AI IN THE COURTROOM

Colin S. Levy

February 2026

Contents

1. Introduction

Artificial intelligence has arrived in the courtroom. Not through a single defining event, but through a series of collisions between an old institution and a new technology, each one exposing gaps in rules that were written for a world in which legal research meant opening a book and judicial drafting meant dictating to a clerk. Lawyers have been experimenting with AI for several years now. The conversation about what it means for judges, court administrators, and the structure of the justice system itself is considerably younger.

This guide is written for lawyers and legal professionals who encounter AI in the judicial branch and need to understand what courts are doing, why they are doing it, and what it means for practice. Whether you are an attorney preparing to file in a jurisdiction with an AI disclosure requirement, a court administrator evaluating an AI transcription vendor, or a legal aid lawyer considering an AI chatbot for client intake, this document is intended to give you the practical and institutional context you need.

The stakes are real. When a lawyer submits a brief riddled with fabricated citations, the filing is struck and the attorney sanctioned. The system corrects itself. But when a judge relies on AI-generated analysis that contains errors, the consequences are structurally different. A flawed judicial opinion can become precedent, bind future courts, alter the rights of individuals who were never parties to the original case, and corrode the institutional legitimacy on which judicial authority depends.[1]

This guide is part of a series on AI and law that includes *AI For Lawyers*, *AI for Legal Teams*, *AI Agents Data Handling and Cybersecurity Guide*, and *Law School in the Age of AI*. Those documents address individual practitioners, organizational adoption, cybersecurity, and legal education. This one turns to the bench and the institution that surrounds it: how courts are confronting AI, what tools they are adopting, what policies have emerged, and what the judiciary's engagement with these technologies reveals about the adaptability of a legal system designed centuries before anyone imagined a machine that could write a brief.

Consider the pace: in less than three years since the *Mata v. Avianca* sanctions, the RAILS AI Use in Courts Tracker has cataloged over 300 AI-related court orders across federal and state jurisdictions, a Southern District of New York ruling has unsettled the boundaries of attorney-client privilege, UNESCO has published the first global set of principles for judicial AI governance, Brazil has launched over 140 judicial AI projects, Indiana's Supreme Court has used AI transcription to cut mental health appeal timelines from 143 days to 31, and two federal judges have retracted opinions contaminated by undisclosed AI assistance. What follows is an effort to examine each of these developments on its own terms.

2. When AI Went Wrong in Court

The judiciary's reckoning with AI was catalyzed not by curiosity but by embarrassment. It began with a case that nobody involved would have chosen to be known for, and it has since spread in directions that implicate not just the bar but the bench itself. Understanding these early failures is essential for any practitioner navigating the current landscape, because the rules and orders that followed were direct responses to what went wrong.

The Brief That Changed the Conversation

In early 2022, Roberto Mata filed a personal injury lawsuit against Avianca Airlines in the Southern District of New York, alleging that a metal serving cart struck his knee during a flight. The claim was unremarkable. What made it notorious was the legal research submitted in his support.[2]

Attorney Steven Schwartz of the firm Levidow, Levidow & Oberman, who handled the substantive legal work but was not admitted to practice in the Southern District, used ChatGPT to identify supporting case law. The brief he drafted cited six judicial decisions: *Varghese v. China Southern Airlines*, *Shaboon v. Egyptair*, *Petersen v. Iran Air*, *Martinez v. Delta Airlines*, *Estate of Durden v. KLM Royal Dutch Airlines*, and *Miller v. United Airlines*. None of them existed.[3]

When opposing counsel at Condon & Forsyth could not locate the authorities, Judge P. Kevin Castel ordered Schwartz to produce copies of the cited opinions. Schwartz returned to ChatGPT, which confirmed the cases were real and could be found on Westlaw and LexisNexis. He submitted what he believed were the opinions. They, too, were fabrications. At the sanctions hearing, Schwartz testified that he had been "operating under the false perception that this website could not possibly be fabricating cases on its own."[4]

On June 22, 2023, Judge Castel imposed a $5,000 sanction on Schwartz, his colleague Peter LoDuca, and the firm. The court found both attorneys had acted with "subjective bad faith" sufficient for sanctions under Federal Rule of Civil Procedure 11. Judge Castel also ordered the attorneys to send individual letters to each judge falsely identified as the author of a fabricated opinion, enclosing the order, the hearing transcript, and the fictitious text.[5]

The opinion was measured in its treatment of the technology itself. Judge Castel acknowledged that "there is nothing inherently improper about using a reliable artificial intelligence tool for assistance," and that "technological advances are commonplace." The fault lay not in using ChatGPT, but in failing to verify what it produced and in deceiving the court about the provenance of the citations.[6] For practitioners, the lesson is direct: an AI tool is a research assistant, not a legal authority. Anything it produces must be verified against primary sources before submission to any court.

When the Problem Reached the Bench

The assumption, after *Mata*, was that the AI problem was a lawyer problem. That assumption held for about two years. In the summer of 2025, two federal judges issued opinions that bore unmistakable signs of AI hallucination: invented quotes, references to parties who did not exist in the underlying record, and misstated case outcomes.

U.S. District Judge Henry T. Wingate of the Southern District of Mississippi issued a temporary restraining order on July 20, 2025, blocking enforcement of Mississippi's HB 437, a law banning diversity, equity, and inclusion programs in public higher education. The order cited nonexistent allegations, named parties who were not in the case, and referenced a decision that did not exist. After the Mississippi Attorney General's office flagged the errors, Judge Wingate replaced the order with a corrected version and wiped the original from the docket. In response to an inquiry from Senate Judiciary Committee Chairman Charles Grassley, Judge Wingate acknowledged that a law clerk had used the AI search tool Perplexity to "analyze publicly available information from the docket" and compose the draft.[7] He announced new procedures requiring all draft opinions to undergo independent secondary review and all cited cases to be printed and attached to the final draft.

In a separate incident, U.S. District Judge Julien Xavier Neals of the District of New Jersey withdrew an opinion in *In re CorMedix Inc. Securities Litigation* on July 23, 2025, after defense counsel at Willkie Farr & Gallagher identified misstated case outcomes, nonexistent quotations, and inaccurate attributions. Judge Neals disclosed that a law school intern had used OpenAI's ChatGPT to conduct legal research, in violation of both a verbal chambers policy and the intern's law school policy.[8] He subsequently implemented a written prohibition on AI use in opinion drafting and a multi-level review process for all chambers output.

The institutional implications are distinct from those of attorney misconduct. A lawyer's flawed brief can be identified and stricken before it influences a decision. A judicial opinion contaminated by AI-generated fabrications may not be caught at all, or may be cited by other courts before the errors surface. For attorneys practicing in these courts, the takeaway is sobering: you cannot assume that a judicial opinion citing specific authorities has been verified by human review. Senator Grassley, in his letters to both judges, argued that the judiciary needed "more decisive, meaningful and permanent AI policies and guidelines."[9]

The Accumulating Sanctions

The disciplinary record has continued to grow. A Connecticut federal judge fined a solo practitioner $500 for filing a brief laden with AI-generated fictitious citations. In *Benjamin v. Costco Wholesale Corp.*, a plaintiff's attorney was ordered to pay a $1,000 penalty for relying on cases fabricated by an AI tool called ChatOn. An Australian database tracking court decisions tainted by

generative AI hallucinations has identified over 150 instances to date.[10] Across these cases, the emerging legal standard is straightforward: the duty to verify citations and factual assertions is nondelegable. The tool that generated the text is irrelevant to the attorney's obligation under Rule 11.

