushing tasks from lawyers – e.g., repetitive contract proofreading, document coding, or endless document review. **76% of lawyers in one study said AI helped decrease burnout** feelings . By automating "grunt work," AI allows lawyers (and paralegals) to focus on higher-value activities that are more engaging – like complex legal analysis, counseling clients, or courtroom strategy. This not only improves mental health and job satisfaction, but also arguably the **ethical lawyering**: lawyers spending more time on strategy and client interaction (and less on copy-pasting or search tasks) can provide better service. Some have observed that integrating AI can *"elevate lawyers to focus on strategic initiatives"* rather than routine processing . Plus, embracing these tools can position lawyers as innovators, improving morale – 46% of legal pros believed AI tools could open more career growth opportunities, seeing them as a chance to learn new skills and roles .

- **Better Client Service:** Faster turnaround, more consistent quality, and potentially lower costs all lead to a better experience for legal clients. AI can also enable new services – for example, law firms might use AI to offer 24/7 contract review hotlines or instant compliance checks, which previously were not feasible. In-house teams can respond to business queries (like "can we do X under GDPR?") faster with AI-assisted research, making the legal department *less of a bottleneck* in enterprise processes. All of this enhances the value that legal teams provide to their stakeholders. Some clients are even expecting it: being able to say "we use AI to ensure we didn't miss anything in this review" can increase client confidence in the work product.

- **Strategic Insights and Data-Driven Decisions:** Through analytics, AI provides macro-level insights that inform strategy. For example, law firms can analyze their own billing and matter data with AI to find inefficiencies or optimal pricing strategies. Litigation prediction tools offer a more **objective assessment of risk**, which helps lawyers advise clients more accurately on whether to settle or fight. In the corporate realm, analyzing thousands of past contracts could reveal which clauses often cause disputes and thus should be negotiated differently – turning anecdotal wisdom into evidence-based policy.

In short, the benefits of AI in law center on the triad of **efficiency, quality, and insight**. AI augments human capabilities, allowing legal professionals to do more in less time, with greater consistency and supported by data. Importantly, this doesn't just benefit the bottom line – it can also enhance **ethical compliance** (e.g., by catching conflicts or errors that could otherwise cause malpractice) and even access to justice in the long run (if routine legal services become cheaper and more available through AI-driven tools for pro se litigants or legal aid attorneys). We are already seeing that those firms

and departments that smartly leverage AI are gaining a competitive edge – a recent survey found **47% of legal organizations call AI "essential" to daily work** and 77% increased AI spending over the past year 【0†L33-L37** (reflecting deepening adoption as benefits accrue). Indeed, a partner at Allen & Overy remarked that not adopting generative AI over time would be *"a serious competitive disadvantage"* for law firms【30†L218-L226】.

Limitations and Risks of Legal AI

Despite its promise, AI in the legal field comes with significant **limitations, risks, and potential downsides** that lawyers must vigilantly manage. Some of the key challenges include:

- **Accuracy and "Hallucinations":** As noted earlier, generative AI can produce incorrect or even fabricated information in a very confident manner. An LLM *guesses* an answer based on patterns, which means if it hasn't seen a fact or case before, it may just create one. This tendency for AI to *"make stuff up"* is obviously dangerous in law, where citing a non-existent case or misquoting a statute can mislead the court and violate ethical duties. The 2023 incident where an attorney's use of ChatGPT led to a brief full of fake cases is a stark warning . Judges have taken notice: the Texas judge's order requiring certification of human-checking of AI content explicitly cites the fact that current AI tools *"are prone to hallucinations … they make stuff up — even quotes and citations."* . Therefore, a major limitation of using AI is the **constant need for human verification**. AI is not a source of truth; it's a tool for drafting or searching the truth that still requires a lawyer's confirmation. This limitation somewhat tempers the efficiency gains – one must budget time to **verify AI outputs** (checking the cases an AI cited, proofreading its contract edits against the original, etc.). Until AI hallucinations are fully solved (if ever), lawyers must treat AI-generated text with initial skepticism, as a *draft* requiring review, not a ready-to-file work product.

- **Bias and Fairness:** AI systems can inherit biases present in their training data or reinforce historical biases. In the legal context, this can have especially grave implications. For example, **COMPAS's racial bias** in risk scoring (flagging black defendants as high risk more often) shows that AI can perpetuate and even obscure discriminatory outcomes under the guise of objectivity. If lawyers or judges rely on such tools unchecked, it could **amplify unfairness** in the justice system. Bias isn't limited to race; a contract review AI might be biased towards vendor-friendly readings if trained mostly on vendor contracts, or an

employment law AI might under-predict discrimination findings if historical cases had biases. Ensuring **algorithmic fairness** is extremely challenging – it requires careful curation of training data, algorithmic adjustments, and ongoing bias audits. In the EU, the AI Act will require such diligence for high-risk AI, but compliance will be non-trivial. Lawyers also face an interpretive challenge: if an AI analytics tool suggests a certain judge is unfavorable to, say, plaintiffs of a certain profile, using that insight might stray into ethically gray territory (could it encourage forum shopping or profiling that runs afoul of anti-discrimination norms?). Overall, bias means **trust in AI outputs should never be automatic** – legal AI needs to be transparently tested, and where possible, results that affect decisions (especially about people) should be cross-checked or supplemented with other reasoning.

