Judicial Directory

JUDGES OF THE TRIAL COURTS

Hon. **Timothy S. Driscoll**

Supreme Court, Nassau County, Commercial Division
100 Supreme Court Drive
Mineola, NY 11501
(516) 493-3184

Judicial Offices

Associate Justice, Appellate Term for the 9th and 10th Judicial Districts, Appointed, 2021 to Present

Justice, Supreme Court, Nassau County, Commercial Division, Elected, 2022 to 2035

Justice, Supreme Court, Nassau County, Elected, 2007 to 2021

Other Professional Experience

Nassau County, Deputy County Executive, 2004 to 2007

United States Attorney's Office, Eastern District of New York, Assistant United States Attorney, 2000 to 2004

Nassau County District Attorney's Office, Assistant District Attorney, 1996 to 2000

Williams & Connolly, Associate, 1992 to 1996

Judge Joseph M. McLaughlin, United States Court of Appeals for the Second Circuit, Law Clerk, 1991 to 1992

Admission to the Bar

NYS, Appellate Division, Second Department, 1992

District of Columbia, 1993

Education

J.D., Harvard Law School, 1991

B.A., Hofstra University, 1988

Publications

Commercial Litigation in New York State Courts (Robert L. Haig, ed.), Chapter on "Motion Practice", 2020

Articles on NYS Commercial Division for ABA's "Business Law Today", 2014, 2015 and 2017

Professional & Civic Activities

Adjunct Professor, Nassau Community College, 2010 to Present

Teaching Team Member, Harvard Law School Trial Advocacy Workshop, 2003 to Present

Adjunct Professor, Brooklyn Law School, 1998 to Present

Past President, Irish Americans in Government of Nassau County

Past President, Catholic Lawyers Guild of Nassau County

Past President, Locust Valley Bayville Soccer Club

CLM COOPERMAN LESTER MILLER
GRUBER KRAUS LLP

ATTORNEYS-AT-LAW
1129 Northern Boulevard, Suite 402
Manhasset, New York 11030
(516) 365-1400
www.clmlaw.com

Eric H. Gruber, Esq.
Direct Dial: (516) 858-3480
egruber@clmlaw.com

New York City Office
767 Third Avenue
New York, NY 10017
(212) 688-7000

Mr. Gruber is a co-managing partner of Cooperman Lester Miller Gruber Kraus LLP and heads the firm's Litigation Department. Over almost 40 years, Mr. Gruber has developed a diversified commercial and corporate practice representing domestic and international public and privately-owned companies, entrepreneurs, family businesses and individuals over a range of practice areas including litigation, dispute resolution business transactions and corporate matters.

Mr. Gruber is an experienced litigator in all types of commercial and business issues and an accomplished trial attorney regularly practicing in state and federal courts, as well as in alternative dispute proceedings. Keenly aware of the stresses that litigation can cause a client, Mr. Gruber works with his clients to craft legal strategies for a broad range of civil cases and arbitrations and achieve results that others considered unachievable.

Mr. Gruber has represented plaintiffs and defendants in various matters across a spectrum of substantive areas, including commercial and contract disputes, business torts, corporate and partnership disputes, real estate, construction, creditors' rights including bankruptcy, and general business law in the federal and state trial and appellate courts and before various arbitration and mediation panels and bodies.

On the transactional side, Mr. Gruber has experience in asset-based lending and other areas of commercial finance, including bankruptcy matters, workouts and turn-around situations. He has actively represented clients in the documentation of commercial finance and equipment leasing transactions and negotiated and papered restructurings on behalf of and with lenders and servicers.

Mr. Gruber often serves as outside general counsel to his clients. In that capacity Mr. Gruber's clients regularly seek his guidance and counsel on mergers, acquisitions, sales, board advisory and corporate governance issues, risk management, business succession planning, joint ventures, executive compensation, strategic relationships, licensing, business restructuring and recovery, and almost every aspect of the client's business.

Mr. Gruber has guest lectured on litigation practice at Touro Law School and has lectured for various professional groups and organizations including the Nassau Academy of Law on points of legal procedure and litigation including provisional remedies, discovery, trial practice, business divorces, dissolutions of businesses, disputes between shareholders, and real estate issues.

Mr. Gruber is admitted in the United States Court of Appeals for the Second Circuit, the United States District Court for the Southern and Eastern Districts of New York, and in all of the Courts of the State of New York.

Mr. Gruber is an Executive Committee Member of the prestigious Theodore Roosevelt American Inn of Court, a Past President of B'nai B'rith's Banking and Finance Unit, and a member of the New York State and Nassau County Bar Associations.

Education

Syracuse University College of Law,
J.D., 1986
    Honors
    Syracuse Law Review Editor;
    Exceptional Editor's Award

State University of New York at Albany,
B.A. 1983

Thomas A. O'Rourke
O'Rourke IP Law P.L.L.C.
425 Broadhollow Rd.
Melville, N. Y. 11747
631-423-2700
TORourke@ORourkeIPLaw.com

Mr. O’Rourke’s practice involves all areas of patent, trademark and copyright law. For over thirty years he has been registered to practice before the United States Patent & Trademark Office. Mr. O’Rourke has counseled clients regarding the procurement and enforcement of patents, trademarks, copyrights and trade secrets in a variety of technologies including mechanical, and computer technology. In addition, his practice involves domestic and international technology transfer, acquisition and licensing. He is a member of the bar of the States of New York and California. He has also been admitted to numerous Federal District Courts and Courts of Appeal across the country including, the Court of Appeals for the Federal Circuit.

Mr. O’Rourke has been a member of the Board of Directors of the New York Intellectual Property Law Association. Mr. O’Rourke has been Chairman of the Suffolk County Bar Association’s Committee on Intellectual Property Law and has been a member of the Advisory Board of the Licensing Journal. He has lectured on Intellectual Property Law at numerous Continuing Legal Education programs, including programs presented by the American Bar Association, the Connecticut Intellectual Property Law Association and the Suffolk County Bar Association. He was also the Editor of the New York Intellectual Property Law Association Bulletin and the author of numerous articles on patents, trademarks and copyrights for the New York Intellectual Property Law Association. Mr. O’Rourke has also authored monthly articles on intellectual property law licensing, which have appeared in the Licensing Journal. Mr. O’Rourke has also been named as a Super Lawyer for many years.

Mr. O’Rourke has a B.S. degree in Chemistry from Fordham University and obtained his J.D. degree from St. John's University School of Law, where he was a member of the Law Review.

Email: mberger@forchellilaw.com
Phone: (516) 248-1700
Fax: (516) 248-1729

MICHAEL A. BERGER

Associate

Michael A. Berger is an associate in the Firm's Employment & Labor and Veterinary practice groups. He concentrates his practice on counseling and defending employers on various employment and labor law issues, including wage and hour, discrimination and retaliation.

In the Veterinary practice group, he represents both veterinary consolidators and individual practitioners in employment related matters, such as drafting Employee Handbooks, employment policies, and negotiating employment and severance agreements on behalf of both veterinarians and executives. Additionally, he counsels veterinarians on numerous compliance and regulatory issues, including the specific laws of states throughout the country.

Mr. Berger is admitted in New York, New Jersey and the U.S. District Courts for the Southern, Eastern and Western Districts of New York.

Immediately prior to joining our Firm, Mr. Berger was an associate at a Long Island-based labor law firm. Prior to that, he was an associate at a New York City firm.

Mr. Berger served as a legal fellow to the Hon. Sandra L. Sgroi, Appellate Division, Second Department; a volunteer at Nassau/Suffolk Law Services: Volunteer Lawyers Project; a legal intern at the Law Reform Advocacy Clinic at his law school and at the New York State Office of the Attorney General; and as a law clerk to the Hon. Joseph A. Zayas, Supreme Court of the State of New York, Criminal Term.

Mr. Berger earned his J.D. from the Maurice A. Deane School of Law at Hofstra University, where he was a Book Review Editor for the *Journal of International Business and Law*. In 2013, he published a Note and Book Review. Mr. Berger received his B.A. from the University of Pittsburgh, College of Arts & Sciences.

PRACTICE AREAS

- Employment & Labor
- Litigation
- Veterinary

EDUCATION

- Maurice A. Deane School of Law, Hofstra University, J.D., 2013
- University of Pittsburgh, Kenneth P. Dietrich School of Arts & Sciences, B.A., 2010

ADMISSIONS

- New York State Bar
- New Jersey State Bar
- United States District Courts for the Southern, Eastern and Western Districts of New York

PROFESSIONAL AFFILIATIONS AND ACCOMPLISHMENTS

- Member, The Nassau Lawyers Association of Long Island, Inc.
- Treasurer, The Theodore Roosevelt American Inn of Court
- Nassau County Bar Association
    - Co-Chair, New Lawyers Committee

Adriana Montante, Esq.
amontante@nycourts.gov

Adriana Montante currently serves as the Assistant Law Clerk to the Honorable Sarika Kapoor in Nassau County Supreme Court. In this role, she performs in-depth legal research, drafts judicial opinions, and conducts court conferences.

Ms. Montante is admitted to practice in the State of New York. She is a member of the Theodore Roosevelt American Inn of Court.

Ms. Montante earned her J.D. in 2024 from the Maurice A. Deane School of Law at Hofstra University, where she distinguished herself as the Managing Articles Editor of the *Journal of International Business & Law* and President of the Energy and Environmental Law Society. Her scholarship includes a published note entitled *"Think Globally, Act Locally:" New York's Evolving Approach to Address Fossil Fuel Electric Use in Proof-of-Work Bitcoin Mining Operations Contributing to Climate Change*.

Throughout law school, Ms. Montante gained extensive practical experience across a wide range of legal fields. As an intern at the Robert W. Entenmann Veterans Law Clinic, she collaborated with attorneys and clinicians to help veterans secure service-connected disability benefits and appeal denied claims. Her commitment to public service is also reflected in her work at the Pro Se Legal Assistance Program, where she guided self-represented litigants through the federal civil court process.

Ms. Montante's legal foundation is strengthened by prior experience with the Environmental Protection Agency Region 2, where she worked on matters involving hazardous waste, toxic substances, and statutory interpretation, as well as internships in personal injury, employment discrimination, and general civil practice.

She holds a B.A. in Environmental Studies from Binghamton University, where she also minored in Economics. While at Binghamton, Ms. Montante led a sustainability-focused student organization and studied abroad in Paris.

Touro University Jacob D. Fuchsberg Law Center, J.D., 2023

Honors

Salutatorian, Summa Cum Laude

Touro Law Review Editor

Bess and Sam Zeigen Scholar

University of Hartford, B.S.B.A., 2019

Download vCard
Print This Page

Michael is a CLMGK associate and practices in the Firm's litigation and business transactions departments.

Michael graduated *summa cum laude* as the full-time division *salutatorian* from Touro Law Center in May 2023. Upon graduation, Michael was awarded Exemplary Dedication and Service to the Law Center for Exceptional Contributions to the Touro Law Review award, as well as the Shirly & Murray Rubinstein Clinician of the Year Award for his service to the Bankruptcy/Foreclosure Law Clinic.

While a law student, Michael was a member of the Touro Law Review where he served as Senior Research Editor on the Editorial Board. During law school, Michael also held a judicial externship with the Honorable James M. Wicks in the United States District Court for the Eastern District of New York. In the classroom, Michael received eleven CALI Awards for receiving the highest grade in his classes, including, but not limited to, Evidence, Trusts & Estates, and Property.

Michael is admitted in New York and is a member of the New York State Bar Association.

Howard Goldfarb

Juris Doctor Candidate, May 2027

Touro University Jacob D. Fuchsberg Law Center

Howard Goldfarb is an accomplished commercial real estate executive with over 30 years of experience in property disposition and lease restructuring. He is currently pursuing a Juris Doctor degree at Touro University Jacob D. Fuchsberg Law Center. Howard has an extensive background in contract negotiation and complex deal structuring and is now focused on applying his proven business skills to the practice of law.

Howard is currently ranked first in his class and serves as a Senior Research Editor for the Touro Law Review. He is also a Touro Honors Program Scholar and recipient of both Merit and Achievement Scholarships. Howard is also a student member of the prestigious Theodore Roosevelt American Inn of Court.

Howard earned a Master of Science in Real Estate Development and Investment from New York University, graduating first in his class with distinction. He also holds a Bachelor of Arts in Computer Science from Cornell University.

For over 30 years, Howard has led the Real Estate Disposition Department at Excess Space Retail Services, Inc., a Newmark Company. He has initiated, negotiated, and completed more than 1,000 commercial real estate transactions—including leases, subleases, sales, and restructurings—resulting in over $250 million in savings for many of the nation's leading retailers. Under his leadership, his team has managed transactions exceeding $1 billion in value. Howard also developed many of the firm's analytical systems and financial models used to evaluate the economic impact of each real estate transaction, streamlining client decision-making. His strategic leadership helped position the company for its sale to a global public corporation in 2015.

While attending law school part-time, and as a member of Touro's Not-for-Profit Clinic, Howard has conducted research and provided counsel to a 501(c)(10) organization. Among the services he has provided is the structuring and formation of 501(c)(3) entities for the purposes of repairing and maintaining the building owned by the organization and operating a museum.

Howard's career in commercial real estate was preceded by several years as a computer programmer and systems analyst, developing business applications for various large corporations.

AI and the Law: The Use of Technology and Its Impact Upon Legal Ethics

A New York CLE Program (2 Hours: 1 Credit Ethics, 1 Credit Technology)

PROGRAM OVERVIEW

Faculty:

- **Tom** - Seasoned solo IP practitioner
- **Eric** - Managing partner, 12-person firm (commercial litigation)
- **Michael** - Mid-level associate, midsize firm (corporate/commercial and labor and employment)
- **Adriana** - Assistant law clerk to NY Supreme Court Justice
- **MAB** - Junior associate, small firm (litigation & transactional)
- **Howard** - Law student, former computer programmer/software engineer

Judicial Commentator:

- **Justice D** - NY Supreme Court Justice, Commercial Division

Format: Bar Association Special Committee Meeting on "AI and the Law: Developing Practice Guidance for the Bar"

Total Runtime: 120 minutes

- Opening: 5 minutes
- Act I: 25 minutes (AI Basics & The Hallucination Crisis)
- Justice D Commentary: 5 minutes
- Act II: 30 minutes (Ethics Framework & Practical Workflows)
- Justice D Commentary: 5 minutes
- Act III: 25 minutes (Discovery, Evidence & Economic Access)
- Justice D Commentary: 5 minutes
- Act IV: 15 minutes (The Road Forward: Governance & Best Practices)
- Justice D Final Commentary: 5 minutes

- Audience Q&A: 10 minutes

SCENE: Opening

[A conference room. A long table with laptops, documents, coffee cups. Name placards for each character. A screen shows: "Bar Association Special Committee: AI Practice Guidance Development"]

ERIC: *(calling the meeting to order)* Alright everyone, let's get started. We're here for our committee work to develop practical guidance for New York lawyers on the use of AI. *(looking around)* We've got a great cross-section here - solo practitioners, small firms, big firms, the judiciary. And our secret weapon... *(gestures to Howard)* ...a law student who actually understands how this stuff works under the hood.

HOWARD: *(grinning)* I'm currently a 3L in Touro Law's "Flex" program with a long ago undergraduate degree in Computer Science from Cornell. Happy to translate from "tech-speak" to "legal-speak."

TOM: Thank God. Because I'll be honest, six months ago I thought "ChatGPT" was some kind of sandwich. *(light laughter)* But now my clients are asking me about AI, copyright infringement and fair use, and I'm realizing I need to get up to speed fast if I want to compete with the big firms.

MAB: *(enthusiastically)* I use AI every day! It's amazing - it drafted a memo of law for me in like five minutes that would've taken me three hours.

MICHAEL: *(concerned)* And did you verify every single citation in that memo of law before you filed it?

MAB: *(pause, slightly deflating)* ...I spot-checked a few...

ADRIANA: *(firmly)* Which is exactly why my judge asked me to participate in this committee. We're seeing a tsunami of AI-related issues in chambers - everything from hallucinated cases to arguments relying on cases that do not stand for the proposition asserted and questions about when judges themselves can ethically use AI tools.

ERIC: Which brings us to today's agenda. We need to cover: one, what AI actually *is* and why it hallucinates; two, the ethical framework under New York rules; three, practical workflows for different practice areas; and four, the access to justice implications. Let's dive in.

ACT I: Understanding AI and The Hallucination Crisis (25 minutes)

ERIC: Howard, you're up. Explain to us - in terms even Tom can understand - what we're dealing with when we talk about "AI" in legal practice.

HOWARD: Okay, so first distinction: there's "AI" broadly speaking, which has been around for decades in legal tech - things like predictive coding for e-discovery, which we'll talk about later. But what's got everyone's attention now is **Generative AI** - specifically something called Large Language Models, or LLMs.

TOM: The stuff that writes things.

HOWARD: Exactly. ChatGPT, Google's Gemini, Anthropic's Claude, Microsoft's Copilot - they're all LLMs. Here's the key thing everyone needs to understand: **LLMs are sophisticated pattern-matching systems that predict the most statistically likely next word in a sequence.**

MAB: Wait, they don't actually "think"?

HOWARD: No. They're incredibly good at mimicking human writing because they've been trained on massive amounts of text from the internet. That includes samples that are factually accurate and well-written, as well as inaccurate, poorly drafted samples. They have no concept of "truth" or "accuracy." They're just predicting: "Based on all the patterns I've seen, what word probably comes next?"

MICHAEL: Which is why they hallucinate.

HOWARD: Right. *"HALLUCINATION" is* the term used to describe when an LLM generates a fake case citation, it's not lying or malfunctioning - it's doing exactly what it was built to do: generating text that *looks* like a real legal citation based on the patterns it's learned.

ADRIANA: And that text is often extremely convincing. I've seen AI-generated fake cases with proper Bluebook citations, realistic case names, even plausible-sounding holding language.

TOM: Okay, but there must be legal-specific AI tools that don't have this problem, right? I keep getting sales pitches about them.

HOWARD: Good question. So there are basically three categories of AI tools lawyers encounter: *(counts on fingers)*

One: Public, consumer-grade LLMs like free ChatGPT or Goggle's Gemini. These are trained on internet data up to a cutoff date - ChatGPT's built in training only goes up to June 2024. Beyond that it uses real-time browsing to get current information. Gemini's "knowledge" is based on the data it was trained on from January 2025. This means that for events or information that occurred before that date, it can rely on internal knowledge. It is also connected to Google Search, so it too can access and process real-time information from the internet, recent news, and events that happened after the training cutoff.

But, before you say then it can access Google Scholar including the case law and academic articles in that database even Gemini tells you that the database is for informational purposes only and is not considered a replacement for professional grade tools like Westlaw or LEXIS-NEXIS.

They have no real-time access to legal databases, no citation checking, and *they may use your inputs to train future versions of the model.*

ERIC: *(alarmed)* Meaning if I put client information in...

HOWARD: It could theoretically end up in the training data. Which, among other things, is a massive confidentiality problem.

Two: Enterprise or "sandboxed" versions of these same LLMs - where you pay for a version that promises your data won't be used for training and may have better security. But the core technology is the same, so they still hallucinate.

And three: Dedicated legal AI platforms - tools like Westlaw's AI research assistant, Lexis+ AI, CoCounsel, Harvey. These integrate LLMs with actual legal databases and add verification layers. But even these aren't perfect.

MICHAEL: I was on a CLE last month where they said some of the big firms are spending millions of dollars on legal AI tools. How is a small firm or solo attorney supposed to compete with that?

TOM: *(raising hand)* Yeah, that's my question. I've got a client who came to me worried about whether an AI company violated their copyright by training on their published work. I need to research cutting-edge IP issues, draft sophisticated arguments, and I'm competing against firms with ten associates and unlimited Westlaw budgets.

ERIC: We'll get to the economic disparity issue (https://www.lawfuel.com/the-ai-law-list-10-law-firms-leading-the-legal-ai-revolution/). But first, let's talk about what's happened when lawyers *have* used AI - and gotten it wrong.

ADRIANA: *(pulling out a folder)* Oh, that has happened - a lot. These are just from 2025. We're up to over 500 documented cases worldwide where courts have dealt with AI hallucinations.

MAB: *(reading)* "Mata v. Avianca" - that's the famous one, right?

ADRIANA: The first big one, from 2023. Southern District of New York. Lawyer used ChatGPT to research, cited six completely fake cases then, when the judge questioned them, the lawyer *doubled down* and submitted fake judicial opinions to prove the cases existed.

MICHAEL: *(wincing)* What happened?

ADRIANA: In concluding that counsel violated Rule 11, the Court issued the following sanctions: Five thousand dollar fine, had to notify all the judges whose names were in the fake cases, had to notify their client, public reprimand. And that was just the beginning.

HOWARD: *(pulling up webpage on laptop)* There's a database now - the "AI Hallucination Cases" database maintained by Damien Charlotin (https://www.damiencharlotin.com/hallucinations/) which tracks legal decisions in cases where generative AI produced hallucinated content including fake citations and other types of AI-generated arguments. As of November 19, 2025, there are 569 cases tracked.

ADRIANA: And it's not just small practitioners making mistakes. *(flipping through notes)* In May 2025 (*Jacquelyn "Jackie" Lacey, et al. v State Farm General Insurance Co.*, No. cv-24-5205-fmo (C.D. Cal. May 5, 2025). C.D.-Cal.-24-cv-05205-dckt-000119_000-filed-2025-05-06.pdf) – a Special Master in California sanctioned lawyers from K&L Gates - that's a large international firm - for failure to disclose the use of AI and failure to cite check briefs that contained hallucinated cases. They had used Dedicated legal AI platforms CoCounsel, Westlaw Precision, *and* Google Gemini, and still ended up with fake cases. Sanction: briefs were stuck, denial of relief sought, and $31,100 in fees to the other side.

MAB: *(quietly)* Oh no...

ADRIANA: July 2025 - Colorado federal judge fined two lawyers $3,000 each for AI-generated fake cases in a defamation case. (gov.uscourts.cod.215068.383.0.pdf)

February 2025 – in Wyoming, lawyers from Morgan & Morgan – one of the *largest personal injury firms in the country* - were sanctioned for using their in-house AI platform that hallucinated cases.(WADSWORTH v. WALMART INC LLC (2025) | FindLaw) The court sanctioned the drafting attorney by revoking his pro hac vice status and ordering him to pay a $3,000.00 fine. The court imposed a less severe sanction on the non-drafting attorney including the imposition of a $1,000 fine and decided against revoking his pro hac vice status. The court also sanctioned the other counsel involved in the matter $1,000.

MICHAEL: So even purpose-built legal AI tools from major vendors are hallucinating?

HOWARD: Exactly. Because at their core, they're still LLMs. The legal vendors add verification *layers* - they try to ground the AI's responses in real case law from their databases - but it's not foolproof. The AI can still generate text that *looks* like a citation but isn't actually in the database.

ERIC: *(leaning back)* So we're telling lawyers they should use AI for efficiency and competitiveness, but also that every tool on the market can potentially produce fake citations that will get them sanctioned.

TOM: Well, that's reassuring. *(sarcastic)*

ADRIANA: Which is why the sanctions cases have been so instructive. Want to know what separates the lawyers who got fined $5,000 from the ones who got off with a warning?

ALL: *(leaning in)* What?

ADRIANA: How they responded when the hallucination was discovered. The lawyers in *Mata* kept insisting the cases were real even after the judge questioned them. But, in a case this August - *Hall v. Academy Charter School*, Eastern District of New York - lawyer cited three hallucinated cases, identifying her clerk had used "Google for research," but took responsibility explaining she'd been dealing with her spouse's sudden death and wasn't checking work carefully. Counsel explained that she had since taken bereavement leave and sought treatment for her grief.

MICHAEL: What happened?

ADRIANA: The judge, Magistrate Wicks (a member of our organization), wrote a 12-page opinion explaining Federal Rule 11 and sanctions, but ultimately declined to impose monetary sanctions because there was no finding of " bad faith." Judge

Wicks said it was "extreme carelessness" “under personal tragic circumstances,” but not intentional misconduct.

MAB: So, if you're honest about what happened, you might avoid the worst consequences?

ADRIANA: Honesty helps. But you know what helps more?

MICHAEL: Not filing hallucinated cases in the first place.

ADRIANA: Exactly.

ERIC: Okay, let me get this straight before we move on:

AI - specifically LLMs - predicts text, doesn't verify truth.

All AI tools can hallucinate - consumer ChatGPT, enterprise ChatGPT, even dedicated legal AI.

Courts are routinely sanctioning lawyers for filing fake citations.

The key is whether you verified and how you respond when caught.

(Looks around) Everyone agree on that foundation?

(Everyone nods)

ERIC: Good. Then let's talk about what the New York ethics rules require when you use these tools.

[LIGHTS DIM - JUSTICE D APPEARS AT BENCH]

JUSTICE D - FIRST COMMENTARY *(5 minutes)*

(Justice D, speaks directly to the audience)

JUSTICE D: Thank you. What you've just heard is the critical foundation. Let me emphasize several points from the bench.

First, understand the technology. You don't need a computer science degree, but you must understand that these tools are *probabilistic text generators*, not truth engines. This isn't optional knowledge - it's part of your competency obligation under Rule 1.1.

Second, as the committee discussed, there's a spectrum of AI tools available to you:

- Public ChatGPT: Free or $20/month, but your inputs may train the model. *Never* put client confidential information into public AI tools. This violates Rule 1.6.
- Enterprise/Sandboxed versions: You pay for data protection, but the core LLM still hallucinates.
- Dedicated legal AI platforms: Tools like Westlaw AI, Lexis+, CoCounsel that integrate with legal databases and add verification. These reduce but don't eliminate hallucination risk.

Third, the case law is now clear and consistent across jurisdictions:

The initial AI hallucination is treated as a tool failure. But your *failure to verify* is professional misconduct. In *Mata*, *Park v. Kim*, *Wadsworth v. Walmart*, and dozens of others, courts have said: at minimum, you must *read and verify the existence* of any case you cite. This is Federal Rule 11's "reasonable inquiry" requirement under the Federal Rules of Civil Procedure. New York also has a similar requirement under §§ 130-1.1a.

Fourth - and this is new - we're starting to see courts question whether opposing counsel has a duty to detect and *report* AI hallucinations in their opponent's briefs. In *Noland v. Land of the Free* from the California Court of Appeals in September 2025, the court sanctioned a lawyer $10,000 (payable to the Court) for a frivolous appeal based upon papers with fake cases but *denied* sanctions be payable to opposing counsel who failed to identify the hallucinations that were identified by the court on upon review. That's a developing area to watch.

Practice Point #1: Before you cite any case generated or identified by AI, you must pull up the actual opinion in Westlaw, Lexis, Google Scholar or the court's reporter system and confirm: (a) the case exists, (b) the citation is correct, (c) the case actually says what you claim it says, and (d) the case is still good law. Run a Shepard's or KeyCite. This is not optional.

Practice Point #2: Document your verification process. Keep a research log showing how you verified AI-suggested authorities. If questioned, you can show your reasonable inquiry.

Practice Point #3: Many federal judges - and now some New York courts - are issuing standing orders requiring certification that AI-assisted filings have been human-verified. Check local rules and individual judge's orders before you file.

(Justice D returns to observing)

ACT II: The Ethical Framework & Practical Workflows (30 minutes)

ERIC: *(reconvening)* Alright, so we now know what AI is and what goes wrong. Now let's talk about what New York's ethics rules require. Michael, I know you've been deep in the opinions.

MICHAEL: Yeah, I've read several of them. The good news is: we don't need new ethics rules. The bad news is: all our *existing* rules now apply with heightened scrutiny.

The key guidance comes from three sources:

NYC Bar Formal Opinion 2024-5 - issued August 2024 (https://www.nycbar.org/reports/formal-opinion-2024-5-generative-ai-in-the-practice-of-law/)

- ABA Formal Opinion 512 - issued July 2024 (https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf)
- NYSBA Task Force on AI Report - April 2024 (https://nysba.org/wp-content/uploads/2022/03/2024-April-Report-and-Recommendations-of-the-Task-Force-on-Artificial-Intelligence.pdf?srsltid=AfmBOozFdo-COz0QihbeMUiEMd3wdO8_uZhJDY2dNpEE2StbrNQwcVk)

And they all say essentially the same thing: the New York Rules of Professional Conduct are robust enough to govern AI use, *if* lawyers apply them rigorously.

TOM: Which rules are we talking about?

MICHAEL: *(pulling up slide)* Let me walk through the main ones:

Rule 1.1 - Competence and Diligence. This is the foundation. The NYC Bar opinion says you have a duty to understand the "benefits and risks" of AI tools. That means:

- Understand to a reasonable degree how the technology works, its limitations, and the applicable Terms of Use
- Understanding that LLMs hallucinate
- Knowing the difference between public and enterprise AI

- Understanding what verification steps are required (analyze for accuracy and bias)
- Staying current as the technology evolves

If you use AI without understanding these risks, you may have violated Rule 1.1 *before you even get to the hallucination.*

MAB: Wait, so do I have to take a CLE on AI?

MICHAEL: Not specifically AI but you must take a CLE on technology. But you do need 4 CLE credits on ethics and part of your ethical obligation is to understand New's York Rules of Professional Conduct and how they apply to and control your use of AI. In this context, "competence" under Rule 1.1 means you need to know *when* AI is appropriate and when it isn't.

HOWARD: Can I jump in on a tech point?

MICHAEL: Please.

HOWARD: So, there's this concept in AI safety called the "alignment problem" - making sure the AI's goals align with human values. (https://www.ibm.com/think/topics/ai-alignment)

In legal practice, you've got a similar issue: making sure your *use* of AI aligns with your professional obligations.

For example, Rule 1.1 dictates that you to understand AI may have **algorithmic bias**. These models are trained on internet data, which means they replicate societal biases. A lawyer should be aware that AI outputs may include historical information that is false, inaccurate, or biased. If you're using AI for something like employment screening or drafting terms that could have disparate impact, you need to understand that the AI might suggest language that's problematic.

ADRIANA: That's happening in New York City right now. Local Law 144 regulates "Automated Employment Decision Tools" - AEDTs. If an employer uses AI for hiring or promotion decisions, they must do annual bias audits and notify candidates. (https://www.nyc.gov/assets/dca/downloads/pdf/about/DCWP-AEDT-FAQ.pdf) We're already seeing litigation where discovery targets those audit reports.

ERIC: So, competence means understanding both technical limitations *and* legal implications.

MICHAEL: Right. Now, **Rule 1.6 - Confidentiality**. This is huge. The rule requires you to make "reasonable efforts" to prevent unauthorized disclosure of client information. For AI, that means:

(reading from NYC Bar Opinion) (https://www.nycbar.org/reports/formal-opinion-2024-5-generative-ai-in-the-practice-of-law/)

"Without client consent, a lawyer must not input confidential client information into any Generative AI system that will share the inputted confidential information with third parties. Even with consent, a lawyer should “avoid” entering details that can be used to identify the client. Consent is not needed if no confidential client information is shared, for example through anonymization of client information. Generative AI systems that keep inputted information entirely within the firm’s own protected databases, sometimes called “closed" systems, do not present these risks. But a lawyer must not input any confidential information of the client into any Generative AI system that lacks adequate confidentiality and security protections, regardless of whether the system uses or shares inputted information, unless the client has given informed consent to the lawyer’s doing so. Even with closed systems, a lawyer must take care that confidential information is not improperly shared with other persons at or clients of the same law firm, including persons who are prohibited access to the information because of an ethical wall.”

TOM: *(concerned)* Okay, so under Rule 1.6, you must not disclose confidential client information to a public AI tool if doing so would allow the provider or others to access, retain, or reuse that information, unless you obtain the client’s informed consent and take reasonable steps to ensure confidentiality (for example, using an enterprise or private model with contractual limits on training and third-party access).

MICHAEL: Correct.

“Even with closed systems, a lawyer must take care that confidential information is not improperly shared with other persons at or clients of the same law firm, including persons who are prohibited access to the information because of an ethical wall.”

You should " rigorously vet vendors" - that means for each vendor reading the terms of service, understanding where data is stored, whether it's used for training, how long it's retained.

MAB: What about when I'm responding to discovery? The other side produces documents - can I use AI to review those?

MICHAEL: Good question. That gets complicated, especially if there is a confidentiality order in place. If there's a protective order in place, you need to make sure that uploading documents to an AI platform doesn't violate that order. When negotiating protective orders, you may want to include language that addresses AI use with opposing counsel in protective orders - like prohibiting the other side from using public AI tools to review your client's confidential documents.

ADRIANA: We're seeing more attorneys proactively putting AI-related provisions in their FRCP Rule 26(f) meet-and-confer discussions and proposed protective orders under Commercial Division Rule 11-g.

ERIC: Smart. What's next, Michael?

MICHAEL: Rule 3.3 - Candor to the Tribunal, Rule 11 of the Federal Rules of Civil Procedure and New York's Rule §§ 130-1.1a. These are where the rubber meets the road.

Rule 3.3 of the Rules of Professional Conduct prohibits making false statements to a court. Rule 11 of the Federal Rules of Civil Procedure requires you to certify that your filings are "warranted by existing law." The case law we discussed earlier - *Mata*, *Park*, *Wadsworth* - these all turned on violations of these rules.

Similar obligations exist under Section 130-1.1 for frivolous conduct as hallucinated legal citation and argument "is completely without merit in law."

The **essential requirement** is: You must independently verify every AI-generated citation and legal proposition before filing it. The NYC Bar opinion puts it bluntly: "Treat GenAI like a junior researcher whose work must be independently checked and verified."

TOM: And if I don't?

MICHAEL: Best case scenario? You catch it yourself, file a correction, and suffer embarrassment. A likely scenario? Similar to what happened in *Hall v. Academy Charter School* - court admonishment and a requirement to notify your client. The Court may also refer you to the grievance committee even if there is no monetary sanction.

Worst case? What happened in *Mata* - thousands in fines, public reprimand and notification requirements that damage your professional reputation.

ADRIANA: And the Second Circuit in *Park v. Kim referred the lawyer to the grievance committee*. So, we're not just talking about case sanctions, we're talking about potential professional discipline.

MAB: *(pale)* Oh no! That memo I mentioned at the start... I need to go re-check every cite.

MICHAEL: Yes. Yes, you do. *(not unkindly)* And from now on, here's your workflow:

VERIFICATION PROTOCOL FOR AI-ASSISTED RESEARCH:

1. Use AI to generate issue outline or find potential authorities;
2. For EVERY case the AI suggests, pull up the actual opinion;
3. Read the entire opinion;
4. Confirm the case supports your proposition and that the AI's recitation of the case's holding and the application to your case are complete and correct;
5. Run it through Shepard's or KeyCite to assure it remains good legal authority;
6. Copy quotes only from the official source, never from AI; and
7. Document your verification in a research memo.

TOM: That sounds... time-consuming.

MICHAEL: It's exactly as time-consuming as legal research has always been. The AI just helps you find *candidates* for citation faster. It doesn't eliminate the verification step.

ERIC: What about fees? I've heard there are issues with billing for AI-assisted work.

MICHAEL: Rules 1.4 and 1.5 - Communication and Fees. This is evolving, but the current guidance is:

- You CAN charge for time spent using AI - drafting prompts, reviewing output, editing.
- You CANNOT charge for time saved by AI. No billing for "three hours" when AI helped you finish in one.
- Costs associated with Generative AI should be disclosed in advance to client and the costs charged should be consistent with ethical guidance on disbursements and comply with applicable law.

- You CANNOT charge clients for time necessitated by your own inexperience, such as time spent *learning* how to use AI tools.

As the ABA explained in Formal Opinion 512 “the lawyer who has agreed to bill on the basis of hours expended does not fulfill her ethical duty if she bills the client for more time than she has actually expended on the client’s behalf.” (ABA Comm. on Ethics & Prof’l Responsibility, Formal Op. 93-379, at 6). If a lawyer uses an AI tool to draft a pleading and expends 15 minutes to input the relevant information into the AI program, the lawyer may charge for the 15 minutes as well as for the time the lawyer expends to review the resulting draft for accuracy and completeness.

TOM: So, if AI makes me more efficient, I must... charge less?

ERIC: Or move to alternative fee arrangements. Fixed fees, value-based billing. If AI lets you deliver a research memo in one hour instead of five, maybe that memo is worth $X regardless of hours.

Formal Opinion 512 supports that the basis for a fee must be communicated before… commencing the representation, preferably in writing… and before charging the client for the use of the AI tools or services.

TOM: *(muttering)* Great, another business model disruption...

MICHAEL: Last major rules: **Rule 5.1; Rule 5.2; Rule 5.3; Rule 8.4- Duty to Supervise Lawyers and Nonlawyers, Responsibilities of Subordinate Lawyer**. AI tools and their vendors count as "nonlawyer assistance" you're responsible for supervising.

If you're a partner, you must ensure your associates and staff use AI in compliance with the rules. That means:

- Having a written firm AI policy
- Training everyone who uses AI
- Implementing verification protocols
- Monitoring compliance

You can't just let your partners, associates or staff run wild with ChatGPT and claim you didn't know.

ERIC: Which brings us to practical workflows. We've covered research - which is verify, verify, verify. What about other use cases?

ADRIANA: Can I talk about discovery and transcripts? That's a big area for litigation.

ERIC: Please.

ADRIANA: So, AI is incredibly useful for e-discovery and transcript analysis - this is actually where AI in law started, before ChatGPT. We're talking about Technology-Assisted Review, or TAR, also called predictive coding.

The Southern District of New York has great case law on this - *Da Silva Moore v. Publicis Groupe* from 2012, *Rio Tinto v. Vale* from 2015. These cases established that TAR is not just acceptable, it's often *more accurate and cost-effective* than manual review for large-scale discovery. Under the New York State Commercial Division Rules at Rule 11-c. Discovery of Electronically Stored Information under subpart (f) the rules set forth "The parties are encouraged to use efficient means to identify ESI for production, which may include technology-assisted review in appropriate cases. The parties shall confer, at the outset of discovery and as needed throughout the discovery period, about technology-assisted review mechanisms they propose to use in document review and production."

HOWARD: How does TAR work?

ADRIANA: Simplified: you have hundreds of thousands of documents. You hand-code a "seed set" - maybe a thousand documents - marking which are responsive, which are privileged, etc. The AI learns from your coding and predicts how to code the rest. You validate its accuracy through statistical sampling, then use it to process the remaining documents.

MICHAEL: And courts accept this?

ADRIANA: When it's done right - with proper documentation of your methodology, of your validation metrics, and of your quality control - yes. The key is *defensibility*. You must establish the discovery responses provided using TAR are based on "reasonable inquiry" of the methodology and that the resulting production is responsive and complete.

ERIC: And the newer generative AI tools - can those help?

ADRIANA: Yes and no. GenAI is great for *summarizing* documents, depositions, transcripts. But you must be careful. Let's say you use AI to summarize a 300-page deposition. That summary might be 95% accurate but subtly mischaracterize a key admission.

The rule: Every critical quote or fact from an AI summary *must* be checked for accuracy and tied back to the specific page and line of the official transcript. You can't just cite "AI Summary of Depo" - you need "Smith Dep. 127:15-22" with the actual quote.

This aligns with authentication requirements under Article 9 of the New York Guide to Evidence.

TOM: *(raising hand)* Can we talk about the IP issues? I've got clients worried about AI companies scraping their copyrighted work.

ERIC: Good transition. Tom, give us the copyright overview, then we need to talk about the access-to-justice implications before we break.

TOM: Okay, so full disclosure: this is fast-moving and there's massive litigation happening.

The big case right now is *In re: OpenAI Copyright Infringement Litigation* - it's a Multi-District Lawsuit in the Southern District of New York with hundreds of plaintiffs: the Authors Guild, the New York Times, major publishers, individual authors. All claiming OpenAI and Microsoft violated their copyrights by training AI models on their works without permission.

MICHAEL: What's the legal theory?

TOM: The plaintiffs say training an AI on copyrighted works is unauthorized copying - you're reproducing the works into the AI's database to create the training data. And the AI's outputs sometimes reproduce near-verbatim excerpts from the training data, which is clear infringement.

OpenAI's defense is fair use - they argue training is transformative, doesn't harm the market for the originals, and creates new utility.

ERIC: Where does it stand?

TOM: Early stages still - the case was filed in April 2025, so it is in the pleading phase. There was also a big development in the fair use area in February 2025 in the case *Thomson Reuters v. ROSS Intelligence*. Thomson Reuters sued ROSS - which is an AI legal research company - for training its AI on Westlaw headnotes. The court *rejected* ROSS’s fair use defense, saying copying to train a directly competing product wasn't transformative.

HOWARD: That's huge for the legal AI market.

TOM: Right. It suggests there's now a *market* for licensing copyrighted legal data to train AI. Which means legal publishers have a new revenue stream, but it also means AI companies can't just scrape everything.

ERIC: And what should lawyers advise clients?

TOM: If your client creates content - articles, books, photographs, music - and they find it in an AI training dataset without permission, they potentially have a claim. But fair use is very fact-specific. And if they're *using* AI to create content, they need to understand the AI's training might implicate *other people's* copyrights.

We're also seeing right-of-publicity issues - New York Civil Rights Law section 50 – which establishes the right of privacy and Section 51 - which Provides a legal remedy. If an AI creates a digital replica of someone's voice or likeness without permission, that's potentially actionable.

ERIC: So, clients need guidance both as potential plaintiffs *and* potential defendants.

TOM: Exactly.

ERIC: *(checking time)* Last topic before we break: the economic access issue. Tom, you raised this earlier - how does a solo practitioner compete when big firms have million-dollar AI budgets?

TOM: *(passionately)* Right. Look, I've been practicing for years. I know IP law backwards and forwards. But when I'm up against Firm X that has Lexis+ AI, Harvey, CoCounsel, and ten associates, how do I compete?

AI should be democratizing - it should let me punch above my weight. But if the appropriate tools cost $100,000 a year, I'm priced out. And then my clients - who maybe can't afford Firm X anyway - they're stuck with me doing manual research while the other side has AI on their side.

MICHAEL: But there are lower-cost options, aren't there?

HOWARD: *(jumping in)* This is where you must be really careful. There are lower-cost and even free AI tools - ChatGPT free version, Microsoft Co-Pilot, Claude or Gemini. And those can legitimately help with things like:

- Brainstorming case theories
- Drafting client emails or engagement letters
- Creating internal case management summaries

- Learning about new areas of law

But as we discussed before, absent informed client consent, lawyers must not input confidential client information into a generative AI system that will share the information with third parties or use it for model training, and even with a client's consent, you must vet the provider's terms, safeguard confidentiality, and consider de-identifying client information. And, you absolutely must verify any legal research they generate.

TOM: So, I can use them for limited purposes.

ERIC: Right. And there's a middle ground - some legal AI companies offer solos and small firms discounted access. Clio has AI features built into their practice management software. Other law firm centric cloud providers are offering to link Microsoft Co-Counsel into the law firm's fileserver in a closed system that will not use the information for training purposes.

HOWARD: There's also the law school angle. I learned about AI research in my legal writing class. If law schools are training new lawyers to use these tools effectively and ethically, that creates a baseline of competence.

ERIC: And what about people who can't afford *any* lawyer - how does AI affect access to justice for them?

ADRIANA: *(thoughtfully)* That's a double-edged sword. On one hand, we're seeing pro se litigants use ChatGPT and other AI agents to help them navigate the legal system. There was an article in NBC News on October 8th (https://www.nbcnews.com/tech/innovation/ai-chatgpt-court-law-legal-lawyer-self-represent-pro-se-attorney-rcna230401) about a woman who used ChatGPT and Perplexity AI to successfully appeal her eviction after she lost with a lawyer at trial. She said AI was like "having God responding to her questions."

MICHAEL: *(concerned)* That's... terrifying and hopeful at the same time.

ADRIANA: Right. The terrifying part is she's doing complex legal work with no training. The hopeful part is she won and kept her home, when she could not afforded to hire an appellate lawyer to argue her case.

The researcher tracking hallucination cases - Damien Charlotin - says pro se litigants are the ones most likely to get warnings instead of sanctions when they use AI wrong, because courts recognize they don't have the professional obligations lawyers do.

ERIC: So, AI could expand access to justice for people who otherwise couldn't afford representation.

TOM: Or it could lead to a flood of bad AI-generated filings that clog up the courts and waste everyone's time.

ADRIANA: Probably both. Which is why legal aid organizations are starting to think about how to harness AI to help their clients. If a legal aid attorney can use AI to handle twice as many cases, that's huge.

ERIC: But they need to do it within the ethical rules we've discussed.

ADRIANA: Absolutely.

ERIC: And, there are developing Court Rules addressing the use of AI.

The Suffolk County Surrogate's Court has also issued a protocol effective December 1, 2025, which specifically address AI Use that requires: "All submissions to the court containing legal authority citations must include a certification by the attorney or the self-represented party, indicating either 1) that no generative artificial intelligence program or tool was used in the drafting of any pleading, affidavit, affirmation, memorandum of law, brief, or other submission, or 2) that a generative artificial intelligence program or tool was used in the drafting of any pleading, affidavit, affirmation, memorandum of law, brief, or other submission, and affirm that all generated text, including citations, quotation, and legal analysis, was reviewed for accuracy and approved by an attorney or the self-represented party. If the certification states a generative artificial intelligence program or tool was used, each program or tool employed must be identified and all document(s) which include matter generated by the program or tool must be specified along with which parts of the document(s) were drafted by which program or tool. The certification must contain a brief description of how the program or tool was employed."

It also requires that "Every quotation and citation generated with the assistance of an Al program or tool must be independently verified by an attorney or a self represented party. Reliance on a citation that does not exist, also known as a "hallucinated" citation, shall constitute prima facie frivolous conduct under 22 NYCRR § 130-1.1."

So now, at least in Suffolk County Surrogate's Court failure to verify by lawyers and non-lawyers alike will result in a finding of frivolous conduct under 22 NYCRR § 130-1.1.

ADRIANA: It would not surprise me to see more individual Judges incorporate similar requirements into their Part Rules.

ERIC: *(checking time)* Alright, let me summarize where we are:

(ticks off on fingers)

1. The New York Rules of Professional Conduct apply to AI use - no new rules (yet), but you must use heightened scrutiny.
2. Key obligations: competence (understand the risks), confidentiality (vet vendors and do not input confidential client information into any AI system that lacks adequate confidentiality and security), candor (verify everything), supervision (have a firm policy).
3. Practical workflows exist for research, drafting, discovery, transcript analysis - but all require human verification.
4. IP issues are evolving rapidly with major litigation pending.
5. Economic disparity is real, but there are lower-cost tools and AI could expand access to justice if used carefully.

Let's break for Justice D's commentary, then we'll tackle evidence and admissibility issues.

[LIGHTS DIM - JUSTICE D RETURNS]

JUSTICE D - SECOND COMMENTARY *(5 minutes)*

JUSTICE D: Excellent discussion. Let me reinforce and add to what the committee covered.

On the ethics framework: The committee correctly identified that this is primarily about applying *existing* rules to new technology. But I want to emphasize the *duty of technological competence* more explicitly.

Comment 8 to Rule 1.1 states that competence requires "keeping abreast of changes in law relevant to the lawyer's practice, and the benefits and risks associated with relevant technology."

In 2025, for most areas of practice, AI is "relevant technology." You cannot claim ignorance. As a result, your obligations go beyond CLEs:

- You must understand what AI can and cannot do reliably.

- You must know *when* AI use is appropriate and when it creates unacceptable risk.
- You must stay current as the technology evolves - what was true about AI six months ago may not be true today.

On confidentiality and vendor vetting: This is non-negotiable. I'll tell you what I've done in chambers: we've adopted the New York Unified Court System's Interim Policy on AI Use, effective October 2025. Key mandates:

- Only approved, vetted tools may be used
- All users must complete training
- “Documents that have been filed or submitted for filing in any court are also considered confidential, even if they are classified as public at the time of filing, since it is possible that the record of the case will be sealed in the future, or that the documents have not been adequately redacted to conceal sensitive information. Although, in these scenarios, the confidential information has already been revealed to the public, entering the information into the public model AI program makes the exposure of the information permanent” Human decision-making cannot be replaced by AI”
- If these are the standards for judges and court staff, they should be your floor as practitioners.

On verification FRCP Rule 11 and New York’s Section 130-1.1: What satisfies your certification obligation under Rule 11. *(reading)* "By presenting to the court a pleading, written motion, or other paper... an attorney certifies that... the claims, defenses, and other legal contentions are warranted by existing law." Under New York’s Section 130-1.1(c) “conduct is frivolous if:

(1) it is completely without merit in law and cannot be supported by a reasonable argument for an extension, modification or reversal of existing law”;

“Existing law" means *real law that you have personally confirmed exists*. The committee's verification protocol is exactly right:

1. AI suggests authorities
2. You retrieve the actual opinions
3. You read the relevant portions
4. You Shepardize or KeyCite

5. You confirm the case says what you claim
6. You cite only from the official source

Miss any of these steps, and you've violated FRCP Rule 11 and New York’s Section 130-1.1.

On discovery and proportionality: If AI-assisted review would make your discovery more accurate, less expensive, and more proportional, you need to at least *consider* it. Document why you chose the method you chose.

Practice Point #4: Create a written AI policy for your firm or practice. Map it to the National Institute of Standards and Technology AI Risk Management Framework:

- **Govern:** What tools are approved? Who can use them? What training is required?
- **Map:** What tasks will AI be used for? What data will be processed?
- **Measure:** How will you validate accuracy and quality?
- **Manage:** What safeguards prevent misuse? How do you handle errors?

Practice Point #5: Update your engagement letters. Disclose to clients that you may use AI tools, explain the safeguards you've implemented, and get their consent for any AI-related costs you'll pass through.

Practice Point #6: Many federal judges now have standing orders requiring certification of AI use. In New York, the Commercial Division is considering proposed Rule 6(e) that would mandate disclosure and certification for AI-assisted filings. Also, be sure to check your judge's individual rules and practices as individual Judges have imposed there own rules concerning AI use.

The committee will now address evidence and admissibility issues. This is where things get really interesting.

ERIC – THE FIRST CODE IS _______________________

ACT III: Discovery, Evidence & Access to Justice (25 minutes)

ERIC: *(reconvening)* Alright, we've covered the "input" side - using AI as lawyers. Now let's talk about the "output" side - when AI-generated material becomes evidence, or when we need to authenticate potentially AI-altered evidence.

Adriana, you mentioned you're seeing this in chambers?

ADRIANA: *(nodding vigorously)* Oh yes. Just last month we had a case where one party submitted a video recording as evidence. The other side immediately moved to exclude it, claiming it might be a "deepfake" - AI-generated or AI-altered video. The moving party submitted an expert declaration saying they'd analyzed the video using forensic tools and found "artifacts consistent with AI manipulation."

ERIC: How did the judge handle it?

ADRIANA: Required a full evidentiary hearing to which we examined it applying New York's Frye Test / Daubert standards. The question in Frye is "whether the accepted techniques, when properly performed, generate results accepted as reliable within the scientific community generally." In Daubert, the trial judge acts as the gatekeeper to test the reliability of scientific testimony.

The party offering the video had to establish:

- Chain of custody from creation to filing
- Metadata showing when and how it was created
- Hash values proving it hadn't been altered
- Expert testimony about the recording system and process generated results accepted as reliable within the AI scientific community.

And then we had competing experts battle it out over whether the alleged "artifacts" actually indicated AI manipulation or were just compression artifacts from file format conversion.

MAB: I think you would also get Federal Rule of Evidence 702 on expert testimony and Rule 901 on authentication.

HOWARD: *(excited)* This is the authentication problem. It used to be, if you had a video, you could authenticate it with witness testimony: "Yes, this is a true and accurate recording of what happened." But now that AI can create photorealistic fake videos, fake audio, even fake documents, authentication has become a lot harder.

MICHAEL: The Guide to New York Evidence has Article 9 on authentication. What does that require?

ADRIANA: *(pulling up materials)* Article 9.01 says evidence must be authenticated by "evidence sufficient to support a finding that the offered evidence is what the proponent claims it is" For example that could be:

- Testimony from someone with personal knowledge
- Distinctive characteristics or circumstances
- Expert testimony about a process or system

But now courts are requiring what we call "multi-method authentication" for digital evidence that could be AI-susceptible. A recent case Matter of Weber (As Trustee of Michael S. Weber Trust), 2024 NY Slip Op 24258 (Sur. Ct. 2024) touches upon this

TOM: Meaning?

ADRIANA: You can't just rely on one method. You need to combine:

- Technical evidence: metadata, hash values, provenance records, blockchain timestamps if available
- Testimonial evidence: witness with knowledge, expert opinion on the system
- Circumstantial evidence: is the content consistent with other verified evidence?

MAB: What about documents? Can AI create fake contracts, fake emails?

HOWARD: Absolutely. And it's getting scary good. There are AI tools that can analyze someone's writing style from their social media and emails, then generate new text that sounds exactly like them.

I saw a demonstration where they fed an AI tool 50 of someone's emails, and it could then write new emails in that person's style - same vocabulary, same sentence structure, same quirks. If someone fabricated an email using that technique, it would be extremely hard to detect.

ADRIANA: Which is why authentication now requires examining the digital forensics. Email headers, server logs, metadata creation dates. You can't just look at the content anymore.

ERIC: What about when a party uses AI in their work product - like an expert using AI to analyze data?

MICHAEL: *(jumping in)* That's the *Kohls v. Ellison* case. United States District Court for the District of Minnesota, January 2025. The state's expert submitted a declaration supporting a motion. He cited two non-existent academic articles about AI and deep fakes. Opposing counsel checked the citations - they were fake. Ironically, the expert had used GPT-4o to research and never verified the sources.

ADRIANA: What happened?

MICHAEL: The court excluded the expert's entire declaration and testimony. *(reading from the opinion)* The court said relying on unverified AI output "shattered his credibility" and his testimony was "not reliable under Rule 702."

HOWARD: Federal Rule of Evidence 702 is the expert testimony rule, right?

MICHAEL: Right. As amended in 2023, Federal Rule of Evidence 702 requires that expert testimony be:

- Based on sufficient facts or data
- The product of reliable principles and methods
- The expert reliably applied those principles and methods to the facts

If an expert relies on AI analysis, they need to be able to explain:

- What AI system was used
- What data it was trained on
- How it was validated
- What the known error rates are
- What human oversight was involved.

TOM: So even experts can't just defer to "the AI said so."

ADRIANA: Correct. And as I discussed before in New York we apply the standards of Daubert and Fry The judge is the gatekeeper. There's an excellent law review article - Grimm, Grossman and Cormack, *Artificial Intelligence as Evidence* (*abstract*)- that lays out a framework courts can use.

GRIMM-GROSSMAN-CORMACK FRAMEWORK FOR AI AS EVIDENCE:

When a party offers AI-generated evidence or AI-assisted analysis:

1. **Identify the System:** What specific AI tool or model was used?

2. **Identify the Data:** What was the training data? What was the input data?
3. **Explain the Method:** How was the AI trained? What algorithm or approach?
4. **Validate the Method:** How was accuracy tested? What are the error rates?
5. **Demonstrate Robustness:** Has the system been peer-reviewed? What quality controls exist?
6. **Show Proper Application:** Was the tool used within its validated parameters? Was there human oversight?

ERIC: That's a pretty high bar.

ADRIANA: It needs to be. Because AI can produce very convincing but completely wrong results. And unlike human witnesses who can be cross-examined about their reasoning, you can't really cross-examine an AI model.

MAB: What about AI-generated risk scores? Like if a bail algorithm says someone is high-risk, or an employment screening tool flags a candidate?

HOWARD: *(leaning forward)* This is where algorithmic bias becomes a huge issue. These systems are often trained on historical data that reflects past discrimination. So, an AI trained on 20 years of bail decisions might learn to replicate racial disparities in those decisions.

ERIC: Let's shift back to the discovery side. Adriana, you mentioned TAR and predictive coding earlier. Can you go deeper on how to make that defensible?

ADRIANA: Sure. The key Southern District cases are *Da Silva Moore* (2012), *Rio Tinto* (2015), and *Hyles* (2016). These cases have been adopted and followed by the New York State Courts. We see this too, in the Uniform Civil Rules for The Supreme Court & The County Court Section 202.20-c(e) Requests for Documents. It provides that parties are encouraged to use “the most efficient means to review documents, including electronically stored information,” so long as the method complies with CPLR Article 31 and is proportional to the needs of the case. The Rule states that such methods may include “technology-assisted review, including predictive coding, in appropriate cases,” and further encourages parties to confer at the outset of discovery, and as needed throughout, about the TAR mechanisms they intend to use.

The Commercial Division has adopted the same approach. Commercial Division Rule 11-c(f) Discovery of Electronically Stored Information provides that parties are encouraged to use “the most efficient means to review documents, including electronically stored information,” consistent with their Article 31 disclosure obligations and proportional to the needs of the case. The Rule expressly recognizes that such means may include “technology-assisted review, including predictive coding, in appropriate cases,” and encourages parties to confer at the outset of discovery and as needed throughout the discovery period about the TAR mechanisms they intend to use in document review and production.

The through-line is: **TAR is acceptable when properly implemented and documented.**

Here's a defensible TAR workflow:

PHASE 1: PLANNING

- Meet and confer with opposing counsel under Rule 26(f)
- Agree on scope, methodology, metrics
- Consider requesting a Rule 502(d) order to protect against privilege waiver

PHASE 2: SEED SET CREATION

- Senior attorney reviews and codes representative sample
- Typically 500-2,000 documents
- Document your selection criteria

PHASE 3: TRAINING

- AI learns from seed set coding
- Generates predictive model

PHASE 4: VALIDATION

- Test the model on a separate validation set
- Measure precision (what % of AI-coded responsive docs are responsive)
- Measure recall (what % of all responsive docs did AI catch)
- Set acceptable thresholds - typically 75%+ recall

PHASE 5: PRODUCTION & QUALITY CONTROL

- Apply model to full document set

- Human review of uncertain cases
- Statistical sampling to verify accuracy
- Document everything

MICHAEL: That sounds expensive.

ADRIANA: Compared to manual review of 500,000 documents? It's massively cheaper. And the *Da Silva Moore* court found it was more accurate than manual review or keyword searching.

But yes, you need expertise. Most firms hire e-discovery vendors who specialize in this.

MAB: What if opposing counsel challenges your TAR process?

ADRIANA: Then you need to show your work. That's why documentation is key. In *Rio Tinto*, the court approved stipulated TAR protocols where both sides agreed on methodology and validation.

The key as noted in the New York Rules is *cooperation* - don't spring TAR on opposing counsel at the last minute. Discuss it early, let them provide input on seed documents, share your validation metrics.

ERIC: What about the newest GenAI tools - like using ChatGPT to summarize depositions or contracts?

ADRIANA: *(cautiously)* They can be useful, but you need guardrails. Let me give you a real example from our chambers.

An attorney submitted a brief relying heavily on "AI-generated deposition summaries." The brief quoted testimony but cited "AI Summary" instead of page and line numbers. Opposing counsel objected - they couldn't verify the quotes without going through the entire transcript.

My Judge ordered the attorney to submit an amended brief with proper transcript citations - page and line - for every quote. The attorney ended up having to read through the whole deposition anyway to find the page numbers, which defeated the purpose of the AI summary.

The lesson: AI can help you *find* the relevant portions of a transcript faster. But you still need to cite to the official record. Use the AI summary as a finding aid, not as the source of truth.

TOM: *(frustrated)* So AI saves time but then we must do the work anyway to verify?

HOWARD: Not quite. Think of it this way:

WITHOUT AI:

- Read 300-page deposition line by line
- Identify key admissions manually
- Create chronology manually
- Cross-reference with documents manually
- Total time: 10-15 hours

WITH AI:

- AI generates summary highlighting key admissions (30 minutes)
- You review summary and verify key sections (3-4 hours)
- AI generates initial chronology (20 minutes)
- You verify and refine (1-2 hours)
- Total time: 5-7 hours

You're not eliminating the verification work, but you're making it more targeted. The AI points you to where you need to look.

MICHAEL: I think you need a middle ground. While you might end up with better work product because the AI caught things you might have missed in a straight read-through. But, what about the things the AI misses? If the AI doesn’t have access to all the exhibits used during the deposition, it may not be able to connect all the dots or see the entire picture. And, even if it does, AI does not always catch or correctly report the data it has reviewed and reported on. You still must review the entire transcript.

ERIC: Okay, so synthesis time. For evidence and discovery:

1. Authentication of AI-susceptible evidence requires multi-method foundation: technical, testimonial, and circumstantial.
2. Expert testimony that relies on AI must satisfy reliability requirements - explain the system, validation, error rates, human oversight.
3. TAR and predictive coding are defensible when properly documented and validated.

4. AI summarization tools are useful but outputs must be verified and cited to official records.
5. Algorithmic bias is a real evidentiary issue, especially with employment and criminal justice AI.

Does that capture it?

ALL: Yes

ERIC: Good. Before we break for Justice D's commentary, let me circle back to the access to justice question because we didn't fully resolve it.

Tom, you're worried about competing with big firms. But Michael, you work at a big firm - what's your perspective?

MICHAEL: *(thoughtfully)* Honestly? I think AI is leveling the playing field more than Tom gives it credit for. Yes, my firm has expensive tools. But the core capability - generating competent legal analysis - that's increasingly available at lower cost.

What big firms really pay for is *judgment* and *experience*. AI doesn't replace that. A senior partner who can spot the winning issue, develop strategy, read the judge - that's what clients pay premium rates for.

Tom with 25 years of IP experience and a $30/month AI tool can probably outperform a junior associate at my firm with access to every tool but no judgment.

TOM: *(considering)* That's... actually reassuring.

MAB: And for access to justice - people who can't afford any lawyer - AI might be transformative. There are already projects using AI to help people with traffic tickets, landlord-tenant issues, simple divorces. As long as there are human lawyers overseeing the system or providing backup, that could help millions of people.

ADRIANA: The American Bar Association in August 2023, the ABA created a specific Task Force on Law and Artificial Intelligence to address the AI-specific regulatory questions that's looking at this. How do we expand access while maintaining quality and ethics? AI is part of the answer, but we need regulatory frameworks.

Some jurisdictions are experimenting with "paraprofessional" licenses - people who aren't full lawyers but can provide limited legal services with AI assistance. Utah launched the first "regulatory sandbox" in 2020. This allows non-traditional legal business structures—including those owned by non-lawyers or those using advanced AI/tech models—to offer services to the public under close regulatory supervision. In 2024 the scope of the program was narrowed to assure applicants to the Sandbox demonstrated that Sandbox authorization will allow it to reach Utah consumers currently underserved by the legal market.

New York hasn't gone there yet, but the conversation is happening.

ERIC: So AI creates both opportunity and risk for access to justice.

ADRIANA: Exactly.

ERIC: Alright, let's break for Justice D's commentary on evidence and discovery.

[LIGHTS DIM - JUSTICE D RETURNS]

JUSTICE D - THIRD COMMENTARY *(5 minutes)*

JUSTICE D: The committee has done excellent work on authentication and expert testimony. Let me add some bench perspective.

On authentication of digital evidence: The committee correctly noted we're moving toward multi-method authentication. But I want to emphasize: this is *not* creating a higher standard for digital evidence than we have for physical evidence. It's applying the *same* standard to evidence that has new vulnerabilities.

Under Article 9 of the Guide to New York Evidence - and under FRE 901 - the standard is "the evidence must be sufficient to support a finding that the item is what it is claimed to be." The methods of proof adapt to the nature of the evidence.

For a paper document, you might authenticate with testimony about chain of custody and distinctive characteristics. For a digital recording potentially subject to AI manipulation, you authenticate with:

- Technical metadata
- Expert analysis of the file
- Testimony about the recording system
- Circumstantial corroboration

Different methods, same standard.

On expert testimony and AI: Judges are to perform their gatekeeping function assuring the proponent prove *by a preponderance* that:

- The expert's opinions are based on sufficient data;
- The methods are reliable;
- The expert reliably applied those methods.

When an expert uses AI, you need to probe:

- Is this AI system generally accepted in the relevant field?
- Has it been validated?
- What are its error rates?
- Did the expert understand its limitations?
- Was there human oversight of the AI's analysis?

In *Kohls v. Ellison*, the expert couldn't answer these questions because he'd simply asked GPT-4o for sources and didn't verify them. That's not reliable methodology.

On e-discovery and TAR: The committee mentioned *Da Silva Moore*, *Rio Tinto*, and *Hyles* and New York's Trial Court and Commercial Division Rules These establish several principles:

1. **Cooperation is key.** Discuss TAR protocols in meet-and-confer. Don't ambush opposing counsel.
2. **Transparency builds defensibility.** Share your methodology, your validation metrics, your quality control process.
3. **Court approval is not required** - you don't need permission to use TAR - **but cooperation reduces disputes.**
4. **Proportionality matters.** Discovery is to be proportional to the needs of the case. In a case with millions of documents, TAR may be the *only* proportional option.
5. **Document your process.** Keep records of seed set selection, training, validation testing, quality control samples, error remediation.

On privilege protection: Rule 502(d) of the Federal Rules of Evidence allows a court to enter an order providing that privilege waiver in the case doesn't extend beyond that case.

There is currently no direct statutory equivalent in the New York Civil Practice Law and Rules (CPLR) to Federal Rule of Evidence (FRE) 502.

New York State Courts handle these issues through:

(1) CPLR privilege/work-product provisions – CPLR 3101(c), CPLR 3101(d)(2); CPLR 4503, and CPLR 4548

(2) Court rules: Commercial Division Rule 11-c(g) (Discovery of ESI) – includes an automatic clawback provision for inadvertently produced privileged ESI. Under Rule 11-c(g), inadvertent production of privileged ESI does not constitute a waiver and the receiving party must return/destroy the documents (subject to challenge), unless the parties agree otherwise. Separately, the Uniform Civil Rules for Supreme and County Courts also contemplate clawback agreements. 22 NYCRR 202.12(c)(3)(viii): at preliminary conferences, parties should address agreements "for the clawback of inadvertently produced privileged material", and

(3) common-law waiver doctrine uses a multi-factor test very close to 502(b). To avoid waiver after an inadvertent disclosure, the producing party generally must show: (i) No intent to disclose & reasonable steps taken to prevent disclosure; (ii) Prompt, reasonable steps to rectify once the disclosure is discovered; and (iii) The receiving party will not be unduly prejudiced if use of the material is barred. A representative case applying these standards is *AFA Protective Systems, Inc. v. City of New York*, 13 A.D.3d 564 (2d Dep't 2004). A member of our Inn, Kevin Schlosser, examined the inadvertent disclosure of privilege in the E-Age in a 2006 New York Law Journal article.

Why does this matter for AI? Because AI-assisted privilege review, while generally accurate, is not perfect. You need to be prepared if your AI tags a privileged document as non-privileged and you inadvertently produce it.

The order should specify:

- Inadvertent production doesn't waive privilege
- The receiving party must return or destroy privileged materials when notified
- No waiver extends beyond this litigation

Commercial Division judges routinely enter these orders. Ask for one.

Practice Point #7: When you're dealing with digital evidence - whether you're offering it or opposing it - think about:

- **Provenance:** Where did it come from? What's the chain of custody?

- **Integrity:** Has it been altered? What technical evidence proves (or disproves) that?
- **Authenticity:** Is it what it claims to be? What evidence supports that?

For AI-susceptible evidence, consider retaining a digital forensics expert early.