3. The Judicial Response: Standing Orders and Certifications

In the weeks following the *Mata v. Avianca* sanctions, judges did not wait for bar associations or legislatures. They began drafting their own rules, court by court, chamber by chamber. For practitioners, the resulting patchwork of orders means that AI disclosure obligations vary depending on which judge you draw, which district you file in, and sometimes which side of the case you represent.

Judge Brantley Starr and the Certification Model

On May 30, 2023, days before Judge Castel's *Mata* sanctions order, U.S. District Judge Brantley Starr of the Northern District of Texas issued what appears to be the first federal standing order specifically addressing generative AI.[11] His Mandatory Certification Regarding Generative Artificial Intelligence requires every attorney appearing before his court to file a certificate either confirming that no portion of any filing was drafted by generative AI, or attesting that any AI-generated language was verified for accuracy using traditional legal databases or print reporters.

The order was blunt in its assessment. It named ChatGPT, Harvey.AI, and Google Bard as examples of covered tools and stated that while these platforms "are very powerful and have many uses in the law, legal briefing is not one of them" because "these platforms in their current states are prone to hallucinations and bias."[12] Attorneys who fail to file the certification risk having their filings struck and face potential Rule 11 sanctions.

Divergent Approaches Across the Federal Bench

Judge Starr established the template, but the orders that followed exposed a wide range of views about what the problem actually is and how best to address it. Practitioners should understand the main categories, because the type of order in your jurisdiction determines your obligations.

At one end of the spectrum, Judge Michael J. Newman of the Southern District of Ohio prohibited AI use entirely, with only a narrow carve-out for traditional legal search engines. At the other, the Fifth Circuit acknowledged in June 2024 that it had deliberately declined to issue a generative AI rule, reasoning that existing obligations under the Federal Rules of Appellate Procedure already require attorneys to stand behind the accuracy of their filings.[13] Under this view, AI does not create new duties; it merely provides a new way to breach old ones.

Between these poles sit orders that focus on disclosure, requiring attorneys to identify which AI tools they used and to certify that no confidential client information was shared with an unauthorized platform. Others apply selectively by party: the Eastern District of Missouri's prohibition covers only pro se litigants, while the Eastern District of Texas bars AI use by pro se filers but permits it for represented parties provided counsel verifies the output.

The RAILS AI Use in Courts Tracker, maintained by a consortium of legal organizations, has cataloged over 300 of these orders across federal and state jurisdictions as of mid-2025.[14] Individual judicial orders are giving ground to district-wide local rules, circuit-level policies, and statewide directives. For attorneys, the practical imperative is to check the local rules and standing orders of every court in which you appear, because AI disclosure requirements are increasingly common and the penalties for noncompliance are real.

4. Judge Scott Schlegel and the Path Forward

The early judicial engagement with AI was reactive: sanctions, hastily drafted standing orders, individual chambers scrambling to address a technology they had not anticipated. What has followed from certain quarters of the bench is more purposeful. Among the most visible figures in that transition is Judge Scott Schlegel of the Louisiana Fifth Circuit Court of Appeal.

Judge Schlegel's credentials in legal technology predate the current AI moment. During a decade on the 24th Judicial District Court in Jefferson Parish, he built what has been described as one of the most technologically advanced trial courts in the country, implementing electronic case management, digital evidence systems, and reentry and swift-probation programs aimed at reducing recidivism.[15] He served as Chair of the Louisiana Supreme Court Technology Commission and President of the Louisiana District Judges Association. In 2021, the National Center for State Courts awarded him the William H. Rehnquist Award for Judicial Excellence, the first Louisiana judge to receive the honor, in a ceremony featuring recorded remarks from Chief Justice John G. Roberts, Jr. and Justice Amy Coney Barrett.[16]

Elected to the appellate bench in 2023, Judge Schlegel has since become a member of the ABA Task Force on the Law and Artificial Intelligence and the National Center for State Courts Joint Technology Commission. He is among a small number of sitting judges who engage publicly and substantively with the operational questions of how AI should function in chambers.

The AI in Chambers Guidelines

In collaboration with four other judges and a professor of law and computer science, Judge Schlegel developed the "AI in Chambers" guidelines, published through the ABA's Working Group on AI and the Courts.[17] They are organized around a single governing principle: "We cannot outsource our judgment."

On verification, the guidelines warn that generative AI tools produce outputs with a tone of confidence that bears no correlation to their factual accuracy. Every citation must be independently checked. On confidentiality, sealed or in camera materials should never be uploaded to a consumer AI platform; enterprise-grade tools with contractual data isolation and audit trails are the minimum acceptable standard. On auditability, any tool whose operation cannot be examined does not belong in judicial chambers.

The most practical guidance concerns the boundary between assistance and abdication. Judge Schlegel draws the line at decision-making. Proofreading, editing, and summarizing case

records are appropriate uses. Anything that touches judicial reasoning is not. His analogy is precise: treat AI the way you would treat a first-year law clerk. Rely on it for preparation, never for judgment.[18]

Predictions and Cautions

In January 2026, Judge Schlegel predicted that AI agents, tools capable of executing multi-step tasks with limited supervision, would begin appearing in judicial chambers within the year. He anticipated that many chambers would adopt at least one agent in a narrow support role: editing, citation-checking, timeline creation, bench memo drafting, or summarization of exhibits and transcripts.[19]

His caution was equally specific. Agents, he argued, should assist in modernizing the courts only if judges "refuse to use them as shortcuts around the safeguards that make the justice system legitimate." The distinction is between delegation and abdication. A judge who asks an AI agent to compile a chronology of filings is delegating a mechanical task. A judge who asks an AI agent to assess the strength of an argument is abdicating a judicial one.

Judge Schlegel has also addressed the challenge of AI-generated evidence, particularly deepfakes. He has pointed to Louisiana's Act 250 (2024), which requires attorneys to exercise "reasonable diligence" to determine whether evidence they submit has been generated by AI, as a legislative model for other states.[20] He is working with the National Center for State Courts and the Thomson Reuters Institute to develop judicial protocols for authenticating digital evidence in an era when the line between genuine and synthetic content is increasingly difficult to detect.

5. Privilege, Evidence, and the New Doctrinal Questions

As AI tools become embedded in legal practice, they are generating doctrinal questions that existing law was not built to resolve. For attorneys, the most immediately consequential of these questions is whether communications with AI tools receive any form of privilege protection. The first major ruling arrived in February 2026.