- **Lack of Explainability:** Many AI models, especially deep learning ones like LLMs, operate as "black boxes." They don't provide easily interpretable reasons for their conclusions. In law, reasoning is paramount – you can't just say "the AI predicts we'd lose this case" without explaining *why*. If an AI tool can't explain its suggestion in a way that a lawyer can incorporate into legal reasoning, its utility is limited. For instance, a judge would not accept "because the algorithm said so" as a reason to rule a certain way. This is a limitation for any AI that goes beyond mechanical tasks. Efforts are underway to improve **AI explainability**, and some tools do offer **rationale features** (like showing similar past cases to justify a prediction), but it remains a concern. In scenarios like e-discovery, lack of explainability is less of an issue (the AI is simply surfacing documents), but in outcomes or strategy, it's a real barrier. Moreover, lack of explainability can complicate **compliance with professional duties** – if a lawyer uses an AI-generated draft, the lawyer must be able to affirm the reasoning and facts in it. If the lawyer doesn't understand how the AI arrived at a certain phrasing or conclusion, that could be problematic. The ABA has implied that lawyers need to understand the basics of how their AI tools work (technological competence) , which is hard if the model is a proprietary black box.

- **Confidentiality and Data Security:** Legal work is bound by confidentiality and attorney-client privilege. A big risk of AI tools, especially cloud-based or third-party ones, is **inadvertent disclosure or misuse of sensitive data**. If a lawyer uploads a client memo to a free AI service, that content might be stored on the provider's servers, potentially accessed by others or used to further train the AI (unless the terms explicitly forbid it, as some now do after public concern). Early on, there were reports of lawyers and other professionals inputting sensitive info into ChatGPT and then that info turning up in other users' outputs (OpenAI has since implemented some protections, but the fear remains). Law firms have responded by **banning or restricting use of public AI tools** for any confidential information, at least until they have vetted, secure alternatives. Even with enterprise contracts, there is a risk – it's essentially outsourcing data to a vendor, which requires due diligence (security audits, ensuring compliance with

data localization laws if applicable, etc.). In Europe, **GDPR** adds another layer: if personal data is processed by an AI, you must have a legal basis and perhaps consent, plus individuals might have rights to not be subject to purely automated decisions. All these mean that deploying AI in legal contexts often needs **privacy impact assessments and contractual safeguards**. The **Embroker survey** cited earlier found data privacy was the number one concern among lawyers regarding AI adoption . Many firms only got comfortable by using **in-house models or vetted providers (like Harvey, which promises not to use client data to train others' models)** . Thus, the confidentiality limitation can slow AI adoption and requires legal teams to negotiate robust confidentiality clauses with vendors. Even beyond external risks, there's an internal one: AI can make it *easy* to search or generate content from internal data, which is great for efficiency but raises **who can access what** questions – firms need to ensure that an AI that indexes all documents doesn't accidentally violate ethical walls or Chinese walls within the firm, or allow a user to see another client's data that they shouldn't. Strong permission controls and data segregation thus have to be part of any legal AI implementation.

- **Ethical and Professional Responsibility Issues:** The use of AI can implicate various professional ethics rules. For instance, under the ABA Model Rules, lawyers have a duty of **competence (Rule 1.1)** which now includes keeping abreast of technology . This means lawyers must understand AI's benefits and risks. Ignorance is not an excuse – if AI could have made a task more efficient or thorough, arguably competence might require its considered use (some have argued that someday failing to use AI for large discovery could be seen as not competent, if it's standard practice). But at the same time, **over-reliance on AI** can breach duties: the duty of **supervision (Rules 5.1 and 5.3)** means a lawyer must supervise nonlawyer assistance – which explicitly covers AI tools now . If an AI makes a mistake, the supervising lawyer is on the hook. There's also the duty of **candor (Rule 3.3)** to courts – knowingly presenting false information (even if from an AI) is unethical, and now judges are wise to ask attorneys to certify their filings are vetted by humans . There's concern about **unauthorized practice of law (UPL)** as well: if AI tools are offered directly to consumers (like DoNotPay's "robot lawyer"), are they practicing law without a license? The DoNotPay saga, where a **class action accused the AI startup of UPL and providing substandard legal services** , shows this risk. DoNotPay had to settle that case and also faced an FTC action for **deceptive marketing as a "AI lawyer"** leading to a fine and requirement to inform customers it's not actually a lawyer . Law firms using AI must avoid any implication that the AI is *giving legal advice* independent of the lawyer. And client consent may be an emerging issue: should clients be informed if significant parts of their work are done by AI rather than human lawyers? Some might not care if quality is high, others might object. Ethically, full transparency is safest. Already, certain courts and bars are trending toward requiring **disclosure of AI usage** in legal processes (as with

Judge Starr's order).