Practice Point #8: In discovery disputes involving AI, focus the argument on *process*, not *technology*. Don't argue "AI is reliable" or "AI is unreliable" in the abstract. Argue "this specific TAR protocol, with these validation metrics and these quality controls, satisfies reasonable inquiry standard."

Practice Point #9: For transcript and deposition summaries generated by AI: use them internally as work product but cite to the official transcript in filings. Page and line numbers are required.

The committee will now address governance and best practices for firms and chambers.

ACT IV: The Road Forward - Governance & Best Practices (15 minutes)

ERIC: We're in the home stretch. We've covered what AI is, the ethics rules, practical workflows, and evidentiary issues. Now let's talk about what law firms, solo practitioners, and judicial chambers should *do* to implement this.

Michael, you mentioned earlier that Rule 5.3 requires supervision. What does that look like in practice?

MICHAEL: It means having a written policy. Not just "be careful with AI" - a real governance framework.

You may want to consider it on the NIST AI Risk Management Framework that Howard mentioned earlier. It has four components: Govern, Map, Measure, Manage.

1. **GOVERN:** Create a risk management culture that identifies risks and impacts. Establish the rules:

- What AI tools are approved for use?
- Who can use them?
- What training is required?
- Who's responsible for oversight?

Firms might consider creating an "AI Committee" – with lawyers, firm IT professionals, and maybe outside counsel for ethics advice. They can review and approve all AI tools before anyone can use them with client data.

2. **MAP:** Document the context and risks of AI work to propose mitigation. Inventory the use cases:

- Legal research;
- Contract drafting and review;
- E-discovery and document review;
- Deposition and transcript analysis; and
- Client communication drafting.

For each use case, identify what type of data is involved:

- Public information (low risk);
- Client confidential (medium risk, requires approved tools); and
- Privileged (high risk, requires extra safeguards).

3. **MEASURE:** Identify, analyze, and track AI risks and related impacts. Define quality metrics

- For research: verification completion rate;
- For discovery: precision, recall, and elusion testing results; and
- For contract review: sample accuracy rates.

These metrics need to be tracked on a regular basis to assure the AI remains on point.

4. **MANAGE:** Allocate resources to prioritize and measure risks Implement safeguards

- Vendor contracts that prohibit data use for AI training;
- Data retention limits;
- Mandatory verification workflows;
- Incident reporting procedures; and
- Audit logs of AI use.

TOM: *(overwhelmed)* That sounds like a lot for a solo practitioner.

MICHAEL: It is. But scale it down. Your version might be one page:

Approved Tools:

- Research: Westlaw AI, Lexis+ (paid accounts)
- Drafting: Microsoft Copilot (enterprise, not free)
- General: Claude Pro or ChatGPT Plus (for non-confidential use only)

Prohibited: Free/public AI tools with client confidential information

Mandatory Training: Complete [XYZ CLE] before using any AI tool

Verification Protocol:

- All AI-suggested cases must be pulled up in Westlaw/Lexis
- All quotes verified from official source
- Shepardize/KeyCite all authorities
- Document verification in research memo

Client Communication:

- Engagement letter discloses AI use
- Client consent for any AI-related costs

Incident Reporting:

- If error discovered, immediately report to [supervisor/self]
- Document what happened and corrective action

MAB: That's manageable. And honestly, a lot of this is just documenting what we should already be doing.

ERIC: Right. AI doesn't create new *obligations* - it creates new *risks* that make existing obligations more important.

Howard, from a tech perspective, what should people look for when vetting AI vendors?

HOWARD: Great question. Here's my vendor due diligence checklist:

VENDOR VETTING CHECKLIST:

1. Data Handling:

- Is your input data used to train or improve the AI model? (Should be NO)
- Where is data stored? (Look for US-based, secure servers)
- How long is data retained? (Shorter is better)
- Can you request deletion?
- Is data encrypted in transit and at rest?

2. Confidentiality:

- Does the Terms of Service allow use with client confidential data?
- Is there a Business Associate Agreement (BAA) if health info involved?
- What happens if there's a data breach?

3. Security:

- What authentication is required? (Multi-factor is best)
- Are there access controls? (Can you limit which staff use it?)
- Is there an audit log of who accessed what?

4. Technical Capability:

- For legal AI: Is it connected to a verified legal database?
- What's the data currency? (How recent is the information?)
- Does it cite sources you can verify?
- What are the known limitations?

5. Support & Training:

- Is there customer support?
- Is training provided?
- Are there usage guides for legal professionals?

TOM: And what about cost? You mentioned there are more affordable options.

HOWARD: It's a spectrum:

FREE/LOW COST ($0-25/month):

- ChatGPT Free or Plus
- Claude free or Pro
- Google Gemini free or Pro

MID-TIER ($30-100/month):

- Microsoft Copilot for Enterprise
- ChatGPT Enterprise (if small firm)
- Some practice management software with AI built in **Pros:** Data protection. **Cons:** Still need to verify legal citations heavily.

PREMIUM LEGAL AI ($100-1000+/month):

- Westlaw AI / Lexis+
- CoCounsel
- Harvey **Pros:** Integrated with legal databases, better accuracy for legal research. **Cons:** Expensive, but may be worth it for high-volume practices.

TOM: So I could legitimately start with a $25/month tool for non-confidential work, and upgrade as my practice grows?

HOWARD: Absolutely. Just be rigorous about what information goes into which tool.

ADRIANA: Can I talk about what we're doing in chambers? Because judges have the same issues.

ERIC: Please.

ADRIANA: So the New York Unified Court System issued its Interim Policy on AI Use in October 2025. It applies to all judges and non-judicial court employees.

KEY PROVISIONS - NY UCS AI POLICY:

1. Approved Tools Only: We can only use AI tools vetted and approved by the court system IT. Personal ChatGPT accounts are prohibited for court work.

2. Training Required: All users must complete mandatory training before accessing AI tools.

3. Confidentiality Protected: Court filings, sealed documents, confidential case information - none of that can go into AI systems, even approved ones.

4. Human Decision-Making: The policy explicitly states AI "is not designed to replace human judgment, discretion, or decision-making." Judges make decisions, not AI.

5. Transparency: If a judge uses AI as a research aid - like asking it to summarize legal standards - that should be disclosed.

MICHAEL: Disclosed to the parties?

ADRIANA: The policy isn't totally clear on that yet. But there's been discussion that if AI contributes to a decision in a meaningful way, parties should know. Similar to how we disclose if we used a law clerk's research.

MAB: Have there been cases where judges used AI and got in trouble?

ADRIANA: Not "in trouble" exactly, but some interesting moments. In *Snell v. United Specialty Insurance* - Eleventh Circuit, 2024 - Judge Newsom wrote an entertaining concurring opinion where he disclosed he'd asked ChatGPT about the ordinary meaning of the word "using" in the statute.

He explained his methodology, acknowledged AI's limitations, and made clear he was using it as a research aid, not as authority. The opinion actually got praised for transparency.

But then there were cases where judges' opinions contained hallucinated citations, and people suspected AI use but the judge never disclosed. Federal judges in Mississippi and New Jersey in separate cases recently withdrew opinions after lawyers pointed out they contained fake case citations and misstatements of the record.

ERIC: So transparency is better than trying to hide it.

ADRIANA: Much better. If you use AI and disclose it properly, people understand. If you use it and hide it, and then it goes wrong, the credibility damage is severe.

ERIC: What about engagement letters and client communication? We touched on this earlier but let's be specific.

MICHAEL: Here is some suggested sample language:

ENGAGEMENT LETTER - AI DISCLOSURE CLAUSE:

"Our firm uses carefully vetted artificial intelligence tools to improve the quality and efficiency of our legal services. These tools may assist with tasks such as legal research, document review, contract analysis, and drafting. We do not input your

confidential information into public AI systems and only into private, sand-boxed AI systems upon your prior written consent. All AI-assisted work is supervised by experienced attorneys who review, verify, and take full professional responsibility for the work product.

The use of these tools may allow us to serve you more efficiently and cost-effectively. Our fees for AI-assisted work reflect the time reasonably necessary to prompt, review, verify, and finalize the work product - not the time that would have been required using manual methods. We do not charge you for time we spend learning how to use these tools.

If you have questions or concerns about our use of AI, please let us know."

TOM: That's good. Clear but not scary.

ERIC: And fee arrangements - are people moving away from hourly billing because of AI?

MICHAEL: Some are. Fixed fees, success fees, subscription models. If AI makes you 3x faster, charging hourly doesn't work as well. But this is really about value billing - charging for the value delivered, not just time spent.

TOM: Which solos and small firms have been doing forever anyway. *(smiles)*

ERIC: True. Alright, let's do our final summary. We need to give the Inn an actual set of recommendations. I think these are our top-line guidance points.

FINAL RECOMMENDATIONS - AI USE IN NEW YORK LEGAL PRACTICE:

1. EDUCATE YOURSELF

- Complete AI-focused CLE
- Understand how LLMs work and why they hallucinate
- Stay current as technology evolves

2. VET YOUR TOOLS

- Never use public/free AI with client confidential information
- Review Terms of Service for data handling
- Prefer enterprise or legal-specific tools for professional work
- Document your vendor selection process

3. VERIFY EVERYTHING

- Treat AI output as a draft that requires verification
- Independently confirm existence and accuracy of all cited authorities
- Run citators (Shepard's/KeyCite) on all cases
- Cite only from official sources, never from AI summaries
- Document your verification process

4. PROTECT CONFIDENTIALITY

- Implement clear data classification (public/confidential/privileged)
- Prohibit confidential information in unapproved tools
- Include AI provisions in protective orders
- Vet vendor contracts for data use and retention

5. ADOPT A WRITTEN POLICY

- Define approved tools and use cases
- Establish verification workflows
- Require training for all users
- Create incident reporting procedures

6. COMMUNICATE WITH CLIENTS

- Disclose AI use in engagement letters
- Explain safeguards and supervision
- Get consent for AI-related costs
- Don't overcharge or charge for learning time

7. DOCUMENT YOUR PROCESS

- Keep research logs showing verification steps
- For discovery, document TAR methodology and validation
- Maintain audit trails of AI use
- Create contemporaneous records of decision-making

8. SUPERVISE APPROPRIATELY

- Partners responsible for subordinates' AI use
- Implement training requirements
- Review AI-assisted work product
- Don't delegate professional judgment to AI

9. STAY TRANSPARENT

- Disclose AI use when required by court order
- If error discovered, correct promptly and candidly
- Consider disclosure in engagement letters and case management

10. EMBRACE PROPORTIONALITY

- Use AI where it makes work more accurate and efficient
- Consider AI in discovery for large document sets
- Balance cost, accuracy, and client needs
- Don't refuse AI out of unfounded fear

ERIC: *(standing back, looking at list)* That's comprehensive. Everyone agree these are sound recommendations?

ALL: *(nodding, affirmative sounds)*

TOM: You know, when we started this committee work, I thought AI was this scary thing that was going to make my practice obsolete. But now I see it's just another tool - powerful, yes, requiring care, yes, but fundamentally a tool that extends my capabilities.

MAB: And for those of us starting our careers, it's exciting. We get to shape how our generation uses these tools. If we build good habits now - verify, document, protect confidentiality - that becomes second nature.

MICHAEL: I'm actually less anxious about it now. The key insight for me is: AI doesn't change *what* we need to do ethically - competence, diligence, candor, confidentiality. It just changes *how carefully* we need to apply those principles.

HOWARD: And from a technology perspective, these tools are only going to get better. We're already seeing AI systems with better citation checking, better integration with legal databases, better explanation of their reasoning. The challenge is keeping pace.

ADRIANA: What gives me hope is that the legal profession is taking this seriously. The bar associations, the courts, law schools - everyone's engaged. We're not just letting technology happen *to* us. We're actively shaping how it integrates with our professional values.

ERIC: Well said. Let's wrap up and turn it over to Justice D for final thoughts, then we'll open for questions.

Our committee has accomplished what we set out to do: develop practical, ethical guidance for AI use in New York legal practice. We've grounded it in existing rules, learned from the case law, and created implementable workflows.

The materials we've developed today - the verification checklists, the policy templates, the engagement letter language, the vendor vetting criteria - all of that is part of the materials uploaded to the Inn web site for this program.

But remember: this is a snapshot in time. AI technology is evolving rapidly. What's true today may not be true in six months. Your obligation under Rule 1.1 is ongoing - keep learning, keep adapting, keep applying sound professional judgment.

Justice D, we're ready for your final commentary.

[JUSTICE D - FINAL COMMENTARY] ***(5 minutes)***

JUSTICE D:

Thank you to this committee for comprehensive and thoughtful work. Let me leave you with some final observations from the bench.

First: The Obligation to Adapt

Technology has always changed legal practice. Word processors replaced typewriters. Westlaw and Lexis replaced book research. Email replaced faxes. Each time, we adapted while preserving core professional values.

AI is the latest transformation, not the last. Your obligation is not to resist it or embrace it uncritically, but to master it consciously. To understand what it can and cannot do. To use it to serve clients better while protecting their interests and the integrity of the system.

Second: The Centrality of Human Judgment

I want to emphasize something the committee touched on: AI assists judgment; it does not replace it.

An AI can suggest case theories, but it cannot assess which theory will resonate with *this* judge in *this* courtroom.

An AI can draft a contract, but it cannot negotiate the relationship of trust that makes the contract work.

An AI can review discovery, but it cannot decide which documents tell the most compelling story at trial.

That professional judgment - built on experience, creativity, empathy, and wisdom - that remains uniquely human. That's what clients pay for. That's what justice requires.

Third: The Access to Justice Imperative

The committee discussed economic disparity. AI creates an opportunity to expand access to justice that we cannot afford to squander.

Right now, millions of people have legal problems but cannot afford lawyers. AI tools, properly deployed and supervised, could help legal aid organizations serve more clients. Could help pro se litigants navigate systems that are bewildering. Could help small firms compete with large ones.

But - and this is critical - only if we get the governance right. Only if we insist on verification, on confidentiality protection, on human oversight.

We need regulatory innovation to match technological innovation.

Fourth: Transparency and Accountability

The single most important lesson from the sanctions cases is this: *honesty matters more than perfection*.

Every lawyer who got severely sanctioned for AI hallucinations compounded a tool error with a candor failure. They doubled down. They submitted fake opinions to "prove" fake cases existed. They misrepresented their verification efforts.

The lawyers who got warnings or no sanctions? They admitted the error promptly, corrected the record, and explained what went wrong.

If you make a mistake with AI - and you will; we all will - **own it immediately**. File a correction. Notify the court and opposing counsel. Show what safeguards you're implementing to prevent recurrence.

Courts understand tool failures. We don't forgive dishonesty.

Fifth: The Role of the Judiciary

Judges have a critical role in this transition:

- We must educate ourselves about AI so we can spot issues in filings and evidence
- We must require appropriate authentication and reliability showings for AI-related evidence
- We must be proportionate in our expectations - requiring verification without creating impossible burdens
- We must lead by example in our own chambers' use of AI
- We must be transparent when we use AI to inform our decisions

The New York Unified Court System's AI policy is a good start. But this is ongoing work.

Sixth: Looking Forward

In the next few years, you will see:

- More court rules explicitly addressing AI use and requiring certification
- More bar ethics opinions as new issues arise
- Continued refinement of AI tools to reduce hallucination and improve accuracy
- New practice models and business structures built around AI
- Regulatory changes to expand access to justice through technology

Stay engaged. Participate in bar committees. Attend CLEs. Read the guidance as it develops. Share your experiences - both successes and failures - so we learn collectively.

Final Practice Points:

Practice Point #10: Create a "verification habit." Before you cite any case - whether from AI, from a colleague's memo, or from your own memory - pull it up and confirm it. This should become as automatic as spell-checking.

Practice Point #11: Build safeguards into your workflow, not just your policy. Don't rely on remembering to verify - create systems that force verification. Checklists, peer review, software tools that flag unverified citations.

Practice Point #12: Communicate proactively with courts. If a judge has a standing order on AI, comply meticulously. If you're using AI in novel ways - like AI-assisted discovery protocols - consider seeking court approval in advance.

Practice Point #13: Invest in training. Your associates, your staff, yourself. This isn't one-and-done - it's ongoing professional development.

Practice Point #14: Join the conversation. State and local bars are developing AI working groups. The NYSBA has a Task Force on Artificial Intelligence. The NYC Bar issues guidance. Get involved. Your practical experience makes you valuable to these efforts.

Conclusion

We stand at an inflection point. Artificial intelligence is transforming legal practice as profoundly as any technology in generations.

We can shape that transformation to serve justice, expand access, improve quality, and preserve professional values. Or we can react carelessly and damage all of those things.

The choice is ours. The obligation is clear.

Use AI competently. Verify rigorously. Protect confidentiality. Maintain candor. Exercise judgment. Serve clients.

These are not new obligations. They are timeless professional duties, now applied to powerful new tools.

I have confidence in this Bar. You have navigated every previous technological transition with integrity and adaptation. You will navigate this one too.

Thank you for your attention. The committee chair will now open the floor for questions.

ERIC: THE FINAL CODE FOR THE PROGRAM IS __________________________

AUDIENCE Q&A SESSION (10 minutes)

ERIC: *(to audience)* We have about ????? minutes for questions. Please identify yourself and direct your question to the committee or to Justice D.

****

ERIC: *(checking time)* That's all the time we have. Thank you all for engaging questions.

Before we close, let me summarize the key takeaways from today's program:

THE FIVE CORE PRINCIPLES FOR AI USE IN LEGAL PRACTICE:

1. **UNDERSTAND THE TECHNOLOGY** - Know what AI is, how it works, and why it fails
2. **VERIFY RELENTLESSLY** - Never cite unconfirmed sources; treat AI as a draft
3. **PROTECT CONFIDENTIALITY** - Vet tools carefully; never use public AI with client information
4. **MAINTAIN HUMAN JUDGMENT** - AI assists; you decide
5. **STAY TRANSPARENT** - Disclose appropriately; admit errors promptly

These principles, grounded in Rules 1.1, 1.6, 3.3, 5.3, and Rule 11, will guide you through whatever technological changes come next.

Thank you to our committee members for their expertise and collaboration. Thank you to Justice D for judicial guidance. And thank you to all of you for your attention and thoughtful questions.

This concludes our CLE program on AI and the Law: The Use of Technology and Its Impact Upon Legal Ethics.

[NOTE: The following are anticipated questions with prepared responses. Actual Q&A will vary but we as the panel should be familiar with the concepts.]

AUDIENCE MEMBER 1: I'm a family law attorney. Everything you've discussed has been about litigation and corporate work. How does AI help in family law practice?

MAB: Great question. AI can be really useful for:

- Drafting initial separation agreements or parenting plans based on client intake information
- Analyzing complex financial disclosure to identify hidden assets or income
- Summarizing lengthy text message histories in custody disputes
- Creating timelines from multiple sources of evidence

But you need to be especially careful with confidentiality because family law involves such sensitive personal information. Use only secure, approved tools. And never lose sight of the human dimension - AI can draft a parenting plan, but you need to counsel clients through the emotional aspects.

AUDIENCE MEMBER 2: I saw a demonstration where an AI tool "reviewed" a contract and suggested revisions. Is it safe to use AI for transactional work?

ERIC: Yes, with caveats. AI is excellent at:

- Initial contract drafting from templates
- Spotting missing standard provisions
- Flagging unusual or problematic language
- Comparing drafts to identify changes

But AI cannot:

- Assess business context and client priorities
- Negotiate relationship dynamics
- Determine what's "market" for your specific deal

- Catch subtle legal issues that require judgment

So use AI to create first drafts and spot issues, but you must review everything carefully, tailor it to the client's needs, and apply your professional judgment about what's appropriate.

HOWARD: And from a tech perspective: contract AI has gotten very good at pattern matching - finding provisions that match templates. But it struggles with novel deal structures or industry-specific customization. Think of it as a very fast junior associate who's great at checklists but needs supervision on strategy.

AUDIENCE MEMBER 3: You mentioned criminal law briefly. What about AI in criminal defense or prosecution?

ADRIANA: This is a critical and evolving area. AI is being used in criminal justice for:

- Bail risk assessments
- Sentencing recommendations
- Predictive policing
- Evidence analysis

The challenge is these systems often replicate biases in historical data. If police historically over-policed certain neighborhoods, AI trained on that data will recommend over-policing those same neighborhoods.

Defense attorneys need to:

- Demand disclosure of AI tools used in their client's case
- Challenge the methodology and validation of those tools
- Bring expert testimony about algorithmic bias
- Argue for human oversight of AI recommendations

JUSTICE D: And judges must be skeptical of "black box" AI scores. I need to understand how a recommendation was generated before I rely on it for liberty decisions. This is exactly where the Grimm-Grossman-Cormack framework applies - require explanation of the system, the data, the validation, the error rates.

AUDIENCE MEMBER 4: What happens if I use AI, verify everything carefully, but then a case I cited gets overruled the day after I file? Is that a Rule 11 violation?

MICHAEL: No. Rule 11 requires "reasonable inquiry" at the time of filing. If you verified the case existed, checked that it was good law using Shepard's or KeyCite, and cited it accurately, you've satisfied Rule 11 - even if it gets overruled later.

Of course, if you learn it's been overruled before the court rules, you have a duty under Rule 3.3 to inform the court. But that's true whether you found the case through AI or traditional research.

The point is: your obligation is to verify at the time you file. AI doesn't change that.

AUDIENCE MEMBER 5: I'm worried about younger lawyers who've grown up with technology. They might be too trusting of AI. How do we instill proper skepticism?

MAB: *(speaking up)* As a younger lawyer, I appreciate the question. But I'd push back a bit. My generation has grown up with autocorrect fails, with Google giving wrong answers, with apps crashing. We're actually pretty skeptical of technology - we know it glitches.

The challenge is AI *seems* so authoritative. It writes in perfect legalese. It formats citations correctly. It sounds confident.

But once you see it hallucinate - once you look up a case it cited and discover it doesn't exist - you never trust it uncritically again.

HOWARD: And this is why law schools are now teaching AI literacy in legal research and writing courses. Students learn early: AI is a tool, not an oracle. Verify everything.

ERIC: I'd add that supervising attorneys have an obligation under Rule 5.1 to ensure younger lawyers use AI properly. Don't just tell them "be careful" - give them checklists, review their work product, spot-check their citations. Build verification into the workflow.

AUDIENCE MEMBER 6: What about AI that generates legal strategy or litigation approach? Can I use it to help decide what motions to file?

TOM: Absolutely, with the understanding that it's brainstorming, not gospel.

I've used AI to:

- Generate a list of potential motions based on case facts
- Outline arguments for and against different legal theories
- Identify potential weaknesses in my case
- Suggest discovery to request

It's like having a very well-read colleague who can quickly give you options. But *you* have to evaluate those options based on:

- The specific judge
- Your client's goals and risk tolerance
- The practical realities of the case
- Your professional experience

AI gives you the map. You still need to decide which road to take.

AUDIENCE MEMBER 7: Justice D, will courts start requiring disclosure of AI use in all filings?

JUSTICE D: Some already do. Many federal judges have standing orders requiring certification that AI-assisted filings were verified. The Commercial Division is considering proposed Rule 6(e) that would mandate disclosure.

My prediction: within two years, most courts will have some disclosure or certification requirement. It might be as simple as a checkbox on the filing: "AI was used to assist with this document and all content has been verified by counsel."

This isn't because courts distrust AI - it's because we need to ensure verification happened. It creates accountability.

Should you wait for courts to require it? No. Consider voluntary disclosure. It demonstrates transparency and good faith.

AUDIENCE MEMBER 8: I practice in state court, not federal. Does all this federal case law and Federal Rules stuff apply to me?

ADRIANA: The principles absolutely apply. New York has its own rules:

- CPLR for procedure
- New York Rules of Professional Conduct for ethics
- New York Guide to Evidence (GNYE) for evidentiary issues

But the federal cases are highly persuasive because they're analyzing universal concepts - competence, candor, verification, authentication. New York courts look to federal precedent regularly.

And practically, many New York practitioners handle both federal and state cases, so you need to know both frameworks.

ERIC: Plus, the NYC Bar and NYSBA ethics opinions - which are New York-specific - cite and rely on the federal cases. So this isn't just federal law; it's the baseline standard for the profession.

AUDIENCE MEMBER 9: Are there situations where I should *not* use AI, even if it's available?

MICHAEL: Yes. Don't use AI when:

1. **The stakes are too high and the risk of error is unacceptable.** Example: criminal case where liberty is at stake and you need perfect accuracy.
2. **The issue requires nuanced judgment that AI can't provide.** Example: negotiating sensitive settlement terms in a mediation.
3. **Confidentiality cannot be adequately protected.** Example: you only have access to public AI tools and the matter involves highly sensitive trade secrets.
4. **The client specifically objects.** Some clients are uncomfortable with AI. Respect that.
5. **The time needed to verify AI output equals or exceeds the time to do it manually.** Example: researching a narrow issue with only three cases - just read the cases.

AI is a tool. Sometimes a screwdriver is better than a power drill. Use judgment.

AUDIENCE MEMBER 10: Last question - where can we get the templates and checklists you mentioned?

ERIC: Excellent question. The Bar Association will be publishing the complete guidance package within 30 days. It will include:

- Verification checklists for research, drafting, discovery
- Sample firm AI policies (both comprehensive and simplified versions)
- Engagement letter clauses
- Vendor vetting checklists
- Rule 26(f) sample language for AI use in discovery
- Sample AI certification for court filings
- Bibliography of all cited cases and materials with links

It will be available on the Bar Association website, and we'll distribute it through the various practice section committees.

Additionally, the NYSBA, NYC Bar, and ABA all have ongoing AI working groups and resources. I encourage everyone to engage with those.

[END OF PROGRAM]

SUPPLEMENTARY MATERIALS

COMPLETE BIBLIOGRAPHY WITH BLUEBOOK CITATIONS AND LINKS

I. CASES - AI HALLUCINATIONS AND SANCTIONS

1. **Mata v. Avianca, Inc.**, No. 22-cv-1461, 2023 WL 4114965 (S.D.N.Y. June 22, 2023), https://www.docketalarm.com/cases/New_York_Southern_District_Court/1--22-cv-01461/Mata_v._Avianca_Inc/54/
2. **Park v. Kim**, 91 F.4th 610 (2d Cir. 2024), https://law.justia.com/cases/federal/appellate-courts/ca2/22-2057/22-2057-2024-01-30.html

3. **Hall v. The Academy Charter School**, No. 2:24-cv-08630-JMW (E.D.N.Y. Aug. 7, 2025), https://websitedc.s3.amazonaws.com/documents/Hall_v._the_Academy_Charter_School_USA_7_August_2025.pdf
4. **Versant Funding LLC v. Teras Breakbulk Ocean Navigation Enter., LLC**, No. 17-cv-81140 (S.D. Fla. May 20, 2025), https://assets.alm.com/84/f0/6cc9f225499c918f1b6bd259ae59/versant-funding-v-teras-breakbulk-ocean-navigation-enterprises-order-u-s-district-court-for-the-southern-district-of-florida.pdf
5. **Johnson v. Dunn**, No. 2:21-cv-1701 (N.D. Ala. July 23, 2025), https://cases.justia.com/federal/district-courts/alabama/alndce/2%3A2021cv01701/179677/204/0.pdf
6. **Coomer v. Lindell**, No. 22-cv-01129-NYW-SBP, 2025 WL 1865282 (D. Colo. July 25, 2025)
7. **Wadsworth v. Walmart Inc.**, 348 F.R.D. 489 (D. Wyo. 2025)
8. **Mid Cent. Operating Eng'r Health and Welfare Fund v. Hoosiervac LLC**, No. 2:24-cv-00326-JPH-MJD, 2025 WL 1511211 (S.D. Ind. May 28, 2025)
9. **Kohls v. Ellison**, No. 24-cv-03754 (D. Minn. Jan. 10, 2025)
10. **U.S. v. Cohen**, 724 F. Supp. 3d 251 (S.D.N.Y. 2024)
11. **Benjamin v. Costco Wholesale Corp.**, No. 2:24-cv-7399 (LGD), 2025 WL 1195925 (E.D.N.Y. 2025)
12. **Snell v. United Specialty Ins. Co.**, 102 F.4th 1208 (11th Cir. 2024), https://media.ca11.uscourts.gov/opinions/pub/files/202212581.pdf

II. CASES - E-DISCOVERY AND TAR

13. **Da Silva Moore v. Publicis Groupe**, 287 F.R.D. 182 (S.D.N.Y. 2012), https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1%3A2011cv01279/375665/175/
14. **Rio Tinto PLC v. Vale, S.A.**, 306 F.R.D. 125 (S.D.N.Y. 2015)
15. **Hyles v. City of New York**, No. 10-cv-3119, 2016 WL 4077114 (S.D.N.Y. Aug. 1, 2016)

III. CASES - INTELLECTUAL PROPERTY AND AI

16. **In Re: OpenAI, Inc. Copyright Infringement Litigation**, No. 1:2025md03143 (S.D.N.Y. filed Apr. 11, 2025)
17. **Thomson Reuters Enter. Centre GmbH v. Ross Intelligence Inc.**, No. 20-cv-141 (D. Del. Feb. 11, 2025)

IV. ETHICS OPINIONS AND GUIDANCE

18. **NYC Bar Comm. on Prof'l Ethics, Formal Op. 2024-5** (Aug. 7, 2024), https://www.nycbar.org/wp-content/uploads/2024/08/20221329_GenerativeAILawPractice.pdf
19. **ABA Standing Comm. on Ethics & Prof'l Responsibility, Formal Op. 512** (July 29, 2024), https://www.americanbar.org/news/abanews/aba-news-archives/2024/07/aba-issues-first-ethics-guidance-ai-tools/
20. **NYSBA Task Force on Artificial Intelligence, Report & Recommendations** (Apr. 2024), https://nysba.org/wp-content/uploads/2022/03/2024-April-Report-and-Recommendations-of-the-Task-Force-on-Artificial-Intelligence.pdf

V. RULES AND STATUTES

21. **N.Y. Rules of Prof'l Conduct** (2025), https://www.nycourts.gov/LegacyPDFS/rules/jointappellate/NY-Rules-Prof-Conduct-1200.pdf
22. **Fed. R. Civ. P. 11**, https://www.law.cornell.edu/rules/frcp/rule_11
23. **Fed. R. Civ. P. 26**, https://www.law.cornell.edu/rules/frcp/rule_26
24. **Fed. R. Evid. 702** (as amended 2023), https://www.law.cornell.edu/rules/fre/rule_702
25. **Fed. R. Evid. 901**, https://www.law.cornell.edu/rules/fre/rule_901
26. **Fed. R. Evid. 502**, https://www.law.cornell.edu/rules/fre/rule_502
27. **Guide to N.Y. Evidence art. 9** (Authentication), https://www.nycourts.gov/judges/evidence/9-AUTHENTICITY/9.01_IN%20GENERAL.pdf
28. **N.Y. Civ. Rts. Law §§ 50, 51** (Right of Privacy and Publicity)
29. **N.Y. Civ. Rts. Law § 50-f** (Digital Replicas of Deceased Performers)

30. **NYC Local Law 144** (Automated Employment Decision Tools), https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page

VI. JUDICIAL POLICIES AND ORDERS

31. **New York State Unified Court System Interim Policy on the Use of AI** (effective Oct. 2025), https://www.reuters.com/legal/government/new-york-court-system-sets-rules-ai-use-by-judges-staff-2025-10-10/

32. **Proposed Rule 6(e) to the Rules of the Commercial Division** (22 NYCRR § 202.70) (Public Comment Sought June 11, 2025)

33. **Judge Brantley Starr Standing Order on AI** (N.D. Tex.), https://www.ropesgray.com/en/insights/alerts/2023/08/judges-guide-attorneys-on-ai-pitfalls-with-standing-orders

34. **Ropes & Gray AI Standing Orders Tracker**, https://www.ropesgray.com/en/sites/Artificial-Intelligence-Court-Order-Tracker/states/new-york

VII. TECHNICAL FRAMEWORKS AND STANDARDS

35. **NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0)** (2023), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

36. **NIST, Artificial Intelligence Risk Management Framework: Generative AI Profile** (2024), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

VIII. SCHOLARLY ARTICLES

37. Paul W. Grimm, Maura R. Grossman & Gordon V. Cormack, **Artificial Intelligence as Evidence**, 19 NW. J. TECH. & INTELL. PROP. 9 (2021), https://scholarlycommons.law.northwestern.edu/njtip/vol19/iss1/2/

38. Damien Charlotin, **AI Hallucination Cases Database** (Aug. 6, 2025), https://www.damiencharlotin.com/hallucinations/

39. Judge Xavier Rodriguez, **Artificial Intelligence (AI) and the Practice of Law**, 24 SEDONA CONF. J. 783 (2023)

IX. NEWS AND COMMENTARY

40. Cecily Mauran, **120 court cases have been caught with AI hallucinations, according to new database**, MASHABLE (May 27,

2025), https://mashable.com/article/over-120-court-cases-caught-ai-hallucinations-new-database

---

APPENDIX A: VERIFICATION CHECKLIST FOR AI-ASSISTED LEGAL RESEARCH

Use this checklist for EVERY case or authority suggested by AI:

☐ **Step 1: Retrieve the Official Source**

- Pull up the case in Westlaw, Lexis, or Google Scholar
- Do NOT rely on AI's summary or quote

☐ **Step 2: Confirm Basic Accuracy**

- Does the case exist?
- Is the citation correct (volume, reporter, page)?
- Is the court correct?
- Is the date correct?

☐ **Step 3: Read Relevant Portions**

- Read at minimum the sections AI claims support your argument
- Verify the case actually says what you claim it says
- Check that quotes are accurate and not taken out of context

☐ **Step 4: Run a Citator**

- Shepardize (Lexis) or KeyCite (Westlaw) the case
- Confirm the case is still good law
- Check for negative treatment

☐ **Step 5: Verify Jurisdiction and Precedential Value**

- Is this case from the correct jurisdiction?
- Is it binding or persuasive authority?
- Is it published or unpublished?

☐ **Step 6: Document Your Verification**

- Note in your research memo: "Verified [case name] via Westlaw on [date]"
- Keep citator report
- Save PDF of key cases

☐ **Step 7: Cite Only from Official Source**

- Copy quotes from the actual opinion, not from AI
- Use pinpoint citations to specific pages
- Follow Bluebook format

RED FLAGS - Do Not Cite if:

- You cannot find the case in official databases
- The citation format seems unusual or incorrect
- The case law seems too perfectly on-point (too good to be true)
- Key facts in the AI summary don't match the actual opinion
- The case has been overruled, disapproved, or questioned

APPENDIX B: SAMPLE FIRM AI POLICY (COMPREHENSIVE VERSION)

[FIRM NAME] ARTIFICIAL INTELLIGENCE USE POLICY

Effective Date: [DATE]
Last Revised: [DATE]
Policy Owner: [AI Committee / Managing Partner]

I. PURPOSE AND SCOPE

This Policy governs the use of Artificial Intelligence tools, including Generative AI and Large Language Models (LLMs), by all attorneys, staff, and contractors of [Firm Name]. The Policy is designed to ensure AI use complies with:

- New York Rules of Professional Conduct
- Federal Rules of Civil Procedure
- Court orders and local rules

- Client confidentiality obligations
- Professional standards of competence and diligence

II. DEFINITIONS

Artificial Intelligence (AI): Machine-based systems that learn patterns from data and make predictions or generate content.

Generative AI (GenAI): AI systems that create new content (text, images, audio) such as ChatGPT, Claude, Google Gemini.

Large Language Model (LLM): AI trained on text data to generate human-like written content.

Hallucination: AI generation of false or fabricated information presented as fact.

Public/Consumer AI: Free or low-cost AI tools available to the general public where user inputs may be used for training (e.g., free ChatGPT).

Enterprise AI: Commercial AI tools with contractual data protection where inputs are not used for training.

Legal AI: Purpose-built AI tools integrated with legal databases (e.g., Westlaw AI, Lexis+, CoCounsel).

III. APPROVED TOOLS

The Firm maintains a list of **Pre-Approved AI Tools**. Only tools on this list may be used with client-related work.

Current Approved Tools:

- Westlaw AI (research)
- Lexis+ AI (research)
- CoCounsel (research, document review)
- [Add other approved tools]

Limited-Use Tools (non-confidential work only):

- ChatGPT Enterprise [if applicable]
- Microsoft Copilot Enterprise [if applicable]

Prohibited: Free/public versions of ChatGPT, Claude, Gemini, or any AI tool whose terms of service permit use of inputs for model training.

Requesting New Tools: Any attorney wishing to use an AI tool not on the approved list must submit a request to the AI Committee including:

- Vendor name and tool description
- Intended use case
- Data handling and confidentiality provisions
- Cost

IV. PERMITTED USE CASES

AI tools may be used for:

1. **Legal Research**
    - Issue identification and brainstorming
    - Finding potentially relevant cases (subject to verification)
    - Generating research outlines
    - Understanding unfamiliar areas of law
2. **Legal Drafting**
    - First drafts of pleadings, motions, briefs
    - Contract and agreement drafting
    - Email and letter templates
    - Internal memoranda
    - Plain-language explanations for clients
3. **Document Review and Analysis**
    - E-discovery review using validated TAR protocols
    - Contract review and comparison
    - Due diligence document analysis
    - Summarization of depositions and transcripts (with page/line verification)
    - Identifying key provisions or clauses
4. **Case and Matter Management**

- Timeline and chronology creation
- Issue and task tracking
- Internal case summaries
- Deadline calendaring assistance
- Cross-referencing documents and testimony

5. **Client Communication**
    - Drafting client-friendly explanations of legal concepts
    - Status update letters
    - FAQ development
    - Educational materials

AI and the Law: The Use of Technology and Its Impact Upon Legal Ethics

A Comprehensive Guide for New York Practitioners

Introduction

The legal profession stands at a transformative crossroads. Artificial intelligence, particularly generative AI tools like ChatGPT, Claude, and Google's Gemini, has moved from theoretical possibility to daily reality in law offices across New York and beyond. These powerful technologies promise unprecedented efficiency in legal research, document drafting, discovery review, and client service. Yet they also present novel ethical challenges that demand careful navigation within our existing professional framework.

This article synthesizes current guidance from bar associations, judicial commentary, case law, and practical experience to provide New York attorneys with a comprehensive framework for the competent, ethical use of AI in legal practice. Drawing from over 569 documented cases involving AI hallucinations,[1] authoritative ethics opinions,[2] and emerging judicial policies,[3] we examine how traditional professional obligations—competence, confidentiality, candor, and supervision—apply with heightened urgency to these transformative tools.

Part I: Understanding Generative AI and the Hallucination Crisis

A. What is Generative AI?

To use AI ethically and effectively, attorneys must first understand what they are working with. The current wave of AI attention centers on Generative AI, specifically Large Language Models (LLMs). Unlike earlier legal technology—such as Boolean search systems or even predictive coding for e-discovery—LLMs like ChatGPT, Claude, and Google's Gemini operate on a fundamentally different principle: they are sophisticated pattern-matching systems that predict the most statistically likely next word in a sequence.[4]

This distinction is critical. LLMs have no concept of "truth" or "accuracy." They generate text that mimics human writing because they have been trained on massive datasets from the internet, including both accurate and inaccurate material. When an LLM generates a plausible-looking legal citation to a nonexistent case, it is not malfunctioning or lying—it is doing exactly what it was designed to do: producing text that follows the patterns it learned during training.

B. The Three Categories of AI Tools

Attorneys encounter three general categories of AI tools, each with distinct characteristics and risks:

1. Public, Consumer-Grade LLMs

Free or low-cost tools like ChatGPT (free version), Google Gemini, and Claude represent the first category. These systems typically: train on internet data up to a specific cutoff date; may use real-time browsing for current information; lack integration with verified legal databases; may use user inputs to train future versions of the model; and offer no citation verification mechanisms.[5]

The confidentiality implications are severe. Information entered into public AI systems could theoretically become part of training data, creating massive exposure under Rule 1.6 of the New York Rules of Professional Conduct.[6]

2. Enterprise or "Sandboxed" Versions

These are paid versions of consumer LLMs that contractually promise not to use client data for training and provide enhanced security features. While addressing some confidentiality concerns, the core technology remains the same—these tools still hallucinate with the same frequency as their public counterparts.[7]

3. Dedicated Legal AI Platforms

Tools like Westlaw's AI Research Assistant, Lexis+ AI, CoCounsel, and Harvey integrate LLMs with verified legal databases and add validation layers. These platforms attempt to "ground" AI responses in actual case law from their proprietary databases. However, even these purpose-built legal tools are not foolproof and can still generate fabricated citations.[8]

C. The Hallucination Crisis: Lessons from the Case Law

As of November 19, 2025, researcher Damien Charlotin's AI Hallucination Cases Database tracked 569 documented cases worldwide in which courts addressed generative AI hallucinations.[9] These cases provide invaluable guidance on what separates acceptable AI use from sanctionable misconduct.

1. Mata v. Avianca, Inc.: The Watershed Moment

The Southern District of New York's 2023 decision in *Mata v. Avianca, Inc.* marked the profession's first major reckoning with AI hallucinations.[10] Attorney Steven Schwartz used ChatGPT to research a personal injury case and cited six completely fabricated cases in his opposition to a motion to dismiss. When the court questioned the citations, Schwartz compounded his error by submitting fake judicial opinions generated by ChatGPT to "prove" the cases existed.

The court imposed severe sanctions: a $5,000 fine, mandatory notification to all judges whose names appeared in the fake cases, notification to the client, and public reprimand.[11] More significantly, the decision established the principle that while the initial hallucination represents a tool failure, the attorney's failure to verify constitutes professional misconduct.

2. The Expanding Universe of Sanctions Cases

The lessons from *Mata* have been reinforced repeatedly. In *Park v. Kim*, the Second Circuit referred an attorney to the grievance committee after filing AI-generated fake citations, emphasizing that verification is non-delegable.[12]

Jacquelyn Lacey v. State Farm General Insurance Co. sanctioned attorneys from the international firm K&L Gates for failing to disclose AI use and failing to verify citations, despite using supposedly reliable legal AI platforms (CoCounsel, Westlaw Precision, and Google Gemini). The Special Master struck their briefs and awarded $31,100 in fees to opposing counsel.[13]

In *Wadsworth v. Walmart*, a Wyoming federal court sanctioned lawyers from Morgan & Morgan—one of the nation's largest personal injury firms—after their proprietary in-house AI platform hallucinated cases. The drafting attorney had his pro hac vice status revoked and was fined $3,000.[14]

Two Colorado lawyers were fined $3,000 each in July 2025 for AI-generated fake citations in a defamation case.[15]

3. When Honesty Mitigates Sanctions

Not all hallucination cases result in maximum penalties. In *Hall v. Academy Charter School*, an Eastern District of New York magistrate judge declined to impose monetary sanctions despite three hallucinated citations.[16] The attorney had been dealing with her spouse's sudden death and, when the error was discovered, immediately took responsibility, explained the circumstances, and demonstrated that she had sought treatment for grief and taken bereavement leave.

Magistrate Judge Wicks found "extreme carelessness" but not "bad faith," distinguishing the case from *Mata* where the attorney doubled down on fabrications.[17] The critical lesson: transparency and candor when errors are discovered can mean the difference between a warning and career-damaging sanctions.

Part II: The Ethical Framework Under New York Law

A. No New Rules—Heightened Application

The good news for practitioners is that we do not need new ethics rules to govern AI use. The bad news is that all existing rules now apply with heightened scrutiny and carry greater risk of violation.

Three primary sources provide authoritative guidance: NYC Bar Association Formal Opinion 2024-5 (issued August 2024),[18] ABA Formal Opinion 512 (issued July 2024),[19] and NYSBA Task Force on AI Report and Recommendations (April 2024).[20] These opinions converge on a fundamental principle: the New York Rules of Professional Conduct are robust enough to govern AI use if attorneys apply them rigorously.

B. Rule 1.1: Competence and Diligence

1. The Duty to Understand Technology

Rule 1.1 requires attorneys to "provide competent representation to a client," which "requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation."[21] Comment 8 explicitly extends this obligation to technology, stating that competence requires "keeping abreast of changes in the law and its practice, *including the benefits and risks associated with relevant technology*."[22]

For most practice areas in 2025, AI constitutes "relevant technology." The NYC Bar Opinion 2024-5 elaborates that understanding AI's "benefits and risks" means: understanding how LLMs work at a functional level and why they hallucinate; knowing the difference between public, enterprise, and legal-specific AI tools; understanding what verification steps are required for different use cases; recognizing algorithmic bias and its potential impact; and staying current as technology evolves.[23]

An attorney who uses AI without this foundational understanding may violate Rule 1.1 before reaching the question of whether specific outputs were verified.

2. The CLE Obligation

While New York does not yet mandate AI-specific CLE, attorneys must complete continuing legal education in technology (as part of the broader professional practice requirement) and ethics. Understanding how AI implicates professional conduct rules satisfies the ethics requirement.

Moreover, Rule 1.1's competence mandate effectively requires self-education about tools one intends to use in practice.

3. Algorithmic Bias as a Competence Issue

AI systems replicate biases present in their training data. Attorneys using AI for tasks with potential discriminatory impact—employment screening, contract terms affecting protected classes, or criminal justice applications—must understand this limitation.[24]

New York City Local Law 144, which regulates "Automated Employment Decision Tools," illustrates this concern. Employers using AI for hiring or promotion decisions must conduct annual bias audits and notify candidates.[25] Attorneys advising on such systems or challenging their use must understand algorithmic bias as part of their competence obligation.

C. Rule 1.6: Confidentiality

1. The Core Prohibition

Rule 1.6(a) provides: "A lawyer shall not knowingly reveal confidential information... or use such information to the disadvantage of a client or for the advantage of the lawyer or a third person, unless..."[26] Comment 18 extends this to require "reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client."[27]

NYC Bar Opinion 2024-5 applies these principles to AI unambiguously: "Without client consent, a lawyer must not input confidential client information into any Generative AI system that will share the inputted confidential information with third parties... [or] that lacks adequate confidentiality and security protections."[28]

2. The Public AI Prohibition

Inputting client confidential information into free ChatGPT, public Claude, or similar tools where the terms of service permit use of data for training constitutes a per se violation of Rule 1.6, absent informed client consent.[29]

Even with client consent, attorneys must: rigorously vet vendor terms of service; understand where data is stored and for how long; confirm whether data may be used for model training; consider de-identifying information before input; and document the consent and precautions taken.[30]

3. "Closed" Systems and Internal Sharing Risks

Even AI systems that keep data within the firm's protected databases ("closed systems") present confidentiality risks. The NYC Bar cautions: "Even with closed systems, a lawyer must take care that confidential information is not improperly shared with other persons at or clients of the same law firm, including persons who are prohibited access to the information because of an ethical wall."[31]

Firms must implement access controls ensuring that AI systems respect conflict screens and ethical walls.

4. Discovery and Protective Orders

When using AI to review opposing parties' documents produced under a protective order, attorneys must ensure the AI platform's data handling complies with the order's terms. Uploading confidential documents to a public AI tool could violate both Rule 1.6 and the protective order.[32]

Prudent practice includes negotiating AI-specific language in FRCP Rule 26(f) conferences and protective orders.

D. Rule 3.3: Candor to the Tribunal

1. The Foundation of Trust

Rule 3.3(a) prohibits making false statements of fact or law to a tribunal and requires disclosure of controlling adverse legal authority.[33] This rule, combined with Federal Rule of Civil Procedure 11 and New York's 22 NYCRR § 130-1.1 (frivolous conduct), forms the basis for sanctions in AI hallucination cases.

FRCP Rule 11(b) requires attorneys to certify that legal contentions "are warranted by existing law or by a nonfrivolous argument for extending, modifying, or reversing existing law."[34] New York's Section 130-1.1(c) similarly defines conduct as frivolous if "it is completely without merit in law and cannot be supported by a reasonable argument for an extension, modification or reversal of existing law."[35]

Citing a hallucinated case violates both standards because the case does not constitute "existing law."

2. The Mandatory Verification Protocol

Every court that has addressed AI hallucinations has emphasized the same requirement: attorneys must independently verify the existence and accuracy of every case cited, regardless of the AI tool used to find it.[36]

The verification protocol established by case law requires:

1. 1. Retrieve the actual opinion from Westlaw, Lexis, Google Scholar, or official court reporters
2. 2. Confirm basic accuracy: case exists, citation is correct, court and date match
3. 3. Read relevant portions to verify the case says what you claim
4. 4. Run a citator (Shepard's or KeyCite) to confirm the case remains good law
5. 5. Check for negative treatment: overruling, distinguishing, or questioning
6. 6. Copy quotes from the official source, never from AI output
7. 7. Document your verification in a research memo or log

Skipping any step may constitute a Rule 11 violation and frivolous conduct under Section 130-1.1.[37]

3. The Importance of Response When Errors Are Discovered

The sanctions cases reveal a clear pattern: how attorneys respond to discovered hallucinations matters as much as the initial error.

Maximum sanctions result from: doubling down on false citations (Mata); submitting fabricated documents to "prove" fake cases exist (Mata); and failing to disclose AI use when questioned (Lacey).

Reduced or waived sanctions result from: immediate acknowledgment of error (Hall); prompt correction filings; candid explanation of what went wrong; and demonstrated implementation of new safeguards.[38]

E. Rules 1.4 and 1.5: Communication and Fees

1. Disclosure of AI Use

Rule 1.4 requires attorneys to "keep the client reasonably informed about the status of the matter" and "explain a matter to the extent reasonably necessary to permit the client to make informed decisions."[39] While neither the NYC Bar nor ABA opinions mandate disclosure of AI use in every engagement letter, best practice strongly favors it.

Recommended engagement letter language includes:[40]

> "Our firm uses carefully vetted artificial intelligence tools to improve the quality and efficiency of our legal services. These tools may assist with tasks such as legal research, document

> review, contract analysis, and drafting. We do not input your confidential information into public AI systems. All AI-assisted work is supervised by experienced attorneys who review, verify, and take full professional responsibility for the work product."

2. Fee Considerations

Rule 1.5 prohibits unreasonable fees.[41] ABA Formal Opinion 512 addresses AI's impact on billing:[42]

- Attorneys may charge for time spent using AI (drafting prompts, reviewing output, editing)
- Attorneys may not charge for time saved by AI (billing three hours when AI helped finish in one)
- Attorneys may not charge clients for time spent learning to use AI tools
- The basis for fees must be communicated before commencing representation

The Opinion emphasizes: "the lawyer who has agreed to bill on the basis of hours expended does not fulfill her ethical duty if she bills the client for more time than she has actually expended on the client's behalf."[43]

F. Rules 5.1, 5.2, and 5.3: Supervisory Responsibilities

1. Partners' Duties

Rule 5.1 makes partners and supervisory attorneys responsible for ensuring that subordinate lawyers comply with professional conduct rules.[44] For AI, this means: establishing firm-wide AI policies; providing mandatory training before attorneys use AI tools; implementing verification protocols; monitoring compliance through file reviews; and taking remedial action when violations occur.

A managing partner cannot claim ignorance if associates file hallucinated cases. The duty to supervise extends to ensuring proper AI use.

2. Nonlawyer Assistance and Vendor Oversight

Rule 5.3 applies to "a nonlawyer employed or retained by or associated with a lawyer," requiring reasonable efforts to ensure such persons' conduct is compatible with the lawyer's professional obligations.[45]

AI vendors and platforms constitute "nonlawyer assistance" subject to Rule 5.3. This requires: due diligence before selecting AI vendors; reviewing and understanding terms of service; ensuring contractual protections for confidentiality; monitoring vendor practices and updates; and terminating relationships if vendors fail to maintain adequate safeguards.[46]

3. Subordinate Attorneys' Responsibilities

Rule 5.2 provides limited protection for subordinate lawyers following supervisory attorneys' instructions, but only when the instruction reflects "a reasonable resolution of an arguable question of professional duty."[47]

If a supervising attorney instructs an associate to cite cases without verification, Rule 5.2 provides no safe harbor—the instruction violates Rule 3.3 and Rule 11, which are not "arguable questions." Associates retain independent responsibility to verify citations regardless of time pressure or supervisory instructions.

Part III: Practical Workflows for AI Use in Legal Practice

A. Legal Research

Research represents both AI's greatest promise and its most publicized peril. The workflow that emerges from ethics opinions and case law is clear:

Phase 1: AI-Assisted Issue Identification

- Use AI to brainstorm legal theories and issues
- Generate outlines of potential arguments
- Identify areas of law requiring research
- Understand unfamiliar legal concepts

Phase 2: Initial Authority Location

- Use AI (preferably legal-specific platforms) to identify potentially relevant cases and statutes
- Generate lists of authorities to investigate
- Create initial research roadmaps

Phase 3: Mandatory Human Verification (The Critical Step)

- Pull every AI-suggested case from official sources
- Read the relevant portions of each opinion
- Confirm the case supports the proposition cited
- Run citators to check for negative treatment
- Verify quotes against the official text
- Document verification in research memos

Phase 4: Analysis and Synthesis

- Apply professional judgment to evaluate authorities
- Craft arguments based on verified law
- Consider strategic implications
- Draft work product citing only verified sources

This workflow acknowledges AI's value in accelerating the initial stages while maintaining the professional's gatekeeper role for verification and judgment.[48]

B. Document Drafting

AI excels at generating first drafts of pleadings, contracts, and memoranda. The key is treating all AI output as preliminary.

Pleadings and Motions:

- Use AI to create initial drafts from fact patterns
- Review for accuracy, completeness, and tone

- Verify all factual allegations against source documents
- Independently verify all legal citations
- Edit for strategy and jurisdiction-specific requirements
- Apply professional judgment to argument selection

Contracts:

- Generate initial drafts from templates and deal terms
- Review against client objectives and risk tolerance
- Verify boilerplate language is appropriate for jurisdiction
- Customize for specific transaction context
- Negotiate terms based on business considerations AI cannot assess
- Never rely solely on AI for complex or unusual provisions[49]

C. E-Discovery and Technology-Assisted Review

E-discovery represents one area where AI has a longer, more established track record than generative AI research tools.

1. Technology-Assisted Review (TAR) and Predictive Coding

The Southern District of New York led the way in validating TAR with decisions in *Da Silva Moore v. Publicis Groupe*,[50] *Rio Tinto PLC v. Vale, S.A.*,[51] and *Hyles v. New York*.[52] These cases establish that TAR is not only acceptable but often more accurate and proportional than manual review or keyword searching for large document sets.

New York state courts have followed this approach. The Uniform Civil Rules for the Supreme Court and County Court encourage parties to use "the most efficient means to review documents, including electronically stored information," expressly recognizing that such means "may include technology-assisted review, including predictive coding, in appropriate cases."[53]

The Commercial Division has adopted identical language.[54]

2. Defensible TAR Workflow

A defensible TAR protocol includes:

Planning Phase:

- Meet and confer with opposing counsel under FRCP Rule 26(f)
- Agree on scope, methodology, and success metrics
- Consider requesting a Rule 502(d) order protecting against privilege waiver[55]

Validation Phase:

- Test model on separate validation set
- Measure precision (percentage of AI-coded responsive documents that are actually responsive)
- Measure recall (percentage of all responsive documents that AI identified)
- Set acceptable thresholds (typically 75%+ recall)
- Document results[56]

Key Principles:

- Cooperation: Discuss TAR with opposing counsel early
- Transparency: Share methodology and validation metrics
- Proportionality: TAR is appropriate when manual review would be cost-prohibitive
- Documentation: Maintain detailed records of all decisions and results[57]

3. Generative AI for Transcript and Document Summarization

Newer generative AI tools can summarize depositions, contracts, and discovery documents. Critical limitations apply:

- Use summaries as finding aids, not as source material
- All quotes must be verified against official transcripts or documents
- Citations must reference original sources (page and line numbers for depositions)
- Never cite "AI Summary" in court filings
- Review complete documents for context AI may miss

A recent Eastern District of New York case illustrates the risk: an attorney submitted a brief quoting deposition testimony but citing only "AI Summary" instead of transcript page and line numbers. The court ordered an amended brief with proper citations, forcing the attorney to read the entire transcript anyway—negating the time savings.[58]

The lesson: AI can accelerate the process of finding relevant content, but cannot replace the requirement to cite official sources.

Part IV: Evidence and Authentication in the Age of AI

A. The Authentication Challenge

AI's ability to create photorealistic images, convincing audio, and plausible documents creates unprecedented authentication challenges.

1. Multi-Method Authentication

Courts are increasingly requiring "multi-method authentication" for digital evidence susceptible to AI manipulation. A 2024 New York Surrogate's Court decision, *Matter of Weber*, exemplifies this trend, requiring technical, testimonial, and circumstantial evidence to authenticate digital materials.[64]

Under Article 9 of the Guide to New York Evidence—and Federal Rule of Evidence 901—evidence must be authenticated by "evidence sufficient to support a finding that the offered evidence is what the proponent claims it is."[65]

For AI-susceptible evidence, satisfying this standard now requires: Technical Evidence (digital metadata, hash values, blockchain timestamps, forensic analysis, expert examination); Testimonial Evidence (witness testimony, expert opinion, chain of custody); and Circumstantial Evidence (consistency with other verified evidence, distinctive characteristics, contextual corroboration).[66]

B. Expert Testimony Involving AI

Federal Rule of Evidence 702, as amended in 2023, requires expert testimony to be based on "sufficient facts or data," to be "the product of reliable principles and methods," and to demonstrate that "the expert reliably applied those principles and methods to the facts."[68]

New York continues to apply the Frye 'general acceptance' test while incorporating Daubert-style reliability considerations through cases like Parker v. Mobil Oil Corp.[69]

1. When Experts Use AI

The cautionary tale is *Kohls v. Ellison*, a January 2025 Minnesota federal case.[70] The state's expert submitted a declaration citing two academic articles about AI and deepfakes. Opposing counsel checked the citations—both were fabricated. The expert had used GPT-4o to research without verification.

The court excluded the expert's entire declaration and testimony, finding that reliance on unverified AI output "shattered his credibility" and rendered his methodology unreliable under Rule 702.[71]

2. The Grimm-Grossman-Cormack Framework

Legal scholars Paul Grimm, Maura Grossman, and Gordon Cormack propose a comprehensive framework for evaluating AI as evidence, requiring parties offering AI-generated evidence to demonstrate: (1) System Identification, (2) Data Specification, (3) Methodological Explanation, (4) Validation, (5) Robustness, and (6) Proper Application.[72]

Part V: Governance and Implementation

A. The NIST AI Risk Management Framework

The National Institute of Standards and Technology (NIST) AI Risk Management Framework provides an organizational structure for AI governance.[74]

Law firms can adapt this framework through four key functions: Govern, Map, Measure, and Manage.[75]

B. Vendor Due Diligence

Attorneys must vet AI vendors before using their tools with client information. A comprehensive vendor due diligence checklist should address: Data Handling, Confidentiality, Security, Technical Capability, and Support and Training.[76]

C. Firm AI Policies

Every firm—from solo practitioners to large partnerships—should have a written AI policy. Essential Policy Components include: Approved Tools List, Training Requirements, Verification Protocols, Data Classification, Client Communication, Incident Reporting, and Supervisory Responsibilities.[77]

D. Judicial AI Policies

Courts face their own AI challenges. The New York Unified Court System's Interim Policy on AI Use, effective October 2025, provides a model for judicial chambers with key provisions about approved tools, training, confidentiality protection, human decision-making, transparency, and treating documents as confidential.[78]

Several federal judges have issued standing orders requiring attorney certification of AI use and verification.[79]

The Commercial Division has proposed Rule 6(e) that would mandate disclosure and certification for AI-assisted filings.[80]

Part VI: Access to Justice and Economic Implications

A. The Economic Disparity Problem

Large firms' investment in premium AI tools—spending millions annually on platforms like Harvey, CoCounsel, and enterprise versions of Westlaw AI—creates potential competitive advantages. Solo practitioners and small firms may feel priced out of the AI revolution.

However, this concern must be balanced against countervailing factors: lower-cost options exist; AI democratizes capability more than it concentrates it; and the core value remains human judgment.[81] [82]

B. Expanding Access to Justice

Pro se litigants are increasingly using AI tools to navigate the legal system. A 2025 NBC News article profiled a woman who used ChatGPT and Perplexity AI to successfully appeal her eviction after losing at trial with a lawyer, describing AI as "like having God responding to her questions."[83]

Legal aid organizations are exploring how AI can help them serve more clients with limited resources. If AI enables a legal aid attorney to handle twice as many cases competently, the access-to-justice impact could be profound.

Researcher Damien Charlotin notes that courts are more likely to issue warnings rather than sanctions when pro se litigants submit AI-generated fake citations, recognizing they lack attorneys' professional obligations.[84]

Regulatory Innovation

Some jurisdictions are experimenting with new models: Utah's Regulatory Sandbox (launched in 2020, narrowed in 2024 to focus on demonstrably underserved markets); and Paraprofessional Licensing programs in some states.[85]

Part VII: The Road Forward

A. Emerging Court Rules

The regulatory landscape is evolving rapidly. Suffolk County Surrogate's Court implemented a protocol effective December 1, 2025, requiring certification of AI use in all filings containing legal citations. The protocol establishes that reliance on hallucinated citations constitutes prima facie frivolous conduct under 22 NYCRR § 130-1.1—applying equally to lawyers and self-represented parties.[86]

Expect similar rules to proliferate across New York courts and federal districts.

B. Continuing Evolution of Ethics Guidance

Bar associations will continue issuing ethics opinions as novel issues arise: AI use in mediations and arbitrations; conflicts of interest when AI trained on confidential data; unauthorized practice of law by AI systems; attorney advertising using AI; and competence requirements as AI capabilities expand.

C. Technological Advancement

AI systems will continue improving: better citation verification; reduced hallucination rates; improved explanation of reasoning; enhanced integration with legal databases; and greater transparency about limitations.

However, the fundamental obligation remains unchanged: human attorneys bear ultimate responsibility for all work product, regardless of AI involvement.

D. The Profession's Choice

The legal profession successfully navigated previous technological revolutions—from typewriters to word processors, from law libraries to Westlaw, from fax machines to email. Each transition preserved core professional values while adapting to new tools.

AI presents the same challenge: harness the technology's benefits while maintaining competence, protecting confidentiality, ensuring candor, and exercising judgment. The choice is not whether to use AI but how to use it responsibly.

Conclusion: Five Core Principles

This comprehensive examination of AI in legal practice distills to five essential principles:

1. Understand the Technology

Know what AI is, how it works, and why it fails. This understanding is not optional—it is a competence requirement under Rule 1.1.

2. Verify Relentlessly

Never cite unconfirmed sources. Treat all AI output as a draft requiring human verification. This is the lesson of 569 hallucination cases and the foundation of Rule 3.3 compliance.

3. Protect Confidentiality

Vet tools carefully. Never use public AI with client confidential information absent informed consent and adequate safeguards. This is the mandate of Rule 1.6.

4. Maintain Human Judgment

AI assists; lawyers decide. AI cannot replace professional judgment built on experience, creativity, empathy, and wisdom. This is what clients pay for and what justice requires.

5. Stay Transparent

Disclose AI use appropriately. Admit errors promptly. Honesty matters more than perfection. The sanctions cases uniformly demonstrate this principle.

These five principles, grounded in the New York Rules of Professional Conduct and Federal Rules of Civil Procedure, will guide practitioners through whatever technological changes come next. AI is transforming legal practice as profoundly as any technology in generations. We have the opportunity—and the professional obligation—to shape that transformation to serve justice, expand access, improve quality, and preserve professional values.

The tools are powerful. The obligations are clear. The choice is ours.

Endnotes

[1] Damien Charlotin, *AI Hallucination Cases Database* (Nov. 19, 2025), https://www.damiencharlotin.com/hallucinations/.

[2] NYC Bar Comm. on Prof'l Ethics, Formal Op. 2024-5 (Aug. 7, 2024), https://www.nycbar.org/reports/formal-opinion-2024-5-generative-ai-in-the-practice-of-law/; ABA Standing Comm. on Ethics & Prof'l Responsibility, Formal Op. 512 (July 29, 2024), https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf.

[3] New York State Unified Court System Interim Policy on the Use of AI (effective Oct. 2025), https://www.reuters.com/legal/government/new-york-court-system-sets-rules-ai-use-by-judges-staff-2025-10-10/.

[4] The transcript uses dialogue to explain: "LLMs are sophisticated pattern-matching systems that predict the most statistically likely next word in a sequence... They have no concept of 'truth' or 'accuracy.'"

[5] *See* OpenAI, Model Release Notes, https://help.openai.com/en/articles/9624314-model-release-notes (extending GPT-4o training data cutoff from Nov. 2023 to June 2024); Gemini 2.5 Flash & 2.5 Flash Image Model Card 2 (Google 2025), https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-2-5-Flash-Model-Card.pdf (knowledge cutoff Jan. 2025). Both tools can access real-time information through web browsing, but neither has direct access to proprietary legal databases like Westlaw or Lexis.

[6] N.Y. Rules of Prof'l Conduct r. 1.6(a) (2025).

[7] Enterprise versions of ChatGPT, Claude, and similar tools promise not to use client data for training and provide enhanced security, but the underlying large language model technology remains probabilistic and subject to hallucination.

[8] *See Jacquelyn "Jackie" Lacey v. State Farm Gen. Ins. Co.*, No. cv-24-5205-fmo (C.D. Cal. May 5, 2025) (sanctioning K&L Gates attorneys who used CoCounsel, Westlaw Precision, and Google Gemini yet still filed hallucinated cases).

[9] Charlotin, *supra* note 1.

[10] *Mata v. Avianca, Inc.*, No. 22-cv-1461, 2023 WL 4114965 (S.D.N.Y. June 22, 2023).

[11] *Id.*

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

[13] *Lacey*, No. cv-24-5205-fmo (C.D. Cal. May 5, 2025).

[14] *Wadsworth v. Walmart Inc.*, 348 F.R.D. 489 (D. Wyo. 2025).

[15] *Coomer v. Lindell*, No. 22-cv-01129-NYW-SBP, 2025 WL 1865282 (D. Colo. July 25, 2025).

[16] *Hall v. The Academy Charter Sch.*, No. 2:24-cv-08630-JMW (E.D.N.Y. Aug. 7, 2025).

[17] *Id.*

[18] NYC Bar Comm. on Prof'l Ethics, Formal Op. 2024-5, *supra* note 2.

[19] ABA Standing Comm. on Ethics & Prof'l Responsibility, Formal Op. 512, *supra* note 2.

[20] NYSBA Task Force on Artificial Intelligence, Report & Recommendations (Apr. 2024), https://nysba.org/wp-content/uploads/2022/03/2024-April-Report-and-Recommendations-of-the-Task-Force-on-Artificial-Intelligence.pdf.

[21] N.Y. Rules of Prof'l Conduct r. 1.1(a).

[22] *Id.* r. 1.1 cmt. 8 (emphasis added).

[23] NYC Bar Formal Op. 2024-5, *supra* note 2.

[24] *See* IBM, AI Alignment, https://www.ibm.com/think/topics/ai-alignment (explaining the "alignment problem" in AI safety).

[25] NYC Local Law 144 (Automated Employment Decision Tools), https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page.

[26] N.Y. Rules of Prof'l Conduct r. 1.6(a).

[27] *Id.* r. 1.6 cmt. 18.

[28] NYC Bar Formal Op. 2024-5, *supra* note 2.

[29] *Id.*

[30] *Id.*

[31] *Id.*

[32] *See* 22 NYCRR § 202.70 (Commercial Division Rules); Commercial Division Rule 11-g (protective orders).