United States v. Heppner

Bradley Heppner was arrested on securities and wire fraud charges in November 2025. During the search of his residence, federal agents seized electronic devices containing approximately thirty-one documents generated using Anthropic's consumer AI tool Claude. Heppner had used the platform independently to prepare reports outlining defense theories and potential legal arguments before retaining counsel.[21]

His defense team argued these materials were shielded by attorney-client privilege and the work product doctrine. On February 10, 2026, Judge Jed S. Rakoff of the Southern District of New York rejected both claims from the bench. A written opinion followed on February 17, in which Judge Rakoff characterized the issue as a "nationwide" matter of first impression.[22]

The privilege analysis turned on a structural point. Attorney-client privilege presupposes a relationship between a client and a licensed professional who owes fiduciary duties and is subject to professional discipline. Judge Rakoff held that no such relationship can exist between a user and an AI platform. All recognized privileges, he wrote, require "a trusting human relationship" with a professional bound by legal and ethical obligations.[23] An AI chatbot satisfies none of these criteria.

Even if the communications could somehow be characterized as privileged, the court found that Heppner had waived any protection by using the consumer version of Claude, whose terms of service permit Anthropic to collect user inputs and outputs, use that data for model training, and disclose it to third parties, including government authorities upon lawful request. Voluntary disclosure to a third party under those conditions, the court reasoned, destroys any claim of confidentiality.

The work product analysis was narrower. Because Heppner conducted the AI research on his own initiative rather than at the direction of counsel, the materials did not embody defense counsel's mental impressions, conclusions, or litigation strategy. They were, in the court's framing, a client's independent research project, not an attorney's work product.

The Kovel Opening and What It Means for Practice

Judge Rakoff's opinion was notable as much for what it left open as for what it decided. Had counsel directed Heppner to use Claude, the court suggested, the Kovel doctrine might have applied. Under *United States v. Kovel* (2d Cir. 1961), attorney-client privilege can extend to non-lawyer professionals retained by counsel to assist in providing legal advice, such as accountants, translators, and forensic experts. Judge Rakoff observed that under attorney direction, "Claude might arguably be said to have functioned in a manner akin to a highly trained professional" serving as counsel's agent.[24]

What emerges is a privilege analysis that turns less on what the AI tool is than on who directed its use and under what confidentiality conditions. Consumer platforms used independently by clients will almost certainly receive no protection. Enterprise tools deployed at counsel's direction, governed by contractual data isolation and explicit confidentiality provisions, occupy different doctrinal ground. For law firms and corporate legal departments, *Heppner* amounts to a strong incentive to route all client AI use through counsel and to adopt enterprise-grade platforms with enforceable data isolation commitments. Practitioners advising clients should consider issuing guidance on AI use before litigation arises, not after.

6. The Algorithm in the Courtroom: Risk Assessment Tools

Generative AI tools like ChatGPT represent the most visible form of AI in the courts, but they are not the most consequential. Algorithmic risk assessment tools have been shaping criminal justice outcomes, largely without public scrutiny, for more than two decades. Practitioners handling criminal matters should understand how these tools work, what courts have said about them, and why their continued use raises unresolved constitutional questions.

Architecture and Adoption

Pretrial risk assessment tools are statistical models that estimate the probability a person accused of a crime will fail to appear at a future court date or be rearrested if released before trial. They analyze variables including age, employment status, housing stability, and prior criminal history to generate a composite score, typically categorizing individuals as high, medium, or low risk. Over sixty U.S. jurisdictions use some form of these tools.[25] They appear not only in pretrial detention decisions but also at sentencing, in prison classification, and in parole determinations.

COMPAS, ProPublica, and the Impossibility Theorem

The most widely examined risk assessment tool is COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), developed in the late 1990s by Northpointe, Inc. (now Equivant). COMPAS calculates recidivism risk scores used in bail, sentencing, and parole decisions across several states, including Wisconsin, Florida, and New York.

In May 2016, ProPublica published an investigation analyzing COMPAS scores for over 7,000 defendants in Broward County, Florida. The findings were stark: among defendants who did not go on to reoffend, Black defendants were misclassified as high-risk at a rate of 44.9%, compared to 23.5% for white defendants. Among those who did reoffend, white defendants were misclassified as low-risk at 47.7%, compared to 28% for Black defendants. ProPublica concluded that the tool's error patterns were sharply skewed by race.[26]

Northpointe issued a 39-page technical rebuttal two months later, titled "COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity." The company argued that ProPublica had conflated different definitions of fairness. COMPAS, Northpointe contended, was calibrated: a score of seven meant roughly the same recidivism probability regardless of race. Among defendants who scored seven, approximately 60% of white defendants and 61% of Black defendants went on to reoffend. By this measure, the tool treated both groups equivalently.[27]

The exchange attracted attention from computer scientists and statisticians, who identified a deeper problem. Researchers including Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan demonstrated mathematically that when two populations have different base rates of the outcome being predicted, it is impossible for a risk score to simultaneously achieve calibration (equal predictive meaning across groups) and error rate balance (equal false positive and false negative rates across groups).[28] The choice between these fairness criteria is not a technical decision. It is a normative one. And it is a choice that no algorithm can make on its own.

State v. Loomis and the Transparency Deficit

The legal challenge to COMPAS reached the Wisconsin Supreme Court in 2016. Eric Loomis had been sentenced to six years in prison on charges related to a drive-by shooting. At sentencing, the trial judge cited Loomis's high COMPAS risk scores as a factor in the decision. Loomis challenged the use of a proprietary algorithm whose methodology his defense team could not examine, arguing that it violated his due process right to be sentenced on the basis of accurate information.[29]

The Wisconsin Supreme Court upheld the use of COMPAS, subject to conditions: sentencing courts must explain the other factors that informed their decisions, and any COMPAS report must be accompanied by disclaimers about the tool's limitations, including its inability to account for individual circumstances and its potential for group-level bias. The U.S. Supreme Court declined to hear the case.

What the court did not resolve may matter more than what it did. When a proprietary algorithm influences a decision about individual liberty, the Fourteenth Amendment's due process guarantee and the Sixth Amendment right to confront adverse evidence are both in play. Christopher Slobogin, director of the criminal justice program at Vanderbilt Law School, called the *Loomis* opinion "one of the most sophisticated judicial treatments of risk assessment instruments to date,"[30] but the central constitutional question, whether a defendant has the right to examine the source code behind a score that contributed to a prison sentence, remains unanswered.

Idaho took the most direct legislative step in 2019, enacting a statute that prevents vendors from asserting trade secret protections over pretrial risk assessment tools and empowers defendants to review the data and calculations underlying their scores.[31] No other state has followed suit. For criminal defense attorneys, the practical implication is clear: when a risk assessment score is cited at sentencing or in a pretrial detention hearing, counsel should demand disclosure of the methodology, the input data, and the error rates, even if prevailing law does not guarantee access.

Public Confidence

A 2025 experimental study of 1,800 participants, stratified by race, ethnicity, and gender, found that judges who rely solely on their own expertise are perceived as more legitimate and trustworthy than those who use AI, whether entirely or as a supplement to personal judgment.[32] The finding has implications for courts considering the adoption of algorithmic tools in settings visible to the public. Institutional legitimacy depends not only on the quality of decisions but on whether the process is perceived as fair.