- **Intellectual Property and Data Ownership:** Another risk is who owns the output of AI or the data used to train it. In the ROSS v. Westlaw case, a legal AI startup got sued for using Westlaw's headnotes (editorial summaries) to train its research AI . In 2025 a court found this **violated Thomson Reuters' copyright** and rejected the startup's fair use defense . ROSS had to shut down under the weight of the lawsuit . This indicates that training an AI on proprietary legal content (like Westlaw or Lexis databases) can be legally risky if done without permission. So law firms have to be careful about **what data they use to fine-tune models** – using client documents is sensitive (for confidentiality) and using third-party treatises or databases could be infringing. The output of AI is also a grey area: by U.S. Copyright Office policy, purely AI-generated text isn't protected by copyright (due to lack of human authorship), so who "owns" an AI-written brief? Likely the client or firm as a compilation with human edits, but it's a new question. There's also risk of inadvertently embedding licensed text: if an AI training included some model contract language that is proprietary, and it regurgitates it verbatim, that could violate license. These IP issues mean legal AI use needs **clear policies on data usage** and vetting of training sources. Providers are now careful – e.g. OpenAI offers options to not use a customer's data for training. Still, attorneys should be mindful when using AI that they don't accidentally *leak* or infringe content. The ROSS case's outcome *"could have significant implications for AI companies… that seek to train their systems using copyrighted materials"* , making it a landmark in this area.

- **Reliability and Over-reliance:** Even when AI is functioning as designed (no bugs, no hallucinations), there's the broader risk of **over-reliance**. Lawyers might become complacent and not develop their own skills. For example, junior lawyers could get used to AI drafting and not learn how to draft from scratch, which in the long run could erode the quality of legal reasoning in the profession. There's also the psychological effect: if an AI consistently gives good suggestions, humans might stop questioning them, leading to blind spots. The Spellbook guide wisely notes that *"AI complements human judgment, it doesn't replace it – legal professionals must validate AI outputs, especially in complex or high-stakes matters"* . Over-relying on AI can also reduce the diversity of thought – if everyone uses the same AI suggestions, legal arguments might become homogenized, potentially stifling creativity or missing novel arguments. Additionally, technical glitches could cause problems: imagine an AI tool going down the night before a filing deadline if the team has become dependent on it. Lawyers need fallback plans and should maintain core competencies to operate without AI if necessary.

- **Regulatory Compliance of AI Itself:** Finally, using AI means navigating new regulations *about AI*. The EU AI Act will require those deploying high-risk AI (like in courts or HR hiring tools at a law firm) to perform risk assessments, ensure

human oversight, and possibly register the AI system. Even lower-risk uses have transparency obligations if they produce content – e.g., if a law firm uses an AI to generate marketing content or chatbot answers on its website, the AI Act might require disclosure that it's AI-generated. In the US, while there isn't a comprehensive law yet, the FTC has been clear that it will go after companies for **deceptive AI claims** or harmful AI practices (as seen with DoNotPay). State laws like those in Illinois (for AI in hiring) or proposed bills in various jurisdictions mean legal teams must keep track of compliance not just with traditional laws but also **meta-compliance (complying with laws about AI usage)**.

In summary, the mantra for lawyers should be: **augment, not automate** – use AI to assist, but maintain human control and responsibility. The limitations and risks of AI underscore that it is not a magic bullet; it is a powerful tool with failure modes that need mitigation. Law is a field where the stakes are high (people's rights, significant sums of money, even lives in criminal justice), so the tolerance for AI error is low. As the ABA Ethics guidance suggests, lawyers must use "independent judgment" on AI outputs and remain "accountable" for its use . Those who treat AI as a junior – not senior – partner will navigate these limitations best. They will double-check, validate, and supervise the AI, thereby harnessing its benefits while guarding against its pitfalls.

Notable Failures and Controversies

The deployment of AI in legal settings has already led to some **high-profile failures, controversies, and unintended consequences** that offer cautionary lessons:

- **Hallucinated Citations Case (2023)** – The *Levidow, Levidow & Oberman* law firm in New York made headlines (and not in a good way) when an attorney used ChatGPT to help write a brief. The AI produced **six phony case citations with bogus quotes**, which the lawyer then submitted to the federal court . The opposing counsel and judge discovered that these cases were nonexistent. In a scathing order, Judge P. Kevin Castel sanctioned the two lawyers, imposing a $5,000 fine and lambasting them for "acts of conscious avoidance and false and misleading statements to the court" . The incident became a media sensation and a wake-up call to the entire legal community about blindly trusting AI. It illustrated that even intelligent lawyers can be lulled by AI's authoritative tone. The lawyers claimed they were *"unaware of the possibility that [ChatGPT's] content could be false"* – an astounding admission that highlighted the need for AI literacy. This case has prompted some judges (like Judge Brantley Starr) to

issue standing orders about AI use in filings , and bar associations to educate members that AI's outputs must be verified independently. It's a prime example of how AI, when misused, can threaten the integrity of legal proceedings and attorneys' reputations.