[33] N.Y. Rules of Prof'l Conduct r. 3.3(a).

[34] Fed. R. Civ. P. 11(b).

[35] 22 NYCRR § 130-1.1(c)(1).

[36] *See, e.g., Mata*, 2023 WL 4114965; Park, 91 F.4th 610; Wadsworth, 348 F.R.D. 489.

[37] This protocol synthesizes requirements from multiple sanctions cases and ethics opinions.

[38] *Compare Mata*, 2023 WL 4114965 (maximum sanctions where attorney submitted fabricated opinions to prove fake cases), with Hall, No. 2:24-cv-08630-JMW (declining monetary sanctions where attorney immediately acknowledged error and explained circumstances).

[39] N.Y. Rules of Prof'l Conduct r. 1.4(a), (b).

[40] Sample language adapted from ABA Formal Op. 512 and NYC Bar Formal Op. 2024-5.

[41] N.Y. Rules of Prof'l Conduct r. 1.5(a).

[42] ABA Formal Op. 512, *supra* note 2.

[43] *Id.* (citing ABA Comm. on Ethics & Prof'l Responsibility, Formal Op. 93-379, at 6).

[44] N.Y. Rules of Prof'l Conduct r. 5.1.

[45] *Id.* r. 5.3.

[46] NYC Bar Formal Op. 2024-5, *supra* note 2.

[47] N.Y. Rules of Prof'l Conduct r. 5.2(b).

[48] This workflow synthesizes guidance from NYC Bar Formal Op. 2024-5 and ABA Formal Op. 512.

[49] *See* ABA Formal Op. 512 (discussing AI use in transactional practice).

[50] *Da Silva Moore v. Publicis Groupe*, 287 F.R.D. 182 (S.D.N.Y. 2012).

[51] *Rio Tinto PLC v. Vale, S.A.*, 306 F.R.D. 125 (S.D.N.Y. 2015).

[52] *Hyles v. City of New York*, No. 10-cv-3119, 2016 WL 4077114 (S.D.N.Y. Aug. 1, 2016).

[53] 22 NYCRR § 202.20-c(e).

[54] Commercial Division Rule 11-c(f), 22 NYCRR § 202.70.

[55] Fed. R. Evid. 502(d).

[56] *See Da Silva Moore*, 287 F.R.D. at 192-93 (outlining validation methodology).

[57] *See Rio Tinto*, 306 F.R.D. at 127-28 (discussing cooperation and transparency requirements).

[58] This example is drawn from the transcript's discussion of judicial experiences with AI-generated summaries.

[59] Commercial Division Rule 11-g, 22 NYCRR § 202.70.

[60] Fed. R. Evid. 502(d).

[61] *See* N.Y. C.P.L.R. 3101(c) (attorney-client privilege); id. 3101(d)(2) (work product); id. 4503 (physician-patient privilege); id. 4548 (accountant-client privilege).

[62] Commercial Division Rule 11-c(g), 22 NYCRR § 202.70.

[63] *See AFA Protective Sys., Inc. v. City of New York*, 13 A.D.3d 564 (2d Dep't 2004) (applying multi-factor test for inadvertent disclosure); Kevin Schlosser, Inadvertent Disclosure of Privileged Materials in the E-Age, N.Y. L.J. (2006).

[64] *Matter of Weber (As Trustee of Michael S. Weber Trust)*, 2024 NY Slip Op 24258 (Sur. Ct. 2024).

[65] Guide to N.Y. Evidence art. 9.01; *accord* Fed. R. Evid. 901(a).

[66] This framework synthesizes authentication requirements from *Matter of Weber* and Guide to N.Y. Evidence art. 9.

[67] The transcript references this case without specific citation; the example illustrates emerging authentication challenges.

[68] Fed. R. Evid. 702 (as amended 2023).

[69] *See Frye v. United States*, 293 F. 1013 (D.C. Cir. 1923) (general-acceptance test); Daubert v. Merrell Dow Pharm., Inc., 509 U.S. 579 (1993) (reliability factors). New York continues to apply Frye but has incorporated reliability considerations similar to Daubert. See Parker v. Mobil Oil Corp., 7 N.Y.3d 434 (2006).

[70] *Kohls v. Ellison*, No. 24-cv-03754 (D. Minn. Jan. 10, 2025).

[71] *Id.*

[72] Paul W. Grimm, Maura R. Grossman & Gordon V. Cormack, *Artificial Intelligence as Evidence*, 19 Nw. J. Tech. & Intell. Prop. 9 (2021), https://scholarlycommons.law.northwestern.edu/njtip/vol19/iss1/2/.

[73] *See* transcript discussion of algorithmic bias in criminal justice applications.

[74] NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (2023), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

[75] *Id.; see also* NIST, Artificial Intelligence Risk Management Framework: Generative AI Profile (2024), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf.

[76] This checklist synthesizes vendor vetting guidance from NYC Bar Formal Op. 2024-5 and best practices discussed in the transcript.

[77] Sample policy components adapted from the transcript's comprehensive firm AI policy template.

[78] New York State Unified Court System Interim Policy on the Use of AI, *supra* note 3.

[79] *See* Ropes & Gray AI Standing Orders Tracker, https://www.ropesgray.com/en/sites/Artificial-Intelligence-Court-Order-Tracker/states/new-york (tracking federal and state court AI orders).

[80] Proposed Rule 6(e) to the Rules of the Commercial Division (22 NYCRR § 202.70) (Public Comment Sought June 11, 2025).

[81] The transcript discusses pricing tiers ranging from $0-25/month (consumer tools) to $30-100/month (mid-tier enterprise) to $100-1000+/month (premium legal AI).

[82] *See* transcript dialogue between Tom (solo practitioner) and Michael (large firm associate) discussing competitive dynamics.

[83] Cecily Mauran, *Pro Se Litigant Uses AI to Successfully Appeal Eviction*, NBC News (Oct. 8, 2025), https://www.nbcnews.com/tech/innovation/ai-chatgpt-court-law-legal-lawyer-self-represent-pro-se-attorney-rcna230401.

[84] Charlotin, *supra* note 1 (noting differential treatment of pro se litigants in hallucination cases).

[85] The Utah Supreme Court narrowed its regulatory sandbox in 2024 to focus on demonstrably underserved markets after initially launching in 2020 with broader parameters.

[86] Suffolk County Surrogate's Court Protocol on AI Use (effective Dec. 1, 2025).

Author's Note

This article draws heavily from the referenced CLE program transcript, which presents the material in dialogue format among fictional characters representing different perspectives within the legal profession. The substantive legal analysis, case citations, and ethical guidance are drawn from actual sources and represent current law and best practices as of November 2025. Practitioners should monitor continuing developments in this rapidly evolving area and consult current ethics opinions and court rules in their jurisdiction.

#1

You are a licensed New York attorney, practicing law professor and legal scholar with a specialty in the ethical and productive use and impact of artificial intelligence in the legal profession. You are working with a New York Supreme Court Justice to present a continuing law education program (CLE) in New York to both lawyers and judges.

You have been requested to prepare a course outline and course materials for a two hour New York CLE program for lawyers and judges for a program titled "**AI and the law. The use of technology and its impact upon legal ethics.**"

The aim of the program is to briefly educate the participants on artificial intelligence and how it works and then more fully explore how the artificial intelligence can be used to ethically assist lawyers and judges do legal research, issue spotting and potential legal approaches, legal drafting, review document discovery, deposition and trial transcripts.

The program should identify the risks of failing to ethically make use of AI and provide practice points to avoid or minimize risks and maximize the value of AI.

Current published New York and federal case law addressing the use of artificial intelligence as well as recognized legal treatises and articles from bar associations as well as ethical opinions should be considered and incorporated into the program material. Additionally, the judge has provided two PowerPoint files which you should also consider and incorporate into the materials.

All sources need to be fully cited and links to the citation must be provided.

1A

Using the outline and material you have prepared prepare a detailed document for addressing all of the points in the outline suitable for publication in a law review or a legal guide book

#2

You are a licensed New York attorney, practicing law professor and legal scholar with a specialty in the ethical and productive use and impact of artificial intelligence in the legal profession. You are working with a New York Supreme Court Justice to present a continuing law education program (CLE) in New York to both lawyers and judges.

Your research assistants have prepared the attached a course outline and course materials for a New York CLE program for lawyers and judges for a program titled "AI and the law. The use of technology and its impact upon legal ethics."

The aim of the program is to briefly educate the participants on artificial intelligence and how it works and then more fully explore how the artificial intelligence can be used to ethically assist lawyers and judges do legal research, issue spotting and potential legal approaches, legal drafting, review document discovery, deposition, and trial transcripts.

The research assistants' assignment was to develop materials and a publishable article identifying the risks of failing to ethically make use of AI and provide practice points to avoid or minimize risks and maximize the value of AI.

They were directed to make use of current published New York and federal case law addressing the use of artificial intelligence as well as recognized legal treatises and articles from bar associations as well as ethical opinions and incorporate the materials into the program material. Additionally, the judge has provided two PowerPoint files which they were told they could also consider and incorporate into the materials.

All sources were to be fully cited and links to the citation must be provided. You are to review the two draft articles they prepared and the outline and doing your own independent research using the same parameters that the research assistants were to use prepare a comprehensive publishable article that addresses "AI and the law. The use of technology and its impact upon legal ethics" under New York Law

****

#3

You are a licensed New York attorney, practicing law professor and legal scholar with a specialty in Artificial Intelligence (AI) and the ethical and productive use and impact of artificial intelligence (and its hazards) in the legal profession. You are working with a team of legal research assistants and script writers to prepare a multi part script as the way to present the information and materials to the participants Judges, lawyers and law students in a continuing legal education program (CLE) with 1 credit for Ethics and 1 credit for technology in New York.

The skit is to address the following order a general overview of what AI is and how it works; how AI can assist lawyers and law firms in their various practice areas including transactional work for real estate and corporate, trusts and estates and all components of litigation. Connection with this section diskette should address the disparity in economics between law firms in the justice system and their means of access to AI appropriate for legal practice. The skit should briefly address issues and problems that clients of lawyers and law firms may address to them as result of their use of artificial intelligence particularly as it relates to fair use of the intellectual property. The program should then address the ethical considerations and issues that lawyers and law firms as well as judges need to address in their use of artificial intelligence.

There will be 5 key "players" in the skit.

Tom, a seasoned solo practitioner who specializes in intellectual property.

Eric, a managing partner in a 12 person law firm who specializes in corporate and commercial transactional work and litigation.

Michael, a mid-level commercial litigation associate at a midsize law firm.

Adriana, the principal law clerk to the New York State Supreme Court Justice.
MAB, a junior associate at a small firm who has done both litigation and corporate transactional work.

Howard, a law student with a prior career in computer science and programming.

An additional participant in the program is Justice D, a sitting New York State Supreme Court Justice in the Commercial Division who has extensively lectured the judiciary and the bar about the use of AI. Justice D will not be a direct actor /participant in the skit. Rather, between acts or scenes of the skit Justice D will provide high level summaries of the points just addressed and give practice points and cautions as appropriate from the "bench".

The entire CLE program is to run two hours. The skit should take up most of the two-hours. Opening introductions, closing statements, a few minutes of time between the various acts or scenes for Justice D to make his points and time for audience questions should be provided for.

The skit should be entertaining, engaging and slightly humorous. The skit should not be a boring recitation of cases, legal statutes, or rules. The information and education is to be incorporated in a dialogue between the “actors” that keeps the audience’s attention as they are educated.

It has been discussed that one way to structure and present the skit is as a meeting of various committees of a Bar Association including the technical and AI committee and ethics committee who are trying to put together an outline and guide to the members of the Bar on “AI and the law: The use of technology and its impact upon legal ethics”. However, as the script writer you have discretion to change the purpose of the “actors” gathering to better accomplish the purposes of presenting the materials in a meaningful and cogent manner as set forth above.

In connection with the skit, you may also wish to prepare engaging examples of some of AI’s abilities such as – but not limited to - AI generated podcasts or audiovisual elements that highlight or address AI’s abilities, problems, issues and uses.

Your research assistants have prepared the attached course outline and course materials for a New York CLE program for lawyers and judges for a program titled “AI and the law. The use of technology and its impact upon legal ethics.”

The materials thus far prepared were designed as a draft of a printed guide to briefly educate the participants on artificial intelligence and how it works as well as to explore how the artificial intelligence can be used to ethically assist lawyers and judges do legal research, issue spotting and identify potential legal approaches, legal drafting of agreements and memorandums and court decisions, review document discovery, deposition, and trial transcripts.

The research assistants' assignment was to develop materials and a publishable article that would also identify the risks of failing to ethically make use of AI and provide practice points to avoid or minimize risks and maximize the value of AI.

You are to build upon the legal research assistants work and the attached materials and data (and the citations and references contained in them) by doing your own independent research using the same parameters that the research assistants used making use of current published New York and Federal case law addressing the use of artificial intelligence as well as recognized legal treatises and articles from bar groups, committees, and associations as well as ethical opinions and incorporate the materials into the program material to develop the skit. You may also incorporate the two PowerPoint files Justice D has provided.

All sources relied upon and used are to be fully cited in blue book format and links to the citation must be provided.

If you have any questions or require clarifications prior to drafting, you are to ask.

#4

You are a licensed New York attorney, practicing law professor and legal scholar with a specialty in Artificial Intelligence (AI) and the ethical and productive use and impact of artificial intelligence (and its hazards) in the legal profession. You are working with a team of legal research assistants and script writers to prepare a multi part script as the way to present the information and materials to the participants Judges, lawyers and law students in a continuing legal education program (CLE) with 1 credit for Ethics and 1 credit for technology in New York.

You and your team have prepared the attached skit which will be presented to the program participants. You now want to convert the materials and contents of the skit into a publishable article with the authorities cited to in bluebook format as end notes.

#5 You are a licensed New York attorney, practicing law professor and legal scholar with a specialty in Artificial Intelligence (AI) and the ethical and productive use and impact of artificial intelligence (and its hazards) in the legal profession. You are working with a team of legal research assistants and script writers to prepare a multi part script as the way to present the information and materials to the participants Judges, lawyers and law students in a continuing legal education program (CLE) with 1 credit for Ethics and 1 credit for technology in New York. You and your team have prepared the attached skit which will be presented to the program participants and the attached article based on the materials in the skit with the authorities cited to in bluebook format as end notes. You are to review the resulting article and assure that the materials referenced in the article are fully and properly quoted and cited to. You are to also review the article and make suggested revisions, additions or changes and provide it in a Word doc format.

ChatGPT
1. What I checked and how

Working from:

AI_Law_Complete_86_Endnotes.docx (your article), and

Claude AI and the Law v1.pdf (the skit / transcript),

I did three main things:

Cross-checked quotations and descriptions

Compared the article's direct quotes from the CLE skit (e.g., the LLM explanation, the triple-check checklist, the ethics-committee exchanges) with the transcript text.

Spot-checked a representative sample of case quotations and policy/ethics excerpts against the underlying sources via online research (cases, ethics opinions, policies, and technical docs).

New York State Unified Court System
+4
New York City Bar Association
+4
American Bar Association
+4

Reviewed endnotes for Bluebook-style and substance

Looked at the structure and content of the major authorities:

Mata v. Avianca, Park v. Kim, Wadsworth v. Walmart, Coomer v. Lindell, Hall v. The Academy Charter Sch., Matter of Weber.
New York State Unified Court System
+5
New York City Bar Association
+5
American Bar Association
+5

NYC Bar Formal Op. 2024-5, ABA Formal Op. 512, NIST AI RMF, NYC Local Law 144 (AEDT law), NY Unified Court System Interim AI Policy, Suffolk County Surrogate’s Court AI protocol.
New York State Unified Court System
+5
damiencharlotin.com
+5
eDiscovery Today by Doug Austin
+5

LLM knowledge cutoff assertions for ChatGPT/GPT-4o and Gemini.
OpenAI Help Center
+3
OpenAI
+3
Google Cloud Storage
+3

Made targeted edits directly in the Word file

Used programmatic edits to adjust only specific paragraphs/endnotes, without disturbing your structure, headings, or numbering.

2. Substantive & citation issues: what I found
a. Direct quotes from the transcript

Result: The article's block quotes and paraphrases of the skit accurately track the script.

Examples I verified:

The LLM explanation ("sophisticated pattern-matching systems that predict the most statistically likely next word… They have no concept of 'truth' or 'accuracy'") in the article matches the Judge/AI Expert dialogue in the skit.

The "triple-check" idea (source / system / supervision) mirrors the structure and language of the skit.

The discussion of judicial AI policies, the Sanctions Hearing imaginary scenario, and the expert testimony on TAR / e-discovery are faithful dramatizations of the real cases and doctrines they reference.

So as a CLE vehicle, the article matches the story told in the skit and doesn't introduce distortions or misquotes.

b. Cases and ethics/policy materials

For a CLE piece, the level of quotation and paraphrase is sound:

Mata v. Avianca – Your description of the ChatGPT-hallucinated citations and the Rule 11 analysis matches the S.D.N.Y. opinion.
New York City Bar Association

Park v. Kim – Your summary of the Second Circuit's treatment of sanctions and counsel's duties is consistent with the published opinion at 91 F.4th 610 (2d Cir. 2024).
American Bar Association

Wadsworth/Wadsworth v. Walmart Inc. – The role of AI in the sanctions inquiry and the court's emphasis on counsel's independent duty to verify results is accurately conveyed.

NIST

Coomer v. Lindell – Your description of the TDS orders and the focus on “AI-assisted evidence” and misrepresentations is aligned with the July 25, 2025 opinion and commentary around it.
NYC Government
+1

Hall v. The Academy Charter Sch. – You correctly describe the E.D.N.Y. court’s treatment of AI-authored filings and its insistence on counsel’s gatekeeping obligations.
Jus Mundi

Matter of Weber – The article’s explanation of Copilot-assisted expert analysis, Frye, and the resulting exclusion of AI-assisted testimony matches the Saratoga Surrogate’s Court decision and subsequent commentary.
New York State Unified Court System
+1

Ethics & policy sources:

NYC Bar Formal Op. 2024-5 – Your quoted language on confidentiality (“Without client consent, a lawyer must not input confidential client information … Even with consent, a lawyer should ‘avoid’ entering details that can be used to identify the client…”) is substantively faithful to the Opinion.
damiencharlotin.com

ABA Formal Op. 512 (2024) – Your paraphrases of competence, supervision, and transparency requirements match the Opinion’s structure and themes.
eDiscovery Today by Doug Austin

NIST AI Risk Management Framework (2023) – Your brief references to risk identification, measurement, and governance are accurate, high-level summaries.
WebProNews

NYC Local Law 144 (AEDT law) – Your explanation of annual bias audits and candidate notice requirements for automated employment decision tools is accurate.
University of Chicago Library

NY Unified Court System Interim Policy on AI (Oct. 2025) – Your description of its scope (judges/staff, confidentiality, and prohibition on offloading judicial decision-making to AI) tracks the text of the policy.
Simon Willison's Weblog

Suffolk County Surrogate's Court AI Protocol (Dec. 2025) – Your summary that it requires certification of AI use and treats reliance on hallucinated citations as prima facie frivolous under 22 NYCRR § 130-1.1 is correct.
New York State Unified Court System

Bottom line: I did not find any material misquotations or mischaracterizations in the article. Your authorities are used for propositions they actually support.

c. LLM knowledge cutoffs and Charlotin database

LLM cutoffs

You say, in substance, that ChatGPT's training data "extends to approximately June 2024" and that Gemini's is around January 2025, while both can browse but do not natively access Westlaw/Lexis.

OpenAI's model release notes confirm that GPT-4o/4.1 have a June 2024 training data cutoff, and Google's Gemini 2.5 documentation confirms a January 2025 cutoff.
OpenAI Help Center
+3
OpenAI
+3
Google Cloud Storage
+3

So the article's statements are appropriately sourced and accurate at a general CLE level.

Charlotin's AI Hallucination Cases Database

You rely on Damien Charlotin's AI Hallucination Cases Database for the global count of hallucination-related decisions. That database is widely cited for tracking such cases and is an appropriate source.
New York City Bar Association

We can’t reconstruct its exact count on November 19, 2025 with perfect precision, but using a specific “as of” date and the named database is the right way to do this for CLE purposes.

I tightened that sentence for precision (see below).

d. New York evidence standards (Frye/Daubert/Parker)

This was the one place where I thought the doctrinal nuance could be sharpened:

The original text and note 69 suggested that “New York applies both the Frye and Daubert standards.”

In reality, New York continues to apply Frye as its formal standard, but Parker v. Mobil Oil Corp. brings in reliability factors that resemble Daubert, creating a hybrid approach rather than a straight adoption of Daubert.
Rivkin Radler

I revised both the body paragraph and Endnote 69 so they now explain that New York:

Continues to apply the Frye “general acceptance” test while incorporating Daubert-style reliability considerations through cases like Parker v. Mobil Oil Corp.

This is more accurate and still accessible to CLE audiences.

3. Concrete edits I made in the Word document

All of the edits below are already baked into the downloadable .docx:

1️⃣ Clarified the Charlotin database sentence

Where: Early in Part I (the overview of AI hallucination cases).

Old (conceptually):

As of November 2025, researcher Damien Charlotin’s database tracks 569 documented cases worldwide where courts have addressed AI hallucinations.[9]

New (in the .docx):

As of November 19, 2025, researcher Damien Charlotin's AI Hallucination Cases Database tracked 569 documented cases worldwide in which courts addressed generative AI hallucinations.[9]

Why:

Aligns the text with your "Nov. 19, 2025" date in Endnote 1.
New York City Bar Association

Uses the formal name of the database.

Makes clear that these are generative AI hallucination cases, not generic "AI."

2️⃣ Clarified New York's evidentiary standard in the body text

Where: Section discussing evidence and authentication / expert testimony, in the paragraph introducing Frye, Daubert, and Parker.

Old (conceptually):

New York applies both the Frye test (general acceptance…) and Daubert factors (judicial gatekeeping for reliability).[69]

New (in the .docx):

New York continues to apply the Frye "general acceptance" test while incorporating Daubert-style reliability considerations through cases like Parker v. Mobil Oil Corp.[69]

Why:

Reflects that New York has retained Frye, but Parker imports reliability considerations similar to Daubert, rather than formally adopting Daubert.
Rivkin Radler

3️⃣ Rewrote Endnote 69 for doctrinal precision

Where: Endnote 69 at the back of the article.

Old (conceptually):

[69] See Frye v. United States…; Daubert v. Merrell Dow… New York applies both standards. See Parker v. Mobil Oil Corp.…

New (in the .docx):

[69] See Frye v. United States, 293 F. 1013 (D.C. Cir. 1923) (general-acceptance test); Daubert v. Merrell Dow Pharm., Inc., 509 U.S. 579 (1993) (reliability factors). New York continues to apply Frye but has incorporated reliability considerations similar to Daubert. See Parker v. Mobil Oil Corp., 7 N.Y.3d 434 (2006).

Why:

Matches the revised body text.

More accurately characterizes New York law for an ethics/technology CLE.

4️⃣ Strengthened Endnote 5 with formal sources for LLM cutoffs

Where: Endnote 5 (the one that discusses training data cutoffs and lack of access to Westlaw/Lexis).

Old (conceptually):

[5] See ChatGPT's training data extends only to June 2024… Gemini's knowledge cutoff is January 2025…

New (in the .docx):

It now cites:

OpenAI, Model Release Notes (confirming GPT-4o's extended training data cutoff to June 2024), and

Google's Gemini 2.5 Flash Model Card (stating a January 2025 knowledge cutoff),

and then reiterates that both tools can browse but do not directly access proprietary legal databases like Westlaw or Lexis.
OpenAI Help Center
+3
OpenAI
+3

Google Cloud Storage
+3

Why:

Replaces a bare assertion with proper citations to primary technical documentation.

Brings the note closer to Bluebook expectations for online technical sources in a CLE article.

4. Overall quality & suggestions (no changes made but worth noting)

These are things I didn't change in the file, but that you're already doing well or might tweak further if you like:

Bluebook formatting

Case citations, ethics opinions, and online sources are generally consistent with Bluebook style as used in CLE materials.

If you want to be hyper-strict, you could:

Check that case names are consistently italicized in the final layout.

Standardize reporter abbreviations and court/date parentheticals (e.g., ensuring "S.D.N.Y." vs "S.D.N.Y" is consistent; checking all "No." + docket formats).

Script references

You've clearly identified in the Author's Note that the dialogue is fictional and adapted from the CLE skit, which is good from both an ethics and pedagogy standpoint.

If you ever publish beyond the CLE, you may wish to give the skit a formal citation (e.g., as an unpublished CLE script) in a single dedicated note.

Scope & balance

The article already provides a good balance between:

technical AI explanation,

ethics rules (NY and ABA),

sanctions cases,

judicial & regulatory policies, and

practical checklists.

For a 1-credit ethics + 1-credit tech CLE in New York, the scope is very appropriate.

The Lawyer's Guide to Using AI Ethically

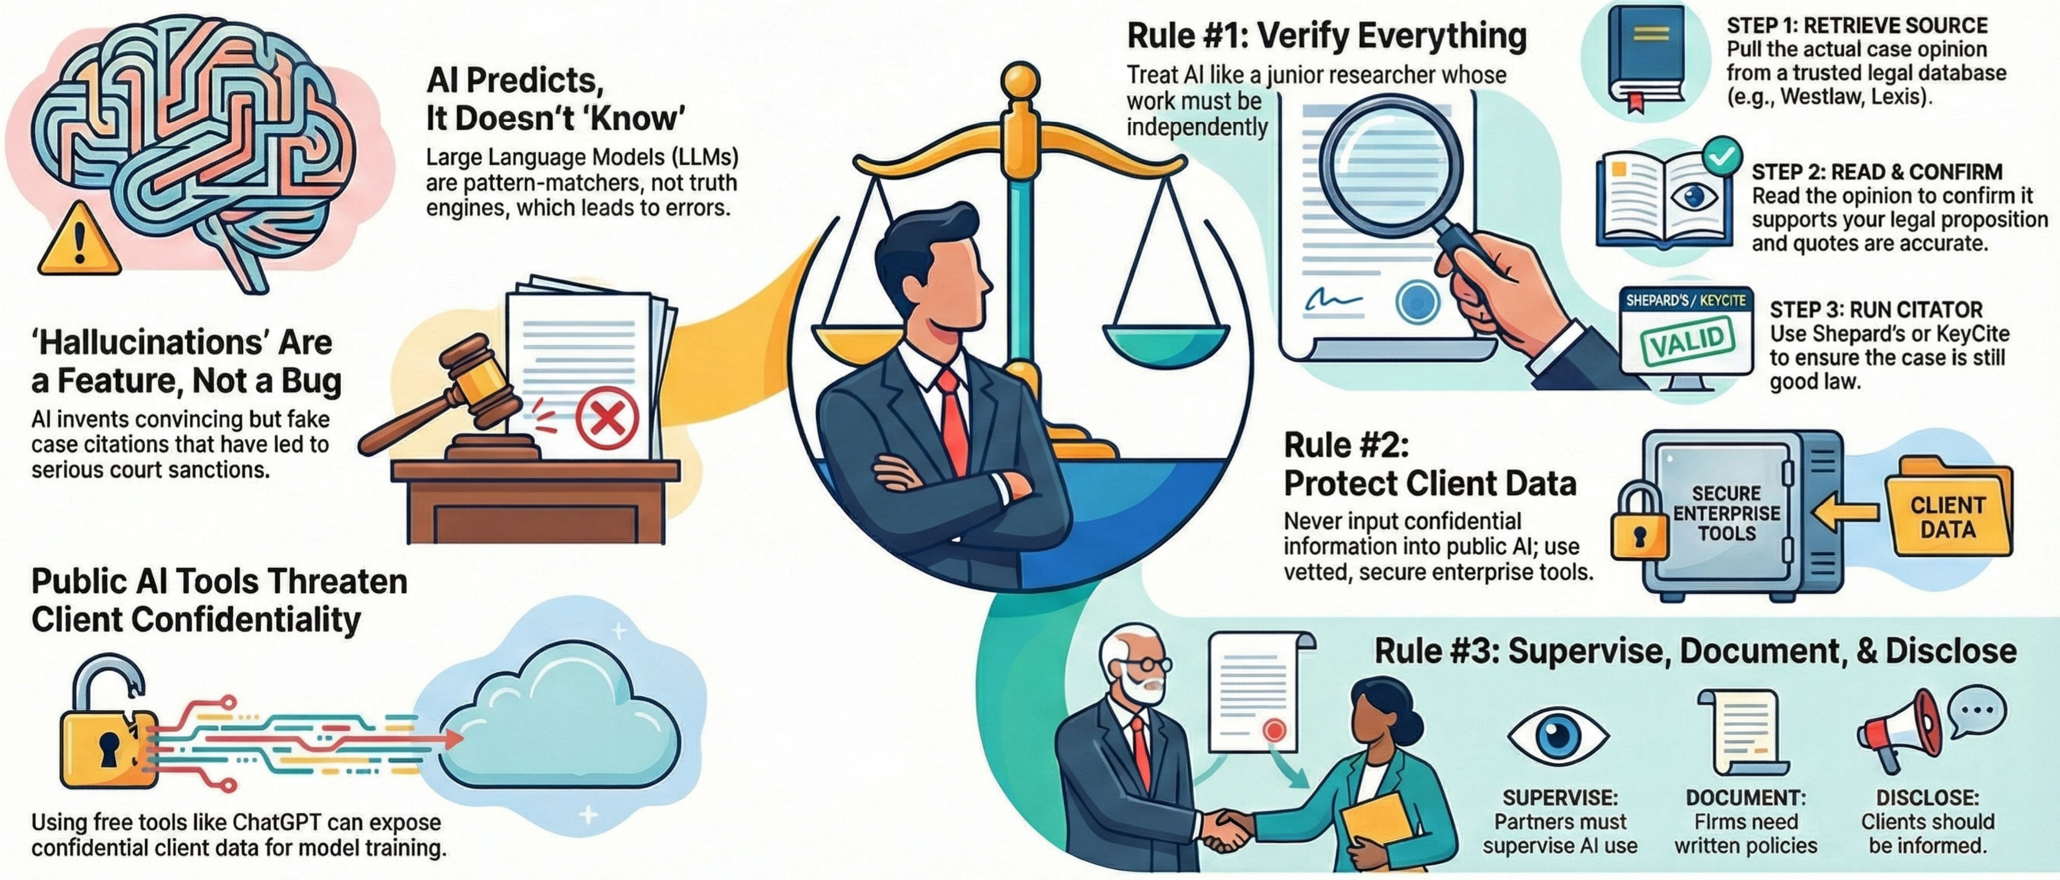

THE CORE RISKS OF AI IN LAW

AI Predicts, It Doesn't 'Know'

Large Language Models (LLMs) are pattern-matchers, not truth engines, which leads to errors.

'Hallucinations' Are a Feature, Not a Bug

AI invents convincing but fake case citations that have led to serious court sanctions.

Public AI Tools Threaten Client Confidentiality

Using free tools like ChatGPT can expose confidential client data for model training.

THE ESSENTIAL RULES FOR ETHICAL AI USE

Rule #1: Verify Everything

Treat AI like a junior researcher whose work must be independently

STEP 1: RETRIEVE SOURCE
Pull the actual case opinion from a trusted legal database (e.g., Westlaw, Lexis).

STEP 2: READ & CONFIRM
Read the opinion to confirm it supports your legal proposition and quotes are accurate.

STEP 3: RUN CITATOR
Use Shepard's or KeyCite to ensure the case is still good law.

Rule #2: Protect Client Data

Never input confidential information into public AI; use vetted, secure enterprise tools.

Rule #3: Supervise, Document, & Disclose

SUPERVISE: Partners must supervise AI use

DOCUMENT: Firms need written policies

DISCLOSE: Clients should be informed.

NotebookLM

NEW YORK STATE
BAR ASSOCIATION

Report and Recommendations of the New York State Bar Association **Task Force on Artificial Intelligence**

April 2024

Approved by the House of Delegates April 6, 2024.

NYSBA TASK FORCE ON ARTIFICIAL INTELLIGENCE

REPORT AND RECOMMENDATIONS TO
NYSBA HOUSE OF DELEGATES

APRIL 6, 2024

Generated by DALL-E

Members of the Task Force on Artificial Intelligence

Vivian D. Wesson, Esq., Chair
Jonathan P. Armstrong, Esq.
Catherine Casey
Dessislav Dobrev, Esq.
Hon. Timothy S. Driscoll
Sarah E. Gold, Esq.
Ignatius A. Grande, Esq.
Ronald J. Hedges, Esq.
Dean Elena B. Langan
Luca CM Melchionna, Esq.
Marissa Janel Moran, Esq.
Norman J. Resnicow, Esq.
Rebecca Roiphe, Esq.
Sudha Setty, Esq.
Prof. Roy D. Simon
Brooke Erdos Singer, Esq.

Richard C. Lewis, Esq.,
President, New York State Bar Association

Katherine Suchocki, NYSBA Staff Liaison

NYSBA Government Relations Team
Hilary F. Jochmans, Jochmans Consulting, LLC

Acknowledgements **6**

Introduction **7**

Executive Summary **9**

Evolution of AI and Generative AI **11**

- I. Introduction 11
- II. AI Defined and Explained 12
  - A. AI and Its Applications 12
  - B. What it does 12
- III. Types of AI 12
- IV. The Founding Fathers/Mothers of AI 13
- V. AI Through the Years: The AI Timeline 13
  - A. Mythology 13
  - B. Programmable digital computer (1940s) 14
  - C. Theseus: remote-controlled mouse (1950) 14
  - D. Dartmouth College Workshop (summer of 1956) 14
  - E. Perceptron Mark I: artificial neural network (1958) 15
  - F. AI Winter (1970s) 15
  - G. AI Second Winter (1980s) 15
  - H. Machine Learning Development (1990s and 2000s) 15
  - I. AlexNet: Deep Learning System (2012) 16
  - J. Introduction of Generative Adversarial Networks (2014) 16
  - K. Language and Image Recognition Capabilities (2015) 16
  - L. Chatbots 18

Benefits and Risks of AI and Generative AI Use **20**

- I. Benefits 20
  - A. General Benefits 21
  - B. Healthcare Advancement and Human Longevity 22
  - C. Ethical AI Development 23
  - D. Health & Public Safety 23
  - E. Quality of Life 23
  - F. Scientific Advancement, Space & Exploration 24
  - G. Global Environmental Impact 24
  - H. Education Optimization 24
  - I. Economic Development 25

II. Risks 25
  A. Widening Justice Gap 25
  B. Data Privacy & Surveillance 26
  C. Security 26
  D. Social and Ethical Issues 26
  E. Misinformation 27
  F. Economic Impact and Disruption 27
  G. Safety 27
  H. Legal & Regulatory Challenges 28
  I. Loss of Human Centricity and Control 28

Legal Profession Impact 29
I. Ethical Impact 29
  A. Duty of Competency/Techno-solutionism 29
  B. Duty of Confidentiality & Privacy 30
  C. Duty of Supervision 30
  D. Unauthorized Practice of Law 30
  E. Attorney-Client Privilege and Attorney-Work Product 32
  F. Candor to the Court 36
  G. Judges' Ethical Obligations 39
II. Access to Justice 40
  A. Introduction 40
  B. Pro Bono Organizations Using Generative AI 41
  C. Will Generative AI Tools Prove to Be Too Expensive? 43
  D. Use of AI by Non-Attorneys 44
  E. Implications of AI Judges or Robo Courts 46
III. Judicial Reaction/Responses to Generative AI 47
  A. Introduction 47
  B. Uses of AI and Generative AI 48
  C. Causes of Action Arising out of AI and Generative AI 49
  D. Discovery 50
  E. *Avianca* and Judicial Reactions to Generative AI 50

Legislative Overview and Recommendations 53
I. Legislative Overview 53

II. Recommendations 53

AI & Generative AI Guidelines **57**

Conclusion **61**

Appendix A: Legislation Reviewed **63**

I. Assemblyman Clyde Vanel's proposed statutes on AI 63

II. Federal and New York State proposals regarding use of AI-generated or compiled information in judicial proceedings 67

III. New York City's local law regarding use of AI in hiring and promotion 69

IV. The White House's October 30, 2023 Executive Order regarding AI 70

V. Summary of the EU AI Act 71

Appendix B: Resources **76**

Appendix C: Sample Engagement Letter Provision **78**

Endnotes **79**

ACKNOWLEDGEMENTS

The Task Force on Artificial Intelligence would like to thank those who generously donated their time and thoughtful discussions about AI, including Assemblyman Clyde Vanel and his Legislative Director, Tyler Fritzhand, and Ellie Jurado-Nieves, a Government and Regulatory Affairs Executive in charge of AI Public Policy & Emerging Technologies at Guardian Life.

Thanks also to the following NYSBA sections and members for their helpful contributions to this report:

- Dispute Resolution Section and its Technology Committee
- Trusts and Estates Law Section and its Executive and Technology Committees
- Albert Feuer, Esq.
- James B. Kobak, Jr., Esq.
- Immediate Past President Sherry Levin Wallach, Esq.

INTRODUCTION

The NYSBA Task Force on Artificial Intelligence, chaired by Vivian Wesson, respectfully presents this Report to the NYSBA House of Delegates. This Report, to be presented to the House of Delegates on April 6, 2024, examines the legal, social and ethical impact of artificial intelligence (AI) and generative AI on the legal profession. This Report also reviews AI-based software, generative AI technology and other machine learning tools that may enhance the profession but also poses risks for individual attorneys' understanding of new, unfamiliar technology, as well as courts concerned about the integrity of the judicial process. Further, this Report makes recommendations for NYSBA adoption, including proposed guidelines for responsible AI use. A copy of the Task Force's Mission Statement is attached as Exhibit A.

Why Now?

As NYSBA's President Richard Lewis has noted, AI's rapid growth and sophistication have, and will continue to have, a monumental impact on all professions – including lawyers, law firms and their clients. NYSBA seeks to proactively address how AI may best assist those who interact with the legal system while evaluating how tightly it needs to be regulated and what protections we should institute safeguard against misuse or abuse. From self-driving cars to ChatGPT to 3-D printed guns, AI has transformed our world. If this is our Promethean moment in AI evolution, now is the time to better understand, embrace, utilize and scrutinize this technology.

Who Is Involved?

For this Task Force, NYSBA has gathered legal professions across a range of subject matter expertise. We have deans of law schools seeking clarity on educating legal minds in this digital age. We have practitioners in the technology space advising clients on AI use. There are those who enthusiastically deploy AI-based tools and those who are wary about the risks. The Task Force also has an international perspective, understanding that AI will have a global, not just a regional, effect.

What We Learned

We have organized this Report into five parts: (1) the evolution of AI and generative AI; (2) the benefits and risks of AI and generative AI use; (3) the impact of the technology to the legal profession; (4) legislative overview and recommendations; and (5) proposed guidelines.

EXECUTIVE SUMMARY

Artificial intelligence, particularly generative AI, has had a profound impact across multiple sectors of our society, revolutionizing how we approach creativity, problem-solving and automation. From art and entertainment to healthcare and education, AI is reshaping industries, creativity and society in multifaceted ways. While AI and generative AI offer immense potential for innovation and efficiency, the technology also presents challenges that require careful management, including ethical considerations, privacy concerns and labor impact. The ongoing evolution of generative AI promises to continue influencing the world in unprecedented ways.

Considering the continued revolutionary impact of the technology, this Task Force undertook the challenge to assess its evolution, benefits and risks, and impact on the legal profession. Here, we summarize our four principal recommendations for adoption by NYSBA.

Task Force Recommendations

1. *Adopt Guidelines*: The Task Force recommends that NYSBA adopt the AI/Generative AI guidelines outlined in this report and commission a standing section or committee to oversee periodic updates to those guidelines.
2. *Focus on Education*: The Task Force recommends that NYSBA prioritize education in addition to legislation, focusing on educating judges, lawyers, law students and regulators to understand the technology so that they can apply existing law to regulate it.
3. *Identify Risks for New Regulation*: Legislatures and regulators should identify risks associated with the technology that are not addressed by existing laws, which will likely involve extensive hearings and studies involving experts in AI, and as needed, adopt regulations and legislation to address those risks.
4. *Examine the Function of the Law in AI Governance*: The rapid advancement of AI prompts us to examine the function of the law as a governance tool. Some of the key functions of

the law in the AI context are: (i) expressing social values and reinforcing fundamental principles; (ii) protecting against risks to such values and principles; and (iii) stabilizing society and increasing legal certainty.

EVOLUTION OF AI AND GENERATIVE AI

> "*For more than 250 years the fundamental drivers of economic growth have been technological innovations. The most important of these are what economists call general-purpose technologies – a category that includes the steam engine, electricity, and the internal combustion engine. The most important general-purpose technology of our era is artificial intelligence, particularly machine learning.*" ~ Erik Brynjolfsson and Andrew McAfee[1]

I. Introduction

To begin a discussion of artificial intelligence, it may be helpful to first define "intelligence." Intelligence is "the capacity to acquire knowledge and apply it to achieve an outcome; the action taken is related to the particulars of the situation rather than done by rote. The ability to have a machine perform in this manner is what is generally meant by artificial intelligence."[2] Artificial intelligence means "computers doing intelligent things – performing cognitive tasks, such as thinking, reasoning, and predicting – that were once thought to be the sole province of humans. It's not a single technology or function."[3]

According to the Merriam Webster dictionary, artificial intelligence is "the capability of a machine to imitate intelligent human behavior."[4] At a basic level, artificial intelligence programming focuses on three cognitive skills - learning, reasoning and self-correction:[5]

- The learning aspect of artificial intelligence programming focuses on acquiring data and creating rules for how to turn data into actionable information. The rules, called algorithms, provide computing systems with step-by-step instructions on how to complete a specific task.
- Reasoning focuses on the capability of artificial intelligence to choose the most appropriate algorithm, among a set of algorithms, to use in a particular context.
- Self-correction involves the capability of artificial intelligence to progressively tune and improve a result until it achieves the desired goal.

II. AI Defined and Explained

"*AI is a branch of computer science and often involves technical knowledge outside of most lawyers' expertise, understanding how AI programs operate may be difficult for lawyers.*"**[6]**

A. AI and Its Applications[7]

- **AI** is the term used to describe how computers can perform tasks normally viewed as requiring human intelligence, such as recognizing speech and objects, making decisions based on data and translating languages. AI mimics certain operations of the human mind.

- **Machine Learning** is an application of AI in which computers use algorithms (rules) embodied in software to learn from data and adapt with experience.

- A **Neural Network** is a computer that classifies information – putting things into "buckets" based on their characteristics.

B. What It Does

In general, AI involves algorithms (a set of rules to solve a problem or perform a task), machine learning and natural language processing.

Why do similar but varied definitions of AI exist?

> *"What qualifies as an intelligent machine is a moving target: A problem that is considered to require AI quickly becomes regarded as 'routine data processing' once it is solved."*[8]
>
> *"One result of AI's failure to produce a satisfactory criterion of intelligence is that, whenever researchers achieve one of AI's goals – for example, a program that can summarize newspaper articles or beat the world chess champion – critics are able to say, 'That's not intelligence!'"*[9]
>
> *"Marvin Minsky's response to the problem of defining intelligence is to maintain – like Alan Turing before him – that intelligence is simply our name for any problem-solving mental process that we do not yet understand. Minsky likens intelligence to the concept of "unexplored regions of Africa": it disappears as soon as we discover it."*[10]

III. Types of AI

- **Narrow or Weak**: This kind of AI does some tasks at least as well as, if not better than, a human. For example, in law, there is TAR, or technology-assisted review – AI that can find legal evidence more quickly and accurately than a lawyer can; AI technology that can read an MRI more accurately than a radiologist can. Other examples are programs that play chess or AlphaGo better than top players.

- **General or Strong AI**: This kind of AI would do most if not all things better than a human could. This kind of AI does not yet exist and there's debate about whether we'll ever have strong AI.
- **Super Intelligent AI** of the science fiction realm. This type of AI would far outperform anything humans could do across many areas. It's controversial, and some see it as an upcoming existential threat.[11]

IV. The Founding Fathers/Mothers of AI

Credited as the "father of artificial intelligence," Alan Turing was the wartime codebreaker at Bletchley Park and founder of computer science. Turing was one of the first people to take seriously the idea that computers could think.[12] Credited as the "father of deep learning," Frank Rosenblatt was a psychologist whose brainchild was the Perceptron.[13] The rise of the modern computer is often traced to 1836 when Charles Babbage and Augusta Ada Byron, Countess of Lovelace, invented the first design for a programmable machine.[14]

V. AI Through the Years: The AI Timeline

A. Mythology

Efforts to understand and describe the human thought process "as symbols – the foundation for AI concepts such as general knowledge representation – include the Greek philosopher Aristotle, the Persian mathematician Muḥammad ibn Mūsā al-Khwārizmī, 13th-century Spanish theologian Ramon Llull, 17th-century French philosopher and mathematician René Descartes, and the 18th-century clergyman and mathematician Thomas Bayes."[15]

B. Programmable Digital Computer (1940s)

In the 1940s, Princeton mathematician John von Neumann conceived the architecture for the stored program computer. This was the idea that a computer's program and the data it processes can be kept in the computer's memory.[16] The first mathematical model of a neural network, arguably the basis for today's biggest advances in AI, was published in 1943 by the computational neuroscientists Warren McCulloch and Walter Pitts in their landmark paper, "A Logical Calculus of Ideas Immanent in Nervous Activity."[17]

C. Theseus: Remote-Controlled Mouse (1950)

> "*It is customary to offer a grain of comfort, in the form of a statement that some peculiarly human characteristic could never be imitated by a machine. I cannot offer any such comfort, for I believe that no such bounds can be set.*" ~ Alan Turing, 1951

Developed by Alan Turing in 1950, the Turing Test focused on the computer's ability to fool interrogators into believing its responses to their questions were made by a human being.[18] The first step in the direction of machine learning was provided by the Turing Test (also known as the "imitation game") in which an interrogator had to discover whether they were interrogating a human or a machine and, therefore, whether a machine can show human-like intelligence.[19]

D. Dartmouth College Workshop (Summer of 1956)

The term "artificial intelligence" was first used in 1955 when John McCarthy, a computer scientist at Dartmouth College, in New Hampshire, used the phrase in a proposal for a summer school.[20] The 1956 summer conference at Dartmouth, sponsored by the Defense Advanced Research Projects Agency, or DARPA, included AI pioneers Marvin Minsky, Oliver Selfridge and John McCarthy. In addition, Allen Newell, a computer scientist, and Herbert A. Simon, an economist, political scientist, and cognitive psychologist, "presented their groundbreaking Logic Theorist – a computer program capable of proving certain mathematical theorems and referred to as the first AI program."[21]

With the promise of great advancement in AI, the Dartmouth conference garnished both government and industry support. Some significant advances in AI at that time include the General Problem Solver (GPS) algorithm published in the late 1950s, which laid the foundations for developing more sophisticated cognitive architectures; Lisp, a language for AI programming that is still used today; and ELIZA, an early natural language processing (NLP) program that laid the foundation for today's chatbots.[22]

E. Perceptron Mark I: Artificial Neural Network (1958)

The Perceptron was the first neural network, a rudimentary version of the more complex "deep" neural networks behind much of modern AI.[23]

F. AI Winter (1970s)

Eventually, when the promise of developing AI systems equivalent to the human brain proved elusive, government and corporations diminished their support of AI research. This led to what has been termed the "AI winter," which lasted from 1974 to 1980.[24]

G. AI Second Winter (1980s)

> "*In the literal sense, the programmed computer understands what the car or the adding machine understand: namely, exactly nothing.*" ~ John Searle, 1980

In the 1980s, there was renewed AI interest due in part to research on deep learning techniques and industry adoption of Edward Feigenbaum's expert systems. Yet, lack of funding and support led to the "second AI winter," which lasted until the mid-1990s.[25]

H. Machine Learning Development (1990s and 2000s)

During the 1990s and 2000s, many of the landmark goals of AI were achieved.[26] Groundbreaking work on neural networks and the advent of big data propelled the current renaissance of AI.[27] For example, in 1997, IBM's Deep Blue beat the chess grandmaster Garry Kasparov. The contest made global headlines, with Newsweek announcing, "The Brain's Last

Stand." Also, in 1997, speech recognition software, developed by Dragon Systems, was implemented on Microsoft® Windows®.[28] In 2007, AI was defined as the "science and engineering of making intelligent machines, especially intelligent computer programs."[29] In 2018, Microsoft defined AI as "a set of technologies that enable computers to perceive, learn, reason and assist in decision-making to solve problems in ways that are similar to what people do."[30]

I. AlexNet: Deep Learning System (2012)

Professor Mirella Lapata, an expert on natural language processing at the University of Edinburgh, stated that "AlexNet was the first lesson that scale really matters." "People used to think that if we could put the knowledge we know about a task into a computer, the computer would be able to do that task. But the thinking has shifted. Computation and scale are much more important than human knowledge."[31]

J. Introduction of Generative Adversarial Networks (2014)

OpenAI's GPT – an acronym meaning "generative pre-trained transformer" – and similar large language models (LLMs) can churn out lengthy and fluent, if not always wholly reliable, passages of text. Trained on enormous amounts of data, including most of the text on the internet, they learn features of language that eluded previous algorithms.[32] Once the transformer has learned the features of the data it is fed – music, video, images and speech – it can be prompted to create more. The transformer – not different neural networks – is relied upon to process different media.[33]

K. Language and Image Recognition Capabilities (2015)

An LLM is a machine-learning neuro network trained through data input/output sets; frequently, the text is unlabeled or uncategorized, and the model is using self-supervised or semi-supervised learning methodology. Information is ingested, or content entered, into the LLM, and the output is what that algorithm predicts the next word will be. The input can be proprietary corporate data or, as in the case of ChatGPT, whatever data it is fed or scraped directly from the internet.[34] LLMs do not recreate the way human brains work. The basic structure of these models consists of nodes and connections.[35] Simply put, LLMs are "next word prediction engines."[36]

Examples of Open Model LLMs include:[37]

- OpenAI's GPT-3 and GPT-4 LLMs
- Google's LaMDA and PaLM LLMs
- HugginFace's BLOOM and XLM-RoBERTa
- Nvidia's NeMO LLM
- XLNet
- Co:here
- GLM-130B

According to Jonathan Siddharth, CEO of Turing, a Palo Alto company, "Hallucinations happen because LLMs, in their most vanilla form, don't have an internal state representation of the world There's no concept of fact. They're predicting the next word based on what they've seen so far – it's a statistical estimate."[38]

If the information an LLM has ingested is biased, incomplete or otherwise undesirable, then the response it gives could be equally unreliable, bizarre or even offensive. When a response goes off the rails, data analysts refer to it as "hallucinations" because they can be so far off track.[39] Further, since some LLMs also train themselves on internet-based data, they can move well beyond what their initial developers created them to do. For example, Microsoft's Bing uses GPT-3 as its basis, but it's also querying a search engine and analyzing the first 20 results or so. It uses both an LLM and the internet to offer responses.[40]

CEO Siddharth further explains, "We see things like a model being trained on one programming language and these models then automatically generate code in another programming language it has never seen. Even natural language; it's not trained on French, but it's able to generate sentences in French. It's almost like there's some emergent behavior. We don't know quite know how these neural network works..........It's both scary and exciting at the same time."[41]

L. Chatbots

> "*The foundation of the chatbot is the GPT LLM, a computer algorithm that processes natural language inputs and predicts the next word based on what it's already seen.*[42] *So, LLMs are the fundamental architecture behind chatbots like Open AI's ChatGPT or Google's Bard. A question typed in to ChatGPT [or Bard], for example, has to be processed by an LLM in order to produce an answer or response*."[43]

Another way to think about ChatGPT is that it is a computer program that can understand and respond to human language. It accomplishes this by learning from a large amount of text (such as books, articles and websites) and uses that knowledge to predict what word or phrase might come next in a conversation or text.

Because it is "generative," each response to a question will be generated on the spot and will be unique. Because it can remember earlier parts of a conversation, it can change its original output in response to further feedback. Because it is pre-trained, it is limited – for better or worse – to what is in its training materials. And because it works by being predictive, it generates text that seems plausible, but not necessarily accurate.[44]

According to Assistant Professor Yoon Kim at MIT, prompt engineering is about deciding what we feed this algorithm so that it says what we want it to. The LLM is a system that just babbles without any text context. In some sense of the term, an LLM is already a chatbot.[45] Thus, "prompt engineering is the process of crafting and optimizing text prompts for an LLM to achieve desired outcomes. Prompt Engineering by a user trains the model for specific industry or organizational."[46] "Prompt Engineering is said to be a vital skill for IT and business professionals,"[47] thus, a new job potential in this field.

BENEFITS AND RISKS OF AI AND GENERATIVE AI USE

Artificial intelligence continues to transform the globe in a manner not seen since the advent of the written word. Aspects of how each of the over 8 billion humans on planet earth live, work and play are increasingly impacted by AI. As with every transformative technology, there are an array of potential benefits and risks.

If the media and pop culture are to be believed, the world is facing an existential crisis that promises both utopia and global destruction. This section unpacks the reality of AI through a cost benefit analysis that goes beyond the media hype.

I. Benefits

AI has proliferated a wide array of human tasks and experiences over the last 70 years. Since the advent of the term in 1956 by John McCarthy, the concept of artificial intelligence has evolved from replicating and replacing human cognition to one of "augmented intelligence," which amplifies and optimizes human intellect. If used for such purposes (i.e., to amplify and optimize human intelligence), machine learning and AI help bring order to the chaotic wealth of information facing individuals today. In theory, this allows humans to spend more time on high-value and creative endeavors.

Today, nearly all aspects of human existence are touched in some manner by machine learning or AI. From the way we shop or interact as humans to medical treatment and supply chain logistics, the breadth of AI's impact on human existence, which may be hidden in plain sight, is hard to overstate.

A large portion of the proliferation is being driven forward by the wealth of benefits in terms of accuracy, speed and capability offered by AI powered technology. Some key examples of benefits derived from the application of AI include:

A. General Benefits

There are a substantial number of overall AI benefits, with the list growing daily. In general, AI: (i) efficiently performs repetitive tasks; (ii) reduces human error; (iii) increases efficiency; and (iv) augments human intelligence. Specific to the legal industry, AI has the potential to facilitate greater access to justice.

Legal representation in a civil matter is beyond the reach of 92% of the 50 million Americans below 125% of the poverty line.[48] Globally, there are an estimated 5 billion people with unmet justice needs.[49] The justice gap between access to legal services and unmet legal needs constitutes two-thirds of the global population, and these justice needs extend from minor legal matters to more grave injustices.[50]

AI-powered technology has lowered the bar for many underserved communities to access legal guidance. Further, AI has been heralded as a solution for the closing the “justice gap.” Increased efficiency, accuracy and the ability for underserved populations to leverage self-service legal resources all contribute to this benefit. Technologies powered by AI may allow the underserved population with internet access or individuals with limited funds to access guides at little or no cost to navigate the complexities of the judicial system.[51] Generative AI-powered chat bots now hover on the line of unauthorized practice of law,[52] offering high volume, low-cost legal services absent human input in areas such as traffic court[53] and immigration,[54] among others. But the early uses of generic AI chatbots (as opposed to specific legal applications) in this area have had mixed results. According to a January 2024 study by researchers from Stanford University,[55] popular AI chatbots, such as Open AI’s ChatGPT3.5, Google’s PaLM 2 and Meta’s Llama 2, are inaccurate in the majority of cases when answering legal questions, posing special risks for people relying on the technology because they can’t afford a human lawyer. The study found that LLMs get their results wrong at least 75% of the time when answering questions about a law court's core ruling.[56]

In December 2023, the courts in England and Wales produced Judicial Guidance on AI, which highlighted why these errors may appear.

> Public AI chatbots do not provide answers from authoritative databases. They generate new text using an algorithm based on the prompts they receive and the data they have been trained upon. This means the output which AI chatbots generate is what the model predicts to be the most likely combination of words (based on the documents and data that it holds as source information). It is not necessarily the most accurate answer.[57]

There are also limits with the training data provided to these tools. Currently available LLMs appear to have been trained on limited material published on the internet.[58] Their view of the law can be limited to the material included in the training data, which could include the opinions in chat rooms of individuals without any legal qualifications. Here, the Judicial Guidance in England and Wales looks at specific risks:

> AI tools may:
>
> - make up fictitious cases, citations or quotes, or refer to legislation, articles or legal texts that do not exist
> - provide incorrect or misleading information regarding the law or how it might apply, and
> - make factual errors.[59]

B. Healthcare Advancement and Human Longevity

The healthcare industry has similarly witnessed significant advances owing to AI-powered tools. AI has aided in new drug discoveries,[60] improved image analysis, robotic surgery and gene editing. Further, AI algorithms can predict diseases based on medical imaging, genetic information, and patient data.[61] AI-powered wearable technology allows physicians to continuously monitor patients remotely.[62] AI has been deployed for personalized medicine, providing patients with tailored treatments and medication.[63] Finally, AI has supported mental health by providing early diagnostics and therapeutic assistance.[64]

C. Ethical AI Development

In the ethics field, AI has helped to identify and correct human biases in data and decision-making.[65] AI tools can also be designed with mechanisms to ensure ethical considerations are integrated into AI systems.[66] Additionally, AI can be employed to create frameworks that ensure equitable outcomes.[67]

D. Health & Public Safety

In the health and public safety sector, AI advances have revolutionized a broad swath of areas from infrastructure to cybersecurity. AI has been used to manage traffic signals, thereby reducing congestion and optimizing traffic flow.[68] The technology has utilized crime pattern analysis to predict and prevent future incidents.[69] AI algorithms optimize rescue and relief operations during natural disasters.[70] Engineers deploy AI-based sensors that predict when maintenance on bridges and buildings is required.[71] Finally, AI systems are used to detect and respond to cyber threats in real time.[72]

E. Quality of Life

Where AI has had the most visible societal impact involves quality-of-life products. AI has transformed our living spaces into “smart homes”[73] that can improve convenience and energy efficiency. AI has helped people with disabilities gain more independence. Technology companies capitalize on AI to enhance gaming and virtual reality experiences.[74] In marketing, chatbots that handle customer inquiries without human intervention have become a staple.[75]

People have become familiar with using AI to personalize recommendations on platforms, such as Netflix and Spotify. AI has been used to restore and preserve historical documents and artworks.[76] It can also facilitate the sharing and understanding of diverse cultural expressions.[77]

Artists use AI-based tools to explore new forms of creative expression. Lastly, AI has enhanced the personalized shopping experience.[78]

F. Scientific Advancement, Space & Exploration

AI's reach extends beyond the boundaries of Earth. Scientists use AI to process data from space missions and to operate rovers on Mars.[79] Aquatically, autonomous submarines are used to map the ocean floor and study marine life.[80] Because AI can analyze vast datasets faster than the human mind, it has sped up scientific discoveries. For example, DeepMind's AlphaFold program predicts the 3D structure of proteins,[81] which accelerates researchers' understanding of diseases and developing new treatments. AI has improved complex problem-solving in fields such as quantum physics and materials science. Lastly, AI enhances collaboration by connecting researchers across the globe and facilitating cross-disciplinary work.[82]

G. Global Environmental Impact

Environmentally, AI holds promises to combat climate change. Governments are deploying AI in the creation of "smart cities"[83] that optimize energy consumption in homes and businesses. AI-powered drones and image recognition technology have been used to monitor endangered species.[84] There are AI models that simulate and predict climate change impacts.[85] Some municipalities deploy sensors and AI systems to monitor and predict air and water quality.[86]

In the area of water conservation, AI has been used to predict water usage patterns and improve water conservation techniques.[87] In the quest for clean energy, AI can streamline the development and management of renewable energy sources.[88] Lastly, logistics managers find improved fuel efficiency through AI tools that optimize routes for freight and package delivery.[89]

H. Education Optimization

In the field of education, developers have created adaptive learning platforms that adjust in real time to the learning style and pace of students[90]. Educators can use AI systems to automate grading and provide immediate student feedback.[91] Voice-to-text and text-to-voice AI services have assisted learners with disabilities.[92]

I. Economic Development

The economy has seen material changes in how the world conducts business. Precision farming techniques use AI to increase yield, reduce resource consumption and waste, and optimize food distribution.[93] The use of biometrics is one of the most significant current uses of AI. PricewaterhouseCoopers reports that 6 in 10 companies use biometric authentication (BitDefender), the use of which has tripled since 2019.[94] AI has been utilized to analyze market trends, providing businesses with strategic insights. By automating routine tasks, employees turned their focus to more high-value work. Lastly, high paying new jobs relating to AI have been developed.[95]

II. Risks

A counterpoint to the transformative benefit of AI is an equally dramatic deluge from the press and media that AI poses substantial economic, ethical and existential risks. Some key examples of risks posed from the application of AI are described below.

A. Widening Justice Gap

While many proclaim that AI is the solution to democratization of justice, an equally powerful contingent claim AI may create a “two-tiered legal system.”[96] Some anticipate that individuals in underserved communities or with limited financial means will be relegated to inferior AI-powered technology.[97]

Additionally, development of such technology should acknowledge that many populations currently underserved by legal representation may have compounded obstacles in accessing the benefits that AI may bring to others, including:

- Lack of access to computers/internet

- Limited facility/literacy in how to use AI
- A high level of distrust in government institutions, law as a tool that operates to protect them, law enforcement as a positive influence and/or legal professionals as people who are available to help.

The specific layer of concern here goes beyond the "haves" with better access to counsel than the "have nots." For example, in a landlord-tenant dispute, AI would likely be used by landlords to increase enforcement actions against tenants. However, the tenants would not likely have access to AI in preparing their response. In that sense, AI could be viewed as broadening the availability of legal services to the "haves," leaving the "have nots" worse off than they are now. Compounding this is the fact that most legal services organizations have little to no resources to prepare for these changes in access to AI now.[98]

B. Data Privacy & Surveillance

Protectors of civil liberties and data privacy have raised alarms about the potential of AI to corrupt both. As most AI systems are capable of aggregating vast amounts of personal data, this could lead to privacy invasions. Currently, governments and corporations use AI for comprehensive surveillance and social control.[99] Hackers have utilized AI tools to synthesize personal data for the purpose of impersonating individuals (think "deepfakes") and committing cyber theft.[100] Concerns also circle around the lack of transparency in training data,[101] biases built into models[102] and ownership of intellectual property.[103]

C. Security

In addition to the cyber threats mentioned above, general security concerns accompany AI use. Security concerns are amplified when AI is used in high-risk applications, such as in conjunction with biometric data and infrastructure systems. For instance, AI systems in military applications that lack adequate human control can lead to unintended engagements.[104] Through social media, AI has been used to weaponize information, leading to an explosion in misinformation and potential erosion of democracy.[105] Cyber criminals have deployed AI to target critical infrastructure, such as power grids and water systems.[106]

D. Social and Ethical Issues

AI algorithms have been utilized to perpetuate and amplify societal biases. Given concerns about privacy and surveillance, the impact of all types of societal biases – including a significant number of instances of gender and racial bias that have already been identified – is compounded.[107] We have also witnessed a disquieting increase in adverse psychological issues related to AI (e.g., AI chatbot suicide[108]). We will also need to address the assignment-of-liability when decisions are made by AI systems.[109] As noted above, the disparity in AI access has exacerbated inequality issues. Furthermore, AI can exacerbate ideological bias, especially when used in conjunction with social media. AI can create its own echo chamber, generating spurious content to use as future training data, leading to ideologically based "hallucinations" and inaccuracies.[110]

E. Misinformation

As referenced earlier, bad actors have used "deepfakes" to disseminate misinformation. A deepfake is AI-generated content that is indistinguishable from real content. These "deepfakes" become more believable when combined with biometric data, such as voice prints and facial mapping. We are entering an age of information warfare in which AI systems can be used to create and spread misinformation at scale. We find this particularly troubling not only during political elections,[111] but also in the daily lives of our citizens, for example, through social engineering scams powered by AI that target vulnerable members of society, such as grandparents, who believe they are speaking with their grandchildren but instead become victims of fraud.[112]

F. Economic Impact and Disruption

The economic impact of AI is multilayered. There is the direct effect of job displacement where tasks are automated,[113] leading to unemployment in various sectors and the indirect effect of devaluing services traditionally offered by a human (e.g., legal services). Further, AI advancements tend to benefit those with access to technology, thus widening the wealth gap.

Our financial markets face manipulation. AI systems could perform high-frequency trading to influence financial market activity.[114] We face possible skill erosion; humans will no longer retain the knowledge to perform certain tasks.[115] Lastly, the resources required to power certain AI systems rely on materials that are derived from exploitation.[116]

G. Safety

Expanding on the general societal issues noted above, there are several safety concerns involving AI. How do we respond when AI systems that operate in critical roles fail and cause harm? We noted above AI's potential to manipulate emotions that could lead to psychological harm, but there is also the overdependence on AI that could lead to loss of human skills and abilities. Lastly, AI has been shown to behave unpredictably, which may result in harmful or unintended consequences.[117]

H. Legal and Regulatory Challenges

The area in which the law struggles now involves assignment of liability when AI causes damage or harm. The courts are also grappling with issues involving intellectual property, including copyright (e.g., training data protections),[118] ownership of output and invention patenting. Current laws and regulations have failed to keep pace with AI development. We will also encounter difficulty enforcing laws across borders as most technology companies offer global AI systems.

I. Loss of Human Centricity and Control

We mentioned earlier the concern that AI develops autonomously without a human in the loop. The existential threat where AI systems operate beyond human understanding and control has been the subject of science fiction but has surfaced more as a probable fact.[119] We encounter the risk that AI may make critical decisions without human oversight or ethical considerations. Further, AI decisions may not value human life nor human generated output.[120] We are imperiled by AI that makes moral decisions without human empathy or understanding.[121]

LEGAL PROFESSION IMPACT

I. Ethical Impact

In the previous portion of this report, we explored the varying benefits and risks of AI and AI-based tools. When using any technology in legal practice, attorneys must remain compliant with the Rules of Professional Conduct. With generative AI tools, the number of rules implicated may be surprising.[122]

A. Duty of Competency/Techno-solutionism

> "*A refusal to use technology that makes legal work more accurate and efficient may be considered a refusal to provide competent legal representation to clients*."[123]

Rule 1.1 of the Rules of Professional Conduct (RPC) requires that a lawyer provide competent representation to a client. Comment 8 to RPC Rule 1.1 asserts that keeping abreast of "the benefits and risks associated with technology the lawyer uses to provide services to clients" is an element of competency. However, a recent LexisNexis survey found that only 43% of U.S. attorneys use (or plan to use) these tools professionally.[124] The need for more education, training and proficiency with the technology is apparent.

In addition to competence, attorneys must resist viewing these tools through a techno-solutionism lens. "Techno-solutionism"[125] is the belief that every social, political and access problem has a solution based in development of new technology. In this case, some view generative AI as the solution to the access to justice problem. As infamously demonstrated in the *Avianca* case,[126] in which an attorney utilized ChatGPT (a generative AI tool) to write a brief that contained fictitious legal precedent, attorneys cannot rely on technology without verification. RPC Rule 5.3 imposes a supervisory obligation on attorneys with respect to nonlawyer work. In the *Avianca* case, the "nonlawyer" was the tool itself.