7. AI Tools for Judicial Research and Case Preparation

Beyond the headline-grabbing failures and policy debates, AI tools are quietly reshaping the daily work of legal research, transcript analysis, and case preparation in courts across the country. Practitioners should be familiar with the tools that judges and their clerks may be using, because these tools increasingly shape the speed and texture of judicial decision-making.

Legal Research Platforms

The two dominant legal research providers have both integrated AI capabilities into their platforms. Thomson Reuters launched Westlaw CoCounsel, which provides AI-assisted legal research, document drafting, and a citation verification feature called "Quick Check Judicial" that analyzes briefs, evaluates the validity of cited authorities, and flags quotations that do not match the underlying source.[33] LexisNexis introduced Lexis+ AI Protege, a supplemental research and drafting tool that uses extractive AI grounded in LexisNexis's verified legal databases rather than general-purpose language models. Its "Brief Analysis" feature identifies unverified quotations and missing citations.

These tools differ from general-purpose chatbots in a critical respect: they are designed to retrieve information from curated, verified legal databases rather than generating text from statistical patterns. The risk of hallucination is not eliminated, but it is substantially reduced because the output is anchored to primary sources. When a judge's clerk uses CoCounsel to draft a bench memo, the resulting product is more likely to cite real authorities than if the same clerk had prompted ChatGPT with the same question.

Transcript Analysis and Court Reporting

AI is also transforming how courts produce and use transcripts. Traditional court reporting relies on stenographers who capture proceedings in real time and produce written transcripts, a process that typically takes days or weeks. Several AI-powered alternatives are now in use.

The Indiana Supreme Court provides the most closely studied example. In September 2024, the court launched an Expedited Mental Health Appeals Pilot Project in Marion County using Anthropic's AI system to generate transcripts within minutes of a hearing's conclusion. Before the pilot, appeals from involuntary commitment proceedings took an average of 143.3 calendar days from the notice of appeal to the Court of Appeals opinion. Under the AI-assisted process, that average dropped to 30.85 days.[34] The pilot demonstrates that AI transcription can produce

measurable, concrete improvements in case timelines, particularly in proceedings where speed has direct consequences for the people involved.

Commercially available tools include Steno's Transcript Genius, which uses generative AI to produce topic-specific summaries and timelines from deposition and hearing transcripts, and CaseMark, a platform launched in February 2026 that transforms raw transcripts into formatted, citation-rich summaries for court reporting firms.[35] SoniClear RealTime 10 uses AI to deliver real-time transcription during proceedings, and court reporters using its technology have reported reducing transcript preparation time by up to 75 percent.

The NCSC AI Sandbox

The National Center for State Courts has developed an AI Sandbox: a secure, hands-on environment where judges and court professionals can experiment with AI tools in simulated court scenarios without risk to actual cases.[36] The sandbox serves as a training platform, allowing court staff to develop familiarity with AI capabilities and limitations before encountering these tools in operational settings. For courts considering AI adoption, the NCSC sandbox offers a practical starting point that avoids the risks of learning on live cases.

Bench Cards for AI-Generated Evidence

The NCSC has also produced bench cards to help judges evaluate AI-generated evidence. These pocket-sized reference guides walk judges through the questions to ask when a party introduces evidence that may have been created or altered by AI: What tool produced the content? What was the chain of custody? Has the output been verified against independent sources? The bench cards address hallucination risks, authentication challenges under Federal Rules of Evidence 901 and 902, and the specific difficulties posed by AI-generated images, audio, and video.

8. AI in Court Administration: From Pilot Programs to Operations

The research and case preparation tools described in the previous chapter assist individual judges and clerks. But courts are also operational institutions: they manage caseloads, process filings, schedule hearings, and communicate with the public. In these administrative functions, AI carries fewer of the risks that complicate its role in deciding cases, and the practical returns are measurable. For court administrators, legal aid organizations, and attorneys who interact with court filing systems, these operational deployments are often the most immediately relevant applications of AI in the judicial branch.

Electronic Filing and Document Processing

The most mature AI applications in court administration involve document processing at scale. Texas provides the clearest illustration. All 254 Texas counties use the eFileTexas system, operated by Tyler Technologies under a $98 million contract with the Texas Office of Court Administration. In August 2023, Tyler acquired Computing System Innovations (CSI), whose AI-powered tools handle automated redaction, document indexing, and data extraction from court filings.[37] The system has processed over 100 million electronic filings since its inception. AI-driven redaction is particularly valuable in jurisdictions with strict rules governing the public availability of personal information in court records: Social Security numbers, dates of birth, and financial account numbers can be identified and removed automatically, reducing the risk of inadvertent disclosure.

CourtDrive, a cloud-based court management platform, offers an AI Docketing Assistant that ingests scheduling orders and automatically populates calendars, calculates filing deadlines, and alerts attorneys and court staff to upcoming obligations. For practitioners, such tools reduce the risk of missed deadlines, though they also introduce a new dependency: if the AI misinterprets a scheduling order, the resulting calendar entries may be wrong. Human verification remains essential.

State Court Initiatives

Several states have moved beyond individual pilot programs to establish institutional AI strategies for their court systems.

In California, Chief Justice Patricia Guerrero established an Artificial Intelligence Task Force in May 2024, prioritizing generative AI as a focus area for the entire judicial branch. The task force is

developing policy recommendations covering document summarization, case research, and public engagement, with a stated commitment to keeping a "human in the loop" at every stage of AI-assisted judicial work.[38]

Illinois released a comprehensive AI policy through its Judicial Conference Task Force on Artificial Intelligence, providing guidance on responsible AI use across all court operations. Connecticut published a Responsible AI Framework for its judicial branch. Maryland adopted Guidelines for the Acceptable Use of Artificial Intelligence Tools and Platforms. New Jersey issued a Statement of Principles for the Judiciary's Ongoing Use of AI. Utah adopted Interim Rules on the Use of Generative AI. Georgia created a Judicial Council Ad Hoc Committee on Artificial Intelligence and the Courts.[39]

The common thread across these state initiatives is a move from ad hoc judicial orders to centralized, statewide governance. For practitioners, this means that AI policies are increasingly determined at the state level rather than by individual judges, creating more uniform expectations within a given state court system.

Scheduling, Translation, and Case Triage

Courts are testing AI in several administrative domains that do not involve adjudication. Automated scheduling for complex multi-party litigation uses predictive analytics to estimate case duration and allocate courtroom time. Real-time translation services help courts that serve multilingual populations, though the NCSC has cautioned that AI should not replace human interpreters for spoken interpretation during hearings because of accuracy limitations with fast speech, accents, and legal terminology.[40] AI-powered document classification routes incoming filings to the appropriate division or judge. Predictive analytics estimate case complexity and duration, helping court administrators allocate resources.