- **Racial Bias in Criminal Risk Assessment (COMPAS)** – Mentioned earlier, the COMPAS algorithm controversy is a landmark in algorithmic bias discussions. In 2016, investigative journalists at ProPublica analyzed COMPAS's predictions for thousands of defendants and found a significant racial skew: **Black defendants who did not reoffend were nearly twice as likely to be misclassified as high risk than white defendants** (45% vs 23%), while white defendants who did reoffend were almost twice as likely to have been mislabeled low risk . This indicated that the algorithm, used in bail and sentencing decisions, could lead to harsher treatment for Black defendants and more lenient treatment for white defendants, all else equal. Northpointe (now Equivant), the company behind COMPAS, disputed some conclusions, but the controversy sparked widespread debate and research into fairness of AI in justice. It also led to at least one legal challenge (in Wisconsin, a defendant argued his sentence using COMPAS violated due process, though the court upheld use of COMPAS with caveats). The COMPAS saga is a cautionary tale of **unintended consequences** – a tool intended to reduce bias and bring consistency ended up reinforcing biases and generating public mistrust in legal AI. It underscores that without transparency and bias mitigation, AI can undermine confidence in the legal system. In response, some jurisdictions have dialed back use of such tools, or at minimum, demanded more transparency (e.g., requiring disclosure of the score to the defendant and how it's used).

- **"Robot Lawyer" Unauthorized Practice and Deception (DoNotPay)** – DoNotPay is a startup that launched as a consumer-friendly AI legal assistant (initially to fight parking tickets via an app). It gained fame as the "world's first robot lawyer," and ambitiously attempted to have an AI **live-argue in traffic court** via an earpiece to a defendant – an idea that was scuttled after several state bar associations warned it would constitute unauthorized practice of law. The company then faced a **class-action lawsuit in 2023** (Faridian v. DoNotPay) alleging that DoNotPay was practicing law without a license and providing poorly prepared legal documents . The plaintiff, a customer, said the services (like AI-generated demand letters and small claims filings) were substandard and that he wouldn't have paid for DoNotPay *"had he known DoNotPay was not actually a lawyer."* DoNotPay reached a settlement in 2024, the terms of which weren't public, effectively conceding the issue . On top of that, in 2024 the U.S. **FTC launched an investigation** and later an enforcement action against DoNotPay for misleading consumers with its marketing . The FTC's order in early 2025 *prohibited DoNotPay from calling itself a "robot lawyer" or implying it is a substitute for a human attorney* and fined the company $193,000 . This

controversy highlights the **regulatory and ethical boundaries** that AI must navigate. DoNotPay's hype overran its capabilities, and it encountered the wall of UPL – only licensed lawyers can provide individualized legal advice in most jurisdictions. The case shows that entrepreneurial zeal in legal AI must be tempered by compliance with existing legal ethics rules. It also suggests a risk of consumer harm when people rely on AI legal tools without understanding their limits – a poorly done demand letter could prejudice a case, for example. The lesson: legal AI companies should avoid overpromising and must likely keep a human in the loop or clearly limit their services to informational, not advisory, roles to stay on the right side of the law.

- **Legal AI Startup vs. Publisher (ROSS Intelligence and Westlaw):** ROSS Intelligence was a much-heralded legal research AI startup – an offshoot of IBM Watson – that aimed to disrupt the duopoly of Westlaw and Lexis. It provided an AI-powered natural language search over case law. But in 2020, Thomson Reuters (owner of Westlaw) sued ROSS, alleging that ROSS illicitly used Westlaw's proprietary headnotes and annotations to train its AI . Westlaw's headnotes (brief summaries of points of law in cases) are copyrighted. ROSS had apparently hired a third-party (LegalEase) to obtain Westlaw content to train its models after TR refused to license it to them . The litigation was devastating for ROSS; even though the case hadn't reached a final verdict, the legal costs forced ROSS to shut down operations by January 2021 . Finally, in 2025, a federal judge ruled that ROSS **did infringe Westlaw's copyrights** and that this use was not fair use . This outcome sent shockwaves in the AI community, because many AI models are trained on vast amounts of data, some of which may be copyrighted. It raised the question: *how to train legal AI without running afoul of IP law?* The ROSS controversy could stifle smaller legal AI entrants (who lack their own dataset of caselaw or capital to license it) and has potentially strengthened incumbents (as Westlaw and Lexis build their own AI, with their data advantage locked up). The case is a prominent example of how **IP and competition issues** intersect in legal AI. It also arguably is a "failure" in that a promising technology (an AI research assistant) was essentially removed from the market not due to its tech failing, but due to legal/IP barriers – a perhaps unintended consequence of how we regulate data and intellectual property.

- **Ban on Judge Analytics (France 2019):** As discussed, France enacted a unique law prohibiting the publication of any analysis of judges' decision data that could identify or profile individual judges . Violating this carries criminal penalties. This can be seen as a *pre-emptive strike* against legal AI analytics in the judiciary. The impetus was that judges feared being rated or having their tendencies exposed in a way that could be used to forum shop or undermine their authority . The unintended consequences of this ban are notable: it arguably stifles an entire segment of legal tech in France (startups like Predictice had to adjust to only provide analytics without naming judges). Critics say it hampers transparency – why shouldn't the public know if a judge is an outlier in,

say, sentencing? The ban itself has been controversial, seen by some as an overreach to protect judges at the expense of public data freedom . This situation underscores a cultural difference: in some jurisdictions, the legal system may reject certain AI uses not because of tech failure, but because of value choices (transparency vs. judicial privacy). It's a reminder that legal AI can clash with tradition and trigger regulatory pushback that limits its deployment.