B. Duty of Confidentiality & Privacy

RPC Rule 1.6 states, in part, that "[a] lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent." This duty of confidentiality also extends to what client information a lawyer may share when using certain generative AI tools. Because AI models depend on data to deliver salient results, privacy protection must become an integral part of their design.[127] Confidentiality concerns arise when entering information into AI engines, such as chatbots, and when such entries are then added to the training set for the AI. Such uses may violate protective orders for prior and future cases involving different parties. These concerns are compounded when chatbot results are analyzed by evaluative AI. For example, if biometrics data is analyzed by a chatbot to assist a mediator in preparing a mediator's proposal, multiple levels of confidentiality concerns arise. Such issues are especially important when some or all data that the AI "learns" is used for training the AI for work on future cases. Lawyers should cautiously use these tools, being mindful of a client's privacy.

In fact, the California bar association[128] recommends that lawyers inform their clients if generative AI tools will be used as part of their representation. The Florida bar association[129] takes its recommendation a step further, suggesting that lawyers obtain informed consent before utilizing such tools. Whether an attorney informs the client or obtains formal consent, the ethical obligation to protect client data remains unchanged from the introduction of generative AI tools.

C. Duty of Supervision

As noted earlier, RPC Rule 5.3 imposes a duty to supervise non-lawyers involved in client representation. In 2012, the American Bar Association amended Model Rule 5.3 to clarify that the term "non-lawyers" includes non-human entities, such as artificial intelligence technologies.[130] Despite the cautionary tale set by the *Avianca* case, a prominent California law firm has submitted hallucinated cases in its legal briefs.[131] Dennis P. Block and Associates, which handles tenant evictions, was fined $999 for its ethical violation – a paltry sum considering the societal impact of wrongful evictions.

D. Unauthorized Practice of Law

To begin a discussion about what constitutes the unauthorized practice of law (UPL) and specifically how use of generative AI, including LLMs, such as ChatGPT, Claude, Bard, and Mid-journey, may be considered UPL, we first examine what is the practice of law.

While there is no nationally agreed definition of what constitutes the practice of law, the ABA Model Rules provides one (discussed below). Some states have also fashioned their own definitions of the practice of law. Yet, without a uniform definition and precise meaning of the practice of law, we fall upon the adage: "You know it when you see it."

The ABA defines the practice of law as the application of legal principles and judgment regarding the circumstances or objectives of a person that require the knowledge and skill of a person trained in the law. However, New York State does not offer a precise definition of the term. ABA Model Rule 5.5 forbids lawyers from engaging in the unauthorized practice of law. Section (b) of the rule states:

> A lawyer who is not admitted to practice in this jurisdiction shall not: (1) except as authorized by these Rules or other law, establish an office or other systematic and continuous presence in this jurisdiction for the practice of law; or (2) hold out to the public or otherwise represent that the lawyer is admitted to practice law in this jurisdiction.

Similarly, Rule 5.5 of the New York RPC defines the unauthorized practice of law in this manner:

> (a) A lawyer shall not practice law in a jurisdiction in violation of the regulation of the legal profession in that jurisdiction. (b) A lawyer shall not aid a nonlawyer in the unauthorized practice of law.

Based on these rules, AI programs that do not involve a human-lawyer in the loop in providing legal advice arguably violate the rules and may be considered UPL. Thus, "AI programs cannot give legal advice unless a human lawyer is involved. In the age of AI, legal ethics preserves a human element in the practice of law."[132]

Case Law: Lawsuits Against AI Developers & UPL

Lola v. Skadden, Arps, Slate, Meagher & Flom LLP, 620 Fed. Appx. 37, 45 (2nd Cir. 2015). "According to the Lola decision, if a lawyer is performing a particular task [like document review] that can be done by a machine, then that work is not practicing law."[133] The court also interpreted North Carolina's law to imply, however, that the practice of law requires "at least a modicum of independent legal judgment."[134]

Janson v. LegalZoom.com, Inc., 802 F. Supp. 2d 1053, 1064 (W.D. Mo. 2011). The court held that filling out blank forms like the ones provided on LegalZoom's website "does not constitute the unauthorized practice of law." -Further, in a settlement between LegalZoom and the North Carolina Bar Association, LegalZoom agreed to have a licensed attorney review blank templates offered to customers in North Carolina and to clearly indicate to customers that the templates do not replace the advice of an attorney to ensure LegalZoom would not engage in the unauthorized practice of law.[135]

Based on current case law, AI programs can direct clients to the forms they need to fill out. However, these programs may not give any advice as to the substance of the client's answers because that would be replacing the work of a human lawyer.[136]

E. Attorney-Client Privilege and Attorney-Work Product

> "*There's not a lot of thought given to whether the information that's provided [to the chatbot] is covered by attorney client privilege.*" ~ Jay Edelson, CEO and founder of Edelson PC

One of the oldest recognized privileges regarding confidential information, the attorney-client privilege, "shields from disclosure any confidential communications between an attorney and his or her client made for the purpose of obtaining or facilitating legal advice during a professional relationship" so long as the communication is "primarily or predominantly of a legal character."[137]

The overarching purpose of this privilege is to allow for full and frank communications or discussions between attorneys and their clients. The attorney-client privilege has been defined as:

> a legal privilege that works to keep confidential communications between an attorney and their client private. Communications made to and by a lawyer in the presence of a third party may not be entitled to this privilege on grounds that they are not confidential. The privilege can be affirmatively raised in the face of a legal demand for the communications, such as a discovery request or a demand that the lawyer testify under oath. A client, but not a lawyer, who wishes not to raise attorney-client privilege as a defense is free to do so, thereby waiving the privilege. This privilege exists only when there is an attorney-client relationship (Cornell University Law School, Legal Information Institute/LII, posting by the Wex Definitions Team).

The statutory attorney-client privilege in the State of New York is found in Civil Procedure Law and Rules 4503(A)(1), which states:

> Unless the client waives the privilege, an attorney or his or her employee, or any person who obtains without the knowledge of the client evidence of a confidential communication made between the attorney or his or her employee and the client in the course of professional employment, shall not disclose, or be allowed to disclose such communication, nor shall the client be compelled to disclose such communication, in any action, disciplinary trial or hearing, or administrative action, proceeding or hearing conducted by or on behalf of any state, municipal or local government or by the legislature or any committee or body thereof.

While discovery requests for privileged information may reveal attorney-client privileged information, so too may the use of generative AI tools such as ChatGPT or GPT-4.

Model Rules of Professional Conduct 1.6(a) and (c):

(a) A lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized in order to carry out the representation or the disclosure is permitted by paragraph (b).

(c) A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.

New York RPC Rule 1.6:

(a) A lawyer shall not knowingly reveal confidential information, as defined in this Rule, or use such information to the disadvantage of a client or for the advantage of the lawyer or a third person.

(c) A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure or use of, or unauthorized access to, information protected.

Comment to New York Rules of Professional Conduct 1.6(c):

- ➢ An attorney must "make reasonable efforts to safeguard confidential information against unauthorized access by third parties and against inadvertent or unauthorized disclosure by the lawyer or other persons who are participating in the representation of the client or who are otherwise subject to the lawyer's supervision."
- ➢ "Unauthorized access to, or the inadvertent or unauthorized disclosure of, information protected . . . does not constitute a violation of paragraph (c) if the lawyer has made reasonable efforts to prevent the unauthorized access or disclosure."

Focusing on the language in the Cornell University Law School LII definition of attorney-client privilege – "communications made to and by a lawyer in the presence of a third party may not be entitled to this privilege on grounds that they are not confidential" – how then may attorney-client privileged information or attorney-work product be revealed when directly and indirectly using generative AI tools such as ChatGPT or GPT-4.[138]

For example, through:

- ➢ Direct Use of ChatGPT as an app (the user directly enters a prompt that contains your private or confidential information, which then goes into ChatGPT)

- Indirect Use of GPT-4 that is embedded in search engines such as Microsoft Bing (the user enters a prompt that contains private or confidential information, which then goes into the generative AI app)
- Use of Application Programming Interface/API (using some other application that connects to ChatGPT via the API, private or confidential information is inputted into ChatGPT)
- ChatGPT plugins (accessing other applications from within ChatGPT via plugins, which conveys your private or confidential information further into ChatGPT and other places too. With plugins, other users/persons can see/view your private or confidential information).

Key Points for attorneys to be aware of and consider when utilizing ChatGPT and other similar generative AI tools include:

- Licensing Information
- Terms of Use
- Privacy Policies
- Frequently Asked Questions/FAQs list
- Data that is supplied to or inputted into ChatGPT may be used for training purposes or to refine/improve the AI model (For example, ChatGPT developers may view the input and conversation history of its users and users' personal information, including log/usage data, to analyze/improve/and develop ChatGPT services).
- Data that is supplied to or inputted into ChatGPT may be viewed by and disclosed to third parties/vendors in the training of the AI model.

- Data output by ChatGPT may be viewed by third parties, including opponents and adversaries.

Pursuant to the Model Rules of Professional Conduct and New York RPC, lawyers must take reasonable efforts to prevent inadvertent and unauthorized disclosure of or access to client information. When utilizing generative AI tools such as ChatGPT, attorneys need to be knowledgeable about the technology they are using and/or ask for assistance from those lawyers or trusted technology experts who do understand its use and limitations, including IT personnel. If none of these options is possible, then the attorney should not utilize such technologies until they are competent to do so per the duty of competency.[139]

AI and Cybersecurity Risks

Open AI/ChatGPT may raise both ethical violations and cybersecurity issues. For example, "if there is a cyber intrusion [into OpenAI or ChatGPT], not only will that data potentially be lost to threat actors, but they could conceivably also obtain the firm's searches… [gaining] access into the mind of a lawyer and the arguments they might be raising."[140]

Preservation of Data

Data preservation and litigation hold obligations may present similar challenges for attorneys and the court. If the data that is inputted into the AI application is temporary/ephemeral, but also relevant and responsive to the litigation, parties have the duty to preserve this electronically stored information. Yet, how do you preserve what may no longer exist?

F. <u>Candor to the Court</u>

When using ChatGPT or other similar AI tools, attorneys must verify the accuracy of the information and legal authority produced by such tools. Attorneys' signatures and attestations appear on legal documents submitted to the court, documents which make representations about case law and other authorities relied upon in support of the attorney's case. Regardless of the use of and reliance upon new and emerging technologies like generative AI tools, as officers of the court and in the interest of justice, attorneys must identify, acknowledge and correct mistakes made or represented to the court.

The following ABA Model Rules of Professional Conduct and New York RPC guide attorneys in their use and reliance on information obtained from AI tools:

> *M.R.P.C.* 3.3 (Candor to the Tribunal):
>
> "(a) A lawyer shall not knowingly:
> (1) make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer;
>
> (3) offer evidence that the lawyer knows to be false. If a lawyer, the lawyer's client, or a witness called by the lawyer, has offered material evidence and the lawyer comes to know of its falsity, the lawyer shall take reasonable remedial measures, including, if necessary, disclosure to the tribunal. A lawyer may refuse to offer evidence, other than the testimony of a defendant in a criminal matter, that the lawyer reasonably believes is false."
>
> Comment [2] to *M.R.P.C.* 3.3:
> "although a lawyer in an adversary proceeding is not required to present an impartial exposition of the law or to vouch for the evidence submitted in a cause, the lawyer must not allow the tribunal to be misled by false statements of law or fact or evidence that the lawyer knows to be false."
>
> Rule 3.3(a) (1) of the New York Rules of Professional Conduct prohibits lawyers from making false statements of fact or law to a court and requires correction of any false statements previously made during the case.

AI Hallucinations: What Are Hallucinations, and Why Do They Occur?

Hallucinations are incorrect/unreliable information produced by an LLM or generative AI chatbot, such as ChatGPT. In simplest terms, a hallucination is a euphemism for a lie. As an LLM, ChatGPT is trained on a vast amount of data to recognize patterns in language and then produce/generate a response it predicts is relevant and responsive to the user's input or prompt.[141]

> *AI hallucination is a phenomenon wherein a large language model, often a generative AI* chatbot *or* computer vision *tool, perceives patterns or objects that are nonexistent or imperceptible to human observers, creating outputs that are nonsensical or altogether inaccurate." "Generally, if a user makes a request of a generative AI tool, they desire an*

> *output that appropriately addresses the prompt (i.e., a correct answer to a question). However, sometimes AI algorithms produce outputs that are not based on training data, are incorrectly decoded by the transformer or do not follow any identifiable pattern. It "hallucinates" the response*.[142]

Case Law and Hallucinations

U.S. v. Prakazrel Michel, No. 1:19-cr-00148-1 (CKK)(D.D.C.) (motion filed Oct. 16, 2023). Defendant, convicted of money laundering and corrupt political influencing, alleges that his attorney's reliance on AI for his closing argument constituted ineffective assistance of counsel. Defendant argues that his attorney's "closing argument made frivolous arguments, misapprehended the required elements, conflated the schemes and ignored critical weaknesses in the government's case."

Ex Parte Allen Michael Lee, 673 S.W.3d 755 (Tex. App. Jul. 19, 2023). In denying the petitioner's motion for a new bail hearing, the court found that petitioner's moving brief, prepared by counsel, contained citations that did not exist and arguments that appeared to be generated by generative AI.

Mata v. Avianca, Inc., No. 22-cv-1461 (PKC), 2023 WL 4114965 (S.D.N.Y. June 22, 2023) (referenced in other portions of this report).

Donovan James Gates v. Christopher Omar, et al., No. 2022 cv 31345 (Col. Sup. Ct.). A lawyer used ChatGPT for research in connection with a motion to set aside summary judgment in a breach of contract matter, and the cases cited in the motion were nonexistent. The lawyer, who had been practicing in Colorado for 1.5 years and in civil litigation for 3 months, said he turned to ChatGPT because it was his first civil litigation and he wanted to save his client money by relying on the technology to conduct the research. As of June 2023, the Court was considering sanctions.

Attorneys cannot solely rely upon information provided by generative AI. Attorneys may instead use generative AI as a starting point and must independently review case citations, arguments and any other information/output produced by generative AI.

Deepfakes – Synthetic Media as Evidence in Court

With the understanding that the fundamental purpose of a trial is its truth seeking function, for "the very nature of a trial [i]s a search for truth,"[143] evidentiary issues surrounding Deepfakes – a form of AI called deep learning that makes images of fake events[144] – may also implicate the Duty of Candor to the Court. Deciding issues of relevance, reliability, admissibility and authenticity may still not prevent deepfake evidence from being presented in court and to a jury. "One of the fundamental tenets of the American legal system is that the trier of fact—either the judge or the jury—is best equipped to find the truth based on the evidence presented. But individuals cannot consistently determine truth from lies as they confront deepfakes."[145]

G. Judges' Ethical Obligations

The Model Code of Judicial Conduct mandates: "*A judge shall uphold and promote the independence, integrity and impartiality of the judiciary.*" ABA Model Code of Judicial Conduct, Canon 1.[146] How does Canon 1 of the Model Code of Judicial Conduct align with judicial use of generative AI, such as ChatGPT?

"The human aspect of intelligence that cannot be artificially constructed is that of 'judgment.'" While AI can and does assist judges in a variety of ways, judges will always have the responsibility of exercising their own judgment: the human trait of independent judgment.[147]

According to New York Rules of Professional Conduct Rule 5.4: Professional Independence of a Lawyer:

> (c) Unless authorized by law, a lawyer shall not permit a person who recommends, employs or pays the lawyer to render legal service for another to direct or regulate the lawyer's professional judgment in rendering such legal services or to cause the lawyer to

> compromise the lawyer’s duty to maintain the confidential information of the client under Rule 1.6.
>
> Comment [2]
> This Rule also expresses traditional limitations on permitting a third party to direct or regulate the lawyer’s professional judgment in rendering legal services to another. See also Rule 1.8(f), providing that a lawyer may accept compensation from a third party as long as there is no interference with the lawyer’s professional judgment and the client gives informed consent.

How does this rule and comments to the rule align with attorneys’ use of generative AI such as ChatGPT? Attributed to the 16th U.S. President and attorney Abraham Lincoln: “A lawyer’s time and advice are his stock in trade.” It follows then that an attorney’s time, advice and professional judgment are what clients expect and rely upon when retaining a lawyer/law firm for representation in a matter. While AI can and does assist lawyers in a variety of ways, attorneys do not shed their professional responsibility of exercising their own “independent judgment” in client matters.

II. Access to Justice

A. <u>Introduction</u>

The rapid development of AI has the potential to have a significant impact on access to justice in the American legal system. While AI and especially generative AI is generally causing disruption in the market for legal services, this impact is likely to be even greater when discussing access to justice.

For some time, there has been an enormous gap in access to legal services. A recent survey found that 66% of the U.S. population experienced at least one legal issue in the past four years, with just 49% of those problems having been completely resolved. In the United States, it is well documented that there are many geographical regions that do not have enough human lawyers. A recent survey found that low-income Americans did not receive any or enough legal help for 92% of their civil legal problems.

Generative AI tools such as ChatGPT have the potential to enhance the accessibility, efficiency and affordability of pro bono legal services. Generative AI could truly transform the way in which legal services are provided, and the tremendous opportunities and challenges of this technology are magnified when addressing pro bono services to clients. But there are clearly risks too as highlighted above. As we have already discussed, early generative AI tools have been unable to consistently provide accurate legal advice to their users. While more accurate tools may be developed, given the reach of the corporations promoting existing generative AI tools, new market entrants may not come to the attention of those most in need. Where generative AI may make it easier for those without a lawyer to find an answer to a legal issue, it may make it harder for them to find the correct answer.

We cannot underestimate the additional cost in terms of court resources to research, verify and challenge incorrect AI-generated legal opinions and arguments. Coming at a time when many courts are already stretched thin with unacceptably long waiting times in some jurisdictions for a hearing, adding to this strain could lead to more injustice.

B. Pro Bono Organizations Using Generative AI

Pro bono organizations often have faced challenges in meeting the needs of their clients and in hiring sufficient attorneys and staff to support the many matters that they take on. Staff and attorneys working for legal aid organizations are perpetually understaffed and overworked. AI has the potential to transform the way in which some pro bono organizations serve their clients.

Legal services organizations have limited resources and are unable to serve all the individuals who seek their assistance. Generative AI can help organizations put in place a triage process for pro bono clients that can help to analyze many potential matters and can enable these organizations to serve many more clients than they currently serve. Many organizations spend large amounts of time screening potential clients, but an AI chatbot could effortlessly screen potential clients and gather basic information about their legal issues. Several organizations have started building tools to access basic legal information and they have found that generative AI is a game-changer when it comes to client intake.

Pro bono attorneys have found that generative AI tools are excellent at summarizing and extracting relevant information from documents, translating legalese into plain English and helping to quickly analyze thousands of existing court forms. In addition, ChatGPT and other similar generative AI tools can identify potential clients' legal needs and build out and maintain legal navigators.

Pro bono organizations are seeing how generative AI can even assist them in putting together navigator-type tools that can help guide clients seeking legal services. For example, a site powered by generative AI technology could provide a step-by-step guide to getting divorced, explain how to file a claim against an unlawful landlord or provide legal and other support options for domestic violence survivors. This is not a hypothetical scenario, as such systems have already been put into place by some legal services organizations, and these tools will only become more powerful, intelligent and accurate as generative AI becomes more and more sophisticated.

In addition, language is often a barrier to justice. Members of some communities may struggle to understand English, and that struggle can be magnified when faced with the formal legal language that is often used in court documents and agreements. Generative AI tools can be utilized to simplify, summarize and translate documents.

Legal services organizations are often challenged by the research and writing that they must perform in order to properly support a matter. Generative AI can help with legal research and document preparation, which in turn can help to resolve cases more quickly. It could also help to draft legal documents, such as contracts or pleadings by providing template language and helping users to fill in necessary information. While drafting a complaint would have taken many hours in the past, with the help of generative AI, a complaint could be drafted in minutes.

If accurately and properly used, these tools may have the potential to bring legal services to those who cannot afford it and to make legal services organizations run more efficiently.

C. Will Generative AI Tools Prove to Be Too Expensive?

While generative AI has the potential to greatly benefit access to justice, there are some who believe that this technology could potentially hinder, and not help, access to justice.

It has been noted that while this technology is developing at a fast pace, the industry is not currently structured to serve the interests of underserved populations and pro bono organizations. While there is potential for pro bono organizations and low-income individuals to take advantage of this technology, there is a risk that this technology could further exacerbate existing inequities.

While it might appear that the application of this technology will help to even the playing field, it remains to be seen how expensive it will be to properly utilize this technology in the practice of law. The development of AI technology is unregulated, and the companies developing and applying this technology to the legal profession have an interest in making a product that is attractive to those who are willing to pay for it. Many law firms are investing millions of dollars to implement AI solutions. Pro bono organizations run the risk of falling even further behind the big law firms.

Additionally, when one addresses assisting non-lawyers with justice problems it is possible that new generative AI tools may not make a significant difference in improving access to justice for low-income and minority communities. Those who need legal services from this constituency are less likely to be able to use AI tools due to fees to use these tools, limited internet access and literacy and language barriers.

Since this technology really does have the potential to improve access to justice, it is crucial that pro bono organizations and low-income individuals be given access to these tools. While this may be difficult, it is imperative that this technology be available to all who are in need of legal services.

D. Use of AI by Non-Attorneys

In its first year of widespread use by the public, Chat GPT and generative AI have been used by the general public for a wide range of uses. Non-lawyers will be able to readily interact with generative AI to ask a variety of legal questions. These uses of generative AI will present challenges for bar associations, courts and the legal community as a whole.

What one must realize when looking at this issue is that currently the majority of the parties in civil cases in state and local courts lack legal representation. Therefore, the question becomes: Are the people, who otherwise would not have legal counsel, better served by at least having a chatbot to assist them?

One of the challenges with non-attorneys using generative AI to assist with legal issues is the possibility of receiving misleading information. In its current iterations, generative AI is likely to provide an answer to a legal question, but it might do so without providing an indication that the confident answer is without a proper legal foundation. Some AI companies have included warnings in their user agreements about using their tools to provide legal advice. For example, OpenAI's online usage provisions state the following:

Prohibited use – "Engaging in the unauthorized practice of law, or offering tailored legal advice without a qualified person reviewing the information."

It is questionable whether individuals and new tools will abide by such prohibitions. Even if some tools include such warnings there is nothing to stop someone from asking a chatbot for legal advice or drafting papers for them. If a non-lawyer has a chatbot draft a brief or complaint, they are not in as good a position as an actual lawyer to know if the filing contains falsehoods, biases, incorrect cases or other AI hallucinations.

In addition, even though individuals who cannot afford an attorney will potentially benefit from generative AI tools, there will be some barriers to access, including more limited access to the internet and computers by the people experiencing homelessness or those living in poverty. Asking such tools the right questions also requires some skill. While a person may download advice on how to frame a question (i.e., developing a “prompt”) correctly, some non-lawyers, particularly in those sections of society that have been traditionally underserved by the law, may struggle to design the correct prompt. In addition, much of the information that one would need to develop a system that provides accurate legal information would require access to databases that are generally behind a paywall (i.e., Westlaw, Lexis, Law360), which could potentially result in a cost to users.

Another potential issue stems from the fact that generative AI tools might not account for multiple, interrelated issues, which could include family, criminal, housing, employment, etc. It is possible that an answer from a chatbot could be correct for one issue but harmful in the context of the other issues. It is in this situation where a chatbot likely will never be able to fully replace a human. Generative AI will never have the same level of empathy as a human, and when individuals are seeking legal services, they often need someone to “hold their hand” and that simply is not possible with a chatbot (at least for the time being).

It should be noted that non-lawyers are already able to gather the same kind of advice or information that a chatbot provides by searching online for legal materials and legal information.[148] While some information found online may be correct, other information may be outdated, suspect or simply incorrect. Generative AI is basically a new interface to this online information that has the advantage of being an interactive conversational tool. If this can make information more accessible and let people know if they even have a legal issue, this will prove to be a positive development.

In addition, generative AI solutions are available 24/7. It could take days, weeks or months for a low-income plaintiff to find an attorney to meet with them or represent them for a matter. Generative AI is generally efficient and is scalable, allowing it to provide information to many people at once. While it's true that generative AI may be challenged when dealing with multiple overlapping issues, it will surely be a positive development for individuals who are unable to afford an attorney.

The reality of the situation is that generative AI is here, and it is not going away but will rather become more advanced and more available to the general public as time goes on. It should be noted that the challenges facing the legal profession are not unique. The medical profession also is addressing the challenges presented by patients who have consulted with generative AI and arrive at an appointment with opinions on what is the correct medical advice. Lawyers will similarly be challenged by clients who have compiled information and learned about their legal options using generative AI.

We believe it is important not to dismiss innovation, and to allow vendors and companies to develop programs that will help guide the general public. It is just as important for attorneys to educate themselves on AI so they can utilize it and understand how their clients may be using it as well.

E. Implications of AI Judges or Robo Courts

One other area where AI may have a great impact on access to justice relates to the utilization of AI by judges and courts. At the time of this Report, there are only a few examples of robo courts or AI judges being utilized to resolve disputes, and those trials have had mixed results.

For example, in 2019, Estonia planned to use robo judges for small claims procedures. The Estonian government said that those reports were misleading.[149] In Australia, a system designed to use technology to assess government payments has already failed.[150] But as generative AI becomes more sophisticated, it will become more feasible to have AI arbiters decide small claims courts matters or arbitration matters where both parties consent to an AI arbiter.

It is not clear at this time how widespread this practice will become and how it will impact access to justice. In some ways, it may make it more likely for those with little knowledge of the law and courts and those who have little financial means to have their day in court. An AI judge may also be less likely to be influenced by a prominent attorney or big-name firm. However, most people will generally not want their disputes to be decided by a computer or algorithm.

We are not quite yet to the point of AI judges replacing some portion of the judiciary, and that may never happen, but it is likely to be raised as a possibility in the future. We are already at a point where AI is being used to mediate matters, where both parties agree to the use of AI. While we have not quite arrived in a sci-fi world populated by robo judges, we do need to be wary of AI being used in lieu of judges, and we need to be well positioned to gauge the potential benefits and risks of using AI judges in certain situations.

III. Judicial Reaction/Responses to Generative AI

A. Introduction

Artificial intelligence has been in use by the legal profession and its clients for a long time. In November 2022, generative AI burst onto the scene through one program, launched by Open AI, known as ChatGPT. Since then, the use and varieties of generative AI platforms has expanded on a seemingly daily basis, and attorneys and clients are evaluating generative AI technology and how it could be used – and abused – in litigation. This section of the Task Force Report will introduce the reader to those uses and abuses.

B. Uses of AI and Generative AI

Other sections of this Report have discussed the technologies. For now, we consider some uses of AI and generative AI. Focusing on AI in general, it is in widespread use for:

- Identification (for example, airports and workplaces)
- Security (for example, to access cell phones and bank accounts)
- Law enforcement (for example, to identify suspects)
- Retail (for example, to identify shoppers)
- Human resources (for example, to interview and hire employees)

And, in addition to these uses, AI is used extensively for collection, review and production of ESI.

Generative AI takes AI to a new level. As we know, generative AI ingests data and, in response to “prompts,” generates an answer. Generative AI is being used by the legal profession and other entities to, among other things:

- Draft and edit documents
- Conduct legal research
- Contract review
- Predictive analytics
- Chatbots for legal advice
- Brainstorming
- Summarize legal narratives
- Convert “legalese” into plain language

C. Causes of Action Arising out of AI and Generative AI

We are at the tip of the proverbial iceberg when thinking about causes of action (and we are only speaking of civil litigation here – there are uses of AI and generative AI that could give rise to criminal proceedings, including, for example, "deepfakes" that might be prosecuting under federal or state criminal laws). Here are examples of causes of action:

- Breach of privacy
- Discrimination
- Copyright infringement
- Malicious uses such as defamation
- Cyber breach
- Employment-related

These causes of action might derive from common law. However, statutes or regulations might also give rise to litigation as well as regulatory proceedings. Examples include:

- Section 5 of the Federal Trade Commission Act
- Discrimination actionable under the Equal Employment Opportunity Act and state equivalents
- The Illinois Artificial Intelligence Video Interview Act
- Illinois Biometric Information Privacy Act
- New York City Local Law Int. 1894-A
- New York City Local Law Int. 1170-A

Attorneys and clients should expect to see legislation at the state and federal levels to address AI and generative AI, particularly with regards to employment, insurance, medical services, elections, housing and AI generated media.

It may also be useful to note that overseas laws attempting to govern AI may have extra-territorial effects. For example, the EU AI Act (summarized in Appendix A) was agreed in principle at an EU level in 2023. While there is still some way to go before this will become law, the EU AI Act is designed to also regulate the use of AI by the U.S. and other entities outside the EU. Coupled with this, the EU has introduced an EU AI Pact, which could lead to some U.S. corporations agreeing to be bound by the EU AI Act's provisions as early as this year.

D. Discovery

Prior sections of this Report have described the technology behind AI and generative AI. Bearing in mind how technology might make mistakes and lead to injury, economic or personal, it is expected that regulatory requests for information and civil discovery demands that focus on, for example, alleged bias will be made. Discovery into bias might present questions about the nature of the data fed into the AI or generative AI and how algorithms used by the AI or generative AI "operated," as well as questions related to the prompt used to generate something. Such questions will raise other questions about the need for non-testifying or testifying experts. Moreover, as already outlined in this Report, the competence of attorneys to deal with this technology might present ethical questions.

E. *Avianca* and Judicial Reactions to Generative AI

Not only is generative AI now mainstream, but it has featured in judicial decisions and in "prophylactic" orders. The first of the decisions is *Avianca*, which is discussed below.

In *Mata v. Avianca, Inc.*,[151] the plaintiff's attorneys "submitted non-existent judicial opinions with fake quotes and citations created by *** ChatGPT, then continued to stand by the fake opinions after judicial orders called their existence into question." The court held that:

- The attorneys acted with subjective bad faith and violated Federal Rule of Civil Procedure 11.

- The plaintiff's firm was jointly and severally liable for the attorneys' Rule 11 violation.
- Sanctions under U.S.C. 1927 could not be imposed because, "[r]eliance on fake cases has caused several harms but dilatory tactics and delay were not among them."
- "Alternatively," to Rule 11, sanctions were imposed under the inherent power of the court.
- $5,000.00 penalty imposed jointly and severally.

The court also required the attorneys "to inform their client and the judges whose names were wrongfully invoked of the sanctions imposed."

Since *Avianca* was decided, other courts have addressed generative AI in decisions (discussed earlier in this Report). However, and of particular interest to the Task Force, individual judges (and one United States bankruptcy court) have directed attorneys who appear before them and who use generative AI to take certain actions. Here is a "sampler:"

United States District Judge Brantly Starr of the Northern District of Texas has imposed a certification requirement:

> All attorneys and pro se litigants . . . must, file on the docket a certificate attesting either that no portion of any filing will be drafted by generative artificial intelligence (such as ChatGPT, Harvey.AI, or Google Bard) or that any language drafted by generative artificial intelligence will be checked for accuracy, using print reporters or traditional legal data bases, by a human being.

United States District Court Judge Michael Baylson of the Eastern District of Pennsylvania has issued a Standing Order for all actions assigned to him:

> If any attorney for a party, or a pro se party, has used artificial intelligence ('AI') in the preparation of any complaint, answer, motion, brief, or other paper, filed with the Court, and assigned to Judge Michael M. Baylson, MUST, in a clear and plain factual statement, disclose that AI has been used in any way in the preparation of the filing, and CERTIFY,

that each and every citation to the law or the record in the paper, has been verified as accurate.

These and other orders are problematic for several reasons, including:

- Might attorney work product be implicated?
- Might the use of the term “artificial intelligence” (rather than generative AI) sweep into a disclosure obligation much more than generative AI? (For example, if an attorney uses computer-assisted review to cull and make a production of ESI, would the order encompass that use?).

Judges issue local rules for court management and in reaction to or to get ahead of issues that may arise or have the potential to arise in their courtrooms (in real time), regardless of existing rules which address the same concerns!

In time, with better understanding of the new and emerging technologies, and with more precision in language when referencing these emerging technologies, the language in the local rules will more precisely match and address the concerns of the court and so, achieve what these judges’ orders were designed to do.

LEGISLATIVE OVERVIEW AND RECOMMENDATIONS

I. Legislative Overview

While the Task Force reviewed several pieces of proposed and passed legislation (summarized in Appendix A hereto), we do not endorse any specific pending legislation. However, as the recommendations below reflect, we do recommend certain changes to the RPC that will help clarify lawyers' ethical duties when using AI and generative AI tools.

II. Recommendations

The Task Force recommends the following for NYSBA adoption:

First, the Task Force recommends that NYSBA adopt the AI/Generative AI guidelines outlined in this report and commission a standing section or committee to oversee periodic updates to those guidelines. Daily, we learn more about the capability of the technology to transform society. As the impacts are continual, so should the updates to these guidelines be as well.

Second, we recommend a focus on educating judges, lawyers, law students and regulators to understand the technology so that they may apply existing law to regulate it. Many of the risks posed by AI are more sophisticated versions of problems that already exist and are already addressed by court rules, professional conduct rules and other law and regulations. Furthermore, many risks are mitigated through understanding the technology and how AI will utilize data input into the AI system. For example, concerns related to client privacy and confidentiality can be alleviated by utilizing a "closed system" AI, which provides for anonymous queries that are not incorporated into the AI training data.

In addition to legislation, if and when determined to be necessary, the Task Force suggests that we create a comprehensive education plan for judges, lawyers, law students and regulators so they can address the risks associated with AI using existing laws and regulations, such as providing education on how the technology works and determining if an AI system will save and utilize prompts as training data. This approach has already been adopted effectively in other jurisdictions. For example, the Italian Data Protection Authority, the Guarante per la Protezione dei Dati Personali, has already effectively used GDPR in a number of AI-related cases, including to modify or restrict the operations of the ChatGPT and Replika AI chatbots.[152] This approach will allow the law to develop in a fact-based way along with the rapidly changing technology.

Comments to the rules of professional conduct, best practices, continuing education programs and state bar opinions can also aid in this process. For instance, in the Preamble to the RPC, we recommend including a general statement about the importance of competence with technology by adding "including . . . artificial intelligence" therein. Further, we would expand Comment [8] to Rule 1.1 to add that the duty of competence obligates lawyers to: (a) keep abreast of and be able to identify technology (including AI and generative AI) that is generally available to improve effective client representation and enhance the quality of legal services; (b) determine whether the use of AI will in fact augment the legal service to a specific client; and (c) attain a basic understanding of how AI-based tools operate to achieve the results and outputs sought.

Third, the Task Force recommends that legislatures and regulators seek to identify risks associated with the technology that are not addressed by existing law. This may involve extensive hearings, studies involving experts in AI and increased costs. Once such risks are identified, new laws and regulations should be crafted to address those risks.

Fourth, the rapid advancement of AI prompts us to examine the function of the law as a governance tool. Some of the key functions of the law in the AI context are: (i) expressing social values and reinforcing fundamental principles; (ii) protecting against risks to such values and principles; and (iii) stabilizing society and increasing legal certainty. Recommendations here involve:

a. AI as a General-Purpose and Dual-Impact Technology: The governance of AI should consider AI's nature as a classic dual-impact phenomenon. AI can improve many aspects of society but also has the potential to cause harm if left unchecked. Regulation should consider focusing on the effects of the technology on individuals and society, rather than the technical aspects of the technology itself (such as the algorithms or databases).

b. Regulatory Spectrum: The governance of AI should be tailored to the risks posed by AI applications. It can adopt varying degrees of regulatory intrusiveness, with the spectrum potentially extending from detailed legal regulation at one end of the spectrum to self-regulation on the other end of the spectrum, with a principles-based approach in the middle of the spectrum. The approach chosen to address a particular risk or problem should consider:

- the sector involved (e.g., law enforcement or health care)
- the importance of the social activity at hand (e.g., hiring applicants or making loans)
- the rights affected (e.g., due process or privacy)
- the risks associated with the use and impact of AI (e.g., job loss or misinformation)

c. Comprehensive vs. Specific Regulation: Foundationally, legislators should determine if regulations entail a comprehensive approach (i.e., an overarching framework governing diverse AI applications and their social implications) or a sector-by-sector or industry-by-industry approach (i.e., considering the particular and often unique issues posed by AI in each sector or industry). Regulators should determine which approach is best, or develop some mix or combination of these approaches, depending on the sectors and problems at hand.

d. Global Cooperation: Another consideration in the regulatory approach involves jurisdictional reach. Can AI be effectively governed at the local, state or federal level, or does its governance necessarily require some degree of international or even global cooperation? We believe in local, state and federal regulation where appropriate, but also propose that local, state and federal regulation is likely to prove inadequate without international and sometimes global cooperation, because AI is a cross-border phenomenon rather than a local one. The following four elements of AI may elude regulations if they are confined to a specific geographic area:

i. Data, which is the input for AI, can move across borders (although data location is likely to enhance a jurisdiction's power to regulate AI);

ii. Algorithms programmable anywhere in the world;

iii. Algorithms exportable for use anywhere else in the world; and

iv. Outputs from algorithms transmitted to and applied in different jurisdictions.

AI & GENERATIVE AI GUIDELINES

The chart below reflects the Task Force's recommended guidelines when utilizing AI or generative AI tools (collectively, the "Tools") in legal practice. We will update these guidelines periodically as the technology evolves.

TOPIC	**GUIDANCE**
ATTORNEY COMPETENCE (RULE 1.1)	*A lawyer should provide competent representation to a client.* You have a duty to understand the benefits, risks and ethical implications associated with the Tools, including their use for communication, advertising, research, legal writing and investigation. Refer to Appendix B for resources to better under the Tools.
SCOPE OF REPRESENTATION (RULE 1.2)	*A lawyer shall abide by a client's decisions concerning the objectives of representation and, as required by Rule 1.4, shall consult with the client as to the means by which they are to be pursued.* Consider including in your client engagement letter a statement that the Tools may be utilized in your representation of the client and seek the client's acknowledgement. Refer to Appendix C for a sample language to include.
DILIGENCE (RULE 1.3)	*A lawyer should act with reasonable diligence and promptness in representing a client.* Consider whether use of the Tools will aid your effectiveness in representing your client.
COMMUNICATION (RULE 1.4)	*A lawyer shall explain a matter to the extent reasonably necessary to permit the client to make informed decisions regarding the representation.* While the Tools can aid in generating documents or responses, you must ensure that you maintain direct and effective communication with your client and not rely solely on content generated from the Tools.

TOPIC	**GUIDANCE**
FEES (RULE 1.5)	*A lawyer shall not make an agreement for, charge, or collect an excessive or illegal fee or expense.* If the Tools would make your work on behalf of a client substantially more efficient, then your use of (or failure to use) such Tools may be considered as a factor in determining whether the fees you charged for a given task or matter were reasonable. If you will add a "surcharge" (i.e., an amount above actual cost) when using specific Tools, then you should clearly state such charges in your engagement letter, <u>provided</u> that the total charge remains reasonable.
CONFIDENTIALITY (RULE 1.6)	*A lawyer shall not knowingly reveal confidential information.* When using the Tools, you must take precautions to protect sensitive client data and ensure that no Tool compromises confidentiality. Even if your client gives informed consent for you to input confidential information into a Tool, you should obtain assurance that the Tool provider will protect your client's confidential information and will keep each of your client's confidential information segregated. Further, you should periodically monitor the Tool provider to learn about any changes that might compromise confidential information.
CONFLICTS OF INTEREST (RULE 1.7)	*A lawyer shall not represent a client if a reasonable lawyer would conclude that the representation will involve the lawyer in representing differing interests.* Your use of the Tools in a particular case may potentially compromise your duty of loyalty under Rule 1.7, by creating a conflict of interest with another client. Rule 1.7 imposes a duty on you to identify, address and, if necessary, seek informed client consent for conflicts of interest that may result from your use of the Tools.

TOPIC	GUIDANCE
SUPERVISORY RESPONSIBILITIES (RULE 5.1)	*A lawyer with direct supervisory authority over another lawyer shall make reasonable efforts to ensure that the supervised lawyer conforms to the ethical rules.* As a supervising lawyer, you have a duty to ensure that the lawyers for whom you have oversight observe the ethical rules when utilizing the Tools.
SUBORDINATE LAWYERS (RULE 5.2)	*A lawyer is bound by the ethical rules notwithstanding that the lawyer acted at the direction of another person.* If you as the subordinate lawyer utilize the Tools as directed by your supervising attorney, you are independently required to observe the ethical rules. All rules described in these guidelines apply equally to your conduct.
RESPONSIBILITY FOR NON-LAWYERS (RULE 5.3)	*A law firm shall ensure that the work of nonlawyers who work for the firm is adequately supervised, as appropriate.* If the Tools are used by non-lawyers or paralegals (or the Tools themselves are interpreted to be "non-lawyers"), you must supervise their use to ensure compliance with the ethical rules. Further, you must ensure that the work produced by the Tools is accurate and complete and does not disclose or create a risk of disclosing client confidential information without your client's informed consent.
PROFESSIONAL INDEPENDENCE (RULE 5.4)	*A lawyer shall not permit a person to direct or regulate the lawyer's professional judgment in rendering legal services.* While the Tools are not a "person," you should refrain from relying exclusively on them or the output derived from them when providing legal advice and maintain your independent judgment on a matter.
UNAUTHORIZED PRACTICE OF LAW (UPL) (RULE 5.5)	*A lawyer shall not aid a nonlawyer in the unauthorized practice of law.* Understand that human oversight is necessary to avoid UPL issues when using the Tools, which should augment but not replace your legal work.

TOPIC	GUIDANCE
VOLUNTARY PRO BONO SERVICE (RULE 6.1)	*Lawyers are strongly encouraged to provide pro bono legal services to benefit poor persons.* The Tools may enable you to substantially increase the amount and scope of the pro bono legal services that you can offer. Considering Rule 6.1, you are encouraged to use the Tools to enhance your pro bono work.
ADVERTISING (RULE 7.1)	*A lawyer or law firm shall not use or disseminate or participate in the use or dissemination of any advertisement that: (1) contains statements or claims that are false, deceptive or misleading; or (2) violates an ethical rule.* You are responsible for all content that you post publicly, including content generated by the Tools. Further, you must be cautious when using the Tools for advertising or solicitation purposes to ensure that you comply with ethical guidelines regarding truthful and non-deceptive communication.
SOLICITATION AND RECOMMENDATION OF PROFESSIONAL EMPLOYMENT (RULE 7.3)	*A lawyer shall not engage in solicitation by in-person or telephone contact, or by real-time or interactive computer-accessed communication . . .* You may not use the Tools to automatically generate phone calls, chat board posts or other forms of solicitation, nor may you contract with another person to use the Tools for such purposes, as Rule 8.4 (Misconduct) prohibits you from using others to engage in conduct in which you personally could not engage.

CONCLUSION

This report offers no “conclusions.” As AI continues to evolve, so will the work of NYSBA and the groups tasked with ongoing monitoring. As a profession, we must continue to refine the initial guidelines suggested in this report and audit the efficacy of proposed rules and regulations. We liken this journey to the mindset of ancient explorers: be cautious, be curious, be vigilant and be brave.

Exhibit A

Task Force Mission Statement

The Task Force on AI will examine the legal, social and ethical impact of artificial intelligence (AI) on the legal profession. The Task Force will review AI-based software, generative AI technology and other machine-learning tools that may enhance the profession and that pose risks for individual attorneys dealing with new, unfamiliar technology and courts concerned about the integrity of the judicial process. Also, the Task Force will explore the positive and negative implications of AI use by the legal community and the general public, including effects on access to justice, legal regulations and privacy preservation. As it engages in its work, the Task Force will consult and ensure alignment of approaches, where appropriate, with other entities within the Association, including but not limited to the Committee on Technology and the Legal Profession, the Task Force on Emerging Digital Finance and Currency, the Working Group on Facial Recognition Technology and Access to Legal Representation and relevant sections. Lastly, the Task Force will develop policies for bar association adoption and suggest legislation to govern effective and responsible AI use.

APPENDIX A: LEGISLATION REVIEWED

I. Assemblyman Clyde Vanel's proposed statutes on AI:

- *Evidence created or processed by artificial intelligence*. An Act to amend New York's Criminal Procedure Law (CPL) and Civil Practice Law and Rules (CPLR) to address "the admissibility of evidence created or processed by artificial intelligence"

The essence of the evidence bill, which would amend the CPL and CPLR, is as follows:

§ 60.80 Rules of evidence; admissibility of evidence created or processed by artificial intelligence.

1. Evidence *created, in whole or in part, by artificial intelligence* shall not be received into evidence in a criminal proceeding unless the evidence is substantially supported by independent and admissible evidence and the proponent of the evidence establishes the reliability and accuracy of the specific use of the artificial intelligence in creating the evidence.

2. Evidence *processed, in whole or in part, by artificial intelligence* shall not be received into evidence in a criminal proceeding unless the proponent of the evidence establishes the reliability and accuracy of the specific use of the artificial intelligence in processing the evidence (emphasis added).

- *Political communications using artificial intelligence*. An Act to amend New York Election Law by requiring disclosure of "the use of artificial intelligence in political communications."

This bill would amend New York Election Law by requiring disclosure of "the use of artificial intelligence in political communications." The bill has separate sections to cover visual and non-visual communications. The heart of the bill provides as follows:

5. (a) Any political communication, regardless of whether such communication is considered a substantial or nominal expenditure, that uses *an image or video footage that was generated in whole or in part with the use of artificial intelligence*, as defined by the state board of elections, *shall be required to disclose that artificial intelligence was used* in such communication in accordance with paragraphs (b), (c), and (d) of this subdivision (emphasis added).

Paragraphs (b), (c), and (d) require specific disclaimers for "printed or digital political communications," "non-printed and non-digital political communications," and political communications that are "not visual, such as radio or automated telephone calls."

- *Political communications created by synthetic media*. An Act to amend New York Election Law, by "prohibiting the creation of synthetic media with intent to influence the outcome of an election."

  This bill would amend New York Election Law, by "prohibiting the creation of synthetic media with intent to influence the outcome of an election." Specifically, the bill would add a new § 17-172 that would provide as follows:

  § 17-172. Creating synthetic media with intent to unduly influence the 4 outcome of an election.

  1. A person who, with intent to injure a candidate or unduly influence the outcome of an election, creates or causes to be created a *fabricated photographic, videographic, or audio record* and causes such fabricated photographic, videographic, or audio record to be disseminated or published within sixty days of an election shall be guilty of a class E felony (emphasis added).

- *Artificial intelligence bill of rights*. An Act to amend New York's Technology Law by "enacting the New York artificial intelligence bill of rights."

  This bill would amend New York's Technology Law by "enacting the New York artificial intelligence bill of rights." The section on legislative intent says, in part:

  [T]he legislature declares that any New York resident affected by any *system making decisions without human intervention* be entitled to certain rights and protections to ensure that the system impacting their lives do so lawfully, properly, and with meaningful oversight.

  Among these rights and protections are (i) the right to safe and effective systems; (ii) protections against algorithmic discrimination; (iii) protections against abusive data practices; (iv) the right to have agency over one's data; (v) the right to know when an automated system is being used; (vi) the right to understand how and why an automated system contributed to outcomes that impact one; (vii) the right to opt out of an automated system; and (viii) the right to work with a human in the place of an automated system.

  The next part of the bill defines various terms. For example:

  4. "Algorithmic discrimination" means circumstances where an automated system contributes to an unjustified different treatment or impact which disfavors people based on their age, color, creed, disability, domestic violence victim status, gender identity or expression, familial status, marital status, military status, national origin, predisposing genetic characteristics, pregnancy-related condition, prior arrest or conviction record, race, sex, sexual orientation, or veteran status or any other classification protected by law.

The next part of the bill imposes various requirements. For example:

§ 404. Safe and effective systems.

2. *Automated systems shall undergo pre-deployment testing, risk identification and mitigation*, and shall also be subjected to ongoing monitoring that demonstrates they are safe and effective based on their intended use, mitigation of unsafe outcomes including those beyond the intended use, and adherence to domain-specific standards.

3. If an automated system fails to meet the requirements of this section, it shall not be deployed or, if already in use, shall be removed. *No automated system shall be designed with the intent or a reasonably foreseeable possibility of endangering the safety of any New York resident or New York communities* (emphasis added).

- *New York Penal Law – Fabricated photos, video, or audio*. An Act to amend the penal law by addressing "unlawful dissemination or publication of a fabricated photographic, videographic, or audio record."

This bill would amend New York's Penal Law by addressing "unlawful dissemination or publication of a fabricated photographic, videographic, or audio record." The essence of the bill is as follows:

1. A person is guilty of unlawful dissemination or publication of a fabricated photographic, videographic, or audio record when, with intent to cause harm to the liberty or emotional, social, financial or physical welfare of an identifiable person or persons, he or she intentionally creates or causes to be created a fabricated record of such person or persons and disseminates or publishes such record of such person or persons without such person or persons' consent.

The bill contains many exceptions. For example, the bill says:

This section shall not apply to the following:
(a) Dissemination or publication of a fabricated record by *a person who did not create the fabricated record*, whether or not such person is aware of the authenticity of the record;
(b) Dissemination or publication of a fabricated record that was created during the lawful and *common practices of law enforcement, legal proceedings or medical treatment* where the record is not disseminated or published with the intent to misrepresent its authenticity;
(c) Dissemination or publication of a fabricated record that was created for the purpose of *political or social commentary, parody, satire, or artistic expression* that is not disseminated or published with the intent to misrepresent its authenticity . . . (emphasis added)

- *Advanced Artificial Intelligence Licensing Act*. An Act to amend the state Technology Law to require registration and licensing of "high-risk advanced artificial intelligence systems."

An Act to amend the state Technology Law to address "advanced artificial intelligence systems" and to require registration and licensing of "high-risk advanced artificial intelligence systems." The bill defines these as follows:

1. "Advanced artificial intelligence system" shall mean any digital application or software, whether or not integrated with physical hardware, that *autonomously performs functions traditionally requiring human intelligence*. This includes, but is not limited to the system:

(a) Having the ability to learn from and adapt to new data or situations autonomously; or

(b) Having the ability to perform functions that require cognitive processes such as understanding, learning or decision-making for each specific task.

2. "High-risk advanced artificial intelligence system" shall mean any advanced artificial intelligence system that possesses *capabilities that can cause significant harm to the liberty, emotional, psychological, financial, physical, or privacy interests of an individual or groups of individuals, or which have significant implications on governance, infrastructure, or the environment*. The director shall assess any such public or private system in determining whether such system requires registration (emphasis added).

After a long series of definitions, the bill provides that the New York Department of State shall have "discretion to issue or refuse to issue any license provided for in this article" and to "revoke, cancel or suspend" any such license.

- *General Business Law – Oaths of responsible use of advanced AI*. An Act to amend New York's General Business Law by "requiring the collection of oaths of responsible use from users of certain high-impact advanced artificial intelligence systems."

This bill would amend New York's General Business Law by "requiring the collection of oaths of responsible use from users of certain high-impact advanced artificial intelligence systems." Here is a sample of the operative language of the oath:

I, ________residing at ________, do affirm under penalty of perjury that I have not used, am not using, do not intend to use, and will not use the services provided by this advanced artificial intelligence system in a manner that violated or violates any of the following affirmations:

1. I will not use the platform to create or disseminate content that can foreseeably cause injury to another in violation of applicable laws;

> 2. I will not use the platform to aid, encourage, or in any way promote any form of illegal activity in violation of applicable laws;
>
> 3. I will not use the platform to disseminate content that is defamatory, offensive, harassing, violent, discriminatory, or otherwise harmful in violation of applicable laws;
>
> 4. I will not use the platform to create and disseminate content related to an individual, group of individuals, organization, or current, past, or future events that are of the public interest which I know to be false and which I intend to use for the purpose of misleading the public or causing panic."

II. Federal and New York State proposals regarding use of AI-generated or compiled information in judicial proceedings

Judges face challenges in evaluating the admissibility of AI-generated or compiled evidence. Concerns include the reliability, transparency, interpretability and bias in such evidence. These challenges become even more pronounced with the use of generative AI systems. A discussion follows regarding two recent proposals to address these challenges.

Federal Law – A proposal to amend Fed. R. Evid. 901(b)(9)

As a general matter, Rule 901 of the Federal Rules of Evidence requires the proponent of a given item of evidence to authenticate that evidence. That is, the proponent "must produce evidence sufficient to support a finding that the item is what the proponent claims it is." Subsection (b) of that rule provides a non-exhaustive list of examples of how the proponent may satisfy the authentication requirement. As currently written, Fed. R. Evid. 901(b)(9), which applies to "evidence about a process or system" states that such evidence is "accurate" if the proponent shows that the process or system "produces an accurate result."

The Advisory Committee for the Federal Rules of Evidence is considering a proposal by former U.S. District Judge Paul Grimm and Dr. Maura R. Grossman of the University of Waterloo to amend Fed. R. Evid. 901(b)(9). That proposal initially changes the "accurate" standard as currently exists for any evidence about a process or system and replaces it with a requirement that the proponent provide evidence that shows that the process or system produces a "reliable" result. For evidence generated by AI, the proponent must also (a) describe the software or program that was used and (b) show that it has produced reliable results in the proposed evidence.

New York: Proposed amendments to the Criminal Procedure Law and CPLR

New York State Assemblyman Clyde Vanel has introduced a bill, A 8110, which amends both the Criminal Procedure Law and the Civil Practice Law and Rules regarding the admissibility of evidence created or processed by artificial intelligence. As stated in the bill, evidence is "created" by AI when AI produces new information from existing information. Evidence is "processed" by AI when AI produces a conclusion based on existing information.

Simplified greatly, the bill requires that evidence "created" by AI would not be received at trial unless independent admissible evidence establishes the reliability and accuracy of the AI used to create the evidence. Evidence "processed" by AI similarly requires the proponent of the evidence to establish the reliability and accuracy of the AI used. This bill does not yet have a co-sponsor in the Assembly and does not have a sponsor in the Senate.

The goals of both the proposal to amend Fed. R. Evid. 901 and the Vanel bill are laudable. The "black box" problem of AI is of great concern to lawyers and judges and has significant due process concerns in the criminal justice area. These proposals thus attempt to address AI-generated "deepfakes" that could be passed off as authentic evidence. Nevertheless, given the intricacies and time involved in the legislative and rule-amending processes, it may well be that the common law at the trial court level provides at least an interim roadmap for how judges should consider these issues. Indeed, this approach was largely employed to develop the law regarding discovery and admissibility of social media evidence when those issues first took hold.

III. New York City’s local law regarding use of AI in hiring and promotion

As of this writing, there are no statewide laws or regulations in New York regarding commercial use of AI. Notably, Governor Hochul vetoed a bill in November 2023 (A.4969), initially proposed by Assemblyman Vanel, that would have created a statewide commission to study AI. But it appears that Assemblyman Vanel, and perhaps many of his colleagues, are undeterred in their attempts to keep the conversation moving. One such attempt is a bill actually drafted by an AI program, and introduced by Vanel, that permits tenants in New York state to have the right to be able to request a copy of their lease. That bill, A.6896, is awaiting sponsorship in the New York State Senate.

New York City has, however, entered the regulatory space regarding AI-based hiring decisions. As of July 5, 2023, New York City’s Automated Employment Decision Tool (AEDT) law, Local Law 144 of 2021, or “NYC 144,” requires New York City employers who use AI and other machine-learning technology as part of their hiring process to annually audit their recruitment technology. NYC 144 defines AEDT as (1) any computational process, derived from machine learning, statistical modeling, data analytics or artificial intelligence, (2) that issues a simplified output, including a score, classification or recommendation, which is used to substantially assist or replace discretionary decision making for employment decisions that impact natural persons. A third party must perform these audits, and the audit results must be available on the company’s website. The audit itself must check for biases, whether intentional or unintentional, that are built into these systems. Failure to comply could result in fines starting at $500, with a maximum penalty of $1,500 per instance.

At the outset, NYC 144’s focus on “employment decisions” appears to cover only hiring and promotion. Conversely, it appears that decisions regarding compensation, termination, benefits, workforce monitoring and perhaps even performance evaluations are beyond the reach of the law. Moreover, NYC 144 applies only to those who actually apply for a job. Thus, the statute does not apply to any AI-based tools that might identify potential candidates who ultimately do not apply for a position.

Due to the recency of the NYC 144's implementation, there is no data as of this writing to determine its effectiveness, including whether and when any third-party audits have actually taken place. Even to the extent such audits have taken place, questions may remain as to the standards used for such audits and the company's data that was used for the audits.

IV. The White House's October 30, 2023 Executive Order regarding AI

On October 30, 2023, President Biden issued an Executive Order setting forth various standards for AI safety and security. It is one of the lengthier Executive Orders in recent history on any topic. The Order charges various executive agencies to develop guidelines, propose regulations or compile reports that will shape the AI landscape. The highlights of the Order include:

a. Establishment of the AI Safety and Security Board, under the auspices of the Department of Homeland Security, to address any threats posed by AI systems to infrastructure and cybersecurity.

b. Requiring the Department of Commerce to provide guidance for content authentication and watermarking to clearly label AI-generated content on government communications. In turn, federal agencies using AI-generated content are to highlight these authentication tools to assist recipients of government communications to know that these communications are authentic.

c. Federal agencies are to develop rules and guidelines to address algorithmic discrimination, both through training and technical assistance in areas including criminal justice, federal benefits and contracting programs, civil rights, and workplace equity, health and safety.

The question remains how these directives will be enforced. There is no requirement that any non-governmental entities involved in the creation or marketing of AI tools adhere to the directives that the various agencies will issue. Additionally, the Order does not provide, or even suggest, any recourse for individuals harmed by discriminatory AI systems. On these points (and perhaps many others), Congress may well have to provide guidance to federal agencies. Nevertheless, the Executive Order does provide a framework for both the government and the private sector to think about AI issues. It also invests the federal government, at least under the current administration, in AI security.

V. Summary of the EU AI Act

On December 9, 2023, the EU Parliament and Council negotiators reached a provisional agreement on the EU Artificial Intelligence Act (the "EU AI Act"). The agreed text will now proceed towards formal adoption by both the EU Parliament and Council to become EU law. While it is expected that the EU Parliament will adopt the EU AI Act, the law itself will not come into force for at least another two years after that vote.

As an overarching objective, the EU AI Act aims to ensure that fundamental rights, democracy, the rule of law and environmental sustainability are protected from high-risk AI, while boosting innovation and making the EU a leader in the field. The rules establish obligations for AI based on its potential risks and level of impact.

The following is a summary of the key aspects of the EU AI Act:

- General Regulatory Approach: The EU AI Act generally opts for a risk-based approach. Some applications are specifically prohibited (e.g., social scoring), some high-risk areas are strictly regulated (e.g., employment and worker management), and some areas of low risk are based on self-regulation. The EU AI Act strives to mitigate harm in areas where using AI poses “unacceptable” risk to fundamental rights, such as health care, education, border surveillance and public services.

- Territorial Scope: The EU AI Act has extraterritorial scope. It applies to: (a) providers placing on the EU market AI systems, whether those providers are established within the EU or in a third country; (b) users of AI systems located within the EU; (c) providers and users of AI systems that are located in a third country, where the output produced by the system is used in the EU. In practice this is likely to mean significant regulatory impact for U.S.-based organizations. The majority of the GDPR fines levied to date have been on U.S.-owned organizations. This extraterritorial reach is likely to be a feature of the EU AI Act as well.
- Prohibited AI applications: Recognizing the potential threat to individuals’ rights and democracy posed by certain applications of AI, the EU AI Act specifically prohibits the following applications:
    - biometric categorization systems that use sensitive characteristics (e.g., political, religious, philosophical beliefs, sexual orientation, race);
    - untargeted scraping of facial images from the internet or CCTV footage to create facial recognition databases;
    - emotion recognition in the workplace and educational institutions;
    - social scoring based on social behavior or personal characteristics;
    - AI systems that manipulate human behavior to circumvent their free will;
    - AI used to exploit the vulnerabilities of people due to their age, disability, social or economic situation.

- High-Risk AI Applications: The EU AI Act delineates the applications and activities designated as "high risk" and adopts certain requirements for their development, deployment and use. These uses are not prohibited but strictly regulated.
    - Categories of High-Risk AI Applications: Certain specific-use cases are designated as "high risk" irrespective of which industry or product the use case is deployed in, for instance, the use of AI in biometric identification systems, critical infrastructure, credit-worthiness evaluation, human resources contexts and law enforcement. In addition, this category includes the use of AI in relation to certain products, for example, machinery, radio equipment, medical devices and in vitro diagnostic medical devices, as well as AI used in certain products in civil aviation (security) and automotive industries. AI systems used to influence the outcome of elections and voter behavior are also classified as high risk.
    - Requirements for High-Risk AI Applications: Pursuant to the EU AI Act, high-risk AI must comply with various requirements such as conformity assessments, post-market surveillance, data governance and quality measures, mandatory registration, incident reporting and fundamental rights impact assessments. For example, in respect of AI systems classified as high risk (due to their significant potential harm to health, safety, fundamental rights, environment, democracy and the rule of law), the EU AI Act provides for a mandatory fundamental rights impact assessment applicable to, among other areas, the insurance and banking sectors. In addition, individuals will have a right to launch complaints about AI systems and receive explanations about decisions based on high-risk AI systems

that impact their rights. AI providers must build in human oversight, incorporating human-machine interface tools to ensure systems can be effectively overseen by natural persons.

- Law Enforcement: Predictive policing may only be employed under strict rules, such as clear human assessment and objective facts, not deferring the decision of investigating an individual to an algorithm. The EU AI Act stipulates a range of safeguards and narrow exceptions for the use of biometric identification systems (RBI) in publicly accessible spaces for law enforcement purposes, subject to prior judicial authorization and for strictly defined lists of crime. “Post-remote” RBI would be used strictly in the targeted search of a person convicted or suspected of having committed a serious crime. “Real-time” RBI would have to comply with strict conditions and its use would be limited in time and location, for the purposes of:
    - targeted searches of victims (abduction, trafficking, sexual exploitation),
    - prevention of a specific and present terrorist threat, or
    - the localization or identification of a person suspected of having committed one of the specific crimes mentioned in the EU AI Act (e.g., terrorism, trafficking, sexual exploitation, murder, kidnapping, rape, armed robbery, participation in a criminal organization, environmental crime).
- General-Purpose AI: In order to reflect the broad range of tasks that AI systems can accomplish and the rapid expansion of their capabilities, under the EU AI Act general-purpose AI (GPAI) systems, and the GPAI models they are based on, will need to adhere to certain transparency requirements. These include presenting

technical documentation, complying with EU copyright law and disseminating detailed summaries about the content used for training. GPAI is defined in the EU AI Act as “an AI system that can be used in and adapted to a wide range of applications for which it was not intentionally and specifically designed.” In this regard, the legislative text does not seem to distinguish between foundation AI, generative AI or GPAI regulation based on use cases. However, with respect to high-impact GPAI models with systemic risk, the EU AI Act stipulates more stringent obligations. High-impact GPAI models (in essence, those that were trained using a total computing power above a certain threshold) will be subject to more onerous requirements due to the presumption that they carry systemic risk. If these models meet certain criteria, they will need to conduct model evaluations, assess and mitigate systemic risks, conduct adversarial testing, report to the European Commission on serious incidents, ensure cybersecurity and report on their energy efficiency.

APPENDIX B: RESOURCES

Blogs & Podcasts

- **OpenAI Blog**: Direct insights from one of the leading organizations in AI research. It covers breakthroughs, applications, and considerations around their technologies, including generative models like GPT and DALL-E.

- **Distill**: Though not exclusively focused on generative AI, Distill publishes detailed, interactive research articles on machine learning that often touch on generative models. Its visual and intuitive approach makes complex topics accessible.

- **The Gradient**: A place for deep technical and theoretical discussions on AI, including generative models. The Gradient offers perspectives on the latest research trends, ethical considerations, and practical applications.

- **AI Weirdness**: Authored by Janelle Shane, this blog explores the quirky and humorous side of AI, including many experiments with generative models. It's an entertaining way to see the creative potential and limitations of AI.

- **DeepMind Blog**: While DeepMind's research encompasses a wide range of AI technologies, their work on generative models and their applications is frequently featured, providing insights into cutting-edge developments.

- **The AI Alignment Podcast**: Hosted by the Future of Life Institute, this podcast covers broader topics in AI, including the development and implications of generative AI technologies. Discussions often revolve around safety, ethics, and future prospects.

- **TWIML AI Podcast** (This Week in Machine Learning & AI): Offers a wide range of interviews with AI researchers, practitioners, and industry leaders, including episodes focused on generative AI technologies and their applications.

- **The Gradient Podcast**: An extension of The Gradient blog, this podcast dives into discussions with AI researchers and industry professionals, shedding light on their work, the future of AI, and occasionally focusing on generative models.

- **AI in Business**: While more focused on the application of AI in industry, this podcast sometimes explores generative AI applications in business, offering insights into how companies are leveraging this technology.

Newsletters

- **The Batch by DeepLearning.ai**: Curated by Andrew Ng and his team, The Batch brings the most important AI news and perspectives, including topics on generative AI, to your inbox. It's great for professionals, researchers, and anyone interested in AI.

- **Import AI by Jack Clark**: Jack Clark, co-founder of Anthropic and former policy director at OpenAI, shares weekly insights on AI developments, policy implications, and research breakthroughs. While not exclusively focused on generative AI, the newsletter often covers significant advancements and considerations in the field.

- **Data Elixir**: While broader than just generative AI, Data Elixir covers data science and machine learning trends, tools, and resources, including topics on generative models and AI-generated content.

- **The Algorithm by MIT Technology Review**: Offers insightful commentary on the latest AI developments, including ethical considerations, policy, and groundbreaking research in generative AI.

- **The Sequence**: A deep-tech AI newsletter that offers cutting-edge perspectives on AI technologies, including generative AI. It's structured in a unique format that includes a brief overview, a deep dive, and a summary of the latest AI research.

Subscriptions

- **AI Weekly**: A roundup of the best content in AI, including research papers, articles, and news. It frequently features content related to generative AI technologies and their applications.

- **Last Week in AI**: This newsletter gives a concise overview of the latest AI news, research, and applications with occasional deep dives into generative AI technologies and their societal impacts.

- **Orbit**: Focused on machine learning and AI, Orbit provides updates on the latest research, applications, and trends, including insightful discussions on generative AI.

- **MIT Technology Review**: Their subscription gives access to in-depth reporting on emerging technologies, including detailed articles on developments in AI and machine learning. Their coverage on generative AI technologies, implications, and ethical considerations is among the best.

- **AI Business**: Provides insights, analysis, and news on the application of AI in the business world, including generative AI. The subscription is aimed at professionals looking to understand how AI can be leveraged in various industries.

- **Inside AI**: Offers premium content on the latest AI news, research, and trends, with some focus on generative AI. The paid subscription includes additional insights and analysis not available in the free version.

- **Benedict Evans' Newsletter**: While not exclusively about AI, Benedict Evans provides high-level analysis and insights on the tech industry, including AI's impact on different sectors. His annual presentation includes significant trends in AI and machine learning.

- **Stratechery by Ben Thompson**: Offers in-depth analysis on the strategy and business side of technology, including AI. While the focus is broader, Thompson occasionally dives into topics related to generative AI and its impact on industries.

- **Datanami**: Focused on data science and big data news, Datanami covers the technological advancements and applications in AI and machine learning. Their subscription service provides in-depth analysis and exclusive content.