The federal judiciary is also planning for AI integration. The Long Range Plan for Information Technology in the Federal Judiciary identifies AI data mining tools for case analysis and management as a priority for the NextGen CM/ECF system, the case management platform used across all federal courts.[41]

Courthouse Security

AI is beginning to appear in courthouse facility management as well. Vendors such as LINEV Systems have developed AI-powered X-ray screening systems designed for high-traffic courthouse entrances, using automated threat detection to reduce queues during peak hours while maintaining security standards. While not yet widespread, these tools illustrate the breadth of AI's potential reach within the judicial branch, extending well beyond the courtroom itself.

9. Access to Justice and Self-Represented Litigants

While courts modernize their internal operations with AI, another pressing application involves the people those courts serve. AI's potential contribution to access to justice rests on a specific, quantifiable deficit. Approximately twenty million civil cases are filed in U.S. courts each year, and in roughly 75% of them, at least one party appears without a lawyer. These cases tend to be small in dollar terms but consequential in lived experience: evictions, debt collection, custody disputes, wage claims.[42] For legal aid organizations and courts serving self-represented litigants, AI tools are beginning to address parts of this deficit.

AI Chatbots and Legal Navigation Tools

Legal Aid of North Carolina launched LIA (Legal Information Assistant), a bilingual English-Spanish generative AI chatbot that has conducted over 21,000 conversations with individuals seeking civil legal help since its July 2024 debut. The Nevada Supreme Court has launched a multilingual AI tool that helps unrepresented litigants understand their procedural options.

Pilot navigator programs are underway in multiple states. Nebraska's Supreme Court is testing a self-help center that provides dynamic, accessible guidance on court procedures. New Mexico is running court navigator projects in both urban and rural settings. Kansas is expanding a navigator pilot from two counties toward statewide coverage.[43] These programs use AI to provide the kind of procedural guidance that a court clerk might offer at a help window, but at scale and outside of business hours.

AI in Default Judgment Review and Legal Aid

A collaboration between Stanford Law School and the Los Angeles Superior Court is developing an automated system to review default judgments for legal errors, catching up to 10% of problematic judgments that would otherwise go unreviewed.[44] The Stanford Legal Design Lab, with funding from the Gates Foundation, is building AI "co-pilots" to support legal aid attorneys in eviction defense. The ABA's Task Force on AI concluded in December 2025 that AI has moved "from experiment to infrastructure" for the legal profession, and pointed to over 100 documented AI use cases in legal aid settings.[45]

Design Constraints for Access-to-Justice Tools

Access-to-justice tools carry their own design constraints that practitioners and court administrators should understand. They must avoid crossing the line into providing legal advice, which raises unauthorized-practice-of-law concerns. They must be usable by people with limited digital literacy. And they must not inadvertently widen the gap for populations without reliable internet access or trust in digital systems. A chatbot that helps a tenant understand the timeline for responding to an eviction notice is providing procedural information. A chatbot that tells the tenant whether to file a particular motion is giving legal advice. The line between these two functions is not always crisp, and getting it wrong carries regulatory consequences.

10. How Courts Around the World Are Deploying AI

The United States is not developing judicial AI policy in isolation. Courts in at least a dozen countries are deploying AI tools, and some have moved considerably further than American courts in both adoption and regulation. Practitioners working in international or comparative law contexts should understand these developments, because they are shaping the expectations that global institutions bring to discussions about AI and the rule of law.

United Kingdom

The UK Ministry of Justice announced a three-year AI Action Plan for Justice in 2024, developed in partnership with Microsoft and OpenAI. The plan's stated goal is to provide enterprise-grade AI tools to all 90,000 justice system staff by December 2025, supported by a dedicated Justice AI Unit.[46] HM Courts & Tribunals Service has begun using AI transcription to help judges prepare written decisions, and has piloted machine learning and computer vision tools to extract and analyze information from paper forms. The court system processes over eight million paper forms annually, and automated extraction is expected to reduce the manual processing burden substantially.

In October 2025, the UK Judicial Office authorized judges to use AI tools like ChatGPT for administrative tasks, including summarizing texts and drafting emails, while cautioning against reliance on AI for legal reasoning.[47]

Canada

The Canadian Judicial Council published Guidelines for the Use of Artificial Intelligence in Canadian Courts in October 2024, applicable to all Canadian courts. The guidelines establish that judges retain exclusive responsibility for decisions and that AI cannot replace judicial decision-making. They require that all AI tools used in courts meet stringent information security standards and that AI-assisted decision-making be explainable.[48] Manitoba's Court of King's Bench has gone a step further, requiring lawyers and self-represented litigants to disclose which AI tools they used and how they were applied in court submissions.

China

China has pursued the most comprehensive judicial AI deployment of any country. The "Smart Justice" initiative, formally launched in 2016, mandates AI implementation across all

Chinese courts. The Supreme People's Court issued binding opinions requiring all courts to implement AI systems by 2025, with a long-term goal of full integration by 2030.[49]

In practice, this has produced a national judicial AI platform containing 320 million pieces of legal information, including court rulings, case files, and legal opinions. The Hangzhou Internet Court, established in 2017 as the country's first court dedicated to digital disputes, operates an entirely online litigation platform with AI capabilities including virtual judges that guide parties through procedures, automatic speech-to-text transcription, case searching, sentencing reference tools, and automated generation of litigation documents.[50] Over 3.1 million legal activities have been completed through this system.

The Chinese model is instructive as a point of comparison because it illustrates what happens when AI adoption is mandated from the top down rather than emerging through individual judicial initiative. The scale is unmatched, but the approach raises questions about transparency, due process, and the degree of human oversight that Western legal systems would consider essential.

Brazil

Brazil has launched over 140 judicial AI projects since 2019, making it one of the most active countries in this space.[51] The projects range from tools that find precedents and categorize cases to systems that draft judicial decisions. SIGMA assists judges by analyzing stored texts and comparing them with procedural documents. Projeto Socrates aims to reduce appellate judgment timelines by 25 percent through automated draft decisions based on prior rulings. APOIA, a generative AI assistant integrating ChatGPT and Gemini, drafts reports, summarizes case files, and identifies applicable law.

Brazil's experience also illustrates the risks. In 2024, a judge signed a draft decision prepared by a court clerk using ChatGPT without being informed of the AI's involvement. The decision contained fabricated facts and fictitious jurisprudence, leading to a wrongful conviction. The incident prompted mandatory practical training courses on AI risks for judicial staff across the country.

France

France's Court of Cassation released a report in April 2025 titled "Preparing the Court of Tomorrow," outlining AI priorities for the judicial branch. The report, prepared by a working group of magistrates, clerks, researchers, and AI experts, recommended the creation of an internal supervisory committee for ethical monitoring, a set of operational guidelines, and an ethical charter specific to the court.[52] The Paris Commercial Court has begun using AI to generate pre-reports that summarize case details from anonymized data, with an expected reduction in preparation time by a factor of six.

India and Other Jurisdictions

India's e-Courts Project Phase III, jointly sponsored by the Supreme Court and the Ministry of Law and Justice, allocates dedicated funding for AI and blockchain integration across the court system. The project includes a Legal Research Analysis Assistant (LegRAA) developed for judges, AI-driven court scheduling optimization, and chatbots that provide litigants with real-time case updates and procedural guidance.[53] Estonia has implemented a semi-automated payment order system for small claims under EUR 8,000, in which algorithms prepare orders and human judges determine jurisdiction and ensure proper service of documents. Australia's federal and state courts have published judicial guidance on AI use, and Queensland has issued specific guidelines for self-represented litigants on the limitations of AI chatbots.