- **Other Controversies:** We can note a few other incidents briefly. In 2020, during the COVID court disruptions, some courts tried **AI-driven remote hearing solutions** and faced challenges (like biases in facial analysis for witness credibility – though largely experimental). There was also the UK's Post Office scandal (not AI but an IT system error leading to wrongful prosecutions) often cited as a warning about over-reliance on automated systems without recourse. In the corporate world, some companies using AI for hiring faced lawsuits under Illinois' Biometric privacy law and NYC's new law requiring bias audits for AI hiring tools – while not legal practice per se, law firms advising on compliance with these show how AI can create new legal liabilities if not carefully managed.

Each of these controversies teaches a lesson: **hallucinations teach verification; bias teaches cautious use and demand for explainability; UPL and IP disputes teach working within legal boundaries and obtaining proper rights; regulatory bans teach aligning technology with societal values**. They collectively have slowed down the most extreme or naive deployments of AI and have led to a more sober, responsible approach in the industry. The good news is that the legal community is learning from these failures. For example, after the ChatGPT citation fiasco, many law firms instituted policies and training on AI, emphasizing that all content must be checked. After ROSS, startups in legal tech have been more careful about data licensing, and Westlaw/Lexis are making their content available in AI-powered formats (so future innovators might partner rather than scrape). The DoNotPay episode has drawn clearer lines around consumer-facing legal AI – we're likely to see more collaboration between tech and licensed lawyers (rather than direct-to-consumer legal advice bots without oversight).

Emerging Trends and Regulatory Outlook

AI in the legal industry is a fast-evolving field. Looking ahead, several **emerging trends, innovations, and regulatory developments** in the EU and US are poised to shape how legal AI is adopted:

1. Generative AI Integration into Mainstream Legal Platforms: 2023–2025 saw an explosion of generative AI (LLM) integration. This trend will deepen. The major legal research and drafting platforms are racing to embed GPT-like capabilities natively. Thomson Reuters' **Westlaw Precision + CoCounsel** and LexisNexis's **Lexis+ AI** are leading examples – these tools let users chat with the database, get draft answers with pinpointed sources, automatically summarize or rewrite documents, etc. By having these in mainstream products, AI will become a seamless part of lawyers' daily workflow (rather than a separate app). Microsoft is also bringing AI to Office 365 in a big way (e.g., *Microsoft 365 Copilot*), which will allow lawyers to do things like generate a first draft of a contract from bullet points directly in Word, or get an email summary in Outlook. Legal tech companies are creating specialized "copilots" for lawyers – indeed, in one survey **47% of legal orgs said generative AI is now essential to daily work** . We can expect *"AI copilots"* for specific practice areas: e.g., an "M&A due diligence assistant" that combines data room analysis, contract AI, and task management; or a "Litigation co-pilot" that drafts deposition outlines, summarizes evidence, and even helps prepare jury instructions based on case law. Law firms might develop their own proprietary AI on their vast data (briefs, memos, outcome databases) – effectively *private LLMs fine-tuned to the firm's knowledge*. One trend is **domain-specific LLMs**: there are already open-source or specialized models like *LegalBERT* for legal language, and we might see an emergence of *"LegalGPT"*-type models that are trained on caselaw and statutes and thus perform better than general models on legal tasks. For instance, in 2024, a project called *CaseGPT* was reportedly training on US court opinions to create an open legal LLM. The EU might foster multilingual legal LLMs for European law (given the need to handle many languages and cross-border directives).

2. Retrieval-Enhanced and RAG-Based Legal Assistants: To combat hallucinations, nearly all upcoming legal AI tools are adopting retrieval-enhancement. We'll see robust development of **vector databases of legal texts**: companies creating embeddings of entire law libraries so that AI can fetch the exact passage of a case or regulation to ground its answer (somewhat like an AI-boosted search engine with instant memo writing). This will make AI much more reliable and acceptable in practice. The user experience will be: ask a complex legal question and get an answer *with citations*, where the AI has essentially done the research for you. Law firms are doing this internally too with their document management systems – e.g., one firm connects an AI to its brief bank and time-entry narratives to answer "Have we handled a similar matter? What was the strategy and outcome?" reliably with references to prior files. This trend leverages firms' *internal big data*. As more legal data gets digitized (court transcripts, regulatory guidance, etc.), RAG-powered AI will have more fuel.