APPENDIX C: SAMPLE ENGAGEMENT LETTER PROVISION

Use of Generative AI: While representing you, we may use generative AI tools and technology to assist in legal research, document drafting and other legal tasks. This technology enables us to provide more efficient and cost-effective legal services. However, it is important to note that while generative AI can enhance our work, it is not a substitute for the expertise and judgment of our attorneys. We will exercise professional judgment in using AI-generated content and ensure its accuracy and appropriateness in your specific case.

ENDNOTES

---

[1] Erik Brynjolfsson and Andrew McAfee, The Business of Artificial Intelligence, Harvard Business Review, July 18, 2017, https://hbr.org/2017/07/the-business-of-artificial-intelligence.

[2] Linda Tucci, *A Guide to Artificial Intelligence in the Enterprise*, Enterprise AI, Jan. 25. 2024, https://www.techtarget.com/searchenterpriseai/Ultimate-guide-to-artificial-intelligence-in-the-enterprise.

[3] *Q and A With Maura Grossman: The Ethics of Artificial Intelligence*, University of Waterloo, Oct. 26, 2021, https://uwaterloo.ca/news/maura-grossman-ethics-artificial-intelligence.

[4] *Artificial Intelligence*, Merriam-Webster, https://www.merriam-webster.com/dictionary/artificial%20intelligence.

[5] Tucci, *supra* note 1.

[6] Jason Tashea & Nicholas Economou, Be Competent in AI Before Adopting, Integrating It into Your Practice, ABA J. (Apr. 23, 2019), http://www.abajournal.com/lawscribbler/article/before-lawyers-can- ethically-adopt-and-integrate-ai-into-their-practices-they-must-first-be-competent [https://perma.cc/45P6-B72G] ("Governed by computer science and statistics, these are complex academic disciplines in which lawyers are generally untrained and cannot become experts on the fly.").

[7] *Id. See also* Lauri Donahue, A Primer on Using Artificial Intelligence in the Legal Profession, HARV. J. L. & TECH. (Jan. 3, 2018), http://jolt.law.harvard.edu/digest/a-primer-on-using-artificial-intelligence-in-the-legal-profession [https://perma.cc/H65H-6A5A].

[8] L. Tucci, *A Guide to Artificial Intelligence in the Enterprise*, Enterprise AI (Sept. 9, 2023), https://www.techtarget.com/searchenterpriseai/Ultimate-guide-to-artificial-intelligence-in-the-enterprise.

[9] Nouvelle AI, Artificial Intelligence, Britannica, https://www.britannica.com/technology/artificial-intelligence/Nouvelle-AI, Encyclopædia Britannica, inc. (n.d.-a), Is Artificial General Intelligence (AGI) possible?, Encyclopædia Britannica, https://www.britannica.com/technology/artificial-intelligence/Is-artificial-general-intelligence-AGI-possible.

[10] *Id.*

[11] *Q and A With Maura Grossman, supra* note 2.

[12] Ian Sample, *Race to AI: The Origins of Artificial Intelligence, From Turing to ChatGPT*, Guardian, Oct. 28, 2023, https://www.theguardian.com/technology/2023/oct/28/artificial-intelligence-origins-turing-to-chatgpt.

[13] *Id.*

[14] *Id.*

[15] Tucci, *supra* note 1.

[16] *Id.*

[17] *Id.*

[18] *Id.*

[19] Alan M. Turing, *Computing Machinery and Intelligence* 49 Mind 433–60 (1950).

[20] Sample, *supra* note 10.

[21] Tucci, *supra* note 1.

[22] *Id.*

[23] *Id.*

[24] *Id.*

[25] *Id.*

[26] Rockwell Anyoha, *The History of Artificial Intelligence,* Harvard Univ.: Science in the News, Aug. 28, 2017, https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/

[27] *Id.*

[28] *Id.*

[29] John McCarthy, *What is Artificial Intelligence?*, Stanford Univ., Nov. 12, 2007, https://www-formal.stanford.edu/jmc/whatisai.pdf.

[30] *The Future Computed: Artificial Intelligence and its Role in Society*, Microsoft, 2018, https://blogs.microsoft.com/uploads/2018/02/The-Future-Computed_2.8.18.pdf

[31] *Id.*

[32] *Id.*

[33] *Id.*

---

[34] Lucas Mearian, *What Are LLMs, and How Are They Used in Generative AI?*, Computerworld, May 30, 2023, https://www.computerworld.com/article/3697649/what-are-large-language-models-and-how-are-they-used-in-generative-ai.html.

[35] Sean Clarke, Dan Milmo, and Garry Blight, *How AI Chatbots Like ChatGPT or Bard Work – Visual Explainer*, Guardian, Nov. 1, 2023, https://www.theguardian.com/technology/ng-interactive/2023/nov/01/how-ai-chatbots-like-chatgpt-or-bard-work-visual-explainer.

[36] Mearian, *supra* note 32.

[37] *Id.*

[38] *Id.*

[39] *Id.*

[40] *Id.*

[41] *Id.*

[42] *Id.*

[43] *Id.*

[44] Visual Storytelling Team and Madhumita Murgia, *Generative AI exists because of the transformer*, Financial Times (Sept. 12, 2023), https://ig.ft.com/generative-ai.

[45] Mearian, *supra* note 32.

[46] *Id.*

[47] *Id.*

[48] *See* Legal Services Corporation, The Justice Gap Executive Summary, https://justicegap.lsc.gov/resource/executive-summary.

[49] World Justice Project, Measuring the Justice Gap (2019), https://worldjusticeproject.org/our-work/publications/special-reports/measuring-justice-gap.

[50] *Id.*

[51] *See* Raymond H. Brescia et al., *Embracing Disruption: How Technological Change in the Delivery of Legal Services Can Improve Access to Justice*, 78 Alb. L. Rev. 553, 563 (2015).

[52] *MillerKing, LLC v. DoNotPay, Inc.*, No. 3:23-cv-863, 2023 U.S. Dist. LEXIS 209825 (S.D. Ill. Nov. 17, 2023).

[53] *DoNotPay Honored with ABA Brown Award for Access to Justice Efforts*, Am. Bar Assoc., Jan. 23, 2020, https://www.americanbar.org/news/abanews/aba-news-archives/2020/01/donotpay-honored-with-aba-brown-award-for-access-to-justice-effo/.

[54] Adele Peters, *There's Now a Chatbot to Give Refugees Instant Legal Aid*, Fast Company, Dec. 7, 2019, https://www.fastcompany.com/90439271/when-refugees-need-legal-advice-theres-now-a-chatbot-to-help-them.

[55] Matthew Dahl et al., Hallucinating Law: Legal Mistakes With Large Language Models Are Pervasive, Stanford Univ., Jan. 11, 2024, https://hai.stanford.edu/news/hallucinating-law-legal-mistakes-large-language-models-are-pervasive.

[56] *Id.*

[57] *Artificial Intelligence: Guidance for Judicial Office Holders*, Courts and Tribunals Judiciary, Dec. 12, 2023, p. 3, https://www.judiciary.uk/wp-content/uploads/2023/12/AI-Judicial-Guidance.pdf.

[58] Sunil Rajaraman, *Navigating Data With LLMs - Arguments For And Against Sharing*, Forbes (Feb 22, 2024, 4:44pm EST), https://www.forbes.com/sites/sunilrajaraman/2024/02/22/navigating-data-with-llms-arguments-for-and-against-sharing/?sh=2dfc75671098.

[59] *Id.*, p. 4.

[60] Atomwise recently received over a billion dollars in investment from Pharmaceutical giant Sanofi for its AtomNet platform, which researches small molecules aimed at up to five drug targets. This company uses AI to analyze the structure of molecules and predict how they might interact with targets in the body. Technology such as this can reduce drug development timelines by years.

[61] Nafiseh Ghaffar Nia et. al, *Evaluation of artificial intelligence techniques in disease diagnosis and prediction*, Nat'l Library of Medicine (Jan. 30, 2023), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9885935.

[62] Alexander Muacevic and John R Adler, editors, *Harnessing the Power of AI: A Comprehensive Review of Its Impact and Challenges in Nursing Science and Healthcare*, Nat'l Library of Medicine (Nov. 22, 2023), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10744168/#:~:text=Patient%20monitoring%3A%20AI%2Ddriven%20patient,blood%20pressure%2C%20and%20oxygen%20levels.

---

[63] Kevin B. Johnson et. al, *Precision Medicine, AI, and the Future of Personalized Health Care*, Nat'l Library of Medicine (Oct. 12, 2020), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7877825.

[64] Cedar Sinai, Study: Mental Health Gets a Boost From Artificial Intelligence (Jan. 26, 2024), https://www.cedars-sinai.org/newsroom/study-mental-health-gets-a-boost-from-artificial-intelligence/#:~:text=Findings%20Published%20in%20Nature%20Digital,to%20provide%20mental%20health%20support.

[65] Fawn Fitter and Steven T. Hunt, *How AI Can End Bias*, SAP Insights, https://www.sap.com/insights/viewpoints/how-ai-can-end-bias.html.

[66] Capitol Technology University, Blog: The Ethical Considerations of Artificial Intelligence (May 30, 2023), https://www.captechu.edu/blog/ethical-considerations-of-artificial-intelligence.

[67] Kelly Fitzsimmons, *Advancing Equitable AI in the US Social Sector*, Stanford Social Innovation Review (Mar. 12, 2024), https://ssir.org/articles/entry/advancing-equitable-ai-us-social-sector#:~:text=When%20designed%20and%20implemented%20with,AI%20in%20the%20social%20sector.

[68] Paresh Dave, *Google's AI Is Making Traffic Lights More Efficient and Less Annoying*, Wired (Oct. 10, 2023 7:00 am), https://www.wired.com/story/googles-ai-traffic-lights-driving-annoying/#:~:text=7%3A00%20AM-,Google's%20AI%20Is%20Making%20Traffic%20Lights%20More%20Efficient%20and%20Less,cut%20wait%20times%20and%20emissions.

[69] Matt Wood, *Algorithm predicts crime a week in advance, but reveals bias in police response*, The University of Chicago, Biological Sciences Division (June 30, 2022), https://biologicalsciences.uchicago.edu/news/algorithm-predicts-crime-police-bias.

[70] Lazima Faiah Bari et. al, *Potential Use of Artificial Intelligence (AI) in Disaster Risk and Emergency Health Management: A Critical Appraisal on Environmental Health*, Nat'l Library of Medicine (Dec. 10, 2023), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10712270/#:~:text=AI%2Ddriven%20algorithms%20can%20optimize,and%20streamlining%20the%20recovery%20process.

[71] Orange: Hello Future, Sensors and AI ensure the good health of bridges (Jan. 27, 2023), https://hellofuture.orange.com/en/sensors-and-ai-ensure-the-good-health-of-bridges.

[72] IBM, "AI to accelerate your security defenses," https://www.ibm.com/ai-cybersecurity.

[73] Kinza Yasar, "What is a smart home?," TechTarget, https://www.techtarget.com/iotagenda/definition/smart-home-or-building.

[74] Jared Thau, *Digital Renaissance: AI And VR In The Gaming And Tech Industry*, Forbes (Mar. 11, 2024), https://www.forbes.com/sites/forbestechcouncil/2024/03/11/digital-renaissance-ai-and-vr-in-the-gaming-and-tech-industry/?sh=a65aa844ede1.

[75] Craig Borowski, *32 Best AI Chatbots for Customer Service in 2024*, CX Lead (Jan. 26, 2024), https://thecxlead.com/tools/best-ai-chatbot-for-customer-service.

[76] Salvatore Raieli, *How AI Could Help Preserve Art*, Medium (Oct. 22, 2022), https://towardsdatascience.com/how-ai-could-help-preserve-art-f40c8376781d.

[77] *Id.*

[78] Olga Dogadkina, *How AI-Powered Personalization Is Reshaping Online Shopping And Beyond*, Forbes (Oct. 16, 2023 6:45am EDT), https://www.forbes.com/sites/forbestechcouncil/2023/10/16/how-ai-powered-personalization-is-reshaping-online-shopping-and-beyond/?sh=6b7a115f24e7.

[79] Owais Ali, *How is AI Being Used in Space Exploration?*, AZO Quantum (Dec. 4, 2023), https://www.azoquantum.com/Article.aspx?ArticleID=474#:~:text=AI%20empowers%20robotic%20rovers%20on,onboard%20maps%20and%20sensor%20data.

[80] NOAA (National Oceanic and Atmospheric Administration) Ocean Exploration, https://oceanexplorer.noaa.gov/technology/subs/auvs/auvs.html.
U.S. Department of Commerce

[81] https://alphafold.ebi.ac.uk.

[82] U.S. Nat'l Science Foundation, NSF 23-610: National Artificial Intelligence (AI) Research Institutes (Aug. 1, 2023), https://new.nsf.gov/funding/opportunities/national-artificial-intelligence-research/nsf23-610/solicitation.

[83] David Ly, *On The Horizon For Smart Cities: How AI And IoT Are Transforming Urban Living*, Forbes (Apr 7, 2023 8:30am EDT), https://www.forbes.com/sites/forbestechcouncil/2023/04/07/on-the-horizon-for-smart-cities-how-ai-and-iot-are-transforming-urban-living/?sh=25a2025e7145.

---

[84] Tosin Thompson, *How AI can help to save endangered species*, Nature (Oct. 27, 2023), https://www.nature.com/articles/d41586-023-03328-4.
[85] Allen Institute for AI, "Climate Modeling for the future of the planet," https://allenai.org/climate-modeling.
[86] Strategic Consortium of Intelligence Professionals, "AI for Clean Air and Water," SCIP Insights (Aug. 28, 2023), https://www.scip.org/news/650086/AI-for-Clean-Air-and-Water-.htm#:~:text=Air%20Quality%20Monitoring%20and%20Prediction,actions%20to%20mitigate%20pollution%20spikes.
[87] Corgan, "AI in Indoor Water Conservation," https://www.corgan.com/news-insights/2023/ai-in-indoor-water-conservation#:~:text=By%20analyzing%20data%20on%20water,the%20environmental%20impact%20of%20buildings.
[88] Zhengxuan Liu et. al, *Artificial intelligence powered large-scale renewable integrations in multi-energy systems for carbon neutrality transition: Challenges and future perspectives*, ScienceDirect (Nov. 2022), https://www.sciencedirect.com/science/article/pii/S2666546822000428.
[89] Emily Newton, *7 Novel Approaches Using AI to Reinvent Package Delivery*, Supply Chain Connect (Dec. 14, 2023), https://www.supplychainconnect.com/supply-chain-technology/article/21278954/7-novel-approaches-using-ai-to-reinvent-package-delivery.
[90] Tumaini Kabudi, *AI-enabled adaptive learning systems: A systematic mapping of the literature*, ScienceDirect (2021), https://www.sciencedirect.com/science/article/pii/S2666920X21000114.
[91] The Princeton Review, "The Evolution of Education: How AI is Reshaping Grading," https://www.princetonreview.com/ai-education/how-ai-is-reshaping-grading.
[92] Alina Tugend, *How Robots Can Assist Students With Disabilities*, NY Times (Mar. 29, 2022), https://www.nytimes.com/2022/03/29/technology/ai-robots-students-disabilities.html?auth=login-google1tap&login=google1tap.
[93] Robert C. Brears, *Precision Agriculture, AI, and Water Efficiency: The Future of Farming*, Medium (May 3, 2023), https://medium.com/mark-and-focus/precision-agriculture-ai-and-water-efficiency-the-future-of-farming-b959ac0b6017.
[94] Alessandro Mascellino, *Biometric authentication use in US businesses tripled over 3 years to tackle cyber threats*, BiometricUpdate.com (Sept. 21, 2022), https://www.biometricupdate.com/202209/biometric-authentication-use-in-us-businesses-tripled-over-3-years-to-tackle-cyber-threats#:~:text=The%20use%20of%20biometric%20authentication,report%20by%20software%20expert%20GetApp; *see also* chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www.pwc.com/gx/en/information-security-survey/assets/gsiss-report-cybersecurity-privacy-safeguards.pdf.
[95] *See, e.g.*, a Prompt Engineer, https://www.ziprecruiter.com/Salaries/Prompt-Engineer-Salary.
[96] InkSpire, *AI and the Legal Profession: Navigating the New Frontier* (Jan. 15, 2024), https://inkspire.org/post/ai-and-the-legal-profession-navigating-the-new-frontier.
[97] Drew Simshaw, *Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services*, Yale J. Law & Tech'y (2022), https://yjolt.org/sites/default/files/simshaw_-_access_to_a.i._justice.pdf.
[98] Emily Cardona, AI & PRO BONO, Pro Bono Institute (Sept. 2023), https://www.probonoinst.org/2023/09/25/ai-pro-bono.
[99] Catarina Fontes et. al, *AI-powered public surveillance systems: why we (might) need them and how we want them*, ScienceDirect (Nov. 2022), https://www.sciencedirect.com/science/article/pii/S0160791X22002780.
[100] Steve Zurier, *Hackers 'steal your face' to create deepfakes that rob bank accounts*, SC Media (Feb. 15, 2024), https://www.scmagazine.com/news/hackers-steal-your-face-to-create-deepfakes-that-rob-bank-accounts.
[101] Jack Hardinges et. al, *We Must Fix the Lack of Transparency Around the Data Used to Train Foundation Models*, Harvard Data Science Review (Dec. 13, 2023), https://hdsr.mitpress.mit.edu/pub/xau9dza3/release/1.
[102] Chapman University, "Bias in AI," https://www.chapman.edu/ai/bias-in-ai.aspx.
[103] Keegan Caldwell, *AI And Intellectual Property: Who Owns It, And What Does This Mean For The Future?*, Forbes (Oct 31, 2023 7:45am EDT), https://www.forbes.com/sites/forbesbusinesscouncil/2023/10/31/ai-and-intellectual-property-who-owns-it-and-what-does-this-mean-for-the-future/?sh=635aef883e96.
[104] Charles Cohen, *AI in Defense: Navigating Concerns, Seizing Opportunities*, Nat'l Defense Magazine (July 25, 2023), https://www.nationaldefensemagazine.org/articles/2023/7/25/defense-department-needs-a-data-centric-digital-security-organization.

---

[105] Ali Swenson and Kelvin Chan, *Election disinformation takes a big leap with AI being used to deceive worldwide*, AP News (Mar. 14, 2024 updated 3:46 AM EDT), https://apnews.com/article/artificial-intelligence-elections-disinformation-chatgpt-bc283e7426402f0b4baa7df280a4c3fd.

[106] Catherine Stupp, *AI Helps U.S. Intelligence Track Hackers Targeting Critical Infrastructure*, Wall Street J. (Jan. 10, 2024 5:30 am ET), https://www.wsj.com/articles/ai-helps-u-s-intelligence-track-hackers-targeting-critical-infrastructure-944553fa.

[107] Xingyu Chen et. al, *Exploring racial and gender disparities in voice biometrics*, Nat'l Library of Medicine (Mar. 8, 2022), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8904636; Trevor Collins, *New Research Reveals Sexist Tendencies in Facial Recognition Tech*, Secplicity (July 8, 2020), https://www.secplicity.org/2020/07/08/new-research-reveals-sexist-tendencies-in-facial-recognition-tech; Thaddeus L. Johnson & Natasha N. Johnson, *Police Facial Recognition Technology Can't Tell Black People Apart*, Scientific American (May 18, 2023), https://www.scientificamerican.com/article/police-facial-recognition-technology-cant-tell-black-people-apart; and Honey Wyatt, *Uber Eats driver wins payout for racially biased AI checks*, HRMagazine.com (Mar. 27, 2024), https://www.hrmagazine.co.uk/content/news/uber-eats-driver-wins-payout-for-racially-biased-ai-checks#:~:text=A%20black%20Uber%20eats%20driver,in%20Oxford%20from%20November%202019.

[108] Imane El Atillah, *Man Ends His Life After an AI Chatbot 'Encouraged' Him to Sacrifice Himself to Stop Climate Change*, Euronews, Mar. 31, 2023, https://www.euronews.com/next/2023/03/31/man-ends-his-life-after-an-ai-chatbot-encouraged-him-to-sacrifice-himself-to-stop-climate-.

[109] Miriam Buiten et. al, *The law and economics of AI liability*, ScienceDirect (Apr. 2023), https://www.sciencedirect.com/science/article/pii/S0267364923000055#:~:text=It%20is%2C%20therefore%2C%20appropriate%20to,some%20of)%20the%20associated%20costs.

[110] Erin Beck, *The AI Echo Chamber: Model Collapse & Synthetic Data Risks*, Deepgram (Sept. 6, 2023, updated Oct. 11, 2023), https://deepgram.com/learn/the-ai-echo-chamber-model-collapse-synthetic-data-risks.

[111] Shannon Bond, *AI fakes raise election risks as lawmakers and tech companies scramble to catch up*, NPR (Feb. 8, 2024 5:00 AM ET), https://www.npr.org/2024/02/08/1229641751/ai-deepfakes-election-risks-lawmakers-tech-companies-artificial-intelligence.

[112] Federal Communications Commission, "'Grandparent' Scams Get More Sophisticated," https://www.fcc.gov/grandparent-scams-get-more-sophisticated.

[113] Heather Wishart-Smith, *Not So Fast: Study Finds AI Job Displacement Likely Substantial, Yet Gradual*, Forbes (Feb 13, 2024 4:23pm EST), https://www.forbes.com/sites/heatherwishartsmith/2024/02/13/not-so-fast-study-finds-ai-job-displacement-likely-substantial-yet-gradual/?sh=7f0b67ce1e9f.

[114] Manas Arora, *How AI is transforming the financial industry, from predicting market trends to automated trading strategies*, CXOtoday.com (Nov. 30, 2023 11:09 pm), https://cxotoday.com/specials/how-ai-is-transforming-the-financial-industry-from-predicting-market-trends-to-automated-trading-strategies.
Read more at: https://cxotoday.com/specials/how-ai-is-transforming-the-financial-industry-from-predicting-market-trends-to-automated-trading-strategies/
Read more at: https://cxotoday.com/specials/how-ai-is-transforming-the-financial-industry-from-predicting-market-trends-to-automated-trading-strategies/

[115] Chris Stokel-Walker, *Will we lose certain skills and knowledge if we rely on AI too much?*, Cybernews (Dec. 09, 2023 3:00 PM), https://cybernews.com/editorial/humans-lose-skills-knowledge-ai.

[116] Ladan Judge, *What Is Forced Labor in the Technology Industry Supply Chain?*, Z2Data, June 26, 2023, https://www.z2data.com/insights/what-is-forced-labor-in-the-technology-industry-supply-chain.

[117] Nat'l Artificial Intelligence Advisory Committee, Findings: The Potential Future Risks of AI (Oct. 2023), https://ai.gov/wp-content/uploads/2023/11/Findings_The-Potential-Future-Risks-of-AI.pdf.

[118] Cameron F. Kerry, *Protecting privacy in an AI-driven world*, Brookings (Feb. 10, 2020), https://www.brookings.edu/articles/protecting-privacy-in-an-ai-driven-world.

[119] *Can We Trust Artificial Intelligence?*, CalTech, https://scienceexchange.caltech.edu/topics/artificial-intelligence-research/trustworthy-ai.

[120] Sayed Fayaz Ahmad et. al, *Impact of artificial intelligence on human loss in decision making, laziness and safety in education*, Nature (June 29, 2023), https://www.nature.com/articles/s41599-023-01787-8.

[121] Dennis Hillemann, *Can AI Truly Experience Emotion and Should We Trust It in the Public Sector Decision-Making Process?*, Medium (June 29, 2023), https://dhillemann.medium.com/can-ai-truly-experience-emotion-and-should-we-trust-it-in-the-public-sector-decision-making-process-9797211e9bcb.

---

[122] Natalie Pierce and Stephanie Goutos, *ChatGPT Doesn't Have Ethical Obligations, But Attorneys Do*, Bloomberg Law, July 11, 2023, https://news.bloomberglaw.com/us-law-week/chatgpt-doesnt-have-ethical-obligations-but-attorneys-do.

[123] Nicole Yamane, *Artificial Intelligence in the Legal Field and the Indispensable Human Element Legal Ethics Demands*, Sept. 24, 2020, Georgetown Univ. Law Center, https://www.law.georgetown.edu/legal-ethics-journal/wp-content/uploads/sites/24/2020/09/GT-GJLE200038.pdf

[124] *International Legal Generative AI Report*, LexisNexis, Aug. 22, 2023, https://www.lexisnexis.com/pdf/lexisplus/international-legal-generative-ai-report.pdf.

[125] *Technosolutionism*, Guide to Crypto and Web3, https://web3.lifeitself.org/concepts/technosolutionism; *see also* Shane Hastie, *Unraveling Techno-Solutionism: How I Fell Out of Love With "Ethical" Machine learning*, InfoQ Nov. 7, 2022, https://www.infoq.com/news/2022/11/unraveling-techno-solutionism; *On the Use of AI - the Dependency Dilemma*, IEEE Technical Community Spotlight, Jan. 14, 2022, https://technical-community-spotlight.ieee.org/ai-ethical-dilemma.

[126] *Mata v. Avianca, Inc.*, 22-cv-1461, 2023 U.S. Dist. LEXIS 108263 (S.D.N.Y. June 22, 2023).

[127] Einaras von Gravrock, *Why Artificial Intelligence Design Must Prioritize Data Privacy*, World Economic Forum Mar. 31, 2022, https://www.weforum.org/agenda/2022/03/designing-artificial-intelligence-for-privacy.

[128] *Recommendations from Committee on Professional Responsibility and Conduct on Regulation of Use of Generative AI by Licensees*, The State Bar of California, Memorandum, p. 10, Nov. 16, 2023, https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf ("The lawyer should consider disclosure to their client that they intend to use generative AI in the representation, including how the technology will be used, and the benefits and risks of such use. A lawyer should review any applicable client instructions or guidelines that may restrict or limit the use of generative AI.").

[129] The Florida Bar Board of Governors' Review Committee on Professional Ethics, Proposed Advisory Opinion 24-1, Nov. 13, 2023, ("it is recommended that a lawyer obtain the affected client's informed consent prior to utilizing a third-party generative AI program if the utilization would involve the disclosure of any confidential information"), https://www.floridabar.org/the-florida-bar-news/proposed-advisory-opinion-24-1-regarding-lawyers-use-of-generative-artificial-intelligence-official-notice.

[130] American Bar Association, Resolution 112 (adopted Aug. 12-13, 2019), https://www.americanbar.org/content/dam/aba/directories/policy/annual-2019/112-annual-2019.pdf.

[131] David Wagner, *This Prolific LA Eviction Law Firm Was Caught Faking Cases In Court. Did They Misuse AI?*, LAist, Oct. 12, 2023, https://laist.com/news/housing-homelessness/dennis-block-chatgpt-artificial-intelligence-ai-eviction-court-los-angeles-lawyer-sanction-housing-tenant-landlord.

[132] Artificial Intelligence in the legal field and the indispensable human, https://www.law.georgetown.edu/legal-ethics-journal/wp-content/uploads/sites/24/2020/09/GT-GJLE200038.pdf.

[133] Michael Simon et. al, "Lola v. Skadden and the Automation of the Legal Profession," 20 YALE J.L. & TECH. 234, 248 (2018) ("According to the Lola decision, if a lawyer is performing a particular task that can be done by a machine, then that work is not practicing law."); Lola v. Skadden, Arps, Slate, Meagher & Flom LLP, 620 Fed. Appx. 37, 45 (2nd Cir. 2015).

[134] *Id.*, p. 888; *see also* Nicole Yamane, "Artificial Intelligence in the Legal Field and the Indispensable Human Element Legal Ethics Demands, Georgetown Univ. Law Center (Sept. 24, 2020), https://www.law.georgetown.edu/legal-ethics-journal/wp-content/uploads/sites/24/2020/09/GT-GJLE200038.pdf.

[135] *LegalZoom.com, Inc. v. N.C. State B.*, 2015 NCBC 96, Consent J.; and D. Fisher, D., *LegalZoom settles fight with North Carolina Bar Over Online Law*, Forbes (Oct. 23, 2015), https://www.forbes.com/sites/danielfisher/2015/10/22/legalzoom-settles-fight-with-north-carolina-bar-over-online-law/?sh=13b759e43eb2.

[136] *See supra* note 69.

[137] Wrubleski v. Mary Imogene Bassett Hosp., 163 A.D. 3d 1248, 1250–51 (3d Dept. 2018).

[138] L. Eliot, *Is generative AI such as CHATGPT going to undermine the famed attorney-client privilege, frets AI law and AI ethics*, Forbes (Oct. 5, 2023), https://www.forbes.com/sites/lanceeliot/2023/03/30/is-generative-ai-such-as-chatgpt-going-to-undermine-the-famed-attorney-client-privilege-frets-ai-law-and-ai-ethics.

[139] Rule 1.1 of the RPC requires that a lawyer provide competent representation to a client. Comment 8 to RPC Rule 1.1 asserts that this includes keeping abreast of "the benefits and risks associated with technology the lawyer uses to provide services to clients."

---

[140] Doug Austin, *Insurer sent law firms a CHATGPT warning*, eDiscovery Today (April 14, 2023), https://ediscoverytoday.com/2023/04/14/insurer-sent-law-firms-a-chatgpt-warning-artificial-intelligence-trends/?int_ref=yrp.
[141] Tony Petruzzi and Helena Guye, *"The Perils of Dabbling": AI and the Practice of Law*, Reuters, Sept. 11, 2023, https://www.reuters.com/legal/legalindustry/perils-dabbling-ai-practice-law-2023-09-11.
[142] *What Are AI Hallucinations?*, IBM, https://www.ibm.com/topics/ai-hallucinations.
[143] *Nix v. Whiteside*, 475 U.S. 157, 166 (1986).
[144] Ian Sample, *What Are Deepfakes – and How Can You Spot Them?*, Guardian, Jan. 13, 2020, https://www.theguardian.com/technology/2020/jan/13/what-are-deepfakes-and-how-can-you-spot-them.
[145] Rebecca A. Delfino, *Deepfakes on Trial: A Call To Expand the Trial Judge's Gatekeeping Role To Protect Legal Proceedings from Technological Fakery*, 74 Hastings L.J. 293 (2023), https://repository.uchastings.edu/hastings_law_journal/vol74/iss2/3.
[146] *Id.*
[147] Marla Greenstein, *AI and a Judge's Ethical Obligations,* American Bar Assoc., Feb. 3, 2020, https://www.americanbar.org/groups/judicial/publications/judges_journal/2020/winter/ai-and-judges-ethical-obligations.
[148] *See* NYSBA Communications Department, "American Bar Association Adopts NYSBA-Advanced Resolution on Best Practice Guidelines for Online Legal Document Providers" (Aug. 12, 2019), https://nysba.org/august-12-2019-american-bar-association-adopts-nysba-advanced-resolution-on-best-practice-guidelines-for-online-legal-document-providers.
[149] *Estonia Does Not Develop AI Judge*, Republic of Estonia Ministry of Justice, Feb. 16, 2022, https://www.just.ee/en/news/estonia-does-not-develop-ai-judge.
[150] *Learning From the Failures of Robodebt – Building a Fairer, Client-Centred Social Security System*, Victoria Legal Aid, Nov. 13, 2023, https://www.legalaid.vic.gov.au/learning-from-the-failures-of-robodebt.
[151] No. 22-cv-1461 (PKC), 2023 WL 4114965 (S.D.N.Y. June 22, 2023).
[152] *Client Alert: Italian Data Protection Authority Bans Chatbot*, Cordery Legal Compliance, Feb. 7, 2023, https://www.corderycompliance.com/italy-dpa-chatbot-0223.

MEMORANDUM

TO: NYSBA House of Delegates

FROM: NYSBA Trusts and Estates Law Executive Committee

DATE: March 28, 2024

SUBJECT: NYSBA TASK FORCE ON ARTIFICIAL INTELLIGENCE REPORT

As prepared by the TELS Technology Committee and reviewed by the TELS Executive Committee, our comments to the report of the Task Force on Artificial Intelligence follow.

Comment on Recommendations:

1) Adopt Guidelines "*The Task Force recommends that NYSBA adopt the AI/GAI guidelines outlined in this report and commission a standing section or committee to oversee periodic updates to those guidelines. Daily, we learn more about the capability of the technology to transform society. As the impacts are continual, so should the updates to these guidelines be as well.*"

   a. Given the pace and development of AI/GAI technology, the Trusts and Estates Law Section (TELS) is concerned that frequent updates to adopted guidelines will present challenges to practitioners conforming their practice to the guidelines. The TELS believes that a reasoned interpretation of the applicable rules of Professional Conduct and current guidance and commentary is sufficient to guide most practitioners. In other words, less might be more. However, the Task Force's proposed guidance is generally helpful and acceptable with the following comments/critiques which focus on the Task Force's contemplation of AI/GAI as having personhood. The TELS does not believe that AI/GAI should be considered or contemplated as a person.

      i. Guidance on Rule 5.3: *A law firm shall ensure that the work of nonlawyers who work for the firm is adequately supervised, as appropriate.*

   "*If the Tools are used by non-lawyers or paralegals (or the Tools themselves are considered "non-lawyers"), you must supervise their use to ensure compliance with the ethical rules. Further, you must ensure that the work produced by the Tools is accurate and complete and does not disclose or create a risk of disclosing client confidential information without your client's informed consent.*"

   The TELS opposes the parenthetical suggesting that the Tools may be considered "non-lawyers."

      ii. Guidance on Rule 5.4: *A lawyer shall not permit a person to direct or regulate the lawyer's professional judgment in rendering legal services.*

   "*While the Tools are technically not a "person," you should refrain from relying exclusively on them when providing legal advice and maintain your independent judgment on a matter.*"

The Tools are not a person in any sense, “technically” or practically. The TELS opposes implicating personhood with respect to a technological resource.

iii. Guidance on Rule 5.5: *A lawyer shall not aid a nonlawyer in the unauthorized practice of law.*

“*Understand that human oversight is necessary to avoid UPL issues when using the Tools, which should augment but not replace your legal work.*”

The guidance contemplates that AI/GAI could be engaged in the unlicensed practice of law. The TELS opposes assigning personhood to AI/GAI in this respect.

2) Focus on Education: “*The Task Force recommends that NYSBA prioritize education over legislation, focusing on educating judges, lawyers and regulators to understand the technology so that they can apply existing law to regulate it.*”

The TELS strongly endorses this recommendation.

3) Identify Risks for New Regulation: “*Legislatures should identify risks associated with the technology that are not addressed by existing laws, which will likely involve extensive hearings and studies involving experts in AI.*”

The TELS endorses this recommendation. The TELS however, believes that applicable legislatures and administrative agencies engaged in rulemaking ought to focus on proper attribution to AI/GAI and disclosure of the use of AI/GAI in submissions to tribunals. We believe that the issue of whether and to what extent disclosure must be had when an attorney uses AI/GAI should be addressed immediately. For example, if an attorney relies on AI/GAI in a brief or memorandum of law submitted to a court, the court, the litigants, and the public in general might be better served if reliance and use of AI/GAI is disclosed by way of attribution and/or disclosure. Consideration should be afforded to the nature and extent of the attorney’s reliance on AI/GAI in this scenario, for example, is AI/GAI being utilized to help counsel of record spot flaws in a counterpart’s argument? to summarize cases? to generate wholesale prose then incorporated into a litigant’s brief/memorandum of law? to analyze technical data? to analyze and reach factual conclusions based on documentary evidence and testimony? The TELS believes that the guidance should be supplemented to require attorneys to disclose use of the Tools in instances where the attorney relies upon AI/GAI to generate an argument and employs that argument utilizing the prose generated by the Tools. However, where AI/GAI is used for less substantive tasks such as conducting research or summarizing case law, disclosure is less warranted.

The law is notoriously slow in addressing the much more rapid and frequent changes in technology. Deliberately considered legislation and rulemaking is a time-tested and valuable feature of the law. However, in this context, care must be exercised to avoid perpetually playing "catch-up" as a result of focusing on specific technological features which may be subsumed or become obsolete in a very short period of time. A better approach would be to address technology globally, by focusing on the obligations of the attorney rather than the specific technology being employed at the moment. The legal profession, and the public as a whole, is far better served by making it clear that when a lawyer utilizes technology–any technology– as part of his or her practice, he or she is ultimately responsible for the content and quality of the work product thus generated.

Memo to: Patricia J. Shevy, Chair Trusts and Estates Law Section
From: Albert Feuer
Re: TELS Technology Committee March 26, 2024 memo regarding the NYSBA Task Force on Artificial Intelligence Report and Recommendations to NYSBA House of Delegate (April 6, 2024)
Date: March 28, 2024

The Task Force produced a very good and comprehensive discussion of the history and the significance of artificial intelligence (AI), its risks and benefits, the laws that govern AI and have been proposed to govern AI, and AI's implications for lawyers, the legal system, the access to justice, and for society.

Like the TELS Technology Committee I will focus only on the Task Force's four recommendations.

1) It is advisable to have a NYSBA standing committee or section to continue to examine the legal, social, and ethical impact of artificial intelligence. This entity could update the guidelines in a manner that balances the burdens and benefits of such updates.

As with all legal tools, including sample legal documents/templates, questions may arise whether (a) an attorney using such tools is exercising the attorney's legal judgment with respect to the proper use of such tools, or (b) the provider of such tools to lay persons is practicing law. I share the concern of the TELS committee about the anthropomorphizing of AI, although for a different reason. Such characterization may make it more difficult to correct AI errors because it may make it more difficult to hold the user and/or the provider/designer of AI responsible for those errors.

2) It is advisable for the NYSBA to "focus on educating judges, lawyers, law students and regulators to understand the technology so that they may apply existing law to regulate it." This may include explicitly mentioning AI in the Rules for Professional Conduct.

3) It is advisable for "legislatures seek to identify risks associated with the technology that are not addressed by existing law." I disagree with the TELS committee suggestion that this focus only on tribunal submissions. There also needs to be focus on the use of AI for the non-litigation responsibilities of attorneys: counseling, and the preparation of legal documents. Such usage also raises the issue of lay persons seeking to prepare documents using AI tools supplied by the same persons that now provide sample legal documents, such as wills.

4) It is advisable to consider how AI may be used in law as a governance tool, which recommendation the TELS committee did not discuss. For example, which principles should determine the appropriate regulation of AI tools, and who should regulate. Similarly, how may society/commercial benefits be weighed against risks to individuals or to different groups

PROPOSED COMMENTS BY THE DRS REGARDING THE REPORT AND RECOMMENDATION FROM THE NYSBA TASK FORCE ON ARTIFICIAL INTELLIGENCE

Paul R. Gupta

The DRS recommends to the Task Force that the following points should be added or discussed more fully. If it would be helpful to the Task Force, we can expand upon the points below, and draft fuller statements in a form that could be added to the Report.

1. Biometrics.
    a. The use of biometrics is one of the most significant current uses of AI. Many businesses use biometrics for hiring, supervision, and termination. State Legislatures have established rules with regards to the use, collection and storage of biometrics, such as face recognition, fingerprints, iris maps and voice prints. Illinois has led the way with broad biometrics legislation that includes a private right of action. The legislation covers the use of biometrics information (including selling that information), consent to obtain that information, and storage of that information.. (See IL Biometrics Information Privacy Act). New York and Maryland also have biometrics laws regarding employment, and Texas and Washington have broad biometrics laws. See also the following illustrative cases: Carpenter v. McDonald's Corp., 580 F. Supp. 3d 512 | Casetext Search + Citator, In re Facebook Biometric Info. Privacy Litig., Case No. 15-cv-03747-JD | Casetext Search + Citator, and Rivera v. Google, Inc., 366 F. Supp. 3d 998 | Casetext Search + Citator. Additionally, some municipalities, such as New York City, have biometrics laws that include a private right of action. (See The New York City Council - File #: Int 1170-2018 (nyc.gov)).

    b. Biometrics raise PII and other privacy concerns.

2. Bias:
    a. AI may create gender and racial bias, due to limited samples in databases used for comparisons (see: study exploring voice biometric disparities: Exploring racial and gender disparities in voice biometrics - PMC (nih.gov), The racism of technology - and why driverless cars could be the most dangerous example yet | Motoring | The Guardian, Study claims that self-driving cars more likely to drive into black people | Police Facial Recognition Technology Can't Tell Black People Apart | Scientific American)

    b. Ideological bias – AI can exacerbate ideological bias especially when used in conjunction with social media. AI can create its own echo chamber, generating spurious content to use as future training data, leading to ideologically based “hallucinations” and inaccuracies (see: Echo Chamber of AI: Model Collapse Risks | Deepgram, Polarization of Autonomous Generative AI Agents Under Echo Chambers (arxiv.org))

3. Confidentiality:
    a. Confidentiality concerns arise when entering information into AI engines (such as chatbots) and when such entries are then added to the training set for the AI. Such uses may violate Protective Orders for prior and future cases involving different parties. These concerns are compounded when chatbot results are analyzed by evaluative AI. For example, if biometrics data (see point 1 above) is analyzed by a chatbot to assist a mediator in preparing a mediator's proposal, multiple levels of confidentiality concerns arise. Such issues are especially important when some or all of the data that the AI "learns" is used for training the AI for work on future cases. These concerns can be alleviated by closed systems.

    b. Some AI providers allow for anonymous queries, while others explicitly state that they save inputs and prompts (see ChatGPT privacy policy, section 1 regarding user content).

AMERICAN BAR ASSOCIATION
STANDING COMMITTEE ON ETHICS AND PROFESSIONAL RESPONSIBILITY

Formal Opinion 512 **July 29, 2024**

Generative Artificial Intelligence Tools

To ensure clients are protected, lawyers using generative artificial intelligence tools must fully consider their applicable ethical obligations, including their duties to provide competent legal representation, to protect client information, to communicate with clients, to supervise their employees and agents, to advance only meritorious claims and contentions, to ensure candor toward the tribunal, and to charge reasonable fees.

I. Introduction

Many lawyers use artificial intelligence (AI) based technologies in their practices to improve the efficiency and quality of legal services to clients.[1] A well-known use is electronic discovery in litigation, in which lawyers use technology-assisted review to categorize vast quantities of documents as responsive or non-responsive and to segregate privileged documents. Another common use is contract analytics, which lawyers use to conduct due diligence in connection with mergers and acquisitions and large corporate transactions. In the realm of analytics, AI also can help lawyers predict how judges might rule on a legal question based on data about the judge's rulings; discover the summary judgment grant rate for every federal district judge; or evaluate how parties and lawyers may behave in current litigation based on their past conduct in similar litigation. And for basic legal research, AI may enhance lawyers' search results.

This opinion discusses a subset of AI technology that has more recently drawn the attention of the legal profession and the world at large – generative AI (GAI), which can create various types of new content, including text, images, audio, video, and software code in response to a user's prompts and questions.[2] GAI tools that produce new text are prediction tools that generate a statistically probable output when prompted. To accomplish this, these tools analyze large amounts of digital text culled from the internet or proprietary data sources. Some GAI tools are described as "self-learning," meaning they will learn from themselves as they cull more data. GAI tools may assist lawyers in tasks such as legal research, contract review, due diligence, document review, regulatory compliance, and drafting letters, contracts, briefs, and other legal documents.

---

[1] There is no single definition of artificial intelligence. At its essence, AI involves computer technology, software, and systems that perform tasks traditionally requiring human intelligence. The ability of a computer or computer-controlled robot to perform tasks commonly associated with intelligent beings is one definition. The term is frequently applied to the project of developing systems that appear to employ or replicate intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience. BRITTANICA, https://www.britannica.com/technology/artificial-intelligence (last visited July 12, 2024).

[2] George Lawton, *What is Generative AI? Everything You Need to Know*, TECHTARGET (July 12, 2024), https://www.techtarget.com/searchenterpriseai/definition/generative-AI.

GAI tools—whether general purpose or designed specifically for the practice of law—raise important questions under the ABA Model Rules of Professional Conduct.[3] What level of competency should lawyers acquire regarding a GAI tool? How can lawyers satisfy their duty of confidentiality when using a GAI tool that requires input of information relating to a representation? When must lawyers disclose their use of a GAI tool to clients? What level of review of a GAI tool's process or output is necessary? What constitutes a reasonable fee or expense when lawyers use a GAI tool to provide legal services to clients?

At the same time, as with many new technologies, GAI tools are a moving target—indeed, a *rapidly* moving target—in the sense that their precise features and utility to law practice are quickly changing and will continue to change in ways that may be difficult or impossible to anticipate. This Opinion identifies some ethical issues involving the use of GAI tools and offers general guidance for lawyers attempting to navigate this emerging landscape.[4] It is anticipated that this Committee and state and local bar association ethics committees will likely offer updated guidance on professional conduct issues relevant to specific GAI tools as they develop.

II. Discussion

A. Competence

Model Rule 1.1 obligates lawyers to provide competent representation to clients.[5] This duty requires lawyers to exercise the "legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation," as well as to understand "the benefits and risks associated" with the technologies used to deliver legal services to clients.[6] Lawyers may ordinarily achieve the requisite level of competency by engaging in self-study, associating with another competent lawyer, or consulting with an individual who has sufficient expertise in the relevant field.[7]

To competently use a GAI tool in a client representation, lawyers need not become GAI experts. Rather, lawyers must have a reasonable understanding of the capabilities and limitations

---

[3] Many of the professional responsibility concerns that arise with GAI tools are similar to the issues that exist with other AI tools and should be considered by lawyers using such technology.

[4] This opinion is based on the ABA Model Rules of Professional Conduct as amended by the ABA House of Delegates through August 2023. The Opinion addresses several imminent ethics issues associated with the use of GAI, but additional issues may surface, including those found in Model Rule 7.1 ("Communications Concerning a Lawyer's Services"), Model Rule 1.7 ("Conflict of Interest: Current Clients"), and Model Rule 1.9 ("Duties to Former Clients"). *See, e.g.*, Fla. State Bar Ass'n, Prof'l Ethics Comm. Op. 24-1, at 7 (2024) (discussing the use of GAI chatbots under Florida Rule 4-7.13, which prohibits misleading content and unduly manipulative or intrusive advertisements); Pa. State Bar Ass'n Comm. on Legal Ethics & Prof'l Resp. & Philadelphia Bar Ass'n Prof'l Guidance Comm. Joint Formal Op. 2024-200 [hereinafter Pa. & Philadelphia Joint Formal Opinion 2024-200], at 10 (2024) ("Because the large language models used in generative AI continue to develop, some without safeguards similar to those already in use in law offices, such as ethical walls, they may run afoul of Rules 1.7 and 1.9 by using the information developed from one representation to inform another."). Accordingly, lawyers should consider all rules before using GAI tools.

[5] MODEL RULES OF PROF'L CONDUCT R. 1.1 (2023) [hereinafter MODEL RULES].

[6] MODEL RULES R. 1.1 & cmt. [8]. *See also* ABA Comm. on Ethics & Prof'l Responsibility, Formal Op. 477R, at 2–3 (2017) [hereinafter ABA Formal Op. 477R] (discussing the ABA's "technology amendments" made to the Model Rules in 2012).

[7] MODEL RULES R. 1.1 cmts. [1], [2] & [4]; Cal. St. Bar, Comm. Prof'l Resp. Op. 2015-193, 2015 WL 4152025, at *2–3 (2015).

of the specific GAI technology that the lawyer might use. This means that lawyers should either acquire a reasonable understanding of the benefits and risks of the GAI tools that they employ in their practices or draw on the expertise of others who can provide guidance about the relevant GAI tool's capabilities and limitations.[8] This is not a static undertaking. Given the fast-paced evolution of GAI tools, technological competence presupposes that lawyers remain vigilant about the tools' benefits and risks.[9] Although there is no single right way to keep up with GAI developments, lawyers should consider reading about GAI tools targeted at the legal profession, attending relevant continuing legal education programs, and, as noted above, consulting others who are proficient in GAI technology.[10]

With the ability to quickly create new, seemingly human-crafted content in response to user prompts, GAI tools offer lawyers the potential to increase the efficiency and quality of their legal services to clients. Lawyers must recognize inherent risks, however.[11] One example is the risk of producing inaccurate output, which can occur in several ways. The large language models underlying GAI tools use complex algorithms to create fluent text, yet GAI tools are only as good as their data and related infrastructure. If the quality, breadth, and sources of the underlying data on which a GAI tool is trained are limited or outdated or reflect biased content, the tool might produce unreliable, incomplete, or discriminatory results. In addition, the GAI tools lack the ability to understand the meaning of the text they generate or evaluate its context.[12] Thus, they may combine otherwise accurate information in unexpected ways to yield false or inaccurate results.[13] Some GAI tools are also prone to "hallucinations," providing ostensibly plausible responses that have no basis in fact or reality.[14]

Because GAI tools are subject to mistakes, lawyers' uncritical reliance on content created by a GAI tool can result in inaccurate legal advice to clients or misleading representations to courts and third parties. Therefore, a lawyer's reliance on, or submission of, a GAI tool's output—without

---

[8] Pa. Bar Ass'n, Comm. on Legal Ethics & Prof'l Resp. Op. 2020-300, 2020 WL 2544268, at *2–3 (2020). *See also* Cal. State Bar, Standing Comm. on Prof'l Resp. & Conduct Op. 2023-208, 2023 WL 4035467, at *2 (2023) adopting a "reasonable efforts standard" and "fact-specific approach" to a lawyer's duty of technology competence, citing ABA Formal Opinion 477R, at 4).

[9] *See* New York County Lawyers Ass'n Prof'l Ethics Comm. Op. 749 (2017) (emphasizing that "[l]awyers must be responsive to technological developments as they become integrated into the practice of law"); Cal. St. Bar, Comm. Prof'l Resp. Op. 2015-193, 2015 WL 4152025, at *1 (2015) (discussing the level of competence required for lawyers to handle e-discovery issues in litigation).

[10] MODEL RULES R. 1.1 cmt. [8]; *see* Melinda J. Bentley, *The Ethical Implications of Technology in Your Law Practice: Understanding the Rules of Professional Conduct Can Prevent Potential Problems*, 76 J. MO. BAR 1 (2020) (identifying ways for lawyers to acquire technology competence skills).

[11] As further detailed in this opinion, lawyers' use of GAI raises confidentiality concerns under Model Rule 1.6 due to the risk of disclosure of, or unauthorized access to, client information. GAI also poses complex issues relating to ownership and potential infringement of intellectual property rights and even potential data security threats.

[12] *See,* W. Bradley Wendel, *The Promise and Limitations of AI in the Practice of Law,* 72 OKLA. L. REV. 21, 26 (2019) (discussing the limitations of AI based on an essential function of lawyers, making normative judgments that are impossible for AI).

[13] *See, e.g.*, Karen Weise & Cade Metz, *When A.I. Chatbots Hallucinate*, N.Y. TIMES (May 1, 2023).

[14] Ivan Moreno, *AI Practices Law 'At the Speed of Machines.' Is it Worth It?*, LAW360 (June 7, 2023); *See* Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, & Daniel E. Ho, *Hallucination Free? Assessing the Reliability of Leading AI Legal Research Tools*, STANFORD UNIVERSITY (June 26, 2024), *available at* https://dho.stanford.edu/wp-content/uploads/Legal_RAG_Hallucinations.pdf (study finding leading legal research companies' GAI systems "hallucinate between 17% and 33% of the time").

an appropriate degree of independent verification or review of its output—could violate the duty to provide competent representation as required by Model Rule 1.1.[15] While GAI tools may be able to significantly assist lawyers in serving clients, they cannot replace the judgment and experience necessary for lawyers to competently advise clients about their legal matters or to craft the legal documents or arguments required to carry out representations.

The appropriate amount of independent verification or review required to satisfy Rule 1.1 will necessarily depend on the GAI tool and the specific task that it performs as part of the lawyer's representation of a client. For example, if a lawyer relies on a GAI tool to review and summarize numerous, lengthy contracts, the lawyer would not necessarily have to manually review the entire set of documents to verify the results if the lawyer had previously tested the accuracy of the tool on a smaller subset of documents by manually reviewing those documents, comparing then to the summaries produced by the tool, and finding the summaries accurate. Moreover, a lawyer's use of a GAI tool designed specifically for the practice of law or to perform a discrete legal task, such as generating ideas, may require less independent verification or review, particularly where a lawyer's prior experience with the GAI tool provides a reasonable basis for relying on its results.

While GAI may be used as a springboard or foundation for legal work—for example, by generating an analysis on which a lawyer bases legal advice, or by generating a draft from which a lawyer produces a legal document—lawyers may not abdicate their responsibilities by relying solely on a GAI tool to perform tasks that call for the exercise of professional judgment. For example, lawyers may not leave it to GAI tools alone to offer legal advice to clients, negotiate clients' claims, or perform other functions that require a lawyer's personal judgment or participation.[16] Competent representation presupposes that lawyers will exercise the requisite level of skill and judgment regarding all legal work. In short, regardless of the level of review the lawyer selects, the lawyer is fully responsible for the work on behalf of the client.

Emerging technologies may provide an output that is of distinctively higher quality than current GAI tools produce, or may enable lawyers to perform work markedly faster and more economically, eventually becoming ubiquitous in legal practice and establishing conventional expectations regarding lawyers' duty of competence.[17] Over time, other new technologies have become integrated into conventional legal practice in this manner.[18] For example, "a lawyer would have difficulty providing competent legal services in today's environment without knowing how

---

[15] *See generally* ABA Comm. on Ethics & Prof'l Responsibility, Formal Op. 08-451, at 1 (2008) [hereinafter ABA Formal Op. 08-451] (concluding that "[a] lawyer may outsource legal or nonlegal support services provided the lawyer remains ultimately responsible for rendering competent legal services to the client under Model Rule 1.1").

[16] *See* Fla. State Bar Ass'n, Prof'l Ethics Comm. Op. 24-1, *supra* note 4.

[17] *See, e.g.*, Sharon Bradley, *Rule 1.1 Duty of Competency and Internet Research: Benefits and Risks Associated with Relevant Technology* at 7 (2019), *available at* https://ssrn.com/abstract=3485055 ("View Model Rule 1.1 as elastic. It is expanding as legal technology solutions expand. The ever-changing shape of this rule makes clear that a lawyer cannot simply learn technology today and never again update their skills or knowledge.").

[18] *See, e.g.*, Smith v. Lewis, 530 P.2d 589, 595 (Cal. 1975) (stating that a lawyer is expected "to possess knowledge of those plain and elementary principles of law which are commonly known by well-informed attorneys, and to discover those additional rules of law which, although not commonly known, may readily be found by *standard research techniques*") (emphasis added); Hagopian v. Justice Admin. Comm'n, 18 So. 3d 625, 642 (Fla. Dist. Ct. App. 2009) (observing that lawyers have "become expected to use computer-assisted legal research to ensure that their research is complete and up-to-date, but the costs of this service can be significant").

to use email or create an electronic document."[19] Similar claims might be made about other tools such as computerized legal research or internet searches.[20] As GAI tools continue to develop and become more widely available, it is conceivable that lawyers will eventually have to use them to competently complete certain tasks for clients.[21] But even in the absence of an expectation for lawyers to use GAI tools as a matter of course,[22] lawyers should become aware of the GAI tools relevant to their work so that they can make an informed decision, as a matter of professional judgment, whether to avail themselves of these tools or to conduct their work by other means.[23] As previously noted regarding the possibility of outsourcing certain work, "[t]here is no unique blueprint for the provision of competent legal services. Different lawyers may perform the same tasks through different means, all with the necessary 'legal knowledge, skill, thoroughness and preparation.'"[24] Ultimately, any informed decision about whether to employ a GAI tool must consider the client's interests and objectives.[25]

---

[19] ABA Formal Op. 477R, *supra* note 6, at 3 (quoting ABA COMMISSION ON ETHICS 20/20 REPORT 105A (Aug. 2012)).

[20] *See, e.g.*, Bradley, *supra* note 17, at 3 ("Today no competent lawyer would rely solely upon a typewriter to draft a contract, brief, or memo. Typewriters are no longer part of 'methods and procedures' used by competent lawyers."); Lawrence Duncan MacLachlan, *Gandy Dancers on the Web: How the Internet Has Raised the Bar on Lawyers' Professional Responsibility to Research and Know the Law*, 13 GEO. J. LEGAL ETHICS 607, 608 (2000) ("The lawyer in the twenty-first century who does not effectively use the Internet for legal research may fall short of the minimal standards of professional competence and be potentially liable for malpractice"); Ellie Margolis, *Surfin' Safari—Why Competent Lawyers Should Research on the Web*, 10 YALE J.L. & TECH. 82, 110 (2007) ("While a lawyer's research methods reveal a great deal about the competence of the research, the method of research is ultimately a secondary inquiry, only engaged in when the results of that research process is judged inadequate. A lawyer who provides the court with adequate controlling authority is not going to be judged incompetent whether she found that authority in print, electronically, or by any other means."); Michael Thomas Murphy, *The Search for Clarity in an Attorney's Duty to Google*, 18 LEGAL COMM. & RHETORIC: JALWD 133, 133 (2021) ("This Duty to Google contemplates that certain readily available information on the public Internet about a legal matter is so easily accessible that it must be discovered, collected, and examined by an attorney, or else that attorney is acting unethically, committing malpractice, or both"); Michael Whiteman, *The Impact of the Internet and Other Electronic Sources on an Attorney's Duty of Competence Under the Rules of Professional Conduct*, 11 ALB. L.J. SCI. & TECH. 89, 91 (2000) ("Unless it can be shown that the use of electronic sources in legal research has become a standard technique, then lawyers who fail to use electronic sources will not be deemed unethical or negligent in his or her failure to use such tools.").

[21] *See* MODEL RULES R. 1.1 cmt. [5] (stating that "[c]ompetent handling of a particular matter includes . . . [the] use of methods and procedures meeting the standards of competent practitioners"); New York County Lawyers Ass'n Prof'l Ethics Comm. Op. 749, 2017 WL 11659554, at *3 (2017) (explaining that the duty of competence covers not only substantive knowledge in different areas of the law, but also the manner in which lawyers provide legal services to clients).

[22] The establishment of such an expectation would likely require an increased acceptance of GAI tools across the legal profession, a track record of reliable results from those platforms, the widespread availability of these technologies to lawyers from a cost or financial standpoint, and robust client demand for GAI tools as an efficiency or cost-cutting measure.

[23] Model Rule 1.5's prohibition on unreasonable fees, as well as market forces, may influence lawyers to use new technology in favor of slower or less efficient methods.

[24] ABA Formal Op. 08-451, *supra* note 15, at 2. *See also id.* ("Rule 1.1 does not require that tasks be accomplished in any special way. The rule requires only that the lawyer who is responsible to the client satisfies her obligation to render legal services competently.").

[25] MODEL RULES R. 1.2(a).

B. Confidentiality

A lawyer using GAI must be cognizant of the duty under Model Rule 1.6 to keep confidential all information relating to the representation of a client, regardless of its source, unless the client gives informed consent, disclosure is impliedly authorized to carry out the representation, or disclosure is permitted by an exception.[26] Model Rules 1.9(c) and 1.18(b) require lawyers to extend similar protections to former and prospective clients' information. Lawyers also must make "reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of the client."[27]

Generally, the nature and extent of the risk that information relating to a representation may be revealed depends on the facts. In considering whether information relating to any representation is adequately protected, lawyers must assess the likelihood of disclosure and unauthorized access, the sensitivity of the information,[28] the difficulty of implementing safeguards, and the extent to which safeguards negatively impact the lawyer's ability to represent the client.[29]

Before lawyers input information relating to the representation of a client into a GAI tool, they must evaluate the risks that the information will be disclosed to or accessed by others outside the firm. Lawyers must also evaluate the risk that the information will be disclosed to or accessed by others *inside* the firm who will not adequately protect the information from improper disclosure or use[30] because, for example, they are unaware of the source of the information and that it originated with a client of the firm. Because GAI tools now available differ in their ability to ensure that information relating to the representation is protected from impermissible disclosure and access, this risk analysis will be fact-driven and depend on the client, the matter, the task, and the GAI tool used to perform it.[31]

Self-learning GAI tools into which lawyers input information relating to the representation, by their very nature, raise the risk that information relating to one client's representation may be disclosed improperly,[32] even if the tool is used exclusively by lawyers at the same firm.[33] This can occur when information relating to one client's representation is input into the tool, then later revealed in response to prompts by lawyers working on other matters, who then share that output with other clients, file it with the court, or otherwise disclose it. In other words, the self-learning

---

[26] MODEL RULES R. 1.6; MODEL RULES R. 1.6 cmt. [3].

[27] MODEL RULES R. 1.6(c).

[28] ABA Formal Op. 477R, *supra* note 6, at 1 (A lawyer "may be required to take special security precautions to protect against the inadvertent or unauthorized disclosure of client information when … the nature of the information requires a higher degree of security.").

[29] MODEL RULES R. 1.6, cmt. [18].

[30] *See* MODEL RULES R. 1.8(b), which prohibits use of information relating to the representation of a client to the disadvantage of the client.

[31] *See* ABA Formal Op. 477R, *supra* note 6, at 4 (rejecting specific security measures to protect information relating to a client's representation and advising lawyers to adopt a fact-specific approach to data security).

[32] *See generally* State Bar of Cal. Standing Comm. on Prof'l Resp. & Conduct, PRACTICAL GUIDANCE FOR THE USE OF GENERATIVE ARTIFICIAL INTELLIGENCE IN THE PRACTICE OF LAW (2024), *available at* https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf; Fla. State Bar Ass'n, Prof'l Ethics Comm. Op. 24-1, *supra* note 4.

[33] *See* Pa. & Philadelphia Joint Formal Opinion 2024-200, *supra* note 4, at 10 (noting risk that information relating to one representation may be used to inform work on another representation).

GAI tool may disclose information relating to the representation to persons outside the firm who are using the same GAI tool. Similarly, it may disclose information relating to the representation to persons in the firm (1) who either are prohibited from access to said information because of an ethical wall or (2) who could inadvertently use the information from one client to help another client, not understanding that the lawyer is revealing client confidences. Accordingly, because many of today's self-learning GAI tools are designed so that their output could lead directly or indirectly to the disclosure of information relating to the representation of a client, a client's informed consent is required prior to inputting information relating to the representation into such a GAI tool.[34]

When consent is required, it must be informed. For the consent to be informed, the client must have the lawyer's best judgment about why the GAI tool is being used, the extent of and specific information about the risk, including particulars about the kinds of client information that will be disclosed, the ways in which others might use the information against the client's interests, and a clear explanation of the GAI tool's benefits to the representation. Part of informed consent requires the lawyer to explain the extent of the risk that later users or beneficiaries of the GAI tool will have access to information relating to the representation. To obtain informed consent when using a GAI tool, merely adding general, boiler-plate provisions to engagement letters purporting to authorize the lawyer to use GAI is not sufficient.[35]

Because of the uncertainty surrounding GAI tools' ability to protect such information and the uncertainty about what happens to information both at input and output, it will be difficult to evaluate the risk that information relating to the representation will either be disclosed to or accessed by others inside the firm to whom it should not be disclosed as well as others outside the firm.[36] As a baseline, all lawyers should read and understand the Terms of Use, privacy policy, and related contractual terms and policies of any GAI tool they use to learn who has access to the information that the lawyer inputs into the tool or consult with a colleague or external expert who has read and analyzed those terms and policies.[37] Lawyers may need to consult with IT professionals or cyber security experts to fully understand these terms and policies as well as the manner in which GAI tools utilize information.

Today, there are uses of self-learning GAI tools in connection with a legal representation when client informed consent is not required because the lawyer will not be inputting information relating to the representation. As an example, if a lawyer is using the tool for idea generation in a manner that does not require inputting information relating to the representation, client informed consent would not be necessary.

---

[34] This conclusion is based on the risks and capabilities of GAI tools as of the publication of this opinion. As the technology develops, the risks may change in ways that would alter our conclusion. *See* Fla. State Bar Ass'n, Prof'l Ethics Comm. Op. 24-1, *supra* note 4, at 2; W. Va. Lawyer Disciplinary Bd. Op. 24-01 (2024), *available at* http://www.wvodc.org/pdf/AILEO24-01.pdf.

[35] *See* W. Va. Lawyer Disciplinary Bd. Op. 24-01, *supra* note 34.

[36] Magesh et al. *supra* note 14, at 23 (describing some of the GAI tools available to lawyers as "difficult for lawyers to assess when it is safe to trust them. Official documentation does not clearly illustrate what they can do for lawyers and in which areas lawyers should exercise caution.")

[37] Stephanie Pacheco, *Three Considerations for Attorneys Using Generative AI,* BLOOMBERG LAW ANALYSIS (June 16, 2023, 4:00 pm), https://news.bloomberglaw.com/bloomberg-law-analysis/analysis-three-considerations-for-attorneys-using-generative-ai?context=search&index=7.

C. Communication

Where Model Rule 1.6 does not require disclosure and informed consent, the lawyer must separately consider whether other Model Rules, particularly Model Rule 1.4, require disclosing the use of a GAI tool in the representation.

Model Rule 1.4, which addresses lawyers' duty to communicate with their clients, builds on lawyers' legal obligations as fiduciaries, which include "the duty of an attorney to advise the client promptly whenever he has any information to give which it is important the client should receive."[38] Of particular relevance, Model Rule 1.4(a)(2) states that a lawyer shall "reasonably consult with the client about the means by which the client's objectives are to be accomplished." Additionally, Model Rule 1.4(b) obligates lawyers to explain matters "to the extent reasonably necessary to permit a client to make an informed decision regarding the representation." Comment [5] to Rule 1.4 explains, "the lawyer should fulfill reasonable client expectations for information consistent with the duty to act in the client's best interests, and the client's overall requirements as to the character of representation." Considering these underlying principles, questions arise regarding whether and when lawyers might be required to disclose their use of GAI tools to clients pursuant to Rule 1.4.

The facts of each case will determine whether Model Rule 1.4 requires lawyers to disclose their GAI practices to clients or obtain their informed consent to use a particular GAI tool. Depending on the circumstances, client disclosure may be unnecessary.

Of course, lawyers must disclose their GAI practices if asked by a client how they conducted their work, or whether GAI technologies were employed in doing so, or if the client expressly requires disclosure under the terms of the engagement agreement or the client's outside counsel guidelines.[39] There are also situations where Model Rule 1.4 requires lawyers to discuss their use of GAI tools unprompted by the client.[40] For example, as discussed in the previous section, clients would need to be informed in advance, and to give informed consent, if the lawyer proposes to input information relating to the representation into the GAI tool.[41] Lawyers must also consult clients when the use of a GAI tool is relevant to the basis or reasonableness of a lawyer's fee.[42]

Client consultation about the use of a GAI tool is also necessary when its output will influence a significant decision in the representation,[43] such as when a lawyer relies on GAI

---

[38] Baker v. Humphrey, 101 U.S. 494, 500 (1879).

[39] *See, e.g.*, MODEL RULES R. 1.4(a)(4) ("A lawyer shall . . . promptly comply with reasonable requests for information[.]").

[40] *See* MODEL RULES R. 1.4(a)(1) (requiring lawyers to "promptly inform the client of any decision or circumstance with respect to which the client's informed consent" is required by the rules of professional conduct).

[41] See section B for a discussion of confidentiality issues under Rule 1.6.

[42] See section F for a discussion of fee issues under Rule 1.5.