11. Global and National Policy Frameworks

Alongside the accumulation of incidents, sanctions, and ad hoc orders, a parallel effort has taken shape: the construction of institutional guidance intended to replace improvisation with coherent, system-wide standards. Practitioners should be aware of these policy documents, because they increasingly shape the rules that courts apply.

UNESCO: Fifteen Principles for the Global Judiciary

On December 4, 2025, UNESCO formally launched its Guidelines for the Use of AI in Courts and Tribunals at the Athens Roundtable on AI and the Rule of Law, held in London. The document represents the first global ethical and operational set of principles for judicial AI governance, organized around fifteen principles including transparency, accountability, human oversight, and the protection of human rights.[54]

UNESCO accompanied the guidelines with survey data drawn from judicial operators across 86 countries. Only 9% had received any form of AI-related training or guidance; 44% reported having already used AI tools in their work. Seventy-three percent stated that mandatory regulations for judicial AI use should be established. Read together, those numbers describe an institution absorbing a technology far faster than it is developing the competence to evaluate what that technology produces.

The ABA and the Sedona Conference

In the United States, the ABA Task Force on Law and Artificial Intelligence published guidelines through its Working Group on AI and the Courts, drafted by a panel of five sitting judges and a professor of law and computer science.[55] The guidelines identify two cognitive distortions that pose particular danger in the judicial context.

The first is automation bias: the tendency to defer to AI-generated outputs without independent verification, a risk compounded by the confident, authoritative tone in which large language models present their answers. The second is confirmation bias: the tendency to accept AI results that align with a pre-existing view and to discount or ignore outputs that do not. Both risks are amplified in judicial settings, where the decision-maker is often working alone and under time pressure.

The guidelines also address the ex parte dimension. When a judge obtains information, analysis, or advice from an AI tool, the guidelines note, the judge risks "relying on extrajudicial information and influences that the parties have not had an opportunity to address or rebut." The

concern is not hypothetical. It maps directly onto Model Code of Judicial Conduct Rule 2.9(A), which prohibits ex parte communications, and Rule 2.9(C), which bars independent investigation of adjudicative facts.

New York's Institutional Approach

On October 10, 2025, the New York Unified Court System released its first official AI policy. The Interim Policy on the Use of Artificial Intelligence applies to all judicial and non-judicial personnel and restricts AI use to a list of pre-approved tools, initially including Microsoft Copilot and OpenAI's ChatGPT.[56] The policy states that AI "must not be treated as a substitute for human judgment, discretion, or decision-making."

New York's approach, which ties permissible use to a curated list of vetted platforms rather than relying on disclosure requirements or general prohibitions, offers a different governance model. It requires institutional investment in evaluating tools before they enter chambers, but in exchange it provides clearer boundaries than a regime that asks each judge or clerk to make independent assessments about which platforms are safe to use.

The European Ethical Charter

The European Commission for the Efficiency of Justice (CEPEJ) adopted an Ethical Charter on the Use of Artificial Intelligence in Judicial Systems in December 2018, predating most of the American judicial orders by five years. The charter establishes principles of respect for fundamental rights, non-discrimination, quality and security, transparency, and user control.[57] While no member state has yet deployed AI at scale in its courts, the charter provides the normative baseline against which European judicial AI deployments will be evaluated. The EU AI Act, adopted in 2024, classifies AI systems used in the administration of justice as "high-risk," subjecting them to mandatory conformity assessments, transparency requirements, and human oversight obligations.

12. Ethical Fault Lines for the Bench

The judicial code of conduct predates AI by decades, but its provisions map onto the problems AI creates with more precision than might be expected. Practitioners should understand these ethical constraints, because they affect what judges can and cannot do with AI tools, and that in turn shapes what attorneys can expect from the bench.

Ex Parte Communication and Independent Investigation

Model Code of Judicial Conduct Rule 2.9(A) prohibits a judge from considering "communications made to the judge outside the presence of the parties or their lawyers" concerning a pending matter. Rule 2.9(C) bars judges from conducting independent investigation of adjudicative facts. When a judge prompts an AI tool with a question about a case, the response constitutes information that originates outside the adversarial process. The parties have had no opportunity to examine, challenge, or rebut it.[58]

The NCSC has noted that AI-generated material "could arguably be viewed as information outside the case that is improperly introduced into the judicial decision-making process."[58] The concern is sharpened by the nature of large language models. Their training data encompasses vast quantities of text, potentially including information about the parties, the legal issues, or related cases that is not part of the record. A judge who consults an AI tool may, without realizing it, be introducing extrajudicial facts into deliberations.

Michigan Judicial Ethics Opinion JI-155, issued in October 2023, took the further step of holding that judges must not only understand the ethical constraints on their own AI use but must also "continually evaluate how they or parties before them are using AI technology tools in their own docket."[59]

Confidentiality of Judicial Deliberations

Judges and their staff handle some of the most sensitive information in the legal system: draft opinions, sealed filings, in camera submissions, jury deliberation notes, and internal case discussions. Consumer AI platforms, by the terms of their own privacy policies, may retain, analyze, and use inputs for model training. Uploading a draft opinion to such a platform is functionally equivalent to sharing it with an unknown third party whose data practices are governed by a terms-of-service agreement, not by judicial ethics rules.

Every major set of guidelines, from the ABA to Judge Schlegel's AI in Chambers document to New York's interim policy, converges on the same remedy: restrict judicial AI use to

enterprise-grade tools operating under contractual data isolation, with audit trails and explicit prohibitions on using judicial inputs for model training.

The Disclosure Question

Whether judges should disclose their use of AI tools in preparing opinions is an open question with strong positions on each side. Proponents of disclosure argue that if AI shaped the analysis in an opinion, the parties and reviewing courts are entitled to know. Opponents counter that judges already use Westlaw, law review articles, treatises, and conversations with clerks without itemizing them, and that requiring disclosure for AI alone draws an unprincipled line.

A workable distinction may track the nature of the AI's contribution. Formatting citations or catching typographical errors occupies a different category than suggesting analytical structures or surfacing arguments the judge had not considered. No court has adopted a formal disclosure rule, but the ABA Working Group has flagged the issue, and several state judicial conferences have placed it on their agendas for 2026.

13. What Comes Next

If 2025 was the year AI ceased to be a hypothetical problem for the judiciary, 2026 is the year it becomes an operational one. For practitioners, the following developments are worth tracking.

Institutional consolidation. Individual standing orders are being superseded by district-wide local rules, circuit-level policies, and statewide directives. New York's interim policy, with its curated list of approved tools, represents one governance model. The Fifth Circuit's reliance on existing rules of practice represents another. The competition between these models will shape how the federal judiciary manages AI for the next decade.