3. Multi-Modal and Advanced Capabilities: AI won't be limited to text. Emerging multi-modal AI can process images, voice, and structured data. In law this means: AI could analyze patent drawings or video evidence. For example, an AI might extract text from a photo of a contract (OCR) and then analyze it, or transcribe a witness video and flag emotional cues or inconsistencies (though the latter is speculative and controversial). We may also see AI helping with **quantitative legal analysis** – e.g., processing large datasets for an antitrust case or damages modeling. Tools like *Litigation Analytics* might expand to incorporate economic data, jury verdict data, etc., and AI will help find patterns. Another budding area is **AI in contract negotiation**: startups are testing AI that can autonomously mark up contracts or even negotiate via email within set parameters. While fully autonomous negotiation is likely far off (humans will supervise, especially for non-standard points), AI could handle back-and-forth on standard terms like NDAs or procurement contracts, effectively **automating the negotiation of routine agreements** up to a point. Some legal AI products already advertise an "auto-negotiation" feature for simple contracts – expect this to mature, perhaps with AI suggesting compromises ("Both sides are stuck on X clause; how about using Y language which is midpoint based on market data?").

4. Law Firm Innovation and New Roles: The rise of AI is changing law firm operations. Firms are hiring **legal technologists and data scientists**, and forming AI committees to evaluate tools. The "augmented lawyer" is becoming the model – lawyers who know how to leverage AI will be more valuable. We might see **billing changes**: if AI makes certain tasks very fast, firms may shift from billable hours to flat fees or value-based billing for those tasks. Conversely, firms might bill for AI-assisted work as a separate line (some already do for e-discovery hosting, etc.). Another trend is **collaboration between firms and AI startups**: e.g., Allen & Overy's partnership with Harvey in 2023, which is now followed by many others. Many big firms are piloting multiple AI tools to find a fit. In-house legal departments too are pressuring law firms to use AI to reduce costs – the 2025 Clio Legal Trends report noted that large majorities of big-firm lawyers are adopting AI, and even solos are doing so in significant numbers . The skill to *"prompt"* AI effectively, check its work, and integrate it into legal workflows is becoming part of lawyer training. Some law schools in the US have started offering courses on legal AI or legal technology to prepare new attorneys.

5. Regulatory Developments – EU: The European Union is leading in AI regulation with the **EU AI Act**. This comprehensive law (expected to take effect around 2026) will influence legal AI significantly, especially in Europe but also globally (since global vendors will adjust products to comply). Under the final AI Act, any AI used in the "administration of justice" is high-risk . This means if a court or even an arbitration uses AI to aid decision-making, that system must meet strict requirements: risk assessments, documentation, human oversight, etc. If a legal AI system were to be

used to predict case outcomes or recommend judicial decisions, it might be subject to prior conformity assessments. Also, AI that provides legal advice might be considered high-risk if it could impact someone's rights (this is debatable in the text, but something to watch). The AI Act will also enforce **transparency**: lawyers in Europe using AI may have to inform clients when a significant output is AI-generated. For example, if a law firm uses an AI to generate a first draft of a memo, they might have to disclose that to the client (some national bars could incorporate that as best practice too, even if not explicitly required by the Act, out of caution). The Act could also require that if AI is used to evaluate people (like in hiring or litigation risk scoring), individuals have the right to a human review. For legal AI vendors, complying with the AI Act will likely involve building **audit logs, explaining model decisions where possible, and ensuring data governance** (no unauthorized personal data usage, etc.).

Another EU angle is **data protection**: AI models often need tons of data. Under GDPR, using personal data to train an AI or to feed prompts requires a lawful basis. We may see clarification from regulators or courts on whether using personal data in an AI model is compatible with the purpose for which data was collected. Already, Italy temporarily banned ChatGPT in 2023 for GDPR violations until OpenAI implemented options to delete data and provided transparency. Legal teams in EU must thus tread carefully, possibly anonymizing data before sending to AI, etc. Additionally, **professional secrecy laws in Europe** (like France's strict rules on lawyer–client confidentiality) might limit cloud use of AI – law firms might push for on-premises AI solutions or EU-based cloud with clear non-access clauses.

6. Regulatory Developments – US: The US likely won't have a broad AI law in the near term, but we will see **sectoral and state-level regulations** that affect legal AI. As mentioned, one immediate development is within the judiciary: some courts could formalize rules about AI in filings (e.g., requiring certification or disclosure of AI use, as Judge Starr did). There's already a proposed amendment to the Federal Rules of Evidence being discussed about the admissibility of AI-generated evidence – to ensure that if an AI output is used as evidence or analysis, it meets reliability standards equivalent to an expert witness . Lawyers and bar associations are debating how to update ethics rules. The ABA issued Formal Opinion 512 in 2023 giving a framework for GenAI use, essentially saying existing rules (competence, confidentiality, supervision, etc.) suffice but must be applied to AI . We might see state ethics opinions requiring, for instance, that clients consent before sensitive info is shared with an AI vendor, or warning that using AI without understanding it could violate competence. On the legislative side, bills like the Algorithmic Accountability Act (proposed) and the NIST AI Risk Management Framework (voluntary) encourage organizations to assess AI for bias and reliability – law firms might start adopting these frameworks to self-

regulate their AI use and to be able to show clients and courts they handle AI responsibly.