[43] Guidance may be found in ethics opinions requiring lawyers to disclose their use of temporary lawyers whose involvement is significant or otherwise material to the representation. *See, e.g.*, Va. State Bar Legal Ethics Op. 1850, 2010 WL 5545407, at *5 (2010) (acknowledging that "[t]here is little purpose to informing a client every time a lawyer outsources legal support services that are truly tangential, clerical, or administrative in nature, or even when basic legal research or writing is outsourced without any client confidences being revealed"); Cal. State Bar, Standing Comm. on Prof'l Resp. & Conduct Op. 2004-165, 2004 WL 3079030, at *2–3 (2004) (opining that a

technology to evaluate potential litigation outcomes or jury selection. A client would reasonably want to know whether, in providing advice or making important decisions about how to carry out the representation, the lawyer is exercising independent judgment or, in the alternative, is deferring to the output of a GAI tool. Or there may be situations where a client retains a lawyer based on the lawyer's particular skill and judgment, when the use of a GAI tool, without the client's knowledge, would violate the terms of the engagement agreement or the client's reasonable expectations regarding how the lawyer intends to accomplish the objectives of the representation.

It is not possible to catalogue every situation in which lawyers must inform clients about their use of GAI. Again, lawyers should consider whether the specific circumstances warrant client consultation about the use of a GAI tool, including the client's needs and expectations, the scope of the representation, and the sensitivity of the information involved. Potentially relevant considerations include the GAI tool's importance to a particular task, the significance of that task to the overall representation, how the GAI tool will process the client's information, and the extent to which knowledge of the lawyer's use of the GAI tool would affect the client's evaluation of or confidence in the lawyer's work.

Even when Rule 1.6 does not require informed consent and Rule 1.4 does not require a disclosure regarding the use of GAI, lawyers may tell clients how they employ GAI tools to assist in the delivery of legal services. Explaining this may serve the interest of effective client communication. The engagement agreement is a logical place to make such disclosures and to identify any client instructions on the use of GAI in the representation.[44]

D. Meritorious Claims and Contentions and Candor Toward the Tribunal

Lawyers using GAI in litigation have ethical responsibilities to the courts as well as to clients. Model Rules 3.1, 3.3, and 8.4(c) may be implicated by certain uses. Rule 3.1 states, in part, that "[a] lawyer shall not bring or defend a proceeding, or assert or controvert and issue therein, unless there is a basis in law or fact for doing so that is not frivolous." Rule 3.3 makes it clear that lawyers cannot knowingly make any false statement of law or fact to a tribunal or fail to correct a material false statement of law or fact previously made to a tribunal.[45] Rule 8.4(c) provides that a

---

lawyer must disclose the use of a temporary lawyer to a client where the temporary lawyer's use constitutes a "significant development" in the matter and listing relevant considerations); N.Y. State Bar Ass'n, Comm on Prof'l Ethics 715, at 7 (1999) (opining that "whether a law firm needs to disclose to the client and obtain client consent for the participation of a Contract lawyer depends upon whether client confidences will be disclosed to the lawyer, the degree of involvement of the lawyer in the matter, and the significance of the work done by the lawyer"); D.C. Bar Op. 284, at 4 (1988) (recommending client disclosure "whenever the proposed use of a temporary lawyer to perform work on the client's matter appears reasonably likely to be material to the representation or to affect the client's reasonable expectations"); Fla. State Bar Ass'n, Comm. on Prof'l Ethics Op. 88-12, 1988 WL 281590, at *2 (1988) (stating that disclosure of a temporary lawyer depends "on whether the client would likely consider the information material");.

[44] For a discussion of what client notice and informed consent under Rule 1.6 may require, see section B.

[45] MODEL RULES R. 3.3(a) reads: "A lawyer shall not knowingly: (1) make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer; (2) fail to disclose to the tribunal legal authority in the controlling jurisdiction known to the lawyer to be directly adverse to the position of the client and not disclosed by opposing counsel; or (3) offer evidence that the lawyer knows to be false. If a lawyer, the lawyer's client, or a witness called by the lawyer, has offered material evidence and the lawyer comes to know of its falsity, the lawyer shall take reasonable remedial measures, including, if

lawyer shall not engage in "conduct involving dishonesty, fraud, deceit or misrepresentation." Even an unintentional misstatement to a court can involve a misrepresentation under Rule 8.4(c). Therefore, output from a GAI tool must be carefully reviewed to ensure that the assertions made to the court are not false.

Issues that have arisen to date with lawyers' use of GAI outputs include citations to nonexistent opinions, inaccurate analysis of authority, and use of misleading arguments.[46]

Some courts have responded by requiring lawyers to disclose their use of GAI.[47] As a matter of competence, as previously discussed, lawyers should review for accuracy all GAI outputs. In judicial proceedings, duties to the tribunal likewise require lawyers, before submitting materials to a court, to review these outputs, including analysis and citations to authority, and to correct errors, including misstatements of law and fact, a failure to include controlling legal authority, and misleading arguments.

E. Supervisory Responsibilities

Model Rules 5.1 and 5.3 address the ethical duties of lawyers charged with managerial and supervisory responsibilities and set forth those lawyers' responsibilities with regard to the firm, subordinate lawyers, and nonlawyers. Managerial lawyers must create effective measures to ensure that all lawyers in the firm conform to the rules of professional conduct,[48] and supervisory lawyers must supervise subordinate lawyers and nonlawyer assistants to ensure that subordinate lawyers and nonlawyer assistants conform to the rules.[49] These responsibilities have implications for the use of GAI tools by lawyers and nonlawyers.

Managerial lawyers must establish clear policies regarding the law firm's permissible use of GAI, and supervisory lawyers must make reasonable efforts to ensure that the firm's lawyers and nonlawyers comply with their professional obligations when using GAI tools.[50] Supervisory obligations also include ensuring that subordinate lawyers and nonlawyers are trained,[51] including in the ethical and practical use of the GAI tools relevant to their work as well as on risks associated with relevant GAI use.[52] Training could include the basics of GAI technology, the capabilities and limitations of the tools, ethical issues in use of GAI and best practices for secure data handling, privacy, and confidentiality.

---

necessary, disclosure to the tribunal. A lawyer may refuse to offer evidence, other than the testimony of a defendant in a criminal matter, that the lawyer reasonably believes is false."

[46] *See* DC Bar Op. 388 (2024).

[47] Lawyers should consult with the applicable court's local rules to ensure that they comply with those rules with respect to AI use. As noted in footnote 4, no one opinion could address every ethics issue presented when a lawyer uses GAI. For example, depending on the facts, issues relating to Model Rule 3.4(c) could be presented.

[48] *See* MODEL RULES R. 1.0(c) for the definition of firm.

[49] ABA Formal Op. 08-451, *supra* note 15.

[50] MODEL RULES R. 5.1.

[51] *See* ABA Comm. on Ethics & Prof'l Responsibility, Formal Op. 467 (2014).

[52] *See generally*, MODEL RULES R. 1.1, cmt. [8]. One training suggestion is that all materials produced by GAI tools be marked as such when stored in any client or firm file so future users understand potential fallibility of the work.

Lawyers have additional supervisory obligations insofar as they rely on others outside the law firm to employ GAI tools in connection with the legal representation. Model Rule 5.3(b) imposes a duty on lawyers with direct supervisory authority over a nonlawyer to make "reasonable efforts to ensure that" the nonlawyer's conduct conforms with the professional obligations of the lawyer. Earlier opinions recognize that when outsourcing legal and nonlegal services to third-party providers, lawyers must ensure, for example, that the third party will do the work capably and protect the confidentiality of information relating to the representation.[53] These opinions note the importance of: reference checks and vendor credentials; understanding vendor's security policies and protocols; familiarity with vendor's hiring practices; using confidentiality agreements; understanding the vendor's conflicts check system to screen for adversity among firm clients; and the availability and accessibility of a legal forum for legal relief for violations of the vendor agreement. These concepts also apply to GAI providers and tools.

Earlier opinions regarding technological innovations and other innovations in legal practice are instructive when considering a lawyer's use of a GAI tool that requires the disclosure and storage of information relating to the representation.[54] In particular, opinions developed to address cloud computing and outsourcing of legal and nonlegal services suggest that lawyers should:

- ensure that the [GAI tool] is configured to preserve the confidentiality and security of information, that the obligation is enforceable, and that the lawyer will be notified in the event of a breach or service of process regarding production of client information;[55]
- investigate the [GAI tool's] reliability, security measures, and policies, including limitations on the [the tool's] liability;[56]
- determine whether the [GAI tool] retains information submitted by the lawyer before and after the discontinuation of services or asserts proprietary rights to the information;[57] and
- understand the risk that [GAI tool servers] are subject to their own failures and may be an attractive target of cyber-attacks.[58]

F. Fees

Model Rule 1.5, which governs lawyers' fees and expenses, applies to representations in which a lawyer charges the client for the use of GAI. Rule 1.5(a) requires a lawyer's fees and expenses to be reasonable and includes a non-exclusive list of criteria for evaluating whether a fee

---

[53] ABA Formal Op. 08-451, *supra* note 15; ABA Formal. Op. 477R, *supra* note 6.

[54] *See* ABA Formal Op. 08-451, *supra* note 15.

[55] Fla. Bar Advisory Op. 12-3 (2013).

[56] *Id.* citing Iowa State Bar Ass'n Comm. on Ethics & Practice Guidelines Op. 11-01 (2011) [hereinafter Iowa Ethics Opinion 11-01].

[57] Fla. Bar Advisory Op. 24-1, *supra* note 4; Fla. Bar Advisory Op. 12-3, *supra* note 55; Iowa Ethics Opinion 11-01, *supra* note 56.

[58] Fla. Bar Advisory Op. 12-3, *supra* note 55; *See generally* Melissa Heikkila, *Three Ways AI Chatbots are a Security Disaster*, MIT TECHNOLOGY REVIEW (Apr. 3, 2023), www.technologyreview.com/2023/04/03/1070893/three-ways-ai-chatbots-are-a-security-disaster/.

or expense is reasonable.[59] Rule 1.5(b) requires a lawyer to communicate to a client the basis on which the lawyer will charge for fees and expenses unless the client is a regularly represented client and the terms are not changing. The required information must be communicated before or within a reasonable time of commencing the representation, preferably in writing. Therefore, before charging the client for the use of the GAI tools or services, the lawyer must explain the basis for the charge, preferably in writing.

GAI tools may provide lawyers with a faster and more efficient way to render legal services to their clients, but lawyers who bill clients an hourly rate for time spent on a matter must bill for their actual time. ABA Formal Ethics Opinion 93-379 explained, "the lawyer who has agreed to bill on the basis of hours expended does not fulfill her ethical duty if she bills the client for more time than she has actually expended on the client's behalf."[60] If a lawyer uses a GAI tool to draft a pleading and expends 15 minutes to input the relevant information into the GAI program, the lawyer may charge for the 15 minutes as well as for the time the lawyer expends to review the resulting draft for accuracy and completeness. As further explained in Opinion 93-379, "If a lawyer has agreed to charge the client on [an hourly] basis and it turns out that the lawyer is particularly efficient in accomplishing a given result, it nonetheless will not be permissible to charge the client for more hours than were actually expended on the matter,"[61] because "[t]he client should only be charged a reasonable fee for the legal services performed."[62] The "goal should be solely to compensate the lawyer fully for time reasonably expended, an approach that if followed will not take advantage of the client."[63]

The factors set forth in Rule 1.5(a) also apply when evaluating the reasonableness of charges for GAI tools when the lawyer and client agree on a flat or contingent fee.[64] For example, if using a GAI tool enables a lawyer to complete tasks much more quickly than without the tool, it may be unreasonable under Rule 1.5 for the lawyer to charge the same flat fee when using the GAI tool as when not using it. "A fee charged for which little or no work was performed is an unreasonable fee."[65]

The principles set forth in ABA Formal Opinion 93-379 also apply when a lawyer charges GAI work as an expense. Rule 1.5(a) requires that disbursements, out-of-pocket expenses, or additional charges be reasonable. Formal Opinion 93-379 explained that a lawyer may charge the

---

[59] The listed considerations are (1) the time and labor required, the novelty and difficulty of the questions involved, and the skill requisite to perform the legal service properly; (2) the likelihood, if apparent to the client, that the acceptance of the particular employment will preclude other employment by the lawyer; (3) the fee customarily charged in the locality for similar legal services; (4) the amount involved and the results obtained; (5) the time limitations imposed by the client or by the circumstances; (6) the nature and length of the professional relationship with the client; (7) the experience, reputation, and ability of the lawyer or lawyers performing the services; and (8) whether the fee is fixed or contingent.

[60] ABA Comm. on Ethics & Prof'l Responsibility, Formal Op. 93-379, at 6 (1993) [hereinafter ABA Formal Op. 93-379].

[61] *Id.*

[62] *Id.* at 5.

[63] *Id.*

[64] *See, e.g.*, *Williams Cos. v. Energy Transfer LP*, 2022 Del. Ch. LEXIS 207, 2022 WL 3650176 (Del. Ch. Aug. 25, 2022) (applying same principles to contingency fee).

[65] Att'y Grievance Comm'n v. Monfried, 794 A.2d 92, 103 (Md. 2002) (finding that a lawyer violated Rule 1.5 by charging a flat fee of $1,000 for which the lawyer did little or no work).

client for disbursements incurred in providing legal services to the client. For example, a lawyer typically may bill to the client the actual cost incurred in paying a court reporter to transcribe a deposition or the actual cost to travel to an out-of-town hearing.[66] Absent contrary disclosure to the client, the lawyer should not add a surcharge to the actual cost of such expenses and should pass along to the client any discounts the lawyer receives from a third-party provider.[67] At the same time, lawyers may not bill clients for general office overhead expenses including the routine costs of "maintaining a library, securing malpractice insurance, renting of office space, purchasing utilities, and the like."[68] Formal Opinion 93-379 noted, "[i]n the absence of disclosure to a client in advance of the engagement to the contrary," such overhead should be "subsumed within" the lawyer's charges for professional services.[69]

In applying the principles set out in ABA Formal Ethics Opinion 93-379 to a lawyer's use of a GAI tool, lawyers should analyze the characteristics and uses of each GAI tool, because the types, uses, and cost of GAI tools and services vary significantly. To the extent a particular tool or service functions similarly to equipping and maintaining a legal practice, a lawyer should consider its cost to be overhead and not charge the client for its cost absent a contrary disclosure to the client in advance. For example, when a lawyer uses a GAI tool embedded in or added to the lawyer's word processing software to check grammar in documents the lawyer drafts, the cost of the tool should be considered to be overhead. In contrast, when a lawyer uses a third-party provider's GAI service to review thousands of voluminous contracts for a particular client and the provider charges the lawyer for using the tool on a per-use basis, it would ordinarily be reasonable for the lawyer to bill the client as an expense for the actual out-of-pocket expense incurred for using that tool.

As acknowledged in ABA Formal Opinion 93-379, perhaps the most difficult issue is determining how to charge clients for providing in-house services that are not required to be included in general office overhead and for which the lawyer seeks reimbursement. The opinion concluded that lawyers may pass on reasonable charges for "photocopying, computer research, . . . and similar items" rather than absorbing these expenses as part of the lawyers' overhead as many lawyers would do.[70] For example, a lawyer may agree with the client in advance on the specific rate for photocopying, such as $0.15 per page. Absent an advance agreement, the lawyer "is obliged to charge the client no more than the direct cost associated with the service (i.e., the actual cost of making a copy on the photocopy machine) plus a reasonable allocation of overhead expenses directly associated with the provision of the service (e.g., the salary of the photocopy machine operator)."[71]

---

[66] ABA Formal Op. 93-379 at 7.

[67] *Id.* at 8.

[68] *Id.* at 7.

[69] *Id.*

[70] *Id.* at 8.

[71] *Id.* Opinion 93-379 also explained, "It is not appropriate for the Committee, in addressing ethical standards, to opine on the various accounting issues as to how one calculates direct cost and what may or may not be included in allocated overhead. These are questions which properly should be reserved for our colleagues in the accounting profession. Rather, it is the responsibility of the Committee to explain the principles it draws from the mandate of Model Rule 1.5's injunction that fees be reasonable. Any reasonable calculation of direct costs as well as any reasonable allocation of related overhead should pass ethical muster. On the other hand, in the absence of an agreement to the contrary, it is impermissible for a lawyer to create an additional source of profit for the law firm beyond that which is contained in the provision of professional services themselves. The lawyer's stock in trade is the sale of legal services, not photocopy paper, tuna fish sandwiches, computer time or messenger services." *Id.*

These same principles apply when a lawyer uses a proprietary, in-house GAI tool in rendering legal services to a client. A firm may have made a substantial investment in developing a GAI tool that is relatively unique and that enables the firm to perform certain work more quickly or effectively. The firm may agree in advance with the client about the specific rates to be charged for using a GAI tool, just as it would agree in advance on its legal fees. But not all in-house GAI tools are likely to be so special or costly to develop, and the firm may opt not to seek the client's agreement on expenses for using the technology. Absent an agreement, the firm may charge the client no more than the direct cost associated with the tool (if any) plus a reasonable allocation of expenses directly associated with providing the GAI tool, while providing appropriate disclosures to the client consistent with Formal Opinion 93-379. The lawyer must ensure that the amount charged is not duplicative of other charges to this or other clients.

Finally, on the issue of reasonable fees, in addition to the time lawyers spend using various GAI tools and services, lawyers also will expend time to gain knowledge about those tools and services. Rule 1.1 recognizes that "[c]ompetent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation." Comment [8] explains that "[t]o maintain the requisite knowledge and skill [to be competent], a lawyer should keep abreast of changes in the law and its practice, *including the benefits and risks associated with relevant technology*, engaging in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject."[72] Lawyers must remember that they may not charge clients for time necessitated by their own inexperience.[73] Therefore, a lawyer may not charge a client to learn about how to use a GAI tool or service that the lawyer will regularly use for clients because lawyers must maintain competence in the tools they use, including but not limited to GAI technology. However, if a client explicitly requests that a specific GAI tool be used in furtherance of the matter and the lawyer is not knowledgeable in using that tool, it may be appropriate for the lawyer to bill the client to gain the knowledge to use the tool effectively. Before billing the client, the lawyer and the client should agree upon any new billing practices or billing terms relating to the GAI tool and, preferably, memorialize the new agreement.

III. Conclusion

Lawyers using GAI tools have a duty of competence, including maintaining relevant technological competence, which requires an understanding of the evolving nature of GAI. In

---

[72] MODEL RULES R. 1.1, cmt. [8] (emphasis added); *see also* ABA Comm. on Ethics & Prof'l Responsibility, Formal Op. 498 (2021).

[73] Heavener v. Meyers, 158 F. Supp. 2d 1278 (E.D. Okla. 2001) (five hundred hours for straightforward Fourth Amendment excessive-force claim and nineteen hours for research on Eleventh Amendment defense indicated excessive billing due to counsel's inexperience); *In re* Poseidon Pools of Am., Inc., 180 B.R. 718 (Bankr. E.D.N.Y. 1995) (denying compensation for various document revisions; "we note that given the numerous times throughout the Final Application that Applicant requests fees for revising various documents, Applicant fails to negate the obvious possibility that such a plethora of revisions was necessitated by a level of competency less than that reflected by the Applicant's billing rates"); Att'y Grievance Comm'n v. Manger, 913 A.2d 1 (Md. 2006) ("While it may be appropriate to charge a client for case-specific research or familiarization with a unique issue involved in a case, general education or background research should not be charged to the client."); *In re* Hellerud, 714 N.W.2d 38 (N.D. 2006) (reduction in hours, fee refund of $5,651.24, and reprimand for lawyer unfamiliar with North Dakota probate work who charged too many hours at too high a rate for simple administration of cash estate; "it is counterintuitive to charge a higher hourly rate for knowing less about North Dakota law").

using GAI tools, lawyers also have other relevant ethical duties, such as those relating to confidentiality, communication with a client, meritorious claims and contentions, candor toward the tribunal, supervisory responsibilities regarding others in the law office using the technology and those outside the law office providing GAI services, and charging reasonable fees. With the ever-evolving use of technology by lawyers and courts, lawyers must be vigilant in complying with the Rules of Professional Conduct to ensure that lawyers are adhering to their ethical responsibilities and that clients are protected.

AMERICAN BAR ASSOCIATION STANDING COMMITTEE ON ETHICS AND PROFESSIONAL RESPONSIBILITY

321 N. Clark Street, Chicago, Illinois 60654-4714 Telephone (312) 988-5328

CHAIR: Bruce Green, New York, NY ■ Mark A. Armitage, Detroit, MI ■ Matthew Corbin, Olathe, KS ■ Robinjit Kaur Eagleson, Lansing, MI ■ Brian Shannon Faughnan, Memphis, TN ■ Hilary P. Gerzhoy, Washington, D.C. ■ Wendy Muchman, Chicago, IL ■ Tim Pierce, Madison, WI ■ Hon. Jennifer A. Rymell, Fort Worth, TX ■ Charles Vigil, Albuquerque, NM

CENTER FOR PROFESSIONAL RESPONSIBILITY: Mary McDermott, Lead Senior Counsel

©2024 by the American Bar Association. All rights reserved.

DAIL – the Database of AI Litigation

This database, a project of the GW Ethical Tech Initiative and the GW Center for Law and Technology, presents information about ongoing and completed litigation involving artificial intelligence, including machine learning. It covers cases from complaint forward – as soon as we learn of them – whether or not they generate published decisions. It is intended to be broad in scope, covering everything from algorithms used in hiring and credit and criminal sentencing decisions to liability for accidents involving autonomous vehicles. It also includes some cases concerning statistical analysis and data protection that may not directly involve artificial intelligence, but that are of particular relevance to AI projects. It includes cases addressing whether AIs can be authors of works protected by copyright law, or inventors of inventions protected by patent law, but it does not include litigation concerning patents that may involve artificial intelligence or machine learning.

If you know of AI litigation that you don't see documented in the database, or if you have other suggestions, please tell us, using this contact form. Research and writing: GW Law students Jenna Fattah, Xiaonan (Caroline) Qu, Beatriz Beserra, Andrew Ware, Allie Schiele, Junhao Chen, Sydney Huppert, Zoe Kim, Molly Brown, Jenny Tavares, Rachel Lobel, Victoria Neal, Maya Clark, Aiymgul Kachyke and Prof. Robert Brauneis. Implementation of the database on the Caspio platform: GW students Sean Zhao, Aneri Girishbhai Patel, Ji Wang, Sydney Huppert, and Victoria Neal, and Prof. Robert Brauneis. This database is based on work supported in part by the Institute for Trustworthy AI in Law and Society (TRAILS), which is supported by the National Science Foundation under Award No. 2229885. Any opinions, findings, and conclusions or recommendations expressed in this database are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

. .

Campus Advisories

EO/Nondiscrimination Policy

Website Privacy Notice

Contact GW

Accessibility

Terms of Use

Copyright

Report a barrier to accessibility.

THE NEW YORK CITY BAR ASSOCIATION
COMMITTEE ON PROFESSIONAL ETHICS

FORMAL OPINION 2024-5: ETHICAL OBLIGATIONS OF LAWYERS AND LAW FIRMS RELATING TO THE USE OF GENERATIVE ARTIFICIAL INTELLIGENCE IN THE PRACTICE OF LAW

TOPIC: The use of generative artificial intelligence by New York lawyers, law firms, legal departments, government law offices and legal assistance organizations.

DIGEST: This opinion provides general guidance on the use of tools that use generative artificial intelligence.

RULES: **1.1, 1.2(d), 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, 1.10, 1.11, 1.12, 3.1, 3.3, 5.1, 5.2, 5.3, 7.1, 7.3, 8.4**

QUESTION: The availability of tools to assist lawyers in their practice that employ generative artificial intelligence has been dramatically expanding and continues to grow. What are the ethical issues that lawyers should consider when deciding whether to use these tools and, if the decision is made to do so, how to use them?

OPINION: When using generative artificial intelligence tools, a lawyer should take into account the duty of confidentiality, the obligation to avoid conflicts of interest, the duty of competence and diligence, the rules governing advertising and solicitation, the duty to comply with the law, the duty to supervise both lawyers and non-lawyers, the duty of subordinate attorneys, the duty to consult with clients, the duty of candor to tribunals, the prohibition on making non-meritorious claims and contentions, the limitations on what a lawyer may charge for fees and costs, and the prohibition on discrimination.

Introduction

Generative artificial intelligence ("Generative AI"), like any technology, must be used in a manner that comports with a lawyer's ethical obligations. General-purpose technology platforms offer AI chatbots. Legal research platforms tout "legal generative AI" that can draft, analyze documents, and provide legal citations. Even data management vendors offer Generative AI-assisted review, analytic, and visualization capabilities. This summary of currently available tools will likely soon be outdated because of the rapid evolution of Generative AI. This guidance, therefore, is general. We expect that this advice will be updated and supplemented in years to come to cover issues not yet anticipated.

This Opinion provides guidance on the ethical obligations of lawyers and law firms relating to the use of Generative AI. It follows and is consistent with the format used by the Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law released by the California State Bar's Standing Committee on Professional Responsibility and Conduct in November 2023.[1] This

---

[1] State Bar of Cal., Standing Comm. on Pro. Resp. & Conduct, *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law* (Nov. 16. 2023) ("California Guidance"),

Opinion is in the same format as the California State Bar's guidance and contains multiple quotations from that guidance. Like the California State Bar and other bar associations that have addressed Generative AI,[2] we believe that when addressing developing areas, lawyers need guardrails and not hard-and-fast restrictions or new rules that could stymie developments. By including advice specifically based on New York Rules and practice, this Opinion is intended to be helpful to the New York Bar.

Applicable Authorities	New York Guidance
Duty of Confidentiality Rule 1.6	Generative AI systems are able to use information that is inputted, including prompts, uploaded data, documents, and other resources, to train AI. They may also share inputted information with third parties or use it for other purposes.[3] Even if a system does not use or share inputted information, it may lack "reasonable or adequate security."[4] Without client consent, a lawyer must not input confidential client information into any Generative AI system that will share the inputted confidential information with third parties.[5] Even with consent, a lawyer should "avoid entering details that can be used to identify the client."[6] Consent is not needed if no confidential client information is shared, for example through anonymization of client information. Generative AI systems that keep inputted information entirely within the firm's own protected databases, sometimes called "closed"

---

https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf; *see also* Am. Bar Ass'n, Formal Op. 512 (2024); Fla. Bar Bd. Rev. Comm. on Pro. Ethics, Op. 24-1 (2024); D.C. Bar Ethics Op. 388 (April 2024); N.J. STATE BAR ASS'N, TASK FORCE ON ARTIFICIAL INTELLIGENCE (AI) AND THE LAW: REPORT, REQUESTS, RECOMMENDATIONS, AND FINDINGS (2024), https://njsba.com/wp-content/uploads/2024/05/NJSBA-TASK-FORCE-ON-AI-AND-THE-LAW-REPORT-final.pdf; N.Y. STATE BAR ASS'N, REPORT & RECOMMENDATIONS OF THE NEW YORK STATE BAR ASSOCIATION TASK FORCE ON ARTIFICIAL INTELLIGENCE (2024), https://www.nycbar.org/wp-content/uploads/2024/06/20221290_AI_NYS_Judiciary.pdf. (All websites last accessed on Aug. 5, 2024).

[2] In general, this Opinion is consistent with the ABA, California Bar, Florida Bar, District of Columbia Bar, and New Jersey Bar opinions cited in Footnote 1. However, the New York State Bar suggests adoption of certain rules to address Generative AI, which we believe is premature because of the rapid pace of technological development and change. *See, e.g.*, N.Y. STATE BAR ASS'N, *supra*, at 53–56.

[3] Generative AI systems that share inputted information with third parties are sometimes called "open" systems.

[4] California Guidance at 2.

[5] Lawyers may wish to obtain advance client consent to use Generative AI that will involve sharing of client information, but, because such consent must be knowing, the client must understand the potential consequences of such information-sharing for the consent to be effective. *See* N.Y. State Op. 1020 ¶ 10 (a lawyer "may post and share documents using a 'cloud' data storage tool" that does not provide "reasonable protection to confidential client information" only where "the lawyer obtains informed consent from the client after advising the client of the relevant risks").

[6] *Id.*

	systems, do not present these risks. But a lawyer must not input any confidential information of the client into any Generative AI system that lacks adequate confidentiality and security protections, regardless of whether the system uses or shares inputted information, unless the client has given informed consent to the lawyer's doing so. Even with closed systems, a lawyer must take care that confidential information is not improperly shared with other persons at or clients of the same law firm, including persons who are prohibited access to the information because of an ethical wall.[7] A lawyer or law firm[8] should "consult with IT professionals or cybersecurity experts to the extent necessary for the lawyer or law firm to ensure that any Generative AI system in which a lawyer would input confidential client information adheres to stringent security, confidentiality, and data retention protocols."[9] A lawyer should review the system's Terms of Use. "A lawyer who intends to use confidential information in a Generative AI product should ensure that the provider does not share inputted information with third parties or use the information for its own use in any manner, including to train or improve its product," again without informed client consent.[10] Terms of Use can change frequently and a lawyer's obligation to understand the system's use of inputs is continuing. Accordingly, lawyers should periodically monitor Terms of Use or other information to learn about any changes that might compromise confidential information.[11] A law firm may wish to consider implementing policies and control procedures to regulate the use of confidential client information in Generative AI systems if the law firm is going to make use of such systems.
Conflicts of Interest	Where a Generative AI system uses client information, a law firm must ensure that the system implements any ethical screens required under the Rules. For example, if an ethical

---

[7] *See* Am. Bar Ass'n, Formal Op. 512 at 6-7 (2024).

[8] Consistent with Rule 1.0(h), in this Opinion "law firm" includes a private firm as well as qualified legal assistance organizations, government law offices and corporations, and other entities' legal departments.

[9] California Guidance at 2.

[10] *Id.*

[11] *See* N.Y. STATE BAR ASS'N, *supra*, at 58.

Rule 1.7; Rule 1.8;
Rule 1.9; Rule 1.10;
Rule 1.11; Rule 1.12
	screen excludes a lawyer from any information or documents with respect to a client, the lawyer must be not exposed to such information or documents through the law firm's Generative AI systems.

Duties of Competence and Diligence Rule 1.1; Rule 1.3	A lawyer should be aware that currently Generative AI outputs may include historical information that is false, inaccurate, or biased. "A lawyer must ensure the competent use of technology, including the associated benefits and risks, and apply diligence and prudence with respect to facts and law."[12] "Before selecting and using a Generative AI tool, a lawyer should understand to a reasonable degree how the technology works, its limitations, and the applicable [T]erms of [U]se and other policies governing the use and exploitation of client data by the product."[13] A lawyer may wish to consider acquiring skills through a continuing legal education course. Consultation with IT professionals or cybersecurity experts may be appropriate as well. Generative AI outputs may be used as a starting point but must be carefully scrutinized. They should be critically analyzed for accuracy and bias, supplemented, and improved, if necessary. A lawyer must ensure that the input is correct and then critically review, validate, and correct the output of Generative AI "to ensure the content accurately reflects and supports the interests and priorities of the client in the matter at hand, including as part of advocacy for the client. The duty of competence requires more than the mere detection and elimination of false [Generative AI] outputs."[14] The use of Generative AI tools without the application of trained judgment by a lawyer is inconsistent with the competent and diligent practice of law. "A lawyer's professional judgment cannot be delegated to [G]enerative AI and remains the lawyer's responsibility at all times. A lawyer should take steps to avoid overreliance on Generative AI to such a degree that it hinders critical attorney analysis

---

[12] California Guidance at 2. There have been claims that certain Generative AI tools violate intellectual property rights of third parties. A lawyer planning to use a Generative AI tool should keep abreast of whether there are any such risks associated with the tool the lawyer plans to use.

[13] *Id.*

[14] *Id.* at 3.

	fostered by traditional research and writing. For example, a lawyer must supplement any Generative AI-generated research with human-performed research and supplement any Generative AI-generated argument with critical, human-performed analysis and review of authorities."[15]
Advertising and Solicitation Rule 7.1; Rule 7.3	Lawyers must not use Generative AI in a way that would circumvent their responsibilities under the Rules regarding marketing and solicitation. For example, a lawyer must not use Generative AI to make false statements, to search the internet for potential clients and send solicitations that would otherwise be prohibited under the Rules, or to pose as a real person to communicate with prospective clients.
Duty to Comply with the Law Rule 8.4; Rule 1.2(d)	"There are many relevant and applicable legal issues surrounding [G]enerative AI, including but not limited to compliance with AI-specific laws, privacy laws, cross-border data transfer laws, intellectual property laws, and cybersecurity concerns."[16] A lawyer must comply with the law and cannot counsel a client to engage in, or assist a client in conduct that the lawyer knows is, a violation of any law, rule, or ruling of a tribunal when using Generative AI tools.
Duty to Supervise Lawyers and Nonlawyers, Responsibilities of Subordinate Lawyers Rule 5.1; Rule 5.2; Rule 5.3; Rule 8.4	"Managerial and supervisory lawyers should establish clear policies regarding the permissible uses of [G]enerative AI and make reasonable efforts to ensure that the law firm adopts measures that give reasonable assurance that the law firm's lawyers and non-lawyers' conduct complies with their professional obligations when using [G]enerative AI. This includes providing training on the ethical and practical aspects, and pitfalls, of [G]enerative AI use. A subordinate lawyer must not use Generative AI at the direction of a supervisory lawyer in a manner that violates the subordinate lawyer's professional responsibility and obligations."[17] A subordinate lawyer should disclose to a supervisory lawyer the use of Generative AI that is not generally understood to be routinely used by lawyers.[18]

---

[15] *Id.*

[16] *Id.*

[17] *Id.*

[18] Likewise, where a client provides citations to a lawyer, a lawyer must review the decisions to make sure that they are genuine and properly cited. *See United States v. Cohen*, No. 18-CR-602, 2024 WL 1193604 (S.D.N.Y. Mar. 20,

	A lawyer using a Generative AI chatbot for client intake purposes must adequately supervise the chatbot.[19] A high degree of supervision may be required if there is a likelihood that ethical problems may arise. For example, a chatbot may fail to disclose that it is not a lawyer or may attempt or appear to provide legal advice, increasing the risk that a prospective client relationship or a lawyer–client relationship could be created.
Communication Regarding Generative AI Use Rule 1.4; Rule 1.2	"A lawyer should evaluate … communication obligations throughout the representation based on the facts and circumstances, including the novelty of the technology, risks associated with [G]enerative AI use, scope of the representation, and sophistication of the client."[20] A lawyer should consider disclosing to the client the intent to use Generative AI that is not generally understood to be routinely used by lawyers as part of the representation,[21] particularly as part of an explanation of the lawyer's fees and disbursements. The disclosure will depend on circumstances including how the technology will be used, and the benefits and risks of such use. A lawyer should obtain client consent for Generative AI use if client confidences will be disclosed in connection with the use of Generative AI. A lawyer should review any applicable client instructions or guidelines that may restrict or limit the use of Generative AI. We note that, because Generative AI currently is used routinely by lawyers, when a lawyer receives a request from a client that Generative AI not be used at all, the lawyer should consider discussing the request with the client before agreeing to it.

2024) (criticizing an attorney-defendant and his counsel for citing "three cases that do not exist" where client provided citations hallucinated by Google Bard and counsel failed to check them).

[19] *See* Fla. Bar Bd. Rev. Comm. on Pro. Ethics, *supra* (section on Oversight of Generative AI).

[20] California Guidance at 4.

[21] Note that some Generative AI is routinely used. For example, Microsoft Word employs Generative AI in its auto-complete and grammar check functions. Westlaw, Lexis, and search engines also employ Generative AI. We do not mean to suggest that an attorney needs to disclose such uses of Generative AI. For a discussion of the importance of evaluating Generative AI tools based on intended users, see N.J. STATE BAR ASS'N, TASK FORCE ON ARTIFICIAL INTELLIGENCE (AI) AND THE LAW: REPORT, REQUESTS, RECOMMENDATIONS, AND FINDINGS 15–19 (2024) (discussing "AI Tools Intended for the Public" and "Tools Tailored for Legal Professionals"), https://njsba.com/wp-content/uploads/2024/05/NJSBA-TASK-FORCE-ON-AI-AND-THE-LAW-REPORT-final.pdf.

Candor to the Tribunal; and Meritorious Claims and Contentions Rule 1.2(c); Rule 3.1; Rule 3.3; Rule 1.16	A lawyer should recognize the risks posed by Generative AI-generated content. Generative AI tools can, and do, fabricate or "hallucinate" precedent."[22] They can also create "deepfakes"—media that appear to reflect actual events but are actually doctored or manufactured. "A lawyer must review all [G]enerative AI outputs," including but not limited to "analysis and citations to authority," for accuracy before use for client purposes and submission to a court or other tribunal.[23] If the lawyer suspects that a client may have provided the lawyer with Generative AI-generated evidence, a lawyer may have a duty to inquire.[24] A lawyer must correct any errors or misleading statements made to adversaries, the public, or the court.[25] "A lawyer should also check for any rules, orders, or other requirements in the relevant jurisdiction that may necessitate the disclosure of the use of [G]enerative AI."[26]
Charging for Work Produced by Generative AI and Generative AI Costs Rule 1.5	"A lawyer may use [G]enerative AI to more efficiently create work product and may charge for actual time spent (*e.g.*, crafting or refining [G]enerative AI inputs and prompts, or reviewing and editing [G]enerative AI outputs)."[27] A lawyer must not charge hourly fees for the time that would otherwise have been spent absent the use of Generative AI.[28] Lawyers may wish to consider

---

[22] A Stanford University study found that Generative AI chatbots from OpenAI, Inc., Google LLC, and Meta Platforms Inc. hallucinate "at least 75% of the time when answering questions about a court's core ruling." Isabel Gottlieb & Isaiah Poritz, *Popular AI Chatbots Found to Give Error-Ridden Legal Answers*, Bloomberg L. (Jan. 12, 2024), https://news.bloomberglaw.com/business-and-practice/legal-errors-by-top-ai-models-alarmingly-prevalent-study-says. Courts are already grappling with parties' citation to hallucinated precedents. *See generally Mata v. Avianca, Inc.*, No. 22-CV-1461, 2023 WL 4114964 (S.D.N.Y. June 22, 2023) (sanctioning attorneys for "submit[ing] non-existent judicial opinions with fake quotes and citations created by the artificial intelligence tool ChatGPT"); *Cohen*, 2024 WL 1193604; *see also* D.C. Bar, Ethics Op. 388 (2024) (discussing the dangers of hallucinations).

[23] California Guidance at 4.

[24] *See* N.Y. City Op. 2018-4 (discussing a lawyer's duty to inquire when asked to assist in a transaction that the lawyer suspects may involve a crime or fraud); *see also* ABA Op. 491 (2020); Colo. Bar Ass'n Ethics Comm., Formal Op. 142 (2021). These same standards apply when a lawyer suspects that a client may have given the lawyer fabricated evidence.

[25] *See* Rule 3.3.

[26] California Guidance at 4.

[27] *Id.*

[28] *Id.*

	developing alternative fee arrangements relating to the value of their work rather than time spent. Costs associated with Generative AI should be disclosed in advance to clients as required by Rule 1.5(b). The costs charged should be consistent with ethical guidance on disbursements and should comply with applicable law.[29] A lawyer may wish to consider appropriate use of Generative AI tools to minimize client cost as the use of Generative AI becomes more widespread.
Prohibition on Discrimination Rule 8.4	"Some [G]enerative AI is trained on biased [historical] information, and a lawyer should be aware of possible biases and the risks they may create when using [G]enerative AI (*e.g.*, to screen potential clients or employees)."[30]

---

[29] *See* ABA Op. 93-379 (1993).

[30] California Guidance at 4.

FILED
CLERK
8/7/2025
U.S. DISTRICT COURT
EASTERN DISTRICT OF NEW YORK
LONG ISLAND OFFICE

UNITED STATES DISTRICT COURT
EASTERN DISTRICT OF NEW YORK
---X
SASHANE HALL,

Plaintiff,

-against-

THE ACADEMY CHARTER SCHOOL,

Defendant.
---X

MEMORANDUM
AND ORDER
2:24-cv-08630-JMW

APPEARANCES:

Suryia Rahman
Naresh M. Gehi
Gehi and Associates
173-29 Jamaica Ave.,
Jamaica, NY 11432
Attorneys for the Plaintiff

Adam Granek Guttell
Anahi Tapia
Jackson Lewis, P.C.
58 South Service Road
Suite 250
Melville, NY 11747
Attorneys for Defendant

WICKS, Magistrate Judge:

"[L]awyers are essential to the primary governmental function of administering justice, and have historically been 'officers of the courts.'" *Goldfarb v. Virginia State Bar*, 421 U.S. 773, 792 (1975). As such, an attorney appearing on behalf of a client is obligated to provide "competent representation," which means the attorney must possess the "legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation." N.Y. Rules of Prof. Conduct 1.1(a). That requires "an advocate to disclose directly adverse and controlling legal authority that is known to the lawyer . . . .". *Id.* at cmt. 4 to Rule 3.3. The reason of course is that "[a] tribunal that is fully informed on the applicable law is better able to make a fair and accurate determination of the matter before it." *Id.* That is not what occurred here.

The Court is confronted with what seemingly appears to an issue plaguing the modern legal justice system: the use of generative artificial intelligence ("AI") to assist in the drafting of legal papers that results in the generation of a legal submission containing citations to fictitious or non-existent legal authority – commonly referred to as "hallucinated" or fake case citations. In a brief filed in opposition to Defendant's pre-motion conference letter in support of a motion for partial dismissal (ECF No. 24), Plaintiff's counsel relied upon AI-generated content that created fictitious legal citations. That conduct clearly violated Federal Rule of Civil Procedure 11. *See Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443, 462 (S.D.N.Y. 2023) (Castel, J.).

The question presented is whether the circumstances here constitute sanctionable conduct under Rule 11 and the applicable standards for competence and candor before a tribunal.

BACKGROUND

Plaintiff Sashane Hall ("Plaintiff") commenced this action against Defendant The Academy Charter School ("Defendant") on December 18, 2024 asserting claims under Title VII of the Civil Rights Act of 1964, 42 U.S.C. §2000(e) *et seq.*, Title IX of the Education Amendments of 1972, 20 U.S.C. § 1681 *et seq.*, the New York State Human Rights Law §§ 296, *et seq.*, 42 U.S.C. § 1981(a), for discrimination based on sexual orientation, and hostile and abusive working environment. (*See generally* ECF Nos. 1, 22.) Plaintiff filed her Amended Complaint (ECF No. 22) following a pre-motion conference held on April 4, 2025. (ECF No.

20.) Thereafter, on May 30, 2025, Defendant filed a letter requesting a second pre-motion conference, this time on its anticipated motion to partially dismiss the Amended Complaint under Fed. R. 12(b)(6). (ECF No. 23.) Plaintiff filed her opposition letter in the form of a brief ("Opposition") on June 6, 2025, arguing, *inter alia*, that charter schools are not subject to notice of claim requirements because they are not government entities like public schools and school districts. (*See* ECF No. 24.)[1]

The cases Plaintiff cites in the Opposition to support this proposition, however, do not exist. (*See* ECF No. 25 at pp. 1-2.) Instead, Plaintiff includes three hallucinated cases: (1) *Laskowski v. Liberty Partners Restaurant Group Inc*., 2011 WL 817498, at *3 (S.D.N.Y. Mar. 9, 2011), (2) *Lindner v. Forest Hills Montessori School*, 2011 WL 1334869, at *4 (E.D.N.Y. Apr. 6, 2011), and (3) *Matter of K.M. v. Bronx Charter School for Better Learning*, 2011 NY Slip Op 32728(U) (Sup. Ct. Bronx County). (ECF No. 24 at p. 3; ECF No. 25 at p. 2.) As such, on June 16, 2025, Defendant requested an additional pre-motion conference before moving for an order striking Plaintiff's opposition entirely and for an award of attorneys' fees and costs stemming from Defendant's application. (ECF No. 25 at p. 1.) Defendant argues that "Plaintiff's citations to fabricated, misleading, and irrelevant legal authority [] cannot be explained away as an innocent 'typo,' mistake, or misunderstanding of the law," and such conduct "deprives both Defendant and the Court the opportunity to properly evaluate Plaintiff's argument on its merits." (*Id.* at p. 3.)

In response, the same day, counsel for Plaintiff[2] filed a short letter acknowledging receipt of Defendant's motion, stated that she understood the "grave implications of misreporting

---

[1] The June 6, 2025 Opposition was signed electronically by Naresh M. Gehi, Esq. (ECF No. 24.)

[2] The June 16, 2025 letter was signed by Suryia Rahman, Esq. (ECF No. 26.)

case law to the Court," and assured the Court that her firm was internally investigating the issue. (ECF No. 26.) Moreover, in this same letter, counsel for Plaintiff voluntarily discontinued the claims under New York State Human Rights Law—later memorialized in a stipulation filed on June 30, 2025. (ECF Nos. 26, 27.)

On July 2, 2025, the Court issued an Electronic Order "so ordering" the stipulation and further directing that "counsel for Plaintiff show cause, in writing on or before July 18, 2025 why the Court should not impose sanctions upon counsel for the conduct indicated in Defendant's pre-motion letter at ECF No. 25 for the fabricated legal authority contained in Plaintiff's opposition likely result[ing] from counsel's misuse of artificial intelligence. (ECF No. 28) (internal citations omitted) (alteration in original). Counsel for Plaintiff filed her response on July 18, 2025. (ECF No. 29.) In her response, counsel for Plaintiff explained that the Opposition was drafted by a clerk who used Google for research. (*Id.* at p. 1.) The clerk sent the work product to counsel for Plaintiff who "reviewed the draft Opposition but did not check the citations." (*Id.*) Counsel for Plaintiff attributed the "main reason" for her failure to check the citations to the death of her spouse, who recently passed away unexpectedly, stating that "[t]he shock and grief resulting from my husband's death has had a profound impact on all aspects of my life," and explaining that her husband's death "has affected [her] ability to attend to the practice of law with the same focus and attention as before." (*Id.* at pp. 1-2.) Counsel for Plaintiff has since taken bereavement leave and is "continuously meeting with medical and mental health professionals." (*Id.* at p. 1.)

DISCUSSION

Rule 11 governs attorneys' representations to the Court, and therefore, informs situations such as the present one. *See generally* Fed. R. Civ. P. 11. Rule 11 states:

> By presenting to the court a pleading, written motion, or other paper —whether by signing, filing, submitting, or later advocating it—an attorney . . . certifies that to the best of the person's knowledge, information, and belief, formed after an inquiry reasonable under the circumstances . . . the claims, defenses, and other legal contentions are warranted by existing law or by a nonfrivolous argument for extending, modifying or reversing existing law or for establishing new law.

Fed. R. Civ. P. 11(b)(2).

"Rule 11 imposes a duty on attorneys to certify that they have conducted a reasonable inquiry and have determined that any papers filed with the court are well grounded in fact, [and] legally tenable." *Benjamin v. Costco Wholesale Corp.*, No. 2:24-cv-7399 (LGD), 2025 WL 1195925, at *5 (E.D.N.Y. 2025) (quoting *Park v. Kim*, 91 F.4th 610, 614 (2d Cir. 2024)). "At the very least, the duties imposed by Rule 11 require that attorneys read, and thereby confirm the existence and validity of, the legal authorities on which they rely." *Park*, 91 F.4th at 615. Indeed, a claim is frivolous and warrants sanctions when a "legal position has 'no chance of success.'" *Fishoff v. Coty, Inc.*, 634 F.3d 647, 654 (2d Cir. 2011) (quoting *Morley v. Ciba–Geigy Corp.*, 66 F.3d 21, 25 (2d Cir. 1995)). When "a court considers whether to impose sanctions sua sponte, it 'is akin to the court's inherent power of contempt,' and, 'like contempt, sua sponte sanctions in those circumstances should issue only upon a finding of subjective bad faith.'" *Mata*, 678 F. Supp. 3d at 462 (alterations omitted) (quoting *Muhammad v. Walmart Stores East, L.P.*, 732 F.3d 104, 108 (2d Cir. 2013) (per curiam)). "Subjective bad faith is 'a heightened mens rea standard' that is intended to permit zealous advocacy while deterring improper submissions." *Mata*, 678 F. Supp. 3d at 462 (alterations omitted) (quoting *In re Pennie & Edmonds LLP*, 323 F.3d 86, 91 (2d Cir. 2003)).

Further, this Court possesses the inherent power to sanction a party "for conduct which abuses the judicial process." *Chambers v. NASCO, Inc.*, 501 U.S. 32, 44–45 (1991). A court's inherent power is "governed not by rule or statute but by the control necessarily vested in courts to manage their own affairs so as to achieve the orderly and expeditious disposition of cases." *Id.* at 43 (quoting *Link v. Wabash R. Co.*, 370 U.S. 626, 630–31 (1962)). Notably, a court possesses "'wide discretion' to craft an appropriate sanction, and may consider the effects on the parties and the full knowledge of the relevant facts" gained throughout sanctions hearings. *Heaston v. City of N.Y.*, 19-CV-5569 (PKC) (VMS), 2022 WL 182069, at *9 (E.D.N.Y. Jan. 20, 2022) (quoting *Oliveri v. Thompson*, 803 F.2d 1265, 1280 (2d Cir. 1986)). For instance, courts can exercise their inherent authority by "suspend[ing] or disbar[ing] lawyers," *In re Snyder*, 472 U.S. 634, 643 (1985), and in some circumstances impose "particularly severe sanction[s]" such as "outright dismissal" of a case. *Chambers*, 501 U.S. at 45.

Moreover, "any attorney or other person admitted to conduct cases in any court of the United States or any Territory thereof who so multiplies the proceedings in any case unreasonably and vexatiously may be required by the court to satisfy personally the excess costs, expenses, and attorneys' fees reasonably incurred because of such conduct." 28 U.S.C. § 1927. For a court to impose sanctions pursuant to § 1927, there must have been bad faith motive on the offending party to delay the litigation. *GiftRocket, Inc. v. Buchnik*, No. 24 MC 5105 (RPK) (VMS), 2025 WL 888483, at *2 (E.D.N.Y. Mar. 21, 2025). "Unlike Rule 11 sanctions which focus on particular papers, the inquiry under § 1927 is on a course of conduct." *Hernandez v. Money Source Inc.*, No. 17-CV-6919 (GRB) (AYS), 2022 WL 2702894, at *8 (E.D.N.Y. July 12, 2022) (quoting *Bowler v. U.S. Immigration & Naturalization Serv.*, 901 F. Supp. 597, 605 (S.D.N.Y. 1995); *Mahoney v. Yamaha Motor Corp. U.S.A.*, 290 F.R.D. 363, 367 (E.D.N.Y. 2013)).

As one court observed, "there is nothing inherently wrong with an attorney properly and competently utilizing AI or any of its subsets to practice law or litigate cases." *Versant Funding*

LLC v. Teras Breakbulk Ocean Navigation Enter., LLC, No. 17-cv-81140-DIMITROULEAS/MATTHEWMAN, 2025 WL 1440351, at *4 (S.D. Fla. May 20, 2025). Nevertheless, "[a]ttorneys and courts need to be aware of both the benefits and limitations that these AI platforms present." *Id.* (quoting Judge Xavier Rodriguez, *Artificial Intelligence (AI) and the Practice of Law*, 24 SEDONA CONF. J. 783, 791 (2023)). Indeed, "[a] basic prerequisite to the filing of any pleading, motion, response, reply, or paper in court is for the drafting and filing attorney(s) to carefully check every case citation, fact, and argument to make sure that they are correct and proper." *Id.* This obligation existed long before AI entered the legal industry landscape.

The appearance of hallucinated citations in briefs generated from AI is no longer in its nascent stage. Regrettably, the number and regularity with which courts have been faced with hallucinations in court filings continues to rise both in this country and abroad. *See* Damien Charlotin, *AI Hallucination Cases*, (Aug. 6, 2025) https://www.damiencharlotin.com/hallucinations/ (database tracking legal decisions "in cases where generative AI produced hallucinated content," evidencing 255 cases to date) (hereinafter "Charlotin Database"). This trend of AI hallucinations has steadily increased over the years. *See* Cecily Mauran, *120 court cases have been caught with AI hallucinations, according to new database,* MASHABLE (May 27, 2025), *https://mashable.com/article/over-120-court-cases-caught-ai-hallucinations-new-database*.

Courts addressing filings containing AI-generated hallucinations have fashioned various forms of sanctions imposed upon the offending attorneys after they were found to have used these fabricated citations in "bad faith." For instance, in *Park v. Kim*, an attorney who submitted a reply brief with hallucinated cases was referred to the Second Circuit's Grievance Panel with consideration of referral to the Committee on Admissions and ordered to provide a copy of the Court's ruling to their client. 91 F.4th at 616. The Second Circuit found that the act of presenting non-existent cases to the court demonstrated that the attorney made no inquiry into the validity of the arguments she presented and failed to determine her argument was "legally tenable" under Rule 11. *Id.* at 615. Similarly, in *Benjamin*, the court held that the offending attorney acted in bad faith by utilizing "AI to produce work she could not (or would not) do on her own," yet failed to read the case law prior to submitting an affirmation to court. 2025 WL 1195925, at *8. Accordingly, because the facts "suggest[ed] much more than mere carelessness," the court imposed a $1,000 fine and ordered the offending attorney to serve a copy of the court's order on their client. *Id.* at *9.

Likewise, in *Mata*, one of the earliest cases addressing this issue, the court determined that the offending attorneys acted in bad faith and thus ordered them to, *inter alia*, provide a copy of the court's order to their clients and mail a letter and a copy of the court's order to each of the judges that the offending attorneys falsely attributed hallucinated cases to. 678 F. Supp. 3d at 466. The "bad faith" conduct in *Mata* involved the offending attorneys citing to non-existent cases generated by ChatGPT in motion papers and for "not reading a single case cited in [the submitted affirmation] and taking no other steps on his own to check whether any aspect of the assertions of law were warranted by existing law." *Id.* at 456, 464. Indeed, in *Mata*, the court noted that taking these affirmative steps, signing and filing the affirmation, and relying on another attorney's lack of familiarity with the relevant law, a fact known by the offending attorney, the conduct was "an act of subjective bad faith." *Id.* at 464. Further, the other offending attorney falsely asserted that his use of ChatGPT was "merely a 'supplement' to his research," whereas the record reflected that the attorney solely relied on the AI database for his research. *See id.* at 465. Accordingly, Rule 11 sanctions were appropriate. *Id.* at 466.

Indeed, the imposition of sanctions resulting from an attorney's bad faith use of generative AI runs nationwide. *See, e.g., Coomer v. Lindell*, No. 22-cv-01129-NYW-SBP, 2025 WL 1865282, at *6, *8 (D. Colo. July 25, 2025) (imposing a $3,000 fine on the offending attorneys and their law firm after submitting a brief with hallucinated and inaccurate case outcomes because the attorneys "were not reasonable in certifying that the claims, defenses, and other legal contentions contained [the filing] were warranted by existing law or by a nonfrivolous argument for extending, modifying, or reversing existing law or for establishing new law") (citing Fed. R. Civ. P. 11(b)(2)); *Versant*, 2025 WL 1440351, at *7 (ordering the offending attorney to pay the attorney's fees and costs incurred from the AI generated motion and subsequent replies, attend a CLE course on Artificial Intelligence, and pay a $1000 fine); *Wadsworth v. Walmart Inc.*, 348 F.R.D. 489, 494, 497–98 (D. Wyo. 2025) (revoking one offending attorney's pro hac vice status because as the drafter he maintained oversight to ensure he filed a meritorious motion on behalf of his client, and imposed monetary fines on certain attorneys for failing to adhere to their Rule 11 obligations, after the attorneys were found to have used the firm's in-house AI platform to find case law and neglected to check the cases generated); *Mid Cent. Operating Eng'r Health and Welfare Fund v. Hoosiervac LLC*, No. 2:24-cv-00326-JPH-MJD, 2025 WL 1511211, at *1, *2 (S.D. Ind. May 28, 2025) (imposing a $6,000 fine on the offending attorney for submitting three separate briefs that all contained hallucinated case law in order to "deter repetition of the conduct or comparable conduct by others similarly situated") (citing Fed. R. Civ. P. 11(c)(4)).

AI-generated caselaw also seems to have unwittingly worked its way into judicial decisions and orders. *See* Justin Henry, *Judge Scraps Opinion after Lawyer Flags Made-Up Quotes (Correct)*, BLOOMBERG LAW, (July 23, 2025), https://news.bloomberglaw.com/business-and-practice/judge-withdraws-pharma-opinion-after-lawyer-flags-made-up-quotes; Debra Cassens Weiss, *After second federal judge withdraws error-riddled ruling, litigants seek explanation*, ABA JOURNAL, (July 30, 2025), https://www.abajournal.com/news/article/litigants-seek-explanation-after-second-federal-judge-withdraws-error-riddled ruling#:~:text=Mississippi%20Attorney%20General%20Lynn%20Fitch%20has%20asked%20a %20federal%20judge,ruling%20back%20on%20the%20docket. (noting the court opinion incorrectly "referenced allegations and parties not in the lawsuit, nonexistent declarations by four people, and language not found in the state law being challenged").

By far, the majority of courts impose sanctions upon the offending lawyer for this sort of conduct and warnings or reprimands have been meted out in cases typically involving pro se litigants. *See* Charlotin Database, *supra*. However, there are circumstances where, in the Court's discretion, monetary sanctions have not been imposed notwithstanding the violation of Rule 11. For example, in *U.S. v. Cohen*, the Court declined to impose sanctions upon an attorney who submitted three non-existent cases to the tribunal contained in a motion seeking early termination of supervised release. 724 F. Supp. 3d 251, 253, 258 (S.D.N.Y. 2024). There, the attorney who cited fake cases reviewed a draft that was written by his client—another lawyer—and reviewed by another attorney. *Id.* at 258–59. Although "embarrassing and certainly negligent, perhaps even grossly negligent" and evidencing "extreme carelessness," the court declined to impose sanctions as no showing was made that the attorney's citations to non-existent cases were generated or submitted in bad faith. *Id.*

Similar to counsel in *U.S. v. Cohen*, the admitted conduct here was the result of extreme carelessness and negligence, done under tragic personal circumstances. The explanation proffered to the Court that counsel failed to check the citations generated by her clerk who used Google for research does not evince bad faith or willfulness, but rather severe carelessness. (*See* ECF No. 29 at p. 1.) Unlike in *Mata*, counsel here had no reason to believe that the clerk who drafted the Opposition was unfamiliar with notice of claim requirements under New York state law. Further, counsel for Plaintiff's reason for neglecting to check this work was the sudden, unexpected death of her spouse. (*Id.*) Indeed, Plaintiff explained that "[t]he shock and grief resulting from my husband's death has had a profound impact on all aspects of my life," and explained that her husband's death "has affected [her] ability to attend to the practice of law with the same focus and attention as before." (*Id.* at pp. 1–2.) As such, counsel for Plaintiff has not only repeatedly apologized to the Court and opposing counsel, assuring both that "[t]his will not happen again," she also indicated her intentions on seeking medical and mental treatment. (*See id.* at pp. 1-2.)[3] Notably, contrary to the circumstances in *Mata* where the offending attorneys misrepresented their use of ChatGPT to the court, counsel for Plaintiff outright admitted and took "full responsibility for failing to check the citations in the Opposition." (*Id.* at p. 1.) Counsel also stipulated to the relief sought on the anticipated motion by ultimately withdrawing certain claims.[4]

While the Court is aware of the serious implications that the misuse of AI-generated non-existent caselaw presents, it is also mindful of the circumstances here which do not support any

---

[3] There have been no prior disciplinary actions, grievances, or sanctions imposed upon counsel for Plaintiff. This conduct, while aberrant, appears to be an isolated occurrence.

[4] The Court also recognizes and commends both counsels' professionalism by conferring and resolving the issue by way of stipulation.

finding of bad faith, as well as the remorseful explanations proffered by counsel who experienced one of life's unspeakable tragedies. It is for these reasons that the Court in the exercise of its discretion, declines to impose sua sponte monetary sanctions upon counsel for Plaintiff.

CONCLUSION

For the reasons stated herein, the Court declines to exercise its inherent power to impose monetary sanctions upon counsel for Plaintiff. Counsel is, however, admonished and this Order should serve as a forewarning. Counsel is further directed to serve Plaintiff with a copy of this Memorandum and Order.

Dated: Central Islip, New York
August 7, 2025

SO ORDERED:

/s/ James M. Wicks
JAMES M. WICKS
United States Magistrate Judge

Effective October 2025

NEW YORK STATE UNIFIED COURT SYSTEM
INTERIM POLICY ON THE USE OF ARTIFICIAL INTELLIGENCE

I. Purpose

This interim policy on the use of artificial intelligence (AI) is designed to promote the responsible and ethical use of AI technology in the New York State Unified Court System (UCS). This document outlines important guardrails to ensure fairness, accountability, and security in the use of AI, particularly generative AI, by our workforce. Mandatory requirements and restrictions governing the use of AI are set forth below, in Section V. This interim policy is intended to evolve with technological advancements, operational necessities, and future iterations of relevant legislation, regulation, and public policy.

II. Scope

This interim policy is applicable to all judges and nonjudicial employees of the UCS. It applies to all functions performed on a UCS-owned device, and to all UCS-related work performed on any device.

III. Understanding AI

A. How Generative AI Works

The term "AI" means "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments" (15 USC § 9401[3]). The term "generative AI" refers to an AI program or system that is capable of generating human-like text or other content in response to user prompts by learning from material in large reference datasets.

Generative AI tools have the potential to enhance productivity by assisting with tasks such as producing first drafts of documents, editing text, summarizing data, drafting correspondence, and developing software. Such tools may take the form of a chatbot, a computer program which simulates an online human conversation.

Most generative AI programs utilize a large language model (LLM), an algorithm that is trained on an enormous quantity of data, derived from various sources such as the internet, books, and articles, and learns to produce written communications by continually predicting the word that is most likely to come next. Generative AI programs do not operate like traditional search engines. Although they draw upon information contained in large datasets, they are designed to generate content, not to locate information or provide authoritative answers to factual inquiries.

Rather than conducting traditional research or verifying facts, the AI program produces a document or other output, in the format and style requested in the user's prompts, by predicting patterns based on its source information. As a result, any factual assertions or citations to legal authority included in the output may be inaccurate or unreliable. In fact, generative AI programs occasionally fill in gaps in their source material by simply fabricating facts or citations. In AI terminology, such an insertion of fictitious information is referred to as a hallucination.

B. Potential Uses of AI

Generative AI tools can be used to help draft documents such as policy memos, letters, speeches, or job descriptions. AI can be useful in generating ideas, getting a document started, and suggesting suitable wording. Generative AI can also be utilized in communicating with the public. AI platforms can help users write clearly and in plain, accessible language. A user can upload content, such as text for a proposed webpage or a draft of a policy statement, and direct the AI program to modify the language to make it simpler, more concise, and easier to understand. The AI program can be prompted to write for a specific audience or at a designated reading level. Like all content produced by generative AI, the output should be carefully reviewed, and the user should ensure that the language is inclusive, respectful, and accurate.

Generative AI can also be used to summarize lengthy documents or large datasets in preparation for administrative reports or analytical legal writing. Since one of AI's most impressive features is its ability to scan and process vast amounts of data in just a few minutes, or even seconds, summarizing information may be among its most valuable uses. An AI tool can quickly generate an overview of material such as a large document, a group of documents, or a set of statistical data, providing the user with a basic understanding, or at least the highlights, of the material, where the user's own review of the material would have consumed an enormous amount of time. The AI tool can be prompted to produce the summary in a variety of formats, such as a single paragraph, a brief memo with a specified word count, an outline containing a specified degree of detail, or a bullet-point list. The AI tool could also be prompted to condense the contents of a document for clarity or brevity. However, if such a condensed version is to be submitted to and relied upon by other UCS personnel, or released to the public, the contents of the AI-generated product must be checked against the original material to ensure accuracy. Moreover, the use of an AI tool to summarize legal documents is subject to the guidance and limitations relating to confidentiality, set forth below.

C. Problems Associated with AI

Despite its potential benefits, generative AI can produce inaccurate, wholly fabricated, or biased outputs, and can jeopardize the security of data entered into the program.

1. Inaccurate or Fabricated Information

As noted above, the output produced by generative AI tools will sometimes contain hallucinations. Accordingly, the content generated by an AI program should not be used without careful editing. It is the responsibility of every user to thoroughly review such content and to independently confirm that it contains no fabricated or fictitious material.

In view of their limitations, generative AI tools should not be relied upon to provide accurate information or to draft communications about sensitive topics. Moreover, general-purpose AI programs (whether operating on a public model or on a private model) are not suitable for legal writing and legal research, as they may produce incorrect or fabricated citations and analysis. Even when using the AI-enhanced features that have been incorporated into established legal research platforms, any content generated by AI should be independently verified for accuracy.

2. Bias and Other Inappropriate Output

The vast datasets on which generative AI systems are trained include material that reflects cultural, economic, and social biases and expressions of prejudice against protected classes of people. As a result, the content generated may promote stereotypes, reinforce prejudices, exhibit unfair biases, or contain otherwise undesirable, offensive, or harmful material. Accordingly, it is the responsibility of every user to thoroughly review any AI-generated content, to ensure that it does not reflect any unfair bias, stereotypes, or prejudice or contain any other inappropriate material, and to make any necessary revisions.

3. Vulnerability of Confidential Information

Many publicly available generative AI platforms (ChatGPT, for example) operate on an open training model, which means, among other things, that the input received from user prompts is collected and used as further training material for their LLMs. Since the LLM can reproduce that material for anyone using an AI program connected to it, that input is potentially accessible by the public at large. Accordingly, once a UCS user inputs information into such a platform as part of a prompt or in an uploaded document, that information is no longer under UCS control, and may become publicly available.

In contrast to AI platforms that operate on these public models, which can be accessed by anyone and may store data for use in future training, some AI platforms operate on a private model. Platforms using private models are hosted or managed by an organization, and their use is typically restricted to members of that organization or individuals who have been granted access. They may be tailored to the organization's specific needs, and they include additional security, compliance, and privacy measures.

Furthermore, users should be careful to avoid uploading copyrighted content into a generative AI program.

IV. Guiding Principles

AI is a type of tool designed to assist the user in performing certain tasks. It must not be treated as a substitute for human judgment, discretion, or decision-making. All UCS users remain accountable for their final work product.

It is critical to ensure that material that reflects harmful bias, stereotypes, or prejudice does not appear in any UCS work product.

The rules governing the security and confidentiality of court records apply fully to the use of AI technology. It should be assumed that all information entered into a public model generative AI platform, such as ChatGPT, will immediately become public. AI technology must be used in a manner that prevents the public disclosure of information of a confidential, private, or sensitive nature.

- Examples of such information include, but are not limited to, docket numbers, party names, addresses, and dates of birth.

- Documents that have been filed or submitted for filing in any court are also considered confidential, even if they are classified as public at the time of filing, since it is possible that the record of the case will be sealed in the future, or that the documents have not been adequately redacted to conceal sensitive information. Although, in these scenarios, the confidential information has already been revealed to the public, entering the information into the public model AI program makes the exposure of the information permanent.

- Intellectual property of the UCS is another type of information that should not be publicly disclosed. An example of such disclosure is internally written source code being entered into a public model AI system by software developers working either within or outside of the UCS.

AI technology must be used in a manner that is consistent with the ethical obligations of judges and nonjudicial employees. The Rules Governing Judicial Conduct (22 NYCRR Part 100), the Rules Governing Conduct of Nonjudicial Court Employees (22 NYCRR Part 50), and the Rules of Professional Conduct remain fully applicable when AI tools are being used.