Agents in chambers. Judge Schlegel's prediction that multi-step AI agents will enter judicial chambers is already beginning to materialize. The distinction that will matter is between agents that perform mechanical tasks, such as citation verification or transcript indexing, and agents that engage with substance, such as evaluating the strength of arguments or proposing analytical structures. The former is a tool. The latter is a collaborator. Courts will need to draw that line clearly.

Deepfakes and the authentication crisis. The ability of AI to generate realistic images, audio, and video at low cost is creating an evidentiary challenge for which existing rules of evidence provide limited guidance. Federal Rules of Evidence 901 and 902, which govern authentication, were drafted for a world in which fabricating convincing documentary evidence was expensive and difficult. That world no longer exists. Louisiana's Act 250 represents an early legislative response, but the broader problem will require federal attention.

The training deficit. UNESCO's survey found that 44% of judicial operators have used AI tools while only 9% have received any training. The NCSC AI Sandbox and the growing number of state-level training programs are early responses, but the gap between adoption and competence remains wide. Closing it will require sustained investment in continuing judicial education at the state and federal levels.

Privilege law after *Heppner*. Judge Rakoff's opinion will be tested by subsequent cases, and the distinction between consumer AI tools and enterprise platforms used under counsel's direction will likely become a routine element of privilege analysis in white-collar criminal defense and corporate investigations. Practitioners should begin advising clients now on the privilege implications of their AI use.

Global convergence and divergence. The UNESCO guidelines, the EU AI Act's classification of judicial AI as high-risk, Canada's judicial council guidelines, and China's mandatory adoption mandate represent four distinct governance philosophies. How these approaches interact will shape the international norms that eventually influence American judicial

policy.

The common thread across these developments is a single institutional question: can the judiciary adapt its rules and culture to a technology that evolves faster than any regulatory process can keep pace with, while preserving the deliberation, transparency, and human judgment that distinguish a court from a processing center? For lawyers and legal professionals who practice before these courts, the answer matters not as an abstraction but as a daily reality that shapes how cases are filed, argued, and decided.

This guide is part of the AI and Law Series by Colin Levy. Other titles include AI For Lawyers, AI for Legal Teams, AI Agents Data Handling and Cybersecurity Guide, and Law School in the Age of AI.

Endnotes

1. Judge Scott Schlegel, Louisiana Fifth Circuit Court of Appeal, has made this observation in multiple public settings. See, e.g., "Bot Overlords, Deepfakes, and the Weight of the Robe: Judge Scott Schlegel on AI in the Courts," 3 Geeks and a Law Blog (Jan. 2026).
2. *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023).
3. The six fictitious cases were *Varghese v. China Southern Airlines*, *Shaboon v. Egyptair*, *Petersen v. Iran Air*, *Martinez v. Delta Airlines*, *Estate of Durden v. KLM Royal Dutch Airlines*, and *Miller v. United Airlines*. See Sanctions Opinion, *Mata v. Avianca*, No. 22-cv-1461, Doc. 54 (S.D.N.Y. June 22, 2023).
4. Schwartz Aff., *Mata v. Avianca*, No. 22-cv-1461 (S.D.N.Y. June 22, 2023), Doc. 54. See also Nate Raymond, "Lawyer Apologizes for Fake Court Citations from ChatGPT," CNN Business (May 27, 2023).
5. Sanctions Opinion, *Mata v. Avianca*, Doc. 54 at 10-11.
6. *Id.* at 2. Judge Castel noted that "existing rules impose a gatekeeping role on attorneys to ensure the accuracy of their filings."
7. Letter from Judge Henry T. Wingate to Director, Administrative Office of the Courts (Oct. 2025), responding to Senate Judiciary Committee inquiry. See "Judge Wingate Admits Mistake in Use of AI in Drafting Order Blocking Mississippi's Anti-DEI Law," Magnolia Tribune (Oct. 23, 2025); "Two Federal Judges Admit to Using AI in Botched Rulings After Months of Silence," Washington Times (Oct. 23, 2025).
8. Letter from Judge Julien Xavier Neals to Sen. Charles Grassley (Oct. 2025). See "Federal Judge Withdraws Opinion After Lawyer Points Out Fake Quotes, Misstated Case Outcomes," ABA Journal (July 24, 2025); "2 Federal Judges Reveal AI Use by Staff Members Led to Error-Riddled Opinions," ABA Journal (Oct. 24, 2025).
9. Press Release, "Grassley Releases Judges' Responses Owning Up to AI Use, Calls for Continued Oversight and Regulation," U.S. Senate Committee on the Judiciary (Oct. 2025).
10. See "Courting Change: Treatment of AI-Generated Content by Australian Courts," LK Law (2024), which describes an ongoing database tracking over 150 instances of AI hallucination in legal decisions.
11. Standing Order, Mandatory Certification Regarding Generative Artificial Intelligence, Judge Brantley Starr, N.D. Tex. (May 30, 2023).
12. *Id.* See also "Federal Judge Seeks to Prevent Generative AI Mistakes in Briefs," Legal Dive (June 1, 2023).
13. Fifth Circuit Court of Appeals, Statement on Generative AI (June 2024). See "Is Disclosure and Certification of the Use of Generative AI Really Necessary?" Judicature, Duke Law (2024).
14. RAILS, "Resource: AI Orders," available at https://rails.legal/resources/resource-ai-orders/.
15. "Judge Scott Schlegel Talks ChatGPT, Transforming the Courtroom, and More with Everlaw," Everlaw Blog (2025). See also Louisiana Record, "24th Judicial District Court Judge Scott U. Schlegel Receives The William H. Rehnquist Award for Judicial Excellence" (2021).
16. National Center for State Courts, 26th Annual William H. Rehnquist Award for Judicial Excellence (2021). Judge Schlegel was the first Louisiana judge to receive the honor. See also Judge Schlegel's biography at https://judgeschlegel.com/bio.
17. ABA Working Group on AI and the Courts, "Guidelines for Judicial Officers: Responsible Use of Artificial Intelligence," Judges' Journal, Spring 2025.
18. *Id.* See also "Judge Scott U. Schlegel: A Call for Education Over Regulation (Artificial Intelligence)," Criminal Law Library Blog.
19. Judge Scott Schlegel, "The 2026 Prediction: From AI to Delegation," judgeschlegel.com (Jan. 2026).
20. Louisiana Act 250 (2024). See "AI-Generated Evidence Showing Up in Court Alarms Judges," NBC News (2025).