The FTC will continue to police deceptive claims around AI – not just DoNotPay but more broadly, if a law firm claims "Our AI will win your case guaranteed!" that would draw ire. Also, if AI is used in consumer-facing legal tools (like legal chatbots on websites), expect the FTC and state AGs to ensure such tools don't mislead consumers or practice law illegally.

7. Cross-border Practice and AI: An interesting trend is that AI might enable more cross-jurisdictional legal practice. For example, an English-speaking lawyer could query an AI about German law and get some guidance (with translation and summary). While that's useful, it raises UPL concerns – would that lawyer be "practicing" German law without a license? If they rely on AI outputs to advise a client, possibly yes. But perhaps AI could also help *bridge* by quickly finding local counsel or local law references. It will be intriguing to see if AI makes multi-jurisdiction practice easier (by handling language and initial research), or if regulators clamp down on that. The EU is somewhat ahead in multi-language capabilities (the EU Commission sponsors projects for legal information systems covering multiple countries). One can foresee an EU-level AI that explains how a given EU directive is implemented in each member state's law – a handy compliance tool for companies and counsel.

8. Increased Client Expectations and Market Pressure: As AI becomes more entrenched, clients will expect their lawyers to use it to be efficient. Already, general counsel are asking outside firms: "Are you using AI or modern tech to keep our costs down?" If not, they may switch to firms that do. Conversely, clients may also be concerned about quality control – some might say "Don't use AI on my sensitive case" (just like some clients previously said "don't outsource my document review overseas"). Law firms will have to educate clients on how they use AI responsibly. Possibly engagement letters might start including a clause about AI (whether it may be used and under what conditions). Another market effect is new entrants: alternative legal service providers (ALSPs) and the Big 4 accounting firms' legal arms could leverage AI to handle legal work more cheaply, increasing competition for traditional firms, especially in high-volume work like contract management and compliance. If an ALSP armed with AI can review contracts at half the cost of a law firm, they'll grab that market share. This will likely push law firms to adopt AI to stay competitive on price and speed. We may also see **legal tech companies partnering directly with corporates** (bypassing law firms) – e.g. companies might license an AI to do first-pass contract negotiation in-house without involving outside counsel as much.

9. AI for Access to Justice: On the innovation side, there's a hopeful trend of AI being used to assist unrepresented individuals and legal aid. For example, there are prototypes of AI chatbots that help fill out forms for asylum applications or expunge criminal records by asking the user questions and auto-filling paperwork. Courts in some places are offering online Q&A systems for self-represented litigants (though typically those are more like guided interviews, not full AI). In the next few years, we might see more AI "legal triage" tools deployed by bar associations or courts to guide people to resources or even draft simple pleadings for them. The challenge is doing this within UPL rules – likely these tools must be very clear they are not a lawyer and ideally have some lawyer oversight in their creation or operation. Some jurisdictions might relax UPL for AI in certain contexts if it demonstrably improves access to justice for those who can't afford a lawyer (this is a contentious area; some argue UPL rules themselves should evolve to allow more tech-enabled help). Internationally, there are experiments like *ODR (Online Dispute Resolution)* platforms that use algorithms to settle disputes (e.g., eBay's automated dispute resolution for buyer-seller issues). The courts might borrow from these to handle small claims or minor disputes with minimal human intervention.

10. Emphasis on AI Ethics and Training: Finally, an emerging theme is the focus on **ethical AI and training lawyers to use AI correctly**. Law firms are developing AI ethics guidelines internally (e.g., requiring human review of AI output, prohibiting input of client data into unapproved tools, etc.). We may see the ABA or international lawyer groups publish AI ethics guidelines specific to law practice. On the technical side, there's likely to be more **model validation studies in legal domain** – similar to how medical AI is tested in trials. For instance, law professors might systematically evaluate how accurate different LLMs are at answering bar-exam style questions or performing legal research tasks. Early results are mixed: GPT-4 famously passed the bar exam, but other studies found legal-specific tasks still challenging for AI. Continuous evaluation will help calibrate trust. Also, AI developers are trying to build in *guardrails* for legal use – like cite-checking modules (to verify that citations are real and relevant) and *hallucination reduction techniques*. OpenAI recently announced working on methods to have models cite sources by training them on citation-rich data. Such innovations could directly benefit legal AI, turning a known risk (hallucination) into a solved or mitigated issue over time.

In terms of regulation, one can predict a likely convergence: the **EU's stricter approach may set a de facto standard** for high-risk uses, while the US approach might remain more laissez-faire but with high legal liability for failures (through malpractice, sanctions, FTC actions, etc., rather than pre-approval of AI systems).

Lawyers in international firms will probably comply with the most stringent applicable rules to be safe – meaning many will adopt EU-style caution globally (for example, always keeping a human in the loop, performing bias audits on their AI tools, documenting how AI is used in matters, etc.).