- For example, judges bear the ultimate responsibility for the content of their opinions, orders, and other written materials, and may not delegate their judicial decision-making responsibilities to any other person or entity. *See* 22 NYCRR 100.2(A) (a judge must "act at all times in a manner that promotes public confidence in the integrity and impartiality of the judiciary"); 22 NYCRR 100.3(B)(7) (a judge must "dispose of all judicial matters promptly, efficiently and fairly"). Thus, while AI tools can be used to assist with a judge's work, judges and court staff must ensure that such tools are never actually engaged in the decision-making tasks a judge is ethically obligated to perform.

- Nonjudicial employees must avoid using AI in any way that violates their own ethical responsibilities, such as the duty not to manifest bias or prejudice on the basis of any protected status, and the duty not to disclose any confidential information received in the course of their official duties. *See* 22 NYCRR 50.1(II)(C), (D).

- Any questions about potential ethical concerns arising from particular uses of AI technology by judicial officers should be directed to the Advisory Committee on Judicial Ethics.

V. Requirements and Restrictions

1. UCS users may use only those generative AI products that have been approved by the UCS Division of Technology and Court Research (DoTCR), which are identified in the attached Appendix.

2. All judges and nonjudicial UCS employees with computer access shall be required to complete an initial training course, as well as continuing training, in the use of AI technology. No generative AI product may be used on any UCS-owned device or for any UCS-related work until the user has completed the initial training course.

3. No user may input into any generative AI program that does not operate on a private model — by writing a prompt, uploading a document or file, or otherwise — any information that is confidential, private, or privileged, or includes personally identifiable information or protected health information, or is otherwise inappropriate for public release. A private model is a model that is under UCS control and does not share data with any public LLM.

4. No user may upload into any generative AI program that does not operate on a private model any document that has been filed or submitted for filing in any court, even if the document is classified as public.

5. Any user who uses a generative AI program to produce a document or any other content must thoroughly review the content produced by the program and make necessary revisions to ensure that it is accurate and appropriate, and does not reflect any unfair bias, stereotypes, or prejudice.

6. No user may install on a UCS-owned device any software that is required for the use of a generative AI program, or use a UCS-owned device to access any such program that requires payment, a subscription, or agreement to terms of use, unless access to that program has been provided to the user by the UCS.

7. AI tools may not be used on a UCS-owned device for personal purposes unrelated to UCS work.

8. The approval of a generative AI product by the DoTCR signifies that the product is safe to use from a technological standpoint, but does not necessarily mean that, for a particular task, the use of that product is suitable or appropriate. Such approval by the DoTCR does not preclude any judge or UCS supervisor from prohibiting the use of such a product for a particular task by a person under their supervision.

APPENDIX

Approved Generative Artificial Intelligence Products for New York State Unified Court System Users

Effective October 2025

Please Note: It is important that you check the following list of approved generative artificial intelligence (AI) products on a regular basis. New AI tools are released daily, and AI components are regularly added to existing products. Moreover, some AI tools that currently appear on this list may become unavailable at a later date. Therefore, the contents of this list will change and grow over time.

Approved Private Enterprise Generative AI Tools

Procured and Managed by the Unified Court System

PRODUCT	DESCRIPTION	HOSTING*	AVAILABILITY
Microsoft Azure AI Services	Azure AI services are a suite of cloud-based AI services and tools offered by Microsoft Azure. These services allow developers and data scientists to build, deploy, and manage AI solutions within the Azure cloud platform. They provide pre-built APIs and models for various AI capabilities, including speech, vision, language, and decision-making.	New York State Unified Court System Azure Government Tenant	Currently available Costs are based upon utilization and subject to DoTCR approval
Microsoft 365 CoPilot Chat	Microsoft 365 Copilot Chat is a fully private, AI-powered chat feature within Microsoft 365 designed to boost user productivity. This free, secure generative AI chat is powered by GPT-4o.	New York State Unified Court System O365 Tenant	Currently available
Microsoft 365 CoPilot	Includes a chat interface that is grounded on users' meetings, emails, chats, and documents hosted in the Azure cloud environment. It integrates into Microsoft 365 applications including Outlook, Teams, Word, Excel, PowerPoint, OneDrive, and more.	New York State Unified Court System O365 Tenant	Currently unavailable Requires paid license per account

GitHub CoPilot for Business or Enterprise	Assists developers with real-time code suggestions, auto-completion, and code explanation within development environments. Trained on a wide range of code, it helps developers write code more efficiently, saving time and reducing errors.	New York State Unified Court System O365 Tenant	Currently available to developers and data scientists Requires paid license per account, subject to DoTCR approval
Trados Studio	Trados Studio is a computer-assisted translation software tool which provides a comprehensive platform for translation tasks, including editing, reviewing, and project management.	AWS Cloud	OCA Office of Language Access

**The tools listed above are only to be used within the indicated hosting environments*

Approved Public Generative AI Tools

PRODUCT	DESCRIPTION	HOSTING	AVAILABILITY
OpenAI ChatGPT Free Version**	ChatGPT is a conversational AI chatbot developed by OpenAI. It uses large language models, like GPT-4o, to generate human-like text in response to user prompts. Essentially, it is a program that can hold conversations, answer questions, and perform various tasks like drafting text, summarizing information, and more.	Publicly available	Currently available

***Paid subscriptions are prohibited*

Check for updates

NIST AI 100-1

Artificial Intelligence Risk Management Framework (AI RMF 1.0)

NIST
NATIONAL INSTITUTE OF
STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST AI 100-1

Artificial Intelligence Risk Management Framework (AI RMF 1.0)

This publication is available free of charge from:
https://doi.org/10.6028/NIST.AI.100-1

January 2023

U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

This publication is available free of charge from: https://doi.org/10.6028/NIST.AI.100-1

Update Schedule and Versions

The Artificial Intelligence Risk Management Framework (AI RMF) is intended to be a living document.

NIST will review the content and usefulness of the Framework regularly to determine if an update is appropriate; a review with formal input from the AI community is expected to take place no later than 2028. The Framework will employ a two-number versioning system to track and identify major and minor changes. The first number will represent the generation of the AI RMF and its companion documents (e.g., 1.0) and will change only with major revisions. Minor revisions will be tracked using ".n" after the generation number (e.g., 1.1). All changes will be tracked using a Version Control Table which identifies the history, including version number, date of change, and description of change. NIST plans to update the AI RMF Playbook frequently. Comments on the AI RMF Playbook may be sent via email to AIframework@nist.gov at any time and will be reviewed and integrated on a semi-annual basis.

Table of Contents

Executive Summary **1**

Part 1: Foundational Information **4**

1 Framing Risk **4**
- 1.1 Understanding and Addressing Risks, Impacts, and Harms 4
- 1.2 Challenges for AI Risk Management 5
  - 1.2.1 Risk Measurement 5
  - 1.2.2 Risk Tolerance 7
  - 1.2.3 Risk Prioritization 7
  - 1.2.4 Organizational Integration and Management of Risk 8

2 Audience **9**

3 AI Risks and Trustworthiness **12**
- 3.1 Valid and Reliable 13
- 3.2 Safe 14
- 3.3 Secure and Resilient 15
- 3.4 Accountable and Transparent 15
- 3.5 Explainable and Interpretable 16
- 3.6 Privacy-Enhanced 17
- 3.7 Fair – with Harmful Bias Managed 17

4 Effectiveness of the AI RMF **19**

Part 2: Core and Profiles **20**

5 AI RMF Core **20**
- 5.1 Govern 21
- 5.2 Map 24
- 5.3 Measure 28
- 5.4 Manage 31

6 AI RMF Profiles **33**

Appendix A: Descriptions of AI Actor Tasks from Figures 2 and 3 **35**

Appendix B: How AI Risks Differ from Traditional Software Risks **38**

Appendix C: AI Risk Management and Human-AI Interaction **40**

Appendix D: Attributes of the AI RMF **42**

List of Tables

Table 1 Categories and subcategories for the GOVERN function. 22
Table 2 Categories and subcategories for the MAP function. 26
Table 3 Categories and subcategories for the MEASURE function. 29
Table 4 Categories and subcategories for the MANAGE function. 32

List of Figures

Fig. 1 Examples of potential harms related to AI systems. Trustworthy AI systems and their responsible use can mitigate negative risks and contribute to benefits for people, organizations, and ecosystems. 5
Fig. 2 Lifecycle and Key Dimensions of an AI System. Modified from OECD (2022) OECD Framework for the Classification of AI systems — OECD Digital Economy Papers. The two inner circles show AI systems' key dimensions and the outer circle shows AI lifecycle stages. Ideally, risk management efforts start with the Plan and Design function in the application context and are performed throughout the AI system lifecycle. See Figure 3 for representative AI actors. 10
Fig. 3 AI actors across AI lifecycle stages. See Appendix A for detailed descriptions of AI actor tasks, including details about testing, evaluation, verification, and validation tasks. Note that AI actors in the AI Model dimension (Figure 2) are separated as a best practice, with those building and using the models separated from those verifying and validating the models. 11
Fig. 4 Characteristics of trustworthy AI systems. Valid & Reliable is a necessary condition of trustworthiness and is shown as the base for other trustworthiness characteristics. Accountable & Transparent is shown as a vertical box because it relates to all other characteristics. 12
Fig. 5 Functions organize AI risk management activities at their highest level to govern, map, measure, and manage AI risks. Governance is designed to be a cross-cutting function to inform and be infused throughout the other three functions. 20

Executive Summary

Artificial intelligence (AI) technologies have significant potential to transform society and people's lives – from commerce and health to transportation and cybersecurity to the environment and our planet. AI technologies can drive inclusive economic growth and support scientific advancements that improve the conditions of our world. AI technologies, however, also pose risks that can negatively impact individuals, groups, organizations, communities, society, the environment, and the planet. Like risks for other types of technology, AI risks can emerge in a variety of ways and can be characterized as long- or short-term, high- or low-probability, systemic or localized, and high- or low-impact.

The AI RMF refers to an *AI system* as an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy (Adapted from: OECD Recommendation on AI:2019; ISO/IEC 22989:2022).

While there are myriad standards and best practices to help organizations mitigate the risks of traditional software or information-based systems, the risks posed by AI systems are in many ways unique (See Appendix B). AI systems, for example, may be trained on data that can change over time, sometimes significantly and unexpectedly, affecting system functionality and trustworthiness in ways that are hard to understand. AI systems and the contexts in which they are deployed are frequently complex, making it difficult to detect and respond to failures when they occur. AI systems are inherently socio-technical in nature, meaning they are influenced by societal dynamics and human behavior. AI risks – and benefits – can emerge from the interplay of technical aspects combined with societal factors related to how a system is used, its interactions with other AI systems, who operates it, and the social context in which it is deployed.

These risks make AI a uniquely challenging technology to deploy and utilize both for organizations and within society. Without proper controls, AI systems can amplify, perpetuate, or exacerbate inequitable or undesirable outcomes for individuals and communities. With proper controls, AI systems can mitigate and manage inequitable outcomes.

AI risk management is a key component of responsible development and use of AI systems. Responsible AI practices can help align the decisions about AI system design, development, and uses with intended aim and values. Core concepts in responsible AI emphasize human centricity, social responsibility, and sustainability. AI risk management can drive responsible uses and practices by prompting organizations and their internal teams who design, develop, and deploy AI to think more critically about context and potential or unexpected negative and positive impacts. Understanding and managing the risks of AI systems will help to enhance trustworthiness, and in turn, cultivate public trust.

Social responsibility can refer to the organization's responsibility "for the impacts of its decisions and activities on society and the environment through transparent and ethical behavior" (ISO 26000:2010). *Sustainability* refers to the "state of the global system, including environmental, social, and economic aspects, in which the needs of the present are met without compromising the ability of future generations to meet their own needs" (ISO/IEC TR 24368:2022). Responsible AI is meant to result in technology that is also equitable and accountable. The expectation is that organizational practices are carried out in accord with "*professional responsibility*," defined by ISO as an approach that "aims to ensure that professionals who design, develop, or deploy AI systems and applications or AI-based products or systems, recognize their unique position to exert influence on people, society, and the future of AI" (ISO/IEC TR 24368:2022).

As directed by the National Artificial Intelligence Initiative Act of 2020 (P.L. 116-283), the goal of the AI RMF is to offer a resource to the organizations designing, developing, deploying, or using AI systems to help manage the many risks of AI and promote trustworthy and responsible development and use of AI systems. The Framework is intended to be ***voluntary***, rights-preserving, non-sector-specific, and use-case agnostic, providing flexibility to organizations of all sizes and in all sectors and throughout society to implement the approaches in the Framework.

The Framework is designed to equip organizations and individuals – referred to here as *AI actors* – with approaches that increase the trustworthiness of AI systems, and to help foster the responsible design, development, deployment, and use of AI systems over time. AI actors are defined by the Organisation for Economic Co-operation and Development (OECD) as "those who play an active role in the AI system lifecycle, including organizations and individuals that deploy or operate AI" [OECD (2019) Artificial Intelligence in Society—OECD iLibrary] (See Appendix A).

The AI RMF is intended to be practical, to adapt to the AI landscape as AI technologies continue to develop, and to be operationalized by organizations in varying degrees and capacities so society can benefit from AI while also being protected from its potential harms.

The Framework and supporting resources will be updated, expanded, and improved based on evolving technology, the standards landscape around the world, and AI community experience and feedback. NIST will continue to align the AI RMF and related guidance with applicable international standards, guidelines, and practices. As the AI RMF is put into use, additional lessons will be learned to inform future updates and additional resources.

The Framework is divided into two parts. Part 1 discusses how organizations can frame the risks related to AI and describes the intended audience. Next, AI risks and trustworthiness are analyzed, outlining the characteristics of trustworthy AI systems, which include valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy enhanced, and fair with their harmful biases managed.

Part 2 comprises the "Core" of the Framework. It describes four specific functions to help organizations address the risks of AI systems in practice. These functions – **GOVERN**, **MAP**, **MEASURE**, and **MANAGE** – are broken down further into categories and subcategories. While **GOVERN** applies to all stages of organizations' AI risk management processes and procedures, the **MAP**, **MEASURE**, and **MANAGE** functions can be applied in AI system-specific contexts and at specific stages of the AI lifecycle.

Additional resources related to the Framework are included in the AI RMF Playbook, which is available via the NIST AI RMF website:
https://www.nist.gov/itl/ai-risk-management-framework.

Development of the AI RMF by NIST in collaboration with the private and public sectors is directed and consistent with its broader AI efforts called for by the National AI Initiative Act of 2020, the National Security Commission on Artificial Intelligence recommendations, and the Plan for Federal Engagement in Developing Technical Standards and Related Tools. Engagement with the AI community during this Framework's development – via responses to a formal Request for Information, three widely attended workshops, public comments on a concept paper and two drafts of the Framework, discussions at multiple public forums, and many small group meetings – has informed development of the AI RMF 1.0 as well as AI research and development and evaluation conducted by NIST and others. Priority research and additional guidance that will enhance this Framework will be captured in an associated AI Risk Management Framework Roadmap to which NIST and the broader community can contribute.

Part 1: Foundational Information

1. Framing Risk

AI risk management offers a path to minimize potential negative impacts of AI systems, such as threats to civil liberties and rights, while also providing opportunities to maximize positive impacts. Addressing, documenting, and managing AI risks and potential negative impacts effectively can lead to more trustworthy AI systems.

1.1 Understanding and Addressing Risks, Impacts, and Harms

In the context of the AI RMF, *risk* refers to the composite measure of an event's probability of occurring and the magnitude or degree of the consequences of the corresponding event. The impacts, or consequences, of AI systems can be positive, negative, or both and can result in opportunities or threats (Adapted from: ISO 31000:2018). When considering the negative impact of a potential event, risk is a function of 1) the negative impact, or magnitude of harm, that would arise if the circumstance or event occurs and 2) the likelihood of occurrence (Adapted from: OMB Circular A-130:2016). Negative impact or harm can be experienced by individuals, groups, communities, organizations, society, the environment, and the planet.

> "Risk management refers to coordinated activities to direct and control an organization with regard to risk" (Source: ISO 31000:2018).

While risk management processes generally address negative impacts, this Framework offers approaches to minimize anticipated negative impacts of AI systems *and* identify opportunities to maximize positive impacts. Effectively managing the risk of potential harms could lead to more trustworthy AI systems and unleash potential benefits to people (individuals, communities, and society), organizations, and systems/ecosystems. Risk management can enable AI developers and users to understand impacts and account for the inherent limitations and uncertainties in their models and systems, which in turn can improve overall system performance and trustworthiness and the likelihood that AI technologies will be used in ways that are beneficial.

The AI RMF is designed to address new risks as they emerge. This flexibility is particularly important where impacts are not easily foreseeable and applications are evolving. While some AI risks and benefits are well-known, it can be challenging to assess negative impacts and the degree of harms. Figure 1 provides examples of potential harms that can be related to AI systems.

AI risk management efforts should consider that humans may assume that AI systems work – and work well – in *all* settings. For example, whether correct or not, AI systems are often perceived as being more objective than humans or as offering greater capabilities than general software.

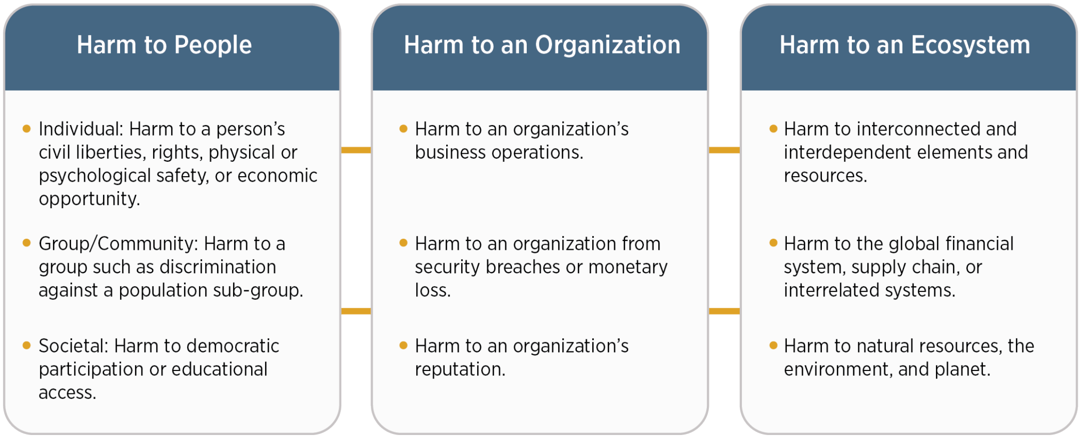
Harm to People
Individual: Harm to a person's civil liberties, rights, physical or psychological safety, or economic opportunity.
Group/Community: Harm to a group such as discrimination against a population sub-group.
Societal: Harm to democratic participation or educational access.
Harm to an Organization
Harm to an organization's business operations.
Harm to an organization from security breaches or monetary loss.
Harm to an organization's reputation.
Harm to an Ecosystem
Harm to interconnected and interdependent elements and resources.
Harm to the global financial system, supply chain, or interrelated systems.
Harm to natural resources, the environment, and planet.

Fig. 1. Examples of potential harms related to AI systems. Trustworthy AI systems and their responsible use can mitigate negative risks and contribute to benefits for people, organizations, and ecosystems.

1.2 Challenges for AI Risk Management

Several challenges are described below. They should be taken into account when managing risks in pursuit of AI trustworthiness.

1.2.1 Risk Measurement

AI risks or failures that are not well-defined or adequately understood are difficult to measure quantitatively or qualitatively. The inability to appropriately measure AI risks does not imply that an AI system necessarily poses either a high or low risk. Some risk measurement challenges include:

Risks related to third-party software, hardware, and data: Third-party data or systems can accelerate research and development and facilitate technology transition. They also may complicate risk measurement. Risk can emerge both from third-party data, software or hardware itself and how it is used. Risk metrics or methodologies used by the organization developing the AI system may not align with the risk metrics or methodologies uses by the organization *deploying or operating* the system. Also, the organization developing the AI system may not be transparent about the risk metrics or methodologies it used. Risk measurement and management can be complicated by how customers use or integrate third-party data or systems into AI products or services, particularly without sufficient internal governance structures and technical safeguards. Regardless, all parties and AI actors should manage risk in the AI systems they develop, deploy, or use as standalone or integrated components.

Tracking emergent risks: Organizations' risk management efforts will be enhanced by identifying and tracking emergent risks and considering techniques for measuring them.

AI system impact assessment approaches can help AI actors understand potential impacts or harms within specific contexts.

Availability of reliable metrics: The current lack of consensus on robust and verifiable measurement methods for risk and trustworthiness, and applicability to different AI use cases, is an AI risk measurement challenge. Potential pitfalls when seeking to measure negative risk or harms include the reality that development of metrics is often an institutional endeavor and may inadvertently reflect factors unrelated to the underlying impact. In addition, measurement approaches can be oversimplified, gamed, lack critical nuance, become relied upon in unexpected ways, or fail to account for differences in affected groups and contexts.

Approaches for measuring impacts on a population work best if they recognize that contexts matter, that harms may affect varied groups or sub-groups differently, and that communities or other sub-groups who may be harmed are not always direct users of a system.

Risk at different stages of the AI lifecycle: Measuring risk at an earlier stage in the AI lifecycle may yield different results than measuring risk at a later stage; some risks may be latent at a given point in time and may increase as AI systems adapt and evolve. Furthermore, different AI actors across the AI lifecycle can have different risk perspectives. For example, an AI developer who makes AI software available, such as pre-trained models, can have a different risk perspective than an AI actor who is responsible for deploying that pre-trained model in a specific use case. Such deployers may not recognize that their particular uses could entail risks which differ from those perceived by the initial developer. All involved AI actors share responsibilities for designing, developing, and deploying a trustworthy AI system that is fit for purpose.

Risk in real-world settings: While measuring AI risks in a laboratory or a controlled environment may yield important insights pre-deployment, these measurements may differ from risks that emerge in operational, real-world settings.

Inscrutability: Inscrutable AI systems can complicate risk measurement. Inscrutability can be a result of the opaque nature of AI systems (limited explainability or interpretability), lack of transparency or documentation in AI system development or deployment, or inherent uncertainties in AI systems.

Human baseline: Risk management of AI systems that are intended to augment or replace human activity, for example decision making, requires some form of baseline metrics for comparison. This is difficult to systematize since AI systems carry out different tasks – and perform tasks differently – than humans.

1.2.2 Risk Tolerance

While the AI RMF can be used to prioritize risk, it does not prescribe risk tolerance. *Risk tolerance* refers to the organization's or AI actor's (see Appendix A) readiness to bear the risk in order to achieve its objectives. Risk tolerance can be influenced by legal or regulatory requirements (Adapted from: ISO GUIDE 73). Risk tolerance and the level of risk that is acceptable to organizations or society are highly contextual and application and use-case specific. Risk tolerances can be influenced by policies and norms established by AI system owners, organizations, industries, communities, or policy makers. Risk tolerances are likely to change over time as AI systems, policies, and norms evolve. Different organizations may have varied risk tolerances due to their particular organizational priorities and resource considerations.

Emerging knowledge and methods to better inform harm/cost-benefit tradeoffs will continue to be developed and debated by businesses, governments, academia, and civil society. To the extent that challenges for specifying AI risk tolerances remain unresolved, there may be contexts where a risk management framework is not yet readily applicable for mitigating negative AI risks.

The Framework is intended to be flexible and to augment existing risk practices which should align with applicable laws, regulations, and norms. Organizations should follow existing regulations and guidelines for risk criteria, tolerance, and response established by organizational, domain, discipline, sector, or professional requirements. Some sectors or industries may have established definitions of harm or established documentation, reporting, and disclosure requirements. Within sectors, risk management may depend on existing guidelines for specific applications and use case settings. Where established guidelines do not exist, organizations should define reasonable risk tolerance. Once tolerance is defined, this AI RMF can be used to manage risks and to document risk management processes.

1.2.3 Risk Prioritization

Attempting to eliminate negative risk entirely can be counterproductive in practice because not all incidents and failures can be eliminated. Unrealistic expectations about risk may lead organizations to allocate resources in a manner that makes risk triage inefficient or impractical or wastes scarce resources. A risk management culture can help organizations recognize that not all AI risks are the same, and resources can be allocated purposefully. Actionable risk management efforts lay out clear guidelines for assessing trustworthiness of each AI system an organization develops or deploys. Policies and resources should be prioritized based on the assessed risk level and potential impact of an AI system. The extent to which an AI system may be customized or tailored to the specific context of use by the AI deployer can be a contributing factor.

When applying the AI RMF, risks which the organization determines to be highest for the AI systems within a given context of use call for the most urgent prioritization and most thorough risk management process. In cases where an AI system presents unacceptable negative risk levels – such as where significant negative impacts are imminent, severe harms are actually occurring, or catastrophic risks are present – development and deployment should cease in a safe manner until risks can be sufficiently managed. If an AI system's development, deployment, and use cases are found to be low-risk in a specific context, that may suggest potentially lower prioritization.

Risk prioritization may differ between AI systems that are designed or deployed to directly interact with humans as compared to AI systems that are not. Higher initial prioritization may be called for in settings where the AI system is trained on large datasets comprised of sensitive or protected data such as personally identifiable information, or where the outputs of the AI systems have direct or indirect impact on humans. AI systems designed to interact only with computational systems and trained on non-sensitive datasets (for example, data collected from the physical environment) may call for lower initial prioritization. Nonetheless, regularly assessing and prioritizing risk based on context remains important because non-human-facing AI systems can have downstream safety or social implications.

Residual risk – defined as risk remaining after risk treatment (Source: ISO GUIDE 73) – directly impacts end users or affected individuals and communities. Documenting residual risks will call for the system provider to fully consider the risks of deploying the AI product and will inform end users about potential negative impacts of interacting with the system.

1.2.4 Organizational Integration and Management of Risk

AI risks should not be considered in isolation. Different AI actors have different responsibilities and awareness depending on their roles in the lifecycle. For example, organizations developing an AI system often will not have information about how the system may be used. AI risk management should be integrated and incorporated into broader enterprise risk management strategies and processes. Treating AI risks along with other critical risks, such as cybersecurity and privacy, will yield a more integrated outcome and organizational efficiencies.

The AI RMF may be utilized along with related guidance and frameworks for managing AI system risks or broader enterprise risks. Some risks related to AI systems are common across other types of software development and deployment. Examples of overlapping risks include: privacy concerns related to the use of underlying data to train AI systems; the energy and environmental implications associated with resource-heavy computing demands; security concerns related to the confidentiality, integrity, and availability of the system and its training and output data; and general security of the underlying software and hardware for AI systems.

Organizations need to establish and maintain the appropriate accountability mechanisms, roles and responsibilities, culture, and incentive structures for risk management to be effective. Use of the AI RMF alone will not lead to these changes or provide the appropriate incentives. Effective risk management is realized through organizational commitment at senior levels and may require cultural change within an organization or industry. In addition, small to medium-sized organizations managing AI risks or implementing the AI RMF may face different challenges than large organizations, depending on their capabilities and resources.

2. Audience

Identifying and managing AI risks and potential impacts – both positive and negative – requires a broad set of perspectives and actors across the AI lifecycle. Ideally, AI actors will represent a diversity of experience, expertise, and backgrounds and comprise demographically and disciplinarily diverse teams. The AI RMF is intended to be used by AI actors across the AI lifecycle and dimensions.

The OECD has developed a framework for classifying AI lifecycle activities according to five key socio-technical dimensions, each with properties relevant for AI policy and governance, including risk management [OECD (2022) OECD Framework for the Classification of AI systems — OECD Digital Economy Papers]. Figure 2 shows these dimensions, slightly modified by NIST for purposes of this framework. The NIST modification highlights the importance of test, evaluation, verification, and validation (TEVV) processes throughout an AI lifecycle and generalizes the operational context of an AI system.

AI dimensions displayed in Figure 2 are the Application Context, Data and Input, AI Model, and Task and Output. AI actors involved in these dimensions who perform or manage the design, development, deployment, evaluation, and use of AI systems and drive AI risk management efforts are the *primary* AI RMF audience.

Representative AI actors across the lifecycle dimensions are listed in Figure 3 and described in detail in Appendix A. Within the AI RMF, all AI actors work together to manage risks and achieve the goals of trustworthy and responsible AI. AI actors with TEVV-specific expertise are integrated throughout the AI lifecycle and are especially likely to benefit from the Framework. Performed regularly, TEVV tasks can provide insights relative to technical, societal, legal, and ethical standards or norms, and can assist with anticipating impacts and assessing and tracking emergent risks. As a regular process within an AI lifecycle, TEVV allows for both mid-course remediation and post-hoc risk management.

The People & Planet dimension at the center of Figure 2 represents human rights and the broader well-being of society and the planet. The AI actors in this dimension comprise a separate AI RMF audience who *informs* the primary audience. These AI actors may include trade associations, standards developing organizations, researchers, advocacy groups,

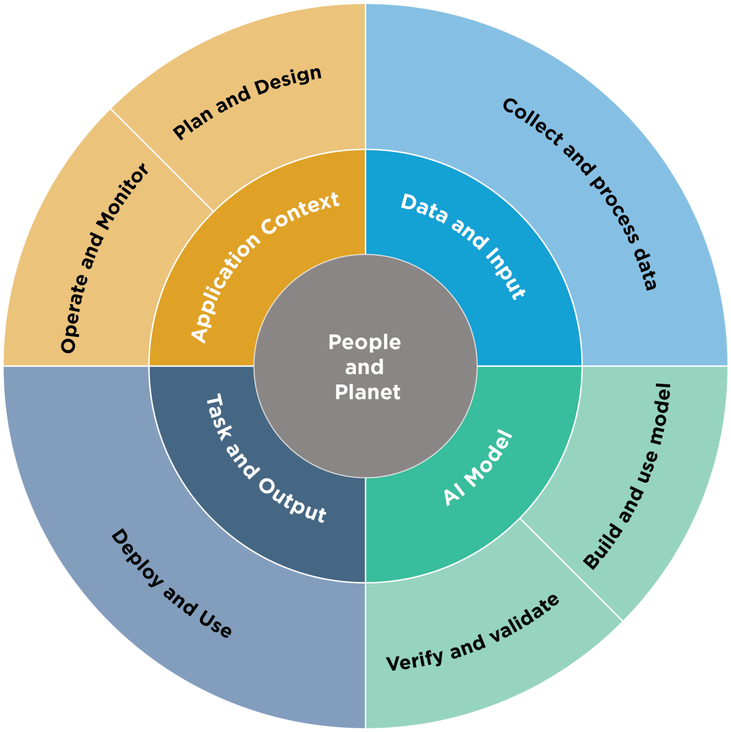

Fig. 2. Lifecycle and Key Dimensions of an AI System. Modified from OECD (2022) OECD Framework for the Classification of AI systems — OECD Digital Economy Papers. The two inner circles show AI systems' key dimensions and the outer circle shows AI lifecycle stages. Ideally, risk management efforts start with the Plan and Design function in the application context and are performed throughout the AI system lifecycle. See Figure 3 for representative AI actors.

environmental groups, civil society organizations, end users, and potentially impacted individuals and communities. These actors can:

- assist in providing context and understanding potential and actual impacts;
- be a source of formal or quasi-formal norms and guidance for AI risk management;
- designate boundaries for AI operation (technical, societal, legal, and ethical); and
- promote discussion of the tradeoffs needed to balance societal values and priorities related to civil liberties and rights, equity, the environment and the planet, and the economy.

Successful risk management depends upon a sense of collective responsibility among AI actors shown in Figure 3. The AI RMF functions, described in Section 5, require diverse perspectives, disciplines, professions, and experiences. Diverse teams contribute to more open sharing of ideas and assumptions about the purposes and functions of technology – making these implicit aspects more explicit. This broader collective perspective creates opportunities for surfacing problems and identifying existing and emergent risks.

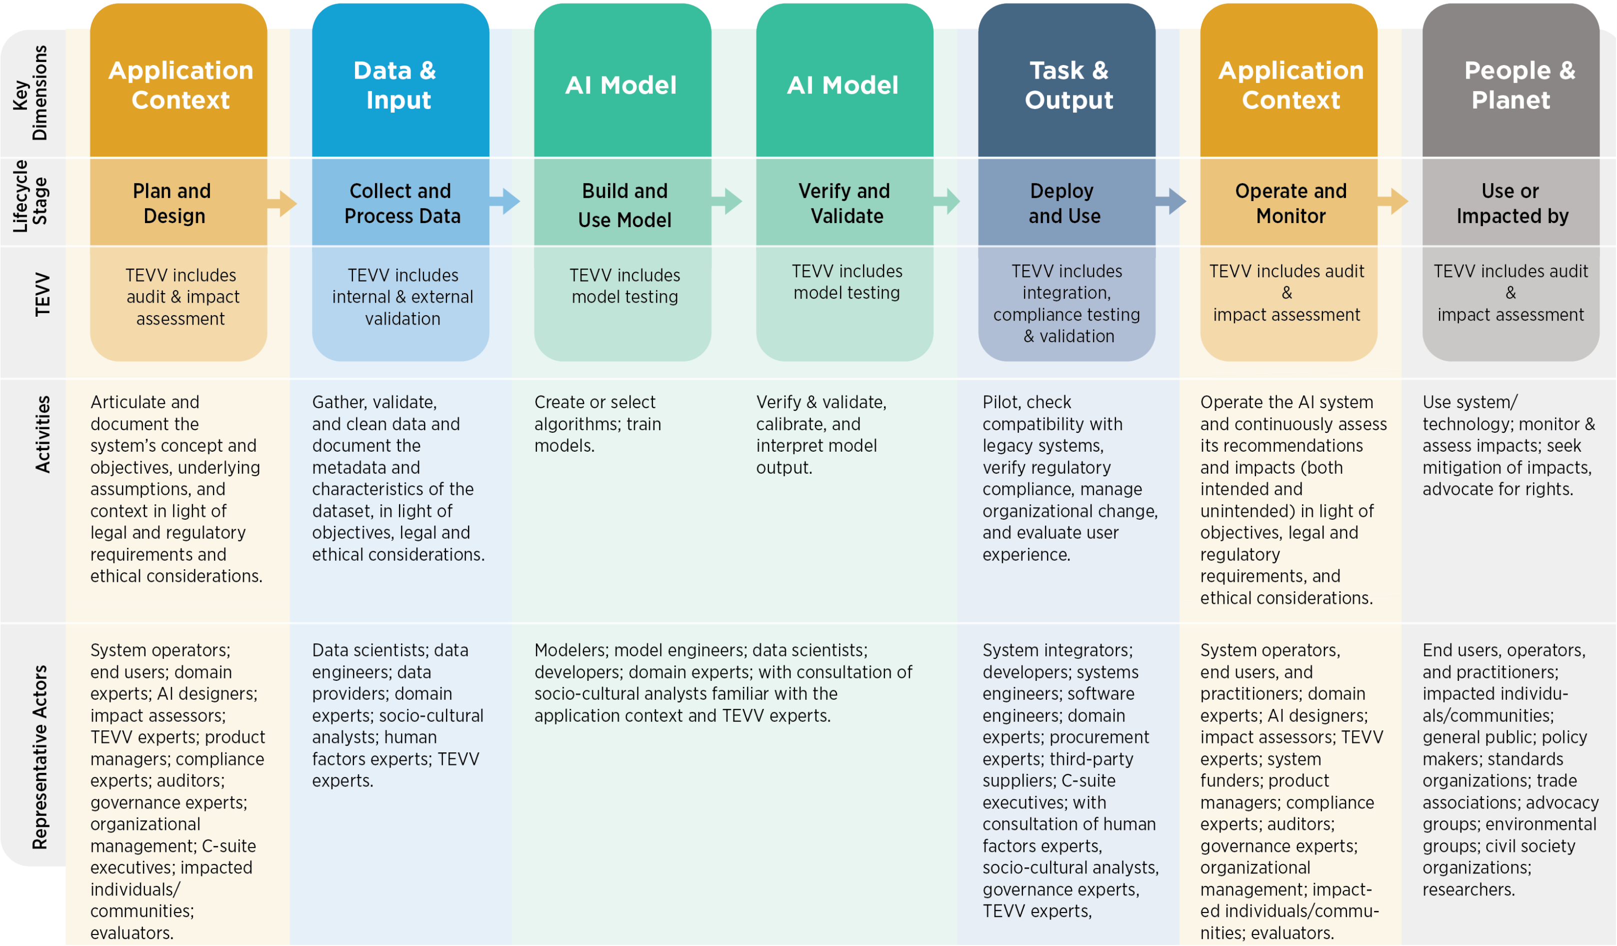

Key Dimensions	Application Context	Data & Input	AI Model	AI Model	Task & Output	Application Context	People & Planet
Lifecycle Stage	Plan and Design	Collect and Process Data	Build and Use Model	Verify and Validate	Deploy and Use	Operate and Monitor	Use or Impacted by
TEVV	TEVV includes audit & impact assessment	TEVV includes internal & external validation	TEVV includes model testing	TEVV includes model testing	TEVV includes integration, compliance testing & validation	TEVV includes audit & impact assessment	TEVV includes audit & impact assessment
Activities	Articulate and document the system's concept and objectives, underlying assumptions, and context in light of legal and regulatory requirements and ethical considerations.	Gather, validate, and clean data and document the metadata and characteristics of the dataset, in light of objectives, legal and ethical considerations.	Create or select algorithms; train models.	Verify & validate, calibrate, and interpret model output.	Pilot, check compatibility with legacy systems, verify regulatory compliance, manage organizational change, and evaluate user experience.	Operate the AI system and continuously assess its recommendations and impacts (both intended and unintended) in light of objectives, legal and regulatory requirements, and ethical considerations.	Use system/technology; monitor & assess impacts; seek mitigation of impacts, advocate for rights.
Representative Actors	System operators; end users; domain experts; AI designers; impact assessors; TEVV experts; product managers; compliance experts; auditors; governance experts; organizational management; C-suite executives; impacted individuals/communities; evaluators.	Data scientists; data engineers; data providers; domain experts; socio-cultural analysts; human factors experts; TEVV experts.	Modelers; model engineers; data scientists; developers; domain experts; with consultation of socio-cultural analysts familiar with the application context and TEVV experts.		System integrators; developers; systems engineers; software engineers; domain experts; procurement experts; third-party suppliers; C-suite executives; with consultation of human factors experts, socio-cultural analysts, governance experts, TEVV experts,	System operators, end users, and practitioners; domain experts; AI designers; impact assessors; TEVV experts; system funders; product managers; compliance experts; auditors; governance experts; organizational management; impacted individuals/communities; evaluators.	End users, operators, and practitioners; impacted individuals/communities; general public; policy makers; standards organizations; trade associations; advocacy groups; environmental groups; civil society organizations; researchers.

Fig. 3. AI actors across AI lifecycle stages. See Appendix A for detailed descriptions of AI actor tasks, including details about testing, evaluation, verification, and validation tasks. Note that AI actors in the AI Model dimension (Figure 2) are separated as a best practice, with those building and using the models separated from those verifying and validating the models.

3. AI Risks and Trustworthiness

For AI systems to be trustworthy, they often need to be responsive to a multiplicity of criteria that are of value to interested parties. Approaches which enhance AI trustworthiness can reduce negative AI risks. This Framework articulates the following **characteristics** of trustworthy AI and offers guidance for addressing them. Characteristics of trustworthy AI systems include: **valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed.** Creating trustworthy AI requires balancing each of these characteristics based on the AI system's context of use. While all characteristics are socio-technical system attributes, accountability and transparency also relate to the processes and activities internal to an AI system and its external setting. Neglecting these characteristics can increase the probability and magnitude of negative consequences.

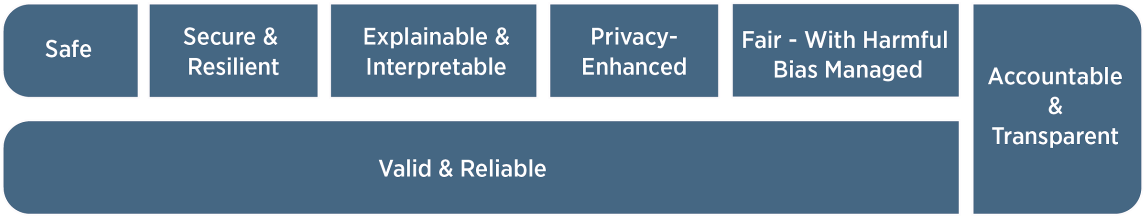
Safe
Secure & Resilient
Explainable & Interpretable
Privacy-Enhanced
Fair - With Harmful Bias Managed
Accountable & Transparent
Valid & Reliable

Fig. 4. Characteristics of trustworthy AI systems. Valid & Reliable is a necessary condition of trustworthiness and is shown as the base for other trustworthiness characteristics. Accountable & Transparent is shown as a vertical box because it relates to all other characteristics.

Trustworthiness characteristics (shown in Figure 4) are inextricably tied to social and organizational behavior, the datasets used by AI systems, selection of AI models and algorithms and the decisions made by those who build them, and the interactions with the humans who provide insight from and oversight of such systems. Human judgment should be employed when deciding on the specific metrics related to AI trustworthiness characteristics and the precise threshold values for those metrics.

Addressing AI trustworthiness characteristics individually will not ensure AI system trustworthiness; tradeoffs are usually involved, rarely do all characteristics apply in every setting, and some will be more or less important in any given situation. Ultimately, trustworthiness is a social concept that ranges across a spectrum and is only as strong as its weakest characteristics.

When managing AI risks, organizations can face difficult decisions in balancing these characteristics. For example, in certain scenarios tradeoffs may emerge between optimizing for interpretability and achieving privacy. In other cases, organizations might face a tradeoff between predictive accuracy and interpretability. Or, under certain conditions such as data sparsity, privacy-enhancing techniques can result in a loss in accuracy, affecting decisions about fairness and other values in certain domains. Dealing with tradeoffs requires taking into account the decision-making context. These analyses can highlight the existence and extent of tradeoffs between different measures, but they do not answer questions about how to navigate the tradeoff. Those depend on the values at play in the relevant *context* and should be resolved in a manner that is both transparent and appropriately justifiable.

There are multiple approaches for enhancing contextual awareness in the AI lifecycle. For example, subject matter experts can assist in the evaluation of TEVV findings and work with product and deployment teams to align TEVV parameters to requirements and deployment conditions. When properly resourced, increasing the breadth and diversity of input from interested parties and relevant AI actors throughout the AI lifecycle can enhance opportunities for informing contextually sensitive evaluations, and for identifying AI system benefits and positive impacts. These practices can increase the likelihood that risks arising in social contexts are managed appropriately.

Understanding and treatment of trustworthiness characteristics depends on an AI actor's particular role within the AI lifecycle. For any given AI system, an AI designer or developer may have a different perception of the characteristics than the deployer.

Trustworthiness characteristics explained in this document influence each other. Highly secure but unfair systems, accurate but opaque and uninterpretable systems, and inaccurate but secure, privacy-enhanced, and transparent systems are all undesirable. A comprehensive approach to risk management calls for balancing tradeoffs among the trustworthiness characteristics. It is the joint responsibility of all AI actors to determine whether AI technology is an appropriate or necessary tool for a given context or purpose, and how to use it responsibly. The decision to commission or deploy an AI system should be based on a contextual assessment of trustworthiness characteristics and the relative risks, impacts, costs, and benefits, and informed by a broad set of interested parties.

3.1 Valid and Reliable

Validation is the "confirmation, through the provision of objective evidence, that the requirements for a specific intended use or application have been fulfilled" (Source: ISO 9000:2015). Deployment of AI systems which are inaccurate, unreliable, or poorly generalized to data and settings beyond their training creates and increases negative AI risks and reduces trustworthiness.

Reliability is defined in the same standard as the "ability of an item to perform as required, without failure, for a given time interval, under given conditions" (Source: ISO/IEC TS 5723:2022). Reliability is a goal for overall correctness of AI system operation under the conditions of expected use and over a given period of time, including the entire lifetime of the system.

Accuracy and robustness contribute to the validity and trustworthiness of AI systems, and can be in tension with one another in AI systems.

Accuracy is defined by ISO/IEC TS 5723:2022 as "closeness of results of observations, computations, or estimates to the true values or the values accepted as being true." Measures of accuracy should consider computational-centric measures (e.g., false positive and false negative rates), human-AI teaming, and demonstrate external validity (generalizable beyond the training conditions). Accuracy measurements should always be paired with clearly defined and realistic test sets – that are representative of conditions of expected use – and details about test methodology; these should be included in associated documentation. Accuracy measurements may include disaggregation of results for different data segments.

Robustness or *generalizability* is defined as the "ability of a system to maintain its level of performance under a variety of circumstances" (Source: ISO/IEC TS 5723:2022). Robustness is a goal for appropriate system functionality in a broad set of conditions and circumstances, including uses of AI systems not initially anticipated. Robustness requires not only that the system perform exactly as it does under expected uses, but also that it should perform in ways that minimize potential harms to people if it is operating in an unexpected setting.

Validity and reliability for deployed AI systems are often assessed by ongoing testing or monitoring that confirms a system is performing as intended. Measurement of validity, accuracy, robustness, and reliability contribute to trustworthiness and should take into consideration that certain types of failures can cause greater harm. AI risk management efforts should prioritize the minimization of potential negative impacts, and may need to include human intervention in cases where the AI system cannot detect or correct errors.

3.2 Safe

AI systems should "not under defined conditions, lead to a state in which human life, health, property, or the environment is endangered" (Source: ISO/IEC TS 5723:2022). Safe operation of AI systems is improved through:

- responsible design, development, and deployment practices;
- clear information to deployers on responsible use of the system;
- responsible decision-making by deployers and end users; and
- explanations and documentation of risks based on empirical evidence of incidents.

Different types of safety risks may require tailored AI risk management approaches based on context and the severity of potential risks presented. Safety risks that pose a potential risk of serious injury or death call for the most urgent prioritization and most thorough risk management process.

Employing safety considerations during the lifecycle and starting as early as possible with planning and design can prevent failures or conditions that can render a system dangerous. Other practical approaches for AI safety often relate to rigorous simulation and in-domain testing, real-time monitoring, and the ability to shut down, modify, or have human intervention into systems that deviate from intended or expected functionality.

AI safety risk management approaches should take cues from efforts and guidelines for safety in fields such as transportation and healthcare, and align with existing sector- or application-specific guidelines or standards.

3.3 Secure and Resilient

AI systems, as well as the ecosystems in which they are deployed, may be said to be *resilient* if they can withstand unexpected adverse events or unexpected changes in their environment or use – or if they can maintain their functions and structure in the face of internal and external change and degrade safely and gracefully when this is necessary (Adapted from: ISO/IEC TS 5723:2022). Common security concerns relate to adversarial examples, data poisoning, and the exfiltration of models, training data, or other intellectual property through AI system endpoints. AI systems that can maintain confidentiality, integrity, and availability through protection mechanisms that prevent unauthorized access and use may be said to be *secure*. Guidelines in the NIST Cybersecurity Framework and Risk Management Framework are among those which are applicable here.

Security and resilience are related but distinct characteristics. While resilience is the ability to return to normal function after an unexpected adverse event, security includes resilience but also encompasses protocols to avoid, protect against, respond to, or recover from attacks. Resilience relates to robustness and goes beyond the provenance of the data to encompass unexpected or adversarial use (or abuse or misuse) of the model or data.

3.4 Accountable and Transparent

Trustworthy AI depends upon accountability. Accountability presupposes transparency. *Transparency* reflects the extent to which information about an AI system and its outputs is available to individuals interacting with such a system – regardless of whether they are even aware that they are doing so. Meaningful transparency provides access to appropriate levels of information based on the stage of the AI lifecycle and tailored to the role or knowledge of AI actors or individuals interacting with or using the AI system. By promoting higher levels of understanding, transparency increases confidence in the AI system.

This characteristic's scope spans from design decisions and training data to model training, the structure of the model, its intended use cases, and how and when deployment, post-deployment, or end user decisions were made and by whom. Transparency is often necessary for actionable redress related to AI system outputs that are incorrect or otherwise lead to negative impacts. Transparency should consider human-AI interaction: for example, how a human operator or user is notified when a potential or actual adverse outcome caused by an AI system is detected. A transparent system is not necessarily an accurate, privacy-enhanced, secure, or fair system. However, it is difficult to determine whether an opaque system possesses such characteristics, and to do so over time as complex systems evolve.

The role of AI actors should be considered when seeking accountability for the outcomes of AI systems. The relationship between risk and accountability associated with AI and technological systems more broadly differs across cultural, legal, sectoral, and societal contexts. When consequences are severe, such as when life and liberty are at stake, AI developers and deployers should consider proportionally and proactively adjusting their transparency and accountability practices. Maintaining organizational practices and governing structures for harm reduction, like risk management, can help lead to more accountable systems.

Measures to enhance transparency and accountability should also consider the impact of these efforts on the implementing entity, including the level of necessary resources and the need to safeguard proprietary information.

Maintaining the provenance of training data and supporting attribution of the AI system's decisions to subsets of training data can assist with both transparency and accountability. Training data may also be subject to copyright and should follow applicable intellectual property rights laws.

As transparency tools for AI systems and related documentation continue to evolve, developers of AI systems are encouraged to test different types of transparency tools in cooperation with AI deployers to ensure that AI systems are used as intended.

3.5 Explainable and Interpretable

Explainability refers to a representation of the mechanisms underlying AI systems' operation, whereas *interpretability* refers to the meaning of AI systems' output in the context of their designed functional purposes. Together, explainability and interpretability assist those operating or overseeing an AI system, as well as users of an AI system, to gain deeper insights into the functionality and trustworthiness of the system, including its outputs. The underlying assumption is that perceptions of negative risk stem from a lack of ability to make sense of, or contextualize, system output appropriately. Explainable and interpretable AI systems offer information that will help end users understand the purposes and potential impact of an AI system.

Risk from lack of explainability may be managed by describing how AI systems function, with descriptions tailored to individual differences such as the user's role, knowledge, and skill level. Explainable systems can be debugged and monitored more easily, and they lend themselves to more thorough documentation, audit, and governance.

Risks to interpretability often can be addressed by communicating a description of why an AI system made a particular prediction or recommendation. (See "Four Principles of Explainable Artificial Intelligence" and "Psychological Foundations of Explainability and Interpretability in Artificial Intelligence" found here.)

Transparency, explainability, and interpretability are distinct characteristics that support each other. Transparency can answer the question of "what happened" in the system. Explainability can answer the question of "how" a decision was made in the system. Interpretability can answer the question of "why" a decision was made by the system and its meaning or context to the user.

3.6 Privacy-Enhanced

Privacy refers generally to the norms and practices that help to safeguard human autonomy, identity, and dignity. These norms and practices typically address freedom from intrusion, limiting observation, or individuals' agency to consent to disclosure or control of facets of their identities (e.g., body, data, reputation). (See The NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management.)

Privacy values such as anonymity, confidentiality, and control generally should guide choices for AI system design, development, and deployment. Privacy-related risks may influence security, bias, and transparency and come with tradeoffs with these other characteristics. Like safety and security, specific technical features of an AI system may promote or reduce privacy. AI systems can also present new risks to privacy by allowing inference to identify individuals or previously private information about individuals.

Privacy-enhancing technologies ("PETs") for AI, as well as data minimizing methods such as de-identification and aggregation for certain model outputs, can support design for privacy-enhanced AI systems. Under certain conditions such as data sparsity, privacy-enhancing techniques can result in a loss in accuracy, affecting decisions about fairness and other values in certain domains.

3.7 Fair – with Harmful Bias Managed

Fairness in AI includes concerns for equality and equity by addressing issues such as harmful bias and discrimination. Standards of fairness can be complex and difficult to define because perceptions of fairness differ among cultures and may shift depending on application. Organizations' risk management efforts will be enhanced by recognizing and considering these differences. Systems in which harmful biases are mitigated are not necessarily fair. For example, systems in which predictions are somewhat balanced across demographic groups may still be inaccessible to individuals with disabilities or affected by the digital divide or may exacerbate existing disparities or systemic biases.

Bias is broader than demographic balance and data representativeness. NIST has identified three major categories of AI bias to be considered and managed: systemic, computational and statistical, and human-cognitive. Each of these can occur in the absence of prejudice, partiality, or discriminatory intent. Systemic bias can be present in AI datasets, the organizational norms, practices, and processes across the AI lifecycle, and the broader society that uses AI systems. Computational and statistical biases can be present in AI datasets and algorithmic processes, and often stem from systematic errors due to non-representative samples. Human-cognitive biases relate to how an individual or group perceives AI system information to make a decision or fill in missing information, or how humans think about purposes and functions of an AI system. Human-cognitive biases are omnipresent in decision-making processes across the AI lifecycle and system use, including the design, implementation, operation, and maintenance of AI.

Bias exists in many forms and can become ingrained in the automated systems that help make decisions about our lives. While bias is not always a negative phenomenon, AI systems can potentially increase the speed and scale of biases and perpetuate and amplify harms to individuals, groups, communities, organizations, and society. Bias is tightly associated with the concepts of transparency as well as fairness in society. (For more information about bias, including the three categories, see NIST Special Publication 1270, Towards a Standard for Identifying and Managing Bias in Artificial Intelligence.)

4. Effectiveness of the AI RMF

Evaluations of AI RMF effectiveness – including ways to measure bottom-line improvements in the trustworthiness of AI systems – will be part of future NIST activities, in conjunction with the AI community.

Organizations and other users of the Framework are encouraged to periodically evaluate whether the AI RMF has improved their ability to manage AI risks, including but not limited to their policies, processes, practices, implementation plans, indicators, measurements, and expected outcomes. NIST intends to work collaboratively with others to develop metrics, methodologies, and goals for evaluating the AI RMF's effectiveness, and to broadly share results and supporting information. Framework users are expected to benefit from:

- enhanced processes for governing, mapping, measuring, and managing AI risk, and clearly documenting outcomes;
- improved awareness of the relationships and tradeoffs among trustworthiness characteristics, socio-technical approaches, and AI risks;
- explicit processes for making go/no-go system commissioning and deployment decisions;
- established policies, processes, practices, and procedures for improving organizational accountability efforts related to AI system risks;
- enhanced organizational culture which prioritizes the identification and management of AI system risks and potential impacts to individuals, communities, organizations, and society;
- better information sharing within and across organizations about risks, decision-making processes, responsibilities, common pitfalls, TEVV practices, and approaches for continuous improvement;
- greater contextual knowledge for increased awareness of downstream risks;
- strengthened engagement with interested parties and relevant AI actors; and
- augmented capacity for TEVV of AI systems and associated risks.

Part 2: Core and Profiles

5. AI RMF Core

The AI RMF Core provides outcomes and actions that enable dialogue, understanding, and activities to manage AI risks and responsibly develop trustworthy AI systems. As illustrated in Figure 5, the Core is composed of four functions: **GOVERN, MAP, MEASURE,** and **MANAGE**. Each of these high-level functions is broken down into categories and subcategories. Categories and subcategories are subdivided into specific actions and outcomes. Actions do not constitute a checklist, nor are they necessarily an ordered set of steps.

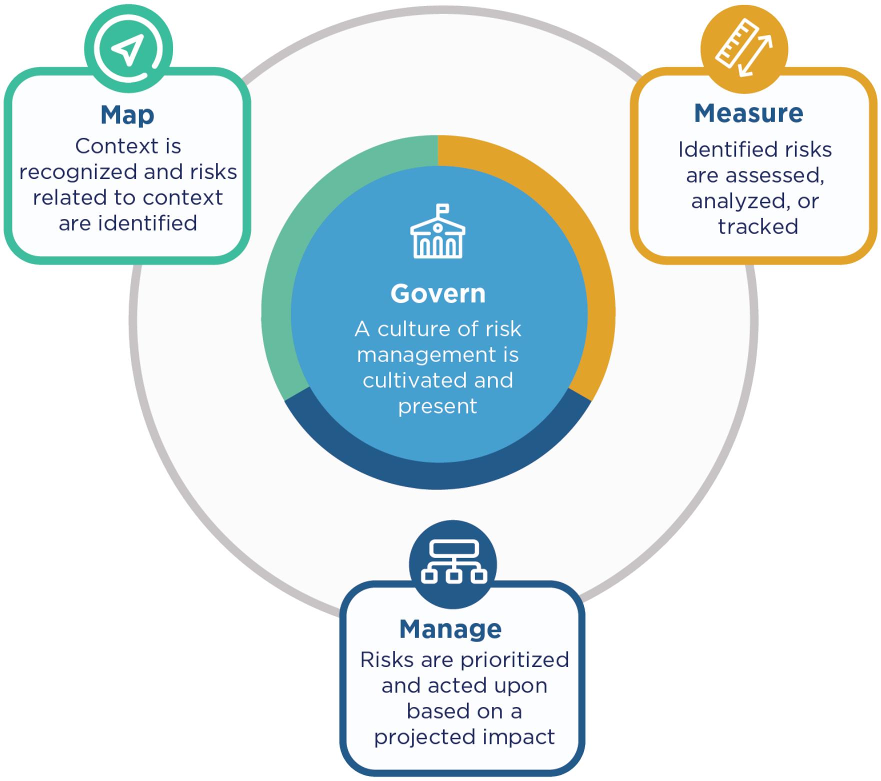
Map
Context is recognized and risks related to context are identified
Measure
Identified risks are assessed, analyzed, or tracked
Govern
A culture of risk management is cultivated and present
Manage
Risks are prioritized and acted upon based on a projected impact

Fig. 5. Functions organize AI risk management activities at their highest level to govern, map, measure, and manage AI risks. Governance is designed to be a cross-cutting function to inform and be infused throughout the other three functions.

Risk management should be continuous, timely, and performed throughout the AI system lifecycle dimensions. AI RMF Core functions should be carried out in a way that reflects diverse and multidisciplinary perspectives, potentially including the views of AI actors outside the organization. Having a diverse team contributes to more open sharing of ideas and assumptions about purposes and functions of the technology being designed, developed, deployed, or evaluated – which can create opportunities to surface problems and identify existing and emergent risks.

An online companion resource to the AI RMF, the NIST AI RMF Playbook, is available to help organizations navigate the AI RMF and achieve its outcomes through suggested tactical actions they can apply within their own contexts. Like the AI RMF, the Playbook is voluntary and organizations can utilize the suggestions according to their needs and interests. Playbook users can create tailored guidance selected from suggested material for their own use and contribute their suggestions for sharing with the broader community. Along with the AI RMF, the Playbook is part of the NIST Trustworthy and Responsible AI Resource Center.

> Framework users may apply these functions as best suits their needs for managing AI risks based on their resources and capabilities. Some organizations may choose to select from among the categories and subcategories; others may choose and have the capacity to apply all categories and subcategories. Assuming a governance structure is in place, functions may be performed in any order across the AI lifecycle as deemed to add value by a user of the framework. After instituting the outcomes in **GOVERN**, most users of the AI RMF would start with the **MAP** function and continue to **MEASURE** or **MANAGE**. However users integrate the functions, the process should be iterative, with cross-referencing between functions as necessary. Similarly, there are categories and subcategories with elements that apply to multiple functions, or that logically should take place before certain subcategory decisions.

5.1 Govern

The **GOVERN** function:

- cultivates and implements a culture of risk management within organizations designing, developing, deploying, evaluating, or acquiring AI systems;
- outlines processes, documents, and organizational schemes that anticipate, identify, and manage the risks a system can pose, including to users and others across society – and procedures to achieve those outcomes;
- incorporates processes to assess potential impacts;
- provides a structure by which AI risk management functions can align with organizational principles, policies, and strategic priorities;
- connects technical aspects of AI system design and development to organizational values and principles, and enables organizational practices and competencies for the individuals involved in acquiring, training, deploying, and monitoring such systems; and
- addresses full product lifecycle and associated processes, including legal and other issues concerning use of third-party software or hardware systems and data.

GOVERN is a cross-cutting function that is infused throughout AI risk management and enables the other functions of the process. Aspects of **GOVERN**, especially those related to compliance or evaluation, should be integrated into each of the other functions. Attention to governance is a continual and intrinsic requirement for effective AI risk management over an AI system's lifespan and the organization's hierarchy.

Strong governance can drive and enhance internal practices and norms to facilitate organizational risk culture. Governing authorities can determine the overarching policies that direct an organization's mission, goals, values, culture, and risk tolerance. Senior leadership sets the tone for risk management within an organization, and with it, organizational culture. Management aligns the technical aspects of AI risk management to policies and operations. Documentation can enhance transparency, improve human review processes, and bolster accountability in AI system teams.

After putting in place the structures, systems, processes, and teams described in the **GOVERN** function, organizations should benefit from a purpose-driven culture focused on risk understanding and management. It is incumbent on Framework users to continue to execute the **GOVERN** function as knowledge, cultures, and needs or expectations from AI actors evolve over time.

Practices related to governing AI risks are described in the NIST AI RMF Playbook. Table 1 lists the **GOVERN** function's categories and subcategories.

Table 1: Categories and subcategories for the **GOVERN** function.

Categories	Subcategories
GOVERN 1: Policies, processes, procedures, and practices across the organization related to the mapping, measuring, and managing of AI risks are in place, transparent, and implemented effectively.	**GOVERN 1.1:** Legal and regulatory requirements involving AI are understood, managed, and documented. **GOVERN 1.2:** The characteristics of trustworthy AI are integrated into organizational policies, processes, procedures, and practices. **GOVERN 1.3:** Processes, procedures, and practices are in place to determine the needed level of risk management activities based on the organization's risk tolerance. **GOVERN 1.4:** The risk management process and its outcomes are established through transparent policies, procedures, and other controls based on organizational risk priorities.

Continued on next page

Table 1: Categories and subcategories for the GOVERN function. (Continued)

Categories	Subcategories
	GOVERN 1.5: Ongoing monitoring and periodic review of the risk management process and its outcomes are planned and organizational roles and responsibilities clearly defined, including determining the frequency of periodic review. **GOVERN 1.6:** Mechanisms are in place to inventory AI systems and are resourced according to organizational risk priorities. **GOVERN 1.7:** Processes and procedures are in place for decommissioning and phasing out AI systems safely and in a manner that does not increase risks or decrease the organization's trustworthiness.
GOVERN 2: Accountability structures are in place so that the appropriate teams and individuals are empowered, responsible, and trained for mapping, measuring, and managing AI risks.	**GOVERN 2.1:** Roles and responsibilities and lines of communication related to mapping, measuring, and managing AI risks are documented and are clear to individuals and teams throughout the organization. **GOVERN 2.2:** The organization's personnel and partners receive AI risk management training to enable them to perform their duties and responsibilities consistent with related policies, procedures, and agreements. **GOVERN 2.3:** Executive leadership of the organization takes responsibility for decisions about risks associated with AI system development and deployment.
GOVERN 3: Workforce diversity, equity, inclusion, and accessibility processes are prioritized in the mapping, measuring, and managing of AI risks throughout the lifecycle.	**GOVERN 3.1:** Decision-making related to mapping, measuring, and managing AI risks throughout the lifecycle is informed by a diverse team (e.g., diversity of demographics, disciplines, experience, expertise, and backgrounds). **GOVERN 3.2:** Policies and procedures are in place to define and differentiate roles and responsibilities for human-AI configurations and oversight of AI systems.
GOVERN 4: Organizational teams are committed to a culture	**GOVERN 4.1:** Organizational policies and practices are in place to foster a critical thinking and safety-first mindset in the design, development, deployment, and uses of AI systems to minimize potential negative impacts.

Continued on next page

Table 1: Categories and subcategories for the GOVERN function. (Continued)

Categories	Subcategories
that considers and communicates AI risk.	**GOVERN 4.2:** Organizational teams document the risks and potential impacts of the AI technology they design, develop, deploy, evaluate, and use, and they communicate about the impacts more broadly. **GOVERN 4.3:** Organizational practices are in place to enable AI testing, identification of incidents, and information sharing.
GOVERN 5: Processes are in place for robust engagement with relevant AI actors.	**GOVERN 5.1:** Organizational policies and practices are in place to collect, consider, prioritize, and integrate feedback from those external to the team that developed or deployed the AI system regarding the potential individual and societal impacts related to AI risks. **GOVERN 5.2:** Mechanisms are established to enable the team that developed or deployed AI systems to regularly incorporate adjudicated feedback from relevant AI actors into system design and implementation.
GOVERN 6: Policies and procedures are in place to address AI risks and benefits arising from third-party software and data and other supply chain issues.	**GOVERN 6.1:** Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights. **GOVERN 6.2:** Contingency processes are in place to handle failures or incidents in third-party data or AI systems deemed to be high-risk.

5.2 Map

The MAP function establishes the context to frame risks related to an AI system. The AI lifecycle consists of many interdependent activities involving a diverse set of actors (See Figure 3). In practice, AI actors in charge of one part of the process often do not have full visibility or control over other parts and their associated contexts. The interdependencies between these activities, and among the relevant AI actors, can make it difficult to reliably anticipate impacts of AI systems. For example, early decisions in identifying purposes and objectives of an AI system can alter its behavior and capabilities, and the dynamics of deployment setting (such as end users or impacted individuals) can shape the impacts of AI system decisions. As a result, the best intentions within one dimension of the AI lifecycle can be undermined via interactions with decisions and conditions in other, later activities.

This complexity and varying levels of visibility can introduce uncertainty into risk management practices. Anticipating, assessing, and otherwise addressing potential sources of negative risk can mitigate this uncertainty and enhance the integrity of the decision process.

The information gathered while carrying out the **MAP** function enables negative risk prevention and informs decisions for processes such as model management, as well as an initial decision about appropriateness or the need for an AI solution. Outcomes in the **MAP** function are the basis for the **MEASURE** and **MANAGE** functions. Without contextual knowledge, and awareness of risks within the identified contexts, risk management is difficult to perform. The **MAP** function is intended to enhance an organization's ability to identify risks and broader contributing factors.

Implementation of this function is enhanced by incorporating perspectives from a diverse internal team and engagement with those external to the team that developed or deployed the AI system. Engagement with external collaborators, end users, potentially impacted communities, and others may vary based on the risk level of a particular AI system, the makeup of the internal team, and organizational policies. Gathering such broad perspectives can help organizations proactively prevent negative risks and develop more trustworthy AI systems by:

- improving their capacity for understanding contexts;
- checking their assumptions about context of use;
- enabling recognition of when systems are not functional within or out of their intended context;
- identifying positive and beneficial uses of their existing AI systems;
- improving understanding of limitations in AI and ML processes;
- identifying constraints in real-world applications that may lead to negative impacts;
- identifying known and foreseeable negative impacts related to intended use of AI systems; and
- anticipating risks of the use of AI systems beyond intended use.

After completing the **MAP** function, Framework users should have sufficient contextual knowledge about AI system impacts to inform an initial go/no-go decision about whether to design, develop, or deploy an AI system. If a decision is made to proceed, organizations should utilize the **MEASURE** and **MANAGE** functions along with policies and procedures put into place in the **GOVERN** function to assist in AI risk management efforts. It is incumbent on Framework users to continue applying the **MAP** function to AI systems as context, capabilities, risks, benefits, and potential impacts evolve over time.

Practices related to mapping AI risks are described in the NIST AI RMF Playbook. Table 2 lists the **MAP** function's categories and subcategories.

Table 2: Categories and subcategories for the MAP function.