21. *United States v. Heppner*, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 10, 2026). See "SDNY Rules AI-Generated Documents Are Not Protected by Privilege," Debevoise Data Blog (Feb. 11, 2026).
22. *Id.* Written opinion issued Feb. 17, 2026. See "Update: Judge Rakoff Issues Written Opinion that AI-Generated Documents Are Not Protected by Privilege," Debevoise Data Blog (Feb. 17, 2026).
23. *Id.* See also "AI Privilege Waivers: SDNY Rules Against Privilege Protection for Consumer AI Outputs," Gibson Dunn (Feb. 2026).
24. *Id.*, citing *United States v. Kovel*, 296 F.2d 918 (2d Cir. 1961).
25. "Are Risk Assessment Tools Setting the Stage for AI Judges?" The Bail Project (2024). See also "AI's Complex Role in Criminal Law: Data, Discretion, and Due Process," ABA GP Solo Magazine (Mar./Apr. 2025).
26. Julia Angwin et al., "Machine Bias," ProPublica (May 23, 2016). The investigation analyzed COMPAS scores for over 7,000 defendants in Broward County, Florida.
27. Northpointe Inc., "COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity" (July 2016). See also Marcello Di Bello, "Algorithmic Fairness: ProPublica v. Northpointe," course materials, Lehman College/CUNY.
28. Jon Kleinberg, Sendhil Mullainathan & Manish Raghavan, "Inherent Trade-Offs in the Fair Determination of Risk Scores," Proceedings of Innovations in Theoretical Computer Science (ITCS 2017). See also Sam Corbett-Davies et al., "Algorithmic Decision Making and the Cost of Fairness," Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2017).
29. *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016), cert. denied, 137 S. Ct. 2290 (2017). See "State v. Loomis," 130 Harv. L. Rev. 1530 (2017).
30. Christopher Slobogin, quoted in Anne L. Washington, "How to Argue with an Algorithm: Lessons from the COMPAS-ProPublica Debate," 17 Colo. Tech. L.J. 131 (2019).
31. Idaho Code Ann. Sec. 19-1910 (2019). See "AI and the Assessment of Risk in Bail, Sentencing and Parole," Law Commission of Ontario (Apr. 2025).
32. "Public Perceptions of Judges' Use of AI Tools in Courtroom Decision-Making: An Examination of Legitimacy, Fairness, Trust, and Procedural Justice," PMC/National Library of Medicine (2025).
33. "New AI Research Tools Now Available: Lexis Protege & Westlaw CoCounsel," USF Zief Law Library Brief (Feb. 2025). Quick Check Judicial analyzes briefs and evaluates the validity of cited authorities.
34. "Involuntary Commitments, Artificial Intelligence & Expedited Appeals," Indiana Court Times (May 5, 2025). See also "Guardrails for the Future: How the Indiana Supreme Court is Preparing for the Age of AI," Indiana Court Times (Aug. 6, 2025).
35. "CaseMark Launches White-Label Platform for Court Reporting Firms," LawNext (Feb. 2026). See also Steno, "Transcript Genius: AI-Powered Analysis Tool" (2025).
36. National Center for State Courts, "AI Sandbox," available at https://www.ncsc.org/resources-courts/ai-sandbox.
37. "Tyler Technologies and Texas Office of Court Administration Announce Decade of Electronic Filing Success," Tyler Technologies Investor Relations (2024). See also "Tyler Buys AI-Backed Court Document Tech Provider CSI," GovTech (Aug. 2023).
38. "Chief Justice Creates Task Force on Use of Generative AI in the California Courts," California Courts Newsroom (May 2024).
39. See Illinois Supreme Court, "Policy on Artificial Intelligence" (2025); Connecticut Judicial Branch, "Responsible AI Framework" (2025); Maryland Judiciary, "Guidelines for the Acceptable Use of Artificial Intelligence Tools and Platforms" (2025); New Jersey Courts, "Statement of Principles for the Judiciary's Ongoing Use of AI" (2025); Utah Courts, "Interim Rules on the Use of Generative AI" (2024); Georgia Judicial Council, "Ad Hoc Committee on Artificial Intelligence and the Courts" (2024).
40. National Center for State Courts, "Navigating AI in Court Translation: Insights for Court Leaders" (2025), cautioning that AI performs poorly with fast speech, filler words, heavy accents, and legal terminology.
41. Administrative Office of the U.S. Courts, "Long Range Plan for Information Technology in the Federal Judiciary, FY 2026," available at https://www.uscourts.gov/.

42. "Access to Justice 2.0: How AI-Powered Software Can Bridge the Gap," ABA Journal (2025).

43. State Justice Institute, Self-Represented Litigation Priority Investment Area (2025). Funded programs in Nebraska, New Mexico, Wyoming, Kansas, and the 11th Judicial Circuit of Florida.

44. "Harnessing AI to Improve Access to Justice in Civil Courts," Stanford Human-Centered AI Institute (2025). The Stanford/LA Superior Court collaboration focuses on default judgments and triage for self-represented litigants.

45. "ABA Task Force: AI Has Moved From Experiment to Infrastructure for the Legal Profession," LawNext (Dec. 2025).

46. "The UK Justice System's AI Plan with Microsoft and OpenAI," AI Magazine (2024). See also "UK Government Unveils AI Strategy To Cut Court Backlogs," Law360 (2024).

47. UK Judicial Office, "Artificial Intelligence (AI) - Judicial Guidance" (Oct. 2025), available at https://www.judiciary.uk/guidance-and-resources/.

48. Canadian Judicial Council, "Guidelines for the Use of Artificial Intelligence in Canadian Courts" (Oct. 2024).

49. Supreme People's Court of China, "Opinions on Regulating and Strengthening the Application of Artificial Intelligence in Judicial Fields" (2022). See also "China's Smart Courts" overview materials.

50. "China Launches Artificial Intelligence Platform to Boost Judicial Efficiency," Supreme People's Court of China English Portal (Dec. 5, 2024). The Hangzhou Internet Court was established in 2017 for digital disputes.

51. "AI Is Helping Judges to Quickly Close Cases, and Lawyers to Quickly Open Them," Rest of World (2025).

52. "AI and Justice: The Court of Cassation Prepares for the Future," ActuIA (Apr. 2025). Report titled "Court of Cassation and Artificial Intelligence: Preparing the Court of Tomorrow."

53. "AI in Judicial Processes: Transforming India's Legal System," IndiaAI (2025). See also e-Courts Project Phase III budget allocation for AI and blockchain under "Future Technological Advancement."

54. UNESCO, "Guidelines for the Use of AI Systems in Courts and Tribunals" (Dec. 4, 2025), launched at the Athens Roundtable on AI and the Rule of Law, London.

55. ABA Working Group on AI and the Courts, "Guidelines for Judicial Officers: Responsible Use of Artificial Intelligence," Judges' Journal, Spring 2025. See also The Sedona Conference, "Navigating AI in the Judiciary" (Feb. 2025).

56. New York State Unified Court System, "Interim Policy on the Use of Artificial Intelligence" (Oct. 10, 2025), available at https://www.nycourts.gov/LegacyPDFS/a.i.-policy.pdf.

57. European Commission for the Efficiency of Justice (CEPEJ), "European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and Their Environment" (Dec. 2018). See also EU AI Act (Regulation 2024/1689), classifying AI in the administration of justice as high-risk.

58. National Center for State Courts, "AI & the Courts: Judicial and Legal Ethics Issues" (2025), available at https://www.ncsc.org/resources-courts/ai-courts-judicial-and-legal-ethics-issues. See also Model Code of Judicial Conduct Rules 2.9(A), 2.9(C).

59. Michigan Judicial Ethics Opinion JI-155 (Oct. 23, 2023). See "From the Bench: AI in the Courts," Illinois State Bar Association AI Section Newsletter (Jan. 2025).