In conclusion, the trajectory is that **AI will become an ever-present assistant in legal work** – not replacing lawyers, but empowering those who learn to use it, and perhaps displacing those who do not. The competitive and economic pressures favor adoption, while the ethical and regulatory pressures shape *how* adoption happens (with oversight, transparency, and accountability). As one Gartner analyst aptly said, *"the ability to harness [AI] effectively for legal work will become table stakes for most employees in the industry."* In other words, much like proficiency with computers and online research is assumed today, **proficiency with AI will be a baseline expectation for the lawyers of tomorrow**. Firms that integrate AI thoughtfully – leveraging its strengths and controlling its weaknesses – stand to gain in productivity and insight. Those who ignore it may find themselves falling behind or facing clients' skepticism.

Both the EU and US are feeling their way through this transformative period. The EU leans toward a **"trustworthy AI"** framework – ensuring AI is lawful, ethical, and robust, especially in law and governance. The US leans toward **innovation with guardrails** imposed post-hoc (through courts and agencies). Ultimately, they share the goal of using AI to improve legal services and justice, without sacrificing the rule of law or public trust. The next few years will likely bring clearer standards (via new laws, bar rules, and precedents) and even more powerful AI tools, making it an exciting but critical time for the legal profession to evolve. By embracing innovation *and* adhering to core legal values, lawyers can turn AI from a disruptor into an indispensable ally in the service of clients and justice.

Sources:

- Global legal AI adoption trends ; Survey of top AI use cases (Ironclad report) .
- Gartner ranking of legal generative AI use cases .
- Example AI tools in research (Westlaw "Ask", CoCounsel) .
- Contract AI capabilities (Kira, LawGeex, Luminance) .
- Litigation AI tools (Lex Machina, Relativity) .
- Compliance AI example (Compliance.ai monitoring) ; Spellbook contract compliance checks .
- Generative AI at Allen & Overy (Harvey partnership) .
- ABA and ethical use of AI .

- Judge requiring AI disclosure .
- ChatGPT fake cases sanction .
- COMPAS bias findings .
- ROSS copyright case outcome .
- DoNotPay class action and FTC order .
- French ban on judge analytics .
- **(Additional citations inline above.)**

References

Allen & Overy. (2023). *Allen & Overy rolls out Harvey AI to lawyers globally*. https://www.allenovery.com

American Bar Association. (2023). *Formal Opinion 512: Generative artificial intelligence tools*. https://www.americanbar.org

Brainspace. (n.d.). *AI-powered analytics for e-discovery*. https://www.brainspace.com

Casetext. (2023). *Casetext launches CoCounsel powered by GPT-4*. https://casetext.com

Clio. (2024). *Legal trends report*. https://www.clio.com

Clio. (2025). *Legal trends report*. https://www.clio.com

Compliance.ai. (n.d.). *Regulatory change management powered by AI*. https://www.compliance.ai

DeepL. (2025). *AI adoption in the legal sector* [Press release]. https://www.deepl.com

DoNotPay, Inc. (2024). *Settlement of consumer litigation* [Court filings].

European Commission. (2024). *Artificial Intelligence Act: Final text*. https://digital-strategy.ec.europa.eu

Everlaw. (n.d.). *Cloud-native e-discovery with AI*. https://www.everlaw.com

Federal Trade Commission. (2025). *FTC takes action against DoNotPay for deceptive AI marketing*. https://www.ftc.gov

Financial Times. (2024). *How AI is removing legal obstacles that slow down business*. https://www.ft.com

Gartner. (2025). *Top use cases for generative AI in legal departments*. https://www.gartner.com

Ironclad. (2024). *State of AI in legal report*. https://ironcladapp.com

Ironclad. (2025). *AI adoption, productivity, and burnout in legal teams*. https://ironcladapp.com

Kira Systems (Litera). (n.d.). *Machine learning contract analysis*. https://kirasystems.com

LawGeex. (n.d.). *Automated contract review and redlining*. https://www.lawgeex.com

Lex Machina. (n.d.). *Litigation analytics and outcome prediction*. https://lexmachina.com

LexisNexis. (2024). *Lexis+ AI: Generative AI for legal research*. https://www.lexisnexis.com

Luminance. (n.d.). *Legal-grade AI for contracts and due diligence*. https://www.luminance.com

New York State Unified Court System. (2023). *Mata v. Avianca, Inc.*, No. 22-cv-1461 (S.D.N.Y.).

OpenAI. (2023). *GPT-4 technical report*. https://openai.com

ProPublica. (2016). *Machine bias: There's software used across the country to predict future criminals—and it's biased against blacks*. https://www.propublica.org

Relativity. (n.d.). *Technology-assisted review and AI for e-discovery*. https://www.relativity.com

Ross Intelligence, Inc. v. Thomson Reuters Corp. (2025). *Copyright infringement decision* (D. Del.).

Spellbook. (n.d.). *AI-powered contract drafting and compliance checks*. https://www.spellbook.legal

Thomson Reuters. (2025). *Generative AI in professional services: Legal edition*. https://www.thomsonreuters.com

Thomson Reuters. (2024). *Ask Westlaw and CoCounsel product overview*. https://legal.thomsonreuters.com

United States District Court for the Northern District of Texas. (2023). *Standing order regarding use of generative AI*.

Westlaw. (n.d.). *Westlaw Precision and AI-assisted legal research*. https://legal.thomsonreuters.com