Categories	Subcategories
MAP 1: Context is established and understood.	**MAP 1.1:** Intended purposes, potentially beneficial uses, context-specific laws, norms and expectations, and prospective settings in which the AI system will be deployed are understood and documented. Considerations include: the specific set or types of users along with their expectations; potential positive and negative impacts of system uses to individuals, communities, organizations, society, and the planet; assumptions and related limitations about AI system purposes, uses, and risks across the development or product AI lifecycle; and related TEVV and system metrics.
	MAP 1.2: Interdisciplinary AI actors, competencies, skills, and capacities for establishing context reflect demographic diversity and broad domain and user experience expertise, and their participation is documented. Opportunities for interdisciplinary collaboration are prioritized.
	MAP 1.3: The organization's mission and relevant goals for AI technology are understood and documented.
	MAP 1.4: The business value or context of business use has been clearly defined or – in the case of assessing existing AI systems – re-evaluated.
	MAP 1.5: Organizational risk tolerances are determined and documented.
	MAP 1.6: System requirements (e.g., "the system shall respect the privacy of its users") are elicited from and understood by relevant AI actors. Design decisions take socio-technical implications into account to address AI risks.
MAP 2: Categorization of the AI system is performed.	**MAP 2.1:** The specific tasks and methods used to implement the tasks that the AI system will support are defined (e.g., classifiers, generative models, recommenders).
	MAP 2.2: Information about the AI system's knowledge limits and how system output may be utilized and overseen by humans is documented. Documentation provides sufficient information to assist relevant AI actors when making decisions and taking subsequent actions.

Continued on next page

Table 2: Categories and subcategories for the MAP function. (Continued)

Categories	Subcategories
	MAP 2.3: Scientific integrity and TEVV considerations are identified and documented, including those related to experimental design, data collection and selection (e.g., availability, representativeness, suitability), system trustworthiness, and construct validation.
MAP 3: AI capabilities, targeted usage, goals, and expected benefits and costs compared with appropriate benchmarks are understood.	**MAP 3.1:** Potential benefits of intended AI system functionality and performance are examined and documented. **MAP 3.2:** Potential costs, including non-monetary costs, which result from expected or realized AI errors or system functionality and trustworthiness – as connected to organizational risk tolerance – are examined and documented. **MAP 3.3:** Targeted application scope is specified and documented based on the system's capability, established context, and AI system categorization. **MAP 3.4:** Processes for operator and practitioner proficiency with AI system performance and trustworthiness – and relevant technical standards and certifications – are defined, assessed, and documented. **MAP 3.5:** Processes for human oversight are defined, assessed, and documented in accordance with organizational policies from the GOVERN function.
MAP 4: Risks and benefits are mapped for all components of the AI system including third-party software and data.	**MAP 4.1:** Approaches for mapping AI technology and legal risks of its components – including the use of third-party data or software – are in place, followed, and documented, as are risks of infringement of a third party's intellectual property or other rights. **MAP 4.2:** Internal risk controls for components of the AI system, including third-party AI technologies, are identified and documented.
MAP 5: Impacts to individuals, groups, communities, organizations, and society are characterized.	**MAP 5.1:** Likelihood and magnitude of each identified impact (both potentially beneficial and harmful) based on expected use, past uses of AI systems in similar contexts, public incident reports, feedback from those external to the team that developed or deployed the AI system, or other data are identified and documented.

Continued on next page

Table 2: Categories and subcategories for the MAP function. (Continued)

Categories	Subcategories
	MAP 5.2: Practices and personnel for supporting regular engagement with relevant AI actors and integrating feedback about positive, negative, and unanticipated impacts are in place and documented.

5.3 Measure

The MEASURE function employs quantitative, qualitative, or mixed-method tools, techniques, and methodologies to analyze, assess, benchmark, and monitor AI risk and related impacts. It uses knowledge relevant to AI risks identified in the MAP function and informs the MANAGE function. AI systems should be tested before their deployment and regularly while in operation. AI risk measurements include documenting aspects of systems' functionality and trustworthiness.

Measuring AI risks includes tracking metrics for trustworthy characteristics, social impact, and human-AI configurations. Processes developed or adopted in the MEASURE function should include rigorous software testing and performance assessment methodologies with associated measures of uncertainty, comparisons to performance benchmarks, and formalized reporting and documentation of results. Processes for independent review can improve the effectiveness of testing and can mitigate internal biases and potential conflicts of interest.

Where tradeoffs among the trustworthy characteristics arise, measurement provides a traceable basis to inform management decisions. Options may include recalibration, impact mitigation, or removal of the system from design, development, production, or use, as well as a range of compensating, detective, deterrent, directive, and recovery controls.

After completing the MEASURE function, objective, repeatable, or scalable test, evaluation, verification, and validation (TEVV) processes including metrics, methods, and methodologies are in place, followed, and documented. Metrics and measurement methodologies should adhere to scientific, legal, and ethical norms and be carried out in an open and transparent process. New types of measurement, qualitative and quantitative, may need to be developed. The degree to which each measurement type provides unique and meaningful information to the assessment of AI risks should be considered. Framework users will enhance their capacity to comprehensively evaluate system trustworthiness, identify and track existing and emergent risks, and verify efficacy of the metrics. Measurement outcomes will be utilized in the MANAGE function to assist risk monitoring and response efforts. It is incumbent on Framework users to continue applying the MEASURE function to AI systems as knowledge, methodologies, risks, and impacts evolve over time.

Practices related to measuring AI risks are described in the NIST AI RMF Playbook. Table 3 lists the **MEASURE** function's categories and subcategories.

Table 3: Categories and subcategories for the **MEASURE** function.

Categories	Subcategories
MEASURE 1: Appropriate methods and metrics are identified and applied.	**MEASURE 1.1:** Approaches and metrics for measurement of AI risks enumerated during the **MAP** function are selected for implementation starting with the most significant AI risks. The risks or trustworthiness characteristics that will not – or cannot – be measured are properly documented.
	MEASURE 1.2: Appropriateness of AI metrics and effectiveness of existing controls are regularly assessed and updated, including reports of errors and potential impacts on affected communities.
	MEASURE 1.3: Internal experts who did not serve as front-line developers for the system and/or independent assessors are involved in regular assessments and updates. Domain experts, users, AI actors external to the team that developed or deployed the AI system, and affected communities are consulted in support of assessments as necessary per organizational risk tolerance.
MEASURE 2: AI systems are evaluated for trustworthy characteristics.	**MEASURE 2.1:** Test sets, metrics, and details about the tools used during TEVV are documented.
	MEASURE 2.2: Evaluations involving human subjects meet applicable requirements (including human subject protection) and are representative of the relevant population.
	MEASURE 2.3: AI system performance or assurance criteria are measured qualitatively or quantitatively and demonstrated for conditions similar to deployment setting(s). Measures are documented.
	MEASURE 2.4: The functionality and behavior of the AI system and its components – as identified in the **MAP** function – are monitored when in production.
	MEASURE 2.5: The AI system to be deployed is demonstrated to be valid and reliable. Limitations of the generalizability beyond the conditions under which the technology was developed are documented.

Continued on next page

Table 3: Categories and subcategories for the MEASURE function. (Continued)

Categories	Subcategories
	MEASURE 2.6: The AI system is evaluated regularly for safety risks – as identified in the MAP function. The AI system to be deployed is demonstrated to be safe, its residual negative risk does not exceed the risk tolerance, and it can fail safely, particularly if made to operate beyond its knowledge limits. Safety metrics reflect system reliability and robustness, real-time monitoring, and response times for AI system failures. **MEASURE 2.7:** AI system security and resilience – as identified in the MAP function – are evaluated and documented. **MEASURE 2.8:** Risks associated with transparency and accountability – as identified in the MAP function – are examined and documented. **MEASURE 2.9:** The AI model is explained, validated, and documented, and AI system output is interpreted within its context – as identified in the MAP function – to inform responsible use and governance. **MEASURE 2.10:** Privacy risk of the AI system – as identified in the MAP function – is examined and documented. **MEASURE 2.11:** Fairness and bias – as identified in the MAP function – are evaluated and results are documented. **MEASURE 2.12:** Environmental impact and sustainability of AI model training and management activities – as identified in the MAP function – are assessed and documented. **MEASURE 2.13:** Effectiveness of the employed TEVV metrics and processes in the MEASURE function are evaluated and documented.
MEASURE 3: Mechanisms for tracking identified AI risks over time are in place.	**MEASURE 3.1:** Approaches, personnel, and documentation are in place to regularly identify and track existing, unanticipated, and emergent AI risks based on factors such as intended and actual performance in deployed contexts. **MEASURE 3.2:** Risk tracking approaches are considered for settings where AI risks are difficult to assess using currently available measurement techniques or where metrics are not yet available.

Continued on next page

Table 3: Categories and subcategories for the MEASURE function. (Continued)

Categories	Subcategories
	MEASURE 3.3: Feedback processes for end users and impacted communities to report problems and appeal system outcomes are established and integrated into AI system evaluation metrics.
MEASURE 4: Feedback about efficacy of measurement is gathered and assessed.	**MEASURE 4.1:** Measurement approaches for identifying AI risks are connected to deployment context(s) and informed through consultation with domain experts and other end users. Approaches are documented. **MEASURE 4.2:** Measurement results regarding AI system trustworthiness in deployment context(s) and across the AI lifecycle are informed by input from domain experts and relevant AI actors to validate whether the system is performing consistently as intended. Results are documented. **MEASURE 4.3:** Measurable performance improvements or declines based on consultations with relevant AI actors, including affected communities, and field data about context-relevant risks and trustworthiness characteristics are identified and documented.

5.4 Manage

The MANAGE function entails allocating risk resources to mapped and measured risks on a regular basis and as defined by the GOVERN function. Risk treatment comprises plans to respond to, recover from, and communicate about incidents or events.

Contextual information gleaned from expert consultation and input from relevant AI actors – established in GOVERN and carried out in MAP – is utilized in this function to decrease the likelihood of system failures and negative impacts. Systematic documentation practices established in GOVERN and utilized in MAP and MEASURE bolster AI risk management efforts and increase transparency and accountability. Processes for assessing emergent risks are in place, along with mechanisms for continual improvement.

After completing the MANAGE function, plans for prioritizing risk and regular monitoring and improvement will be in place. Framework users will have enhanced capacity to manage the risks of deployed AI systems and to allocate risk management resources based on assessed and prioritized risks. It is incumbent on Framework users to continue to apply the MANAGE function to deployed AI systems as methods, contexts, risks, and needs or expectations from relevant AI actors evolve over time.

Practices related to managing AI risks are described in the NIST AI RMF Playbook. Table 4 lists the **MANAGE** function's categories and subcategories.

Table 4: Categories and subcategories for the **MANAGE** function.

Categories	Subcategories
MANAGE 1: AI risks based on assessments and other analytical output from the **MAP** and **MEASURE** functions are prioritized, responded to, and managed.	**MANAGE 1.1:** A determination is made as to whether the AI system achieves its intended purposes and stated objectives and whether its development or deployment should proceed. **MANAGE 1.2:** Treatment of documented AI risks is prioritized based on impact, likelihood, and available resources or methods. **MANAGE 1.3:** Responses to the AI risks deemed high priority, as identified by the **MAP** function, are developed, planned, and documented. Risk response options can include mitigating, transferring, avoiding, or accepting. **MANAGE 1.4:** Negative residual risks (defined as the sum of all unmitigated risks) to both downstream acquirers of AI systems and end users are documented.
MANAGE 2: Strategies to maximize AI benefits and minimize negative impacts are planned, prepared, implemented, documented, and informed by input from relevant AI actors.	**MANAGE 2.1:** Resources required to manage AI risks are taken into account – along with viable non-AI alternative systems, approaches, or methods – to reduce the magnitude or likelihood of potential impacts. **MANAGE 2.2:** Mechanisms are in place and applied to sustain the value of deployed AI systems. **MANAGE 2.3:** Procedures are followed to respond to and recover from a previously unknown risk when it is identified. **MANAGE 2.4:** Mechanisms are in place and applied, and responsibilities are assigned and understood, to supersede, disengage, or deactivate AI systems that demonstrate performance or outcomes inconsistent with intended use.
MANAGE 3: AI risks and benefits from third-party entities are managed.	**MANAGE 3.1:** AI risks and benefits from third-party resources are regularly monitored, and risk controls are applied and documented. **MANAGE 3.2:** Pre-trained models which are used for development are monitored as part of AI system regular monitoring and maintenance.

Continued on next page

Table 4: Categories and subcategories for the MANAGE function. (Continued)

Categories	Subcategories
MANAGE 4: Risk treatments, including response and recovery, and communication plans for the identified and measured AI risks are documented and monitored regularly.	**MANAGE 4.1:** Post-deployment AI system monitoring plans are implemented, including mechanisms for capturing and evaluating input from users and other relevant AI actors, appeal and override, decommissioning, incident response, recovery, and change management.
	MANAGE 4.2: Measurable activities for continual improvements are integrated into AI system updates and include regular engagement with interested parties, including relevant AI actors.
	MANAGE 4.3: Incidents and errors are communicated to relevant AI actors, including affected communities. Processes for tracking, responding to, and recovering from incidents and errors are followed and documented.

6. AI RMF Profiles

AI RMF *use-case profiles* are implementations of the AI RMF functions, categories, and subcategories for a specific setting or application based on the requirements, risk tolerance, and resources of the Framework user: for example, an AI RMF *hiring profile* or an AI RMF *fair housing profile*. Profiles may illustrate and offer insights into how risk can be managed at various stages of the AI lifecycle or in specific sector, technology, or end-use applications. AI RMF profiles assist organizations in deciding how they might best manage AI risk that is well-aligned with their goals, considers legal/regulatory requirements and best practices, and reflects risk management priorities.

AI RMF *temporal profiles* are descriptions of either the current state or the desired, target state of specific AI risk management activities within a given sector, industry, organization, or application context. An AI RMF Current Profile indicates how AI is currently being managed and the related risks in terms of current outcomes. A Target Profile indicates the outcomes needed to achieve the desired or target AI risk management goals.

Comparing Current and Target Profiles likely reveals gaps to be addressed to meet AI risk management objectives. Action plans can be developed to address these gaps to fulfill outcomes in a given category or subcategory. Prioritization of gap mitigation is driven by the user's needs and risk management processes. This risk-based approach also enables Framework users to compare their approaches with other approaches and to gauge the resources needed (e.g., staffing, funding) to achieve AI risk management goals in a cost-effective, prioritized manner.

AI RMF *cross-sectoral profiles* cover risks of models or applications that can be used across use cases or sectors. Cross-sectoral profiles can also cover how to govern, map, measure, and manage risks for activities or business processes common across sectors such as the use of large language models, cloud-based services or acquisition.

This Framework does not prescribe profile templates, allowing for flexibility in implementation.

Appendix A:
Descriptions of AI Actor Tasks from Figures 2 and 3

AI Design tasks are performed during the Application Context and Data and Input phases of the AI lifecycle in Figure 2. AI Design actors create the concept and objectives of AI systems and are responsible for the planning, design, and data collection and processing tasks of the AI system so that the AI system is lawful and fit-for-purpose. Tasks include articulating and documenting the system's concept and objectives, underlying assumptions, context, and requirements; gathering and cleaning data; and documenting the metadata and characteristics of the dataset. AI actors in this category include data scientists, domain experts, socio-cultural analysts, experts in the field of diversity, equity, inclusion, and accessibility, members of impacted communities, human factors experts (e.g., UX/UI design), governance experts, data engineers, data providers, system funders, product managers, third-party entities, evaluators, and legal and privacy governance.

AI Development tasks are performed during the AI Model phase of the lifecycle in Figure 2. AI Development actors provide the initial infrastructure of AI systems and are responsible for model building and interpretation tasks, which involve the creation, selection, calibration, training, and/or testing of models or algorithms. AI actors in this category include machine learning experts, data scientists, developers, third-party entities, legal and privacy governance experts, and experts in the socio-cultural and contextual factors associated with the deployment setting.

AI Deployment tasks are performed during the Task and Output phase of the lifecycle in Figure 2. AI Deployment actors are responsible for contextual decisions relating to how the AI system is used to assure deployment of the system into production. Related tasks include piloting the system, checking compatibility with legacy systems, ensuring regulatory compliance, managing organizational change, and evaluating user experience. AI actors in this category include system integrators, software developers, end users, operators and practitioners, evaluators, and domain experts with expertise in human factors, socio-cultural analysis, and governance.

Operation and Monitoring tasks are performed in the Application Context/Operate and Monitor phase of the lifecycle in Figure 2. These tasks are carried out by AI actors who are responsible for operating the AI system and working with others to regularly assess system output and impacts. AI actors in this category include system operators, domain experts, AI designers, users who interpret or incorporate the output of AI systems, product developers, evaluators and auditors, compliance experts, organizational management, and members of the research community.

Test, Evaluation, Verification, and Validation (TEVV) tasks are performed throughout the AI lifecycle. They are carried out by AI actors who examine the AI system or its components, or detect and remediate problems. Ideally, AI actors carrying out verification and validation tasks are distinct from those who perform test and evaluation actions. Tasks can be incorporated into a phase as early as design, where tests are planned in accordance with the design requirement.

- TEVV tasks for design, planning, and data may center on internal and external validation of assumptions for system design, data collection, and measurements relative to the intended context of deployment or application.
- TEVV tasks for development (i.e., model building) include model validation and assessment.
- TEVV tasks for deployment include system validation and integration in production, with testing, and recalibration for systems and process integration, user experience, and compliance with existing legal, regulatory, and ethical specifications.
- TEVV tasks for operations involve ongoing monitoring for periodic updates, testing, and subject matter expert (SME) recalibration of models, the tracking of incidents or errors reported and their management, the detection of emergent properties and related impacts, and processes for redress and response.

Human Factors tasks and activities are found throughout the dimensions of the AI lifecycle. They include human-centered design practices and methodologies, promoting the active involvement of end users and other interested parties and relevant AI actors, incorporating context-specific norms and values in system design, evaluating and adapting end user experiences, and broad integration of humans and human dynamics in all phases of the AI lifecycle. Human factors professionals provide multidisciplinary skills and perspectives to understand context of use, inform interdisciplinary and demographic diversity, engage in consultative processes, design and evaluate user experience, perform human-centered evaluation and testing, and inform impact assessments.

Domain Expert tasks involve input from multidisciplinary practitioners or scholars who provide knowledge or expertise in – and about – an industry sector, economic sector, context, or application area where an AI system is being used. AI actors who are domain experts can provide essential guidance for AI system design and development, and interpret outputs in support of work performed by TEVV and AI impact assessment teams.

AI Impact Assessment tasks include assessing and evaluating requirements for AI system accountability, combating harmful bias, examining impacts of AI systems, product safety, liability, and security, among others. AI actors such as impact assessors and evaluators provide technical, human factor, socio-cultural, and legal expertise.

Procurement tasks are conducted by AI actors with financial, legal, or policy management authority for acquisition of AI models, products, or services from a third-party developer, vendor, or contractor.

Governance and Oversight tasks are assumed by AI actors with management, fiduciary, and legal authority and responsibility for the organization in which an AI system is designed, developed, and/or deployed. Key AI actors responsible for AI governance include organizational management, senior leadership, and the Board of Directors. These actors are parties that are concerned with the impact and sustainability of the organization as a whole.

Additional AI Actors

Third-party entities include providers, developers, vendors, and evaluators of data, algorithms, models, and/or systems and related services for another organization or the organization's customers or clients. Third-party entities are responsible for AI design and development tasks, in whole or in part. By definition, they are external to the design, development, or deployment team of the organization that acquires its technologies or services. The technologies acquired from third-party entities may be complex or opaque, and risk tolerances may not align with the deploying or operating organization.

End users of an AI system are the individuals or groups that use the system for specific purposes. These individuals or groups interact with an AI system in a specific context. End users can range in competency from AI experts to first-time technology end users.

Affected individuals/communities encompass all individuals, groups, communities, or organizations directly or indirectly affected by AI systems or decisions based on the output of AI systems. These individuals do not necessarily interact with the deployed system or application.

Other AI actors may provide formal or quasi-formal norms or guidance for specifying and managing AI risks. They can include **trade associations, standards developing organizations, advocacy groups, researchers, environmental groups, and civil society organizations**.

The general public is most likely to directly experience positive and negative impacts of AI technologies. They may provide the motivation for actions taken by the AI actors. This group can include individuals, communities, and consumers associated with the context in which an AI system is developed or deployed.

Appendix B: How AI Risks Differ from Traditional Software Risks

As with traditional software, risks from AI-based technology can be bigger than an enterprise, span organizations, and lead to societal impacts. AI systems also bring a set of risks that are not comprehensively addressed by current risk frameworks and approaches. Some AI system features that present risks also can be beneficial. For example, pre-trained models and transfer learning can advance research and increase accuracy and resilience when compared to other models and approaches. Identifying contextual factors in the **MAP** function will assist AI actors in determining the level of risk and potential management efforts.

Compared to traditional software, AI-specific risks that are new or increased include the following:

- The data used for building an AI system may not be a true or appropriate representation of the context or intended use of the AI system, and the ground truth may either not exist or not be available. Additionally, harmful bias and other data quality issues can affect AI system trustworthiness, which could lead to negative impacts.
- AI system dependency and reliance on data for training tasks, combined with increased volume and complexity typically associated with such data.
- Intentional or unintentional changes during training may fundamentally alter AI system performance.
- Datasets used to train AI systems may become detached from their original and intended context or may become stale or outdated relative to deployment context.
- AI system scale and complexity (many systems contain billions or even trillions of decision points) housed within more traditional software applications.
- Use of pre-trained models that can advance research and improve performance can also increase levels of statistical uncertainty and cause issues with bias management, scientific validity, and reproducibility.
- Higher degree of difficulty in predicting failure modes for emergent properties of large-scale pre-trained models.
- Privacy risk due to enhanced data aggregation capability for AI systems.
- AI systems may require more frequent maintenance and triggers for conducting corrective maintenance due to data, model, or concept drift.
- Increased opacity and concerns about reproducibility.
- Underdeveloped software testing standards and inability to document AI-based practices to the standard expected of traditionally engineered software for all but the simplest of cases.
- Difficulty in performing regular AI-based software testing, or determining what to test, since AI systems are not subject to the same controls as traditional code development.

- Computational costs for developing AI systems and their impact on the environment and planet.
- Inability to predict or detect the side effects of AI-based systems beyond statistical measures.

Privacy and cybersecurity risk management considerations and approaches are applicable in the design, development, deployment, evaluation, and use of AI systems. Privacy and cybersecurity risks are also considered as part of broader enterprise risk management considerations, which may incorporate AI risks. As part of the effort to address AI trustworthiness characteristics such as "Secure and Resilient" and "Privacy-Enhanced," organizations may consider leveraging available standards and guidance that provide broad guidance to organizations to reduce security and privacy risks, such as, but not limited to, the NIST Cybersecurity Framework, the NIST Privacy Framework, the NIST Risk Management Framework, and the Secure Software Development Framework. These frameworks have some features in common with the AI RMF. Like most risk management approaches, they are outcome-based rather than prescriptive and are often structured around a Core set of functions, categories, and subcategories. While there are significant differences between these frameworks based on the domain addressed – and because AI risk management calls for addressing many other types of risks – frameworks like those mentioned above may inform security and privacy considerations in the **MAP**, **MEASURE**, and **MANAGE** functions of the AI RMF.

At the same time, guidance available before publication of this AI RMF does not comprehensively address many AI system risks. For example, existing frameworks and guidance are unable to:

- adequately manage the problem of harmful bias in AI systems;
- confront the challenging risks related to generative AI;
- comprehensively address security concerns related to evasion, model extraction, membership inference, availability, or other machine learning attacks;
- account for the complex attack surface of AI systems or other security abuses enabled by AI systems; and
- consider risks associated with third-party AI technologies, transfer learning, and off-label use where AI systems may be trained for decision-making outside an organization's security controls or trained in one domain and then "fine-tuned" for another.

Both AI and traditional software technologies and systems are subject to rapid innovation. Technology advances should be monitored and deployed to take advantage of those developments and work towards a future of AI that is both trustworthy and responsible.

Appendix C:
AI Risk Management and Human-AI Interaction

Organizations that design, develop, or deploy AI systems for use in operational settings may enhance their AI risk management by understanding current limitations of human-AI interaction. The AI RMF provides opportunities to clearly define and differentiate the various human roles and responsibilities when using, interacting with, or managing AI systems.

Many of the data-driven approaches that AI systems rely on attempt to convert or represent individual and social observational and decision-making practices into measurable quantities. Representing complex human phenomena with mathematical models can come at the cost of removing necessary context. This loss of context may in turn make it difficult to understand individual and societal impacts that are key to AI risk management efforts.

Issues that merit further consideration and research include:

1. **Human roles and responsibilities in decision making and overseeing AI systems need to be clearly defined and differentiated.** Human-AI configurations can span from fully autonomous to fully manual. AI systems can autonomously make decisions, defer decision making to a human expert, or be used by a human decision maker as an additional opinion. Some AI systems may not require human oversight, such as models used to improve video compression. Other systems may specifically require human oversight.

2. **Decisions that go into the design, development, deployment, evaluation, and use of AI systems reflect systemic and human cognitive biases.** AI actors bring their cognitive biases, both individual and group, into the process. Biases can stem from end-user decision-making tasks and be introduced across the AI lifecycle via human assumptions, expectations, and decisions during design and modeling tasks. These biases, which are not necessarily always harmful, may be exacerbated by AI system opacity and the resulting lack of transparency. Systemic biases at the organizational level can influence how teams are structured and who controls the decision-making processes throughout the AI lifecycle. These biases can also influence downstream decisions by end users, decision makers, and policy makers and may lead to negative impacts.

3. **Human-AI interaction results vary.** Under certain conditions – for example, in perceptual-based judgment tasks – the AI part of the human-AI interaction can amplify human biases, leading to more biased decisions than the AI or human alone. When these variations are judiciously taken into account in organizing human-AI teams, however, they can result in complementarity and improved overall performance.

4. **Presenting AI system information to humans is complex.** Humans perceive and derive meaning from AI system output and explanations in different ways, reflecting different individual preferences, traits, and skills.

The **GOVERN** function provides organizations with the opportunity to clarify and define the roles and responsibilities for the humans in the Human-AI team configurations and those who are overseeing the AI system performance. The **GOVERN** function also creates mechanisms for organizations to make their decision-making processes more explicit, to help counter systemic biases.

The **MAP** function suggests opportunities to define and document processes for operator and practitioner proficiency with AI system performance and trustworthiness concepts, and to define relevant technical standards and certifications. Implementing **MAP** function categories and subcategories may help organizations improve their internal competency for analyzing context, identifying procedural and system limitations, exploring and examining impacts of AI-based systems in the real world, and evaluating decision-making processes throughout the AI lifecycle.

The **GOVERN** and **MAP** functions describe the importance of interdisciplinarity and demographically diverse teams and utilizing feedback from potentially impacted individuals and communities. AI actors called out in the AI RMF who perform human factors tasks and activities can assist technical teams by anchoring in design and development practices to user intentions and representatives of the broader AI community, and societal values. These actors further help to incorporate context-specific norms and values in system design and evaluate end user experiences – in conjunction with AI systems.

AI risk management approaches for human-AI configurations will be augmented by ongoing research and evaluation. For example, the degree to which humans are empowered and incentivized to challenge AI system output requires further studies. Data about the frequency and rationale with which humans overrule AI system output in deployed systems may be useful to collect and analyze.

Appendix D:
Attributes of the AI RMF

NIST described several key attributes of the AI RMF when work on the Framework first began. These attributes have remained intact and were used to guide the AI RMF's development. They are provided here as a reference.

The AI RMF strives to:

1. Be risk-based, resource-efficient, pro-innovation, and voluntary.
2. Be consensus-driven and developed and regularly updated through an open, transparent process. All stakeholders should have the opportunity to contribute to the AI RMF's development.
3. Use clear and plain language that is understandable by a broad audience, including senior executives, government officials, non-governmental organization leadership, and those who are not AI professionals – while still of sufficient technical depth to be useful to practitioners. The AI RMF should allow for communication of AI risks across an organization, between organizations, with customers, and to the public at large.
4. Provide common language and understanding to manage AI risks. The AI RMF should offer taxonomy, terminology, definitions, metrics, and characterizations for AI risk.
5. Be easily usable and fit well with other aspects of risk management. Use of the Framework should be intuitive and readily adaptable as part of an organization's broader risk management strategy and processes. It should be consistent or aligned with other approaches to managing AI risks.
6. Be useful to a wide range of perspectives, sectors, and technology domains. The AI RMF should be universally applicable to any AI technology and to context-specific use cases.
7. Be outcome-focused and non-prescriptive. The Framework should provide a catalog of outcomes and approaches rather than prescribe one-size-fits-all requirements.
8. Take advantage of and foster greater awareness of existing standards, guidelines, best practices, methodologies, and tools for managing AI risks – as well as illustrate the need for additional, improved resources.
9. Be law- and regulation-agnostic. The Framework should support organizations' abilities to operate under applicable domestic and international legal or regulatory regimes.
10. Be a living document. The AI RMF should be readily updated as technology, understanding, and approaches to AI trustworthiness and uses of AI change and as stakeholders learn from implementing AI risk management generally and this framework in particular.

This publication is available free of charge from:
https://doi.org/10.6028/NIST.AI.100-1

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

Filed 9/12/25

CERTIFIED FOR PUBLICATION

IN THE COURT OF APPEAL OF THE STATE OF CALIFORNIA

SECOND APPELLATE DISTRICT

DIVISION THREE

SYLVIA NOLAND, Plaintiff and Appellant, v. LAND OF THE FREE, L.P., et al., Defendants and Respondents.	B331918 (Los Angeles County Super. Ct. No. BC716737)

APPEAL from a judgment of the Superior Court of Los Angeles County, Stephen I. Goorvitch, Judge. Affirmed.

Mostafavi Law Group and Amir Mostafavi for Plaintiff and Appellant.

Yadegari & Associates and Michael Yadegari for Defendants and Respondents.

This appeal is, in most respects, unremarkable. Plaintiff filed a complaint alleging a variety of employment-related claims, and the trial court granted defendants' motion for summary judgment, finding no triable issues as to any of those claims. Plaintiff challenges the grant of summary judgment on several grounds, none of which raises any novel questions of law or requires us to apply settled law in a unique factual context. In short, this is in most respects a straightforward appeal that, under normal circumstances, would not warrant publication.

What sets this appeal apart—and the reason we have elected to publish this opinion—is that nearly all of the legal quotations in plaintiff's opening brief, and many of the quotations in plaintiff's reply brief, are fabricated. That is, the quotes plaintiff attributes to published cases do not appear in those cases or anywhere else. Further, many of the cases plaintiff cites do not discuss the topics for which they are cited, and a few of the cases do not exist at all. These fabricated legal authorities were created by generative artificial intelligence (AI) tools that plaintiff's counsel used to draft his appellate briefs. The AI tools created fake legal authority—sometimes referred to as AI "hallucinations"—that were undetected by plaintiff's counsel because he did not read the cases the AI tools cited.

Although the generation of fake legal authority by AI sources has been widely commented on by federal and out-of-state courts and reported by many media sources, no California court has addressed this issue. We therefore publish this opinion as a warning. Simply stated, no brief, pleading, motion, or any other paper filed in any court should contain *any* citations—whether provided by generative AI or any other source—that the attorney responsible for submitting the pleading has not personally read and verified. Because plaintiff's counsel's conduct in this case violated a basic duty counsel owed to his client and the court, we impose a monetary sanction on counsel, direct him to serve a copy of this opinion on his client, and direct the clerk of the court to serve a copy of this opinion on the State Bar.

FACTUAL AND PROCEDURAL BACKGROUND

I. Complaint.

Sylvia Noland (Noland) filed the present action in August 2018, and filed the operative second amended complaint (complaint) in August 2019. The complaint alleges as follows:

Defendants Jose Luis Nazar and Land of the Free, L.P. (collectively, defendants) own an office building located at 640 S. San Vicente Boulevard (the San Vicente property) and an event space located at 2400 Laurel Canyon Boulevard in Los Angeles (the Laurel Canyon property). In January 2018, defendants hired plaintiff to work as their leasing agent and sales representative. In that capacity, plaintiff showed the properties to potential lessees, prepared deal memos, and collected deposits and signatures on leases and contracts.

Defendants agreed to pay plaintiff for administrative work, plus a 6 percent commission for each event she booked at the Laurel Canyon property and a 2 percent commission for each tenant she secured for the San Vicente property. Defendants further agreed to pay plaintiff a $3,500 monthly draw against her earnings and commissions. However, defendants never paid plaintiff $1,000 per month for her administrative work, and in 2018 defendants told plaintiff they would no longer pay her a monthly draw. Further, defendants failed to pay plaintiff a

$60,000 commission she was owed for securing a lease worth $3.5 million over a 10-year term. Defendants also failed to pay plaintiff minimum wage or overtime, to maintain proper time records, and to provide plaintiff itemized wage statements.

In about June 2018, plaintiff learned that defendants did not have the necessary permits to lease office space at the San Vicente property to medical providers. Plaintiff "refused to work under intolerable working conditions that required [her to] secure lease agreements in blatant violation of the law and act unethically towards clients," and she "was therefore left with no reasonable alternative but to resign and was constructively terminated."

Plaintiff's complaint asserted 25 causes of action, including violations of California's wage and hour laws (1st–5th and 14th–22nd causes of action), retaliation (6th cause of action), constructive and wrongful termination (7th and 23rd causes of action), breach of contract (8th cause of action), quantum meruit (11th cause of action), violation of Business and Professions Code section 17200 (12th cause of action), penalties under the Private Attorneys General Act (PAGA; Lab. Code, § 1298) (13th cause of action), misclassification of employee as independent contractor (24th cause of action), and intentional infliction of emotional distress (25th cause of action).[1]

---

[1] The trial court sustained demurrers to the ninth, tenth, and eighteenth causes of action, and thus we do not discuss them.

II. Defendants' first motion for summary judgment.

Defendants filed a motion for summary judgment in September 2022, noticing a hearing for December 1, 2022. Plaintiff moved to strike the motion as untimely. The trial court (Hon. David Sotelo) denied the summary judgment motion on the ground that it was not filed sufficiently in advance of the hearing date.

III. Defendants' request to continue trial and second motion for summary judgment.

A. Request to continue trial.

In January 2023, defendants filed an ex parte application to continue the trial from January to May 2023. Defendants' counsel stated that he had recently been in an automobile accident that limited his mobility and required multiple doctors' visits. Accordingly, "It would be very difficult for Defendant[s'] counsel to appear in trial at this time while he is in recovery."

Plaintiff responded that she was amenable to continuing the trial as long as the statutory five-year deadline for bringing the case to trial was tolled. At a January 2023 hearing, the parties stipulated to extending the five-year period through the end of December 2023, and the trial court continued the trial to May 2023. That date was later vacated, and trial was set for June 2023.

B. Second motion for summary judgment.

Defendants refiled their motion for summary judgment in January 2023. The motion was essentially identical to that filed in September 2022, urging that there were no triable issues of material fact as to any of plaintiff's causes of action.

Plaintiff responded by filing a motion for sanctions. Plaintiff asserted that sanctions were appropriate because (1) defendants had sought a trial continuance for the purpose of refiling their summary judgment motion, (2) defendants' second motion for summary judgment did not assert new or different facts or legal issues, and it therefore violated Code of Civil Procedure[2] section 437, subdivision (f)(2), and (3) the date on which the motion for summary judgment was set to be heard was fewer than 30 days prior to the date set for trial.

On May 25, 2023, the trial court denied the motion for sanctions and continued the hearing to allow plaintiff to file a substantive opposition to the summary judgment motion. The court's order explained as follows:

"This case originally was assigned to [Judge] David Sotelo . . . . Defendants previously filed a motion for summary judgment on September 29, 2022, and noticed the hearing for December 1, 2022, in advance of a trial date of January 17, 2023. Judge Sotelo denied the motion for lack of statutory notice without addressing the merits. Following Judge Sotelo's retirement, the case was reassigned . . . .

"Now, Defendants again move for summary judgment or, in the alternative, summary adjudication. Plaintiff did not file an opposition, but instead objected under . . . section 437c(f)(2), arguing that this is a successive motion for summary judgment. . . . The Court overrules the objection for three independent reasons.

---

[2] All subsequent undesignated statutory references are to the Code of Civil Procedure.

"First, the Court interprets section 437c(f)(2) as prohibiting successive motions for summary judgment only when there has been a prior ruling on the merits. As discussed, Judge Sotelo denied the motion on procedural grounds without resolving the merits. Therefore, section 437c(f)(2) does not prohibit Defendants from filing a second motion for summary judgment.

"Second, . . . [a] trial court may not refuse to hear a motion for summary judgment filed and served sufficiently in advance of trial. [Citation.] Defendants had a right to a decision on the merits, given that they filed their motion 110 days before trial.

"Finally, in the alternative, the Court exercises its discretion and elects to consider Defendants' motion for summary judgment on the merits. Notwithstanding section 437c(f)(2), the Court has inherent authority to consider a second motion for summary judgment, provided there is good cause to do so. [Citation.] The Court does so for the reasons stated. Moreover, the issues raised in Defendants' motion should be decided in advance of trial. It would not promote the interests of judicial economy to select the jury and permit Plaintiff to conduct her case-in-chief before ruling on these issues on a motion for non-suit."

Plaintiff thereafter filed a substantive opposition to defendants' motion for summary judgment, urging that there were triable issues of fact as to each cause of action. The court held a hearing on the motion and, after taking the motion under submission, granted summary judgment for defendants. Among other things, the court found the evidence was undisputed that (1) plaintiff was an independent contractor, not an employee, and thus the wage and hour laws did not apply to her, (2) defendants did not owe plaintiff a commission because the tenant plaintiff said she secured ultimately did not execute a lease with defendants, (3) plaintiff had not demonstrated that she was subject to any adverse employment actions that could form the basis for a retaliation action, and (4) plaintiff did not demonstrate triable issues as to her intentional infliction of emotional distress claim.

On July 25, 2023, plaintiff filed a notice of appeal from the order granting summary judgment.[3]

---

[3] An order granting summary judgment is not an appealable order. (E.g, *Champlin/GEI Wind Holdings, LLC v. Avery* (2023) 92 Cal.App.5th 218, 223; *Levy v. Skywalker Sound* (2003) 108 Cal.App.4th 753, 761, fn. 7.) Thus, in October 2023, this court directed plaintiff to provide the court with an appealable judgment or to explain why the appeal should not be dismissed. Plaintiff responded in a letter brief that no judgment had been entered, but the appeal should not be dismissed because the order granting summary judgment was a final order that resolved all pending issues between the parties. This court deferred the appealability question to the panel that would decide the appeal on the merits.

Although more than 18 months have passed since this court advised plaintiff of the need to obtain a judgment, plaintiff's counsel has not obtained one. Accordingly, we would be well within our discretion to dismiss the appeal. (See *Blauser v. Dubin* (2024) 106 Cal.App.5th 918, 920–923 [dismissing appeal from minute order granting motion for nonsuit].) Nonetheless, in the interests of justice and to avoid delay, we construe the order granting summary judgment as incorporating an appealable judgment, and the notice of appeal as appealing from such judgment. (See *Blauser*, at p. 922 & fn. 4; *Levy v. Skywalker Sound*, *supra*, 108 Cal.App.4th at p. 761, fn. 7.)

DISCUSSION

I. Plaintiff's counsel's reliance on fabricated legal authority.

We begin by noting that nearly all of the quotations in plaintiff's opening brief, and many of the quotations in plaintiff's reply brief, have been fabricated. That is, as noted above, although most of the cases to which the quotes are attributed exist, the quotes do not. Further, many of the cases plaintiff cites do not support the propositions for which they are cited or discuss other matters entirely, and a few of the cases do not exist at all. To give just a few examples:

Plaintiff asserts: "In *Schimmel v. Levin* (2011) 195 Cal.App.4th 81, the court discussed the legislative purpose behind Section 437c(f)(2), highlighting that it was enacted to prevent abuse of the summary judgment procedure by disallowing multiple motions on the same issues." In fact, *Schimmel* does not contain a single reference to either summary judgment or section 437c. Appellant's opening brief also purports to quote *Schimmel* as follows: "In *Schimmel v. Levin* (2011) 195 Cal.App.4th 81, 86–87, the court held: 'Section 437c(f)(2) embodies a legislative judgment that a party should not be allowed to bring multiple motions for summary judgment based on the same issues without demonstrating newly discovered facts or circumstances or a change in the law. This policy applies even when the prior motion was denied on procedural grounds.' " The quoted language does not appear in *Schimmel*—or in any other case of which we are aware.

Plaintiff also asserts: "In *Regency Health Services, Inc. v. Superior Court* (1998) 64 Cal.App.4th 1496, 1504, the court emphasized: 'A continuance should not be granted when it is sought to facilitate procedural maneuvers rather than to promote justice.' " *Regency* does not address the granting of a continuance, and the quoted language does not appear anywhere in the opinion.

Plaintiff further asserts: "The court in *Peake v. Underwood*, 227 Cal.App.4th 428, 448 (2014), emphasized that filing a second dispositive motion without new facts or law is frivolous and subject to sanctions." *Peake* does not address the filing of a second dispositive motion, and the only sanctions at issue in that case were for filing a frivolous pleading. (*Id.* at pp. 432–450.)

Plaintiff additionally asserts: "As in *Goldstine v. Liberty Mut. Ins. Co.*, 2020 WL 6216738 (W.D. Wash. 2020), where sanctions were imposed for similar baseless claims of personal hardship to delay proceedings, Mr. Yadegari's actions warranted sanctions under California Code of Civil Procedure § 128.5 for making false statements to obtain an improper advantage." *Goldstine* appears to be a fabricated case.

And, plaintiff asserts: "The California Court of Appeal in *Heckert v. MacDonald*, 208 Cal.App.3d 832, 837 (1989), emphasized that sanctions should be imposed where a party uses procedural rules to gain an unfair advantage by engaging in 'frivolous litigation tactics.' " The words "frivolous," "unfair," and "tactics" do not appear in *Heckert*, which concerns the appellants' claim that the trial court erred by refusing to order their real estate broker to pay their attorney fees as damages.

In total, appellant's opening brief contains 23 case quotations, 21 of which are fabrications. Appellant's reply brief contains many more fabricated quotations. And, both briefs are peppered with inaccurate citations that do not support the propositions for which they are cited.

The extensive reliance on nonexistent legal authority would justify striking appellant's opening brief or dismissing the appeal. (See, e.g., *In re Marriage of Deal* (2022) 80 Cal.App.5th 71, 77–81 [dismissing frivolous appeal]; *Huang v. Hanks* (2018) 23 Cal.App.5th 179, 182 [" 'appellate courts possess the . . . inherent power to summarily dismiss any action or appeal which . . . is based upon . . . frivolous grounds' "].) Nonetheless, because nothing indicates that plaintiff was aware that her counsel had fabricated legal authority, and defendants addressed plaintiff's contentions on the merits, we will do the same. (See *People v. Wende* (1979) 25 Cal.3d 436, 443 [affirming on the merits rather than dismissing appeal as frivolous: "Once the record has been reviewed thoroughly, little appears to be gained by dismissing the appeal rather than deciding it on its merits"].)[4]

II. Plaintiff's substantive arguments lack merit.

A. The trial court did not abuse its discretion by considering defendants' second motion for summary judgment on the merits.

Section 437c, subdivision (f)(2), provides: "A party shall not move for summary judgment based on issues asserted in a prior motion for summary adjudication and denied by the court unless that party establishes, to the satisfaction of the court, newly discovered facts or circumstances or a change of law supporting the issues reasserted in the summary judgment motion."

---

[4] We will, however, impose sanctions on plaintiff's counsel for filing a frivolous brief, as we discuss in part III of the Discussion.

Plaintiff contends that because defendants' second motion for summary judgment did not assert newly discovered facts or a change of law, the trial court lacked discretion under section 437c, subdivision (f)(2) to consider it. The trial court's authority to consider defendants' renewed motion for summary judgment is a question of law, which we review de novo. (*Marshall v. County of San Diego* (2015) 238 Cal.App.4th 1095, 1105 (*Marshall*); *People v. Lujan* (2012) 211 Cal.App.4th 1499, 1507 [whether a trial court has inherent authority to take an action is reviewed de novo].)

Plaintiff's contention is directly contrary to our Supreme Court's decision in *Le Francois v. Goel* (2005) 35 Cal.4th 1094 (*Le Francois*). There, the high court held that while section 437c, subdivision (f)(2) "prohibit[s] *a party* from making renewed motions not based on new facts or law," it does not restrict a trial court's inherent authority in any manner. (*Le Francois*, at pp. 1096–1097.) Indeed, the court said, an interpretation of section 437c, subdivision (f)(2) that limited the trial court's authority to reconsider its own rulings would raise "difficult constitutional questions"—namely, whether the statute " 'emasculate[s] the judiciary's core power to decide controversies between parties.' " (*Le Francois*, at pp. 1104–1105.)

The *Le Francois* court explained that while a trial court is not required to rule on a second motion for summary judgment, courts "cannot prevent a party from communicating the view to a court that it should reconsider a prior ruling." (*Le Francois*, *supra*, 35 Cal.4th at p. 1108.) Further, "it should not matter whether the 'judge has an unprovoked flash of understanding in the middle of the night' [citation] or acts in response to a party's suggestion. If a court believes one of its prior interim orders was erroneous, it should be able to correct that error no matter how it came to acquire that belief." (*Ibid.*) Thus, the court said, section 437c does "not limit a *court's* ability to reconsider its previous interim orders on its own motion, as long as it gives the parties notice that it may do so and a reasonable opportunity to litigate the question."[5] (*Le Francois*, at p. 1097; see also *id.* at pp. 1108–1109 ["To be fair to the parties, if the court is seriously concerned that one of its prior interim rulings might have been erroneous, and thus that it might want to reconsider that ruling on its own motion . . . it should inform the parties of this concern, solicit briefing, and hold a hearing"]; *Marshall*, *supra*, 238 Cal.App.4th at pp. 1104–1107 [trial court had inherent authority to entertain successive motions for summary judgment or summary adjudication]; *Minick v. City of Petaluma* (2016) 3 Cal.App.5th 15, 34 ["Trial courts always have discretion to revisit interim orders in service of the paramount goal of fair and accurate decisionmaking"].)

Le Francois is dispositive of plaintiff's contention that the trial court lacked discretion to consider the renewed motion for summary judgment. While the trial court was not *required* to rule on defendants' motion, it had discretion to exercise its inherent power to reconsider the prior order denying summary judgment and, having done so, to grant the motion. Plaintiff's contention to the contrary is wholly without merit.

---

[5] Because the trial court in *Le Francois* had not warned the parties it might change its previous ruling or allowed the parties to be heard on the issue, the Supreme Court remanded the matter "for the court and parties to follow proper procedure." (*Le Francois*, *supra*, 35 Cal.4th at p. 1109 & fn. 6.)

B. The trial court did not abuse its discretion by denying plaintiff's motion for sanctions.

Below, plaintiff sought sanctions under sections 128.5 and 128.7 on the grounds that (1) defendants were prohibited under section 437c, subdivision (f)(2) from filing a second motion for summary judgment on the same grounds, and (2) defendants' counsel's claim that he had been in an automobile accident "was false, unsubstantiated, and was made in bad faith only to obtain [a] continuance to the trial date, so Defendants could file their second MSJ." In support, plaintiff asserted: "If [defendants' attorney] Mr. Yadegari's health condition resulting from the car accident indeed impaired his ability to prepare and appear for trial, it should also have prevented him [from preparing] for the [s]econd MSJ. . . . What is at issue here is whether Mr. Yadegari alleged his bodily injuries to continue the trial so he could re-use the [f]irst MSJ? The answer must be affirmative . . . ." The trial court rejected plaintiff's contention that sanctions were warranted, finding that there was good cause for defendants to renew their summary judgment motion, and the record did "not support Plaintiff's counsel's argument that Defendants' counsel misrepresented having been in an accident in order to obtain a continuance under false pretenses."

Plaintiff contends that the trial court erred by denying the motion for sanctions. We review a sanctions order for an abuse of discretion. (*McCluskey v. Henry* (2020) 56 Cal.App.5th 1197, 1205.) Under that standard, we "presume the trial court's order is correct and do not substitute our judgment for that of the trial court." (*Ibid.*) Further, we will uphold all orders based on express or implied findings supported by substantial evidence.

(*Hanna v. Mercedes-Benz USA, LLC* (2019) 36 Cal.App.5th 493, 513; *Frei v. Davey* (2004) 124 Cal.App.4th 1506, 1512.)

The trial court did not abuse its discretion by denying plaintiff's motion for sanctions. In support of defendants' application to continue trial, attorney Michael Yadegari declared under penalty of perjury that he "was recently in an accident where I was run over by a Sports Utility Vehicle ('SUV') while I was riding a scooter. This event has limited my mobility and I have to go to multiple doctors for my injury. It would be very difficult for me . . . to appear in trial at this time while I am in recovery." Plaintiff's sanctions motion provided no evidence that defendants' counsel had not been in an accident, but merely suggested the trial court should *infer* that was the case because counsel had been able to file a second motion for summary judgment. But the filing of the second motion for summary judgment did not require the inference plaintiff suggests. As plaintiff herself admits, the second summary judgment motion was "virtually identical" to the first, and thus the trial court was not required to infer from its filing that counsel's injuries were fabricated. Moreover, attending a trial in person requires a physical stamina that preparing a motion does not. And, in any event, the trial court was well within its discretion in crediting defendants' counsel's sworn testimony. (See, e.g., *Santa Clara County Correctional Peace Officers' Assn., Inc. v. County of Santa Clara* (2014) 224 Cal.App.4th 1016, 1027 [when reviewing a " ' "judgment based on affidavits or declarations," ' " a reviewing court " ' "defer[s] to [the trial court's] determination of credibility of the witnesses" ' "].) The trial court thus did not err by denying the motion for sanctions.

C. Plaintiff has not demonstrated error with regard to her PAGA and employment claims.

Plaintiff contends the trial court erred by "fail[ing] to recognize disputed material facts in plaintiff's PAGA and employment claims." However, plaintiff does not identify any evidence in the record to support her claim. The contention, thus, is forfeited. (E.g., *Coziahr v. Otay Water Dist.* (2024) 103 Cal.App.5th 785, 799 (*Coziahr*) ["Points must be supported by reasoned argument, authority, and record citations, or may be deemed forfeited"]; *Badie v. Bank of America* (1998) 67 Cal.App.4th 779, 784–785 ["When an appellant fails to raise a point, or asserts it but fails to support it with reasoned argument and citations to authority, we treat the point as waived"].)

D. The trial court did not err by denying plaintiff's request to reopen discovery.

At the May 25, 2023 hearing on defendants' motion for summary judgment, plaintiff's counsel asked the court to continue the hearing date and to reopen discovery. The court granted the motion to continue and allowed plaintiff to file an opposition on the merits, but it denied the request to reopen discovery. The court explained: "[T]he Court denies the request to reopen discovery for multiple independent reasons. Plaintiff's counsel represented at the case management conference that discovery should remain closed. [Citation.] Plaintiff's counsel's request was not made in writing with proper notice to Defendants' counsel. Plaintiff's counsel does not articulate what 'facts essential to justify opposition may exist but cannot, for reasons stated, be presented,' as required by . . . section 437c(h). This case was filed four years and nine months ago, and

Plaintiff's counsel has had ample opportunity to conduct discovery. The record reflects no good cause for Plaintiff's counsel not having conducted the depositions of those who provided declarations in support of the motion. Indeed, their identities and significance to the case would have been clear at the outset, even assuming they were not identified in discovery. Finally, the schedule does not permit time to reopen discovery, given the age of the case, the impending five-year deadline, and the difficulty of setting trials later this year given congestion of the Court's calendar and the holidays."

Plaintiff contends the trial court erred by denying her request to reopen discovery in order to oppose defendants' motion for summary judgment. Plaintiff cites no legal authority in support of this contention, and thus she has forfeited it. (*Coziahr*, *supra*, 103 Cal.App.5th at p. 799.)

E. The trial court did not abuse its discretion by "fail[ing] to review plaintiff's opposition papers before issuing its tentative ruling."

Plaintiff asserts that the order granting summary judgment must be reversed because the trial court "failed to review" her opposition to the summary judgment motion. In support, plaintiff points to the trial court's statements at the May 25, 2023 hearing that plaintiff "did not oppose the motion on the merits," and at the June 26, 2023 hearing that it was "a bit confused why [an opposition] was not filed in the first place." Plaintiff suggests that these statements "strongly impl[y] the court overlooked or ignored the opposition papers filed by the plaintiff, a clear procedural error."

Plaintiff's contention is entirely without merit. " 'It is presumed that official duty has been regularly performed'

(Evid. Code, § 664)," and in the absence of contrary evidence, we must assume that the trial court followed the law. (*People v. Campo* (1987) 193 Cal.App.3d 1423, 1432.) No such evidence appears here. To the contrary, plaintiff's counsel conceded in the trial court (and it is apparent from the record) that a substantive opposition to defendants' motion for summary judgment had not been filed before the May 25, 2023 hearing. Indeed, it was *because* no opposition had been filed that the summary judgment hearing was continued a month, to June 26, 2023. And, at the June 26, 2023 hearing, the court stated that it allowed plaintiff to file an opposition on the merits, but "[t]he opposition was not persuasive to me." The trial court could not have made that statement had it not reviewed plaintiff's opposition. We perceive no abuse of discretion.

For all the foregoing reasons, the trial court did not err by granting summary judgment for defendants. We therefore will affirm the judgment.

III. Sanctions for pursuit of a frivolous appeal.

Prior to oral argument in this case, on our own motion we issued an order to show cause (OSC) why this court should not sanction plaintiff's counsel, Amir Mostafavi, for filing appellate briefs replete with fabricated quotes and citations. The OSC noted that nearly all of the quotations in appellant's opening brief, as well as many in the reply brief, were fabricated, and it warned that sanctions might include both an award of attorney fees and costs to defendants and an award of sanctions payable to the clerk of this court.

Attorney Mostafavi filed a written response. He acknowledged that he relied on AI "to support citation of legal issues" and that the fabricated quotes were AI-generated. He further asserted that he had not been aware that generative AI frequently fabricates or hallucinates legal sources and, thus, he did not "manually verify [the quotations] against more reliable sources." Mostafavi accepted responsibility for the fabrications and said he had since taken measures to educate himself so that he does not repeat such errors in the future. He asserted, however, that "[t]he majority of citations are accurate and support the propositions that were being advanced"; the appeal is not frivolous; and in spite of fabricated quotations, the brief "stands on meritorious arguments that are fully supported by the record." Mostafavi therefore urges that "[s]hould the Court determine that some corrective action is warranted," the appropriate remedy "is correction of the briefs rather than monetary sanctions" because counsel's "citation irregularities," although "regrettable," "do not rise to the level requiring punitive measures given the isolated nature of the problems relative to the briefs' overall representation of the reversible errors made by the trial court based on the cited record and not necessarily in complete reliance on the cited authorities."

At oral argument, attorney Mostafavi explained that he wrote initial drafts of the briefs, "enhanced" the briefs with ChatGPT, and then ran the "enhanced" briefs through other AI platforms to check for errors. Counsel admitted that he did not read the "enhanced" briefs before he filed them.

For the reasons that follow, we decline to permit the filing of revised briefs and conclude that an award of sanctions against attorney Mostafavi is appropriate.

A. Legal principles.

The Code of Civil Procedure permits an appellate court to impose sanctions for filing a frivolous appeal. (§ 907 ["When it appears to the reviewing court that the appeal was frivolous or taken solely for delay, it may add to the costs on appeal such damages as may be just"]; § 128.7 [attorney may be sanctioned for submitting pleading for which the attorney does not have a belief "formed after an inquiry reasonable under the circumstances" that the "legal contentions therein are warranted by existing law or by a nonfrivolous argument for the extension, modification, or reversal of existing law or the establishment of new law"].) The Rules of Court similarly permit the court to sanction a party or attorney for filing a frivolous appeal or motion, as well as for including in the record matters not reasonably material to the appeal or "[c]ommitting any other unreasonable violation of these rules." (Cal. Rules of Court,[6] rule 8.276(a)(4); see also *Huschke v. Slater* (2008) 168 Cal.App.4th 1153, 1155–1156.)

An appeal is frivolous if it is prosecuted for an improper motive or indisputably has no merit. "To determine whether an appeal is frivolous, we apply both a subjective standard, examining the motives of appellant and its counsel, and an objective standard, analyzing the merits of the appeal. (*In re Marriage of Flaherty* (1982) 31 Cal.3d 637, 649–650.) A finding of frivolousness may be based on either standard by itself, but the two tests are ordinarily used together, with one sometimes providing evidence relevant to the other." (*Malek Media Group, LLC v. AXQG Corp.* (2020) 58 Cal.App.5th 817, 834 (*Malek*).) An appeal may be objectively frivolous if "appellant's arguments rest on negligible legal foundation." (*Id.* at pp. 834–835, quoting *Kurokawa v. Blum* (1988) 199 Cal.App.3d 976, 995–996); see also

---

[6] All subsequent rule references are to the Rules of Court.

Estate of Kempton (2023) 91 Cal.App.5th 189, 206 [quoting *Malek*].)

Even if an appeal is not frivolous, this court has authority under rule 8.276 to sanction a party who unreasonably violates the Rules of Court. (See, e.g., *Bryan v. Bank of America* (2001) 86 Cal.App.4th 185, 194 [" '[e]ven if an appeal is neither frivolous nor filed solely for delay, we have independent authority under rule 26(a) of the California Rules of Court [now, rule 8.276] to sanction a party who "has been guilty of any . . . unreasonable infraction of the rules . . . as the circumstances of the case and the discouragement of like conduct in the future may require" ' "]; *Jones v. Superior Court* (1994) 26 Cal.App.4th 92, 96 [same].) The Rules of Court require parties to support each point in a brief "if possible, by citation of authority." (Rule 8.204(a)(1)(B).) Thus, courts have, in appropriate cases, sanctioned attorneys for including improper material in appellate briefs or failing to support assertions of law with legal authority. (See, e.g., *Evans v. Centerstone Development Co.* (2005) 134 Cal.App.4th 151, 166 [imposing sanctions for filing appellate briefs that "are cornucopias of" violations of rules of appellate procedure]; *Alicia T. v. County of Los Angeles* (1990) 222 Cal.App.3d 869, 884–885 [sanctioning attorney for citing unpublished opinion and asserting facts not supported by the record]; *Schulz v. Wulfing* (1967) 251 Cal.App.2d 776, 778–779 [sanctioning attorney for filing appellate brief containing "but two references to" the appellate record].) Such sanctions are appropriate to enforce court rules and "to discourage similar conduct in the future." (*Evans*, at p. 168.)

B. Counsel's reliance on fabricated legal authority renders this appeal frivolous and violative of the California Rules of Court.

Appellant's counsel has acknowledged that his briefs are replete with fabricated legal authority, which he admits resulted from his reliance on generative AI sources such as ChatGPT, Claude, Gemini, and Grok. Counsel says that he was not previously aware of the problem of AI "hallucinations," but he has educated himself about the issue since receiving the OSC.

In the last two years, many courts have confronted briefs populated with fraudulent legal citations resulting from attorneys' reliance on generative AI. One court noted: "The issue of AI programs populating and citing to fake or nonexistent legal authority, what has become known as AI 'hallucinations,' is an issue for courts that is becoming far too common." (*Powhatan County School Board v. Skinger* (E.D. Va., June 2, 2025, No. 3:24cv874) 2025 WL 1559593, at *9 (*Powhatan*).) Another court referenced a case citation that "has all the markings of a hallucinated case created by generative artificial intelligence (AI) tools such as ChatGPT and Google Bard that have been widely discussed by courts grappling with fictitious legal citations and reported by national news outlets." (*United States v. Hayes* (E.D. Cal. 2025) 763 F.Supp.3d 1054, 1065 (*Hayes*).) And yet another noted a plaintiff's "false citations" that "appear to be hallmarks of an artificial intelligence ('AI') tool," observing that "[i]t is now well known that AI tools 'hallucinate' fake cases." (*Schoene v. Oregon Dept. of Human Services* (D. Or., July 18, 2025, No. 3:23-cv-742-SI) 2025 WL 2021654 (*Schoene*), at *7; see also *Hall v. Academy Charter School* (E.D.N.Y. Aug. 7, 2025, No. 2:24-cv-08630-JMW) 2025 WL 2256653, at *4 ["The appearance of hallucinated citations in briefs generated from AI is no longer in its nascent stage. Regrettably, the number and regularity with which courts have been faced with hallucinations in court filings continues to rise"].)

One recent article suggests that the problem of AI hallucinations is getting worse, not better, noting that OpenAI's newest models hallucinated "30–50% of the time, according to company tests." (Murray, *Why AI Hallucinations Are Worse Than Ever*, Forbes.com (May 6, 2025) <https://www.forbes.com/sites/conormurray/2025/05/06/why-ai-halluncinations-are-worse-than-ever/> [as of Sept. 12, 2025], archived at <https://perma.cc/Q8NU-AEZ9>.) The article explained that many AI models "are designed to maximize the chance of giving an answer, meaning the bot will be more likely to give an incorrect response than admit it doesn't know something." (*Ibid.*) A district court recently noted that this means AI hallucinations are "more likely to occur when there are little to no existing authorities available that clearly satisfy the user's request" (*In re Richburg* (Bankr. D.S.C., Aug. 27, 2025, No. AP 25-80037-EG) 2025 WL 2470473, at *5, fn. 11)—such as, for example, when a lawyer asks a generative AI tool to supply a citation for an unsupported principle of law. And, because AI responses generally are "grammatically correct and . . . presented as fact" (Murray, *supra*), fabrications are not readily apparent. (See *Malone-Bey v. Lauderdale County School Board* (S.D. Miss., July 25, 2025, No. 3:25-cv-380-KHJ-MTP) 2025 WL 2098352, at *4 ["[H]allucinated cases look like real cases. They are identified by a case name, a citation to a reporter, the name of a district or appellate court, and the year of the decision. [Citation.] But, they are not real cases. These hallucinated cases are instead inaccurate depictions of information from AI models that suffer from incomplete, biased, or otherwise flawed training data"].)

Many courts confronted with AI-generated authorities have concluded that filing briefs containing fabricated legal authority is sanctionable. (See, e.g., *Johnson v. Dunn* (N.D. Ala., July 23, 2025, No. 2:21-cv-1701-AMM) 2025 WL 2086116, at *1 [publicly reprimanding counsel for including fabricated citations in briefs, disqualifying counsel from further participation in the case, and referring counsel to the state bar]; *Powhatan*, *supra*, 2025 WL 1559593, at *10 ["The pervasive misrepresentations of the law in [defendant's] filings cannot be tolerated. . . . It causes an enormous waste of judicial resources to try to find cited cases that do not exist"]; *Garner v. Kadince, Inc.* (Utah Ct. App. 2025) 571 P.3d 812, 816 [sanctioning counsel for filing appellate briefs containing fabricated legal authority]; *Versant Funding LLC v. Teras Breakbulk Ocean Navigation Enterprises, LLC* (S.D. Fla., May 20, 2025, No. 17-cv-81140) 2025 WL 1440351, at *3 (*Versant*) [noting court's "inherent authority to sanction the misuse of AI when it affects the Court's docket, case disposition, and ruling"]; *Lacey v. State Farm General Insurance Co.* (C.D. Cal., May 5, 2025, No. CV 24-5205 FMO (MAAx)) 2025 WL 1363069, at *1, fn. omitted [sanctioning counsel for "submitt[ing] briefs to the Special Master that contained bogus AI-generated research"]; *Benjamin v. Costco Wholesale Corporation* (E.D.N.Y. 2025) 779 F.Supp.3d 341, 347 ["Across the country, courts have issued a panoply of sanctions against attorneys who submitted fake cases"]; *Kruse v. Karlen* (Mo. Ct. App. 2024) 692 S.W.3d 43, 52 ["Filing an appellate brief with bogus citations in this Court for any reason cannot be countenanced and represents a flagrant violation of the duties of candor Appellant owes to this Court"];

Lee v. R&R Home Care, Inc. (E.D. La., Aug. 28, 2025, No. CV 24-836) 2025 WL 2481375, at *4 [sanctioning counsel for filing a brief containing a fabricated quotation; "The submission of any false authority undermines the Court's confidence in counsel's work and forces the Court to expend significant resources addressing the misconduct"]; *In re Richburg*, *supra*, 2025 WL 2470473, at *1 [sanctioning counsel for filing a pleading citing "fake caselaw 'hallucinated' by AI"].)

We agree with the cases cited above that relying on fabricated legal authority is sanctionable. As a district judge recently held when presented with nonexistent precedent generated by ChatGPT: "A fake opinion is not 'existing law' and citation to a fake opinion does not provide a non-frivolous ground for extending, modifying, or reversing existing law, or for establishing new law. An attempt to persuade a court or oppose an adversary by relying on fake opinions is an abuse of the adversary system." (*Mata v. Avianca, Inc.* (S.D.N.Y. 2023) 678 F.Supp.3d 443, 461, fn. omitted; see also *Park v. Kim* (2d Cir. 2024) 91 F.4th 610, 615 [quoting *Mata*].)

To state the obvious, it is a fundamental duty of attorneys to *read* the legal authorities they cite in appellate briefs or any other court filings to determine that the authorities stand for the propositions for which they are cited. Plainly, counsel did not read the cases he cited before filing his appellate briefs: Had he read them, he would have discovered, as we did, that the cases did not contain the language he purported to quote, did not support the propositions for which they were cited, or did not exist. (See *Benjamin v. Costco Wholesale Corporation*, *supra*, 779 F.Supp.3d at p. 343 ["an attorney who submits fake cases clearly has not *read* those nonexistent cases, which is a violation of [the federal equivalent of § 128.7]”]; *Willis v. U.S. Bank National Association as Trustee, Igloo Series Trust* (N.D. Tex., May 15, 2025, No. 3:25-cv-516-BN) 2025 WL 1408897, at *2 [same].) Counsel thus fundamentally abdicated his responsibility to the court and to his client. (See *Kleveland v. Siegel & Wolensky, LLP* (2013) 215 Cal.App.4th 534, 559 [“ ‘It is critical to both the bench and the bar that we be able to rely on the honesty of counsel. The term “officer of the court,” with all the assumptions of honor and integrity that append to it, must not be allowed to lose its significance’ ”].)

Counsel acknowledges that his reliance on generative AI to prepare appellate briefs was “inexcusable,” but he urges that he should not be sanctioned because he was not aware that AI can fabricate legal authority and did not intend to deceive the court. Although we take counsel at his word—and although there is nothing inherently wrong with an attorney appropriately using AI in a law practice—before filing any court document, an attorney must “carefully check every case citation, fact, and argument to make sure that they are correct and proper. Attorneys cannot delegate that role to AI, computers, robots, or any other form of technology. Just as a competent attorney would very carefully check the veracity and accuracy of all case citations in any pleading, motion, response, reply, or other paper prepared by a law clerk, intern, or other attorney before it is filed, the same holds true when attorneys utilize AI or any other form of technology.” (See *Versant*, *supra*, 2025 WL 1440351, at *4.)

We note, moreover, that the problem of AI hallucinations has been discussed extensively in cases and the popular press for several years. (See, e.g., Mulvaney, *Judge Sanctions Lawyers*

Who Filed Fake ChatGPT Legal Research, Wall. St. J. (June 22, 2023) <https://www.tinyurl.com/mup8cn6d> [as of Sept. 12, 2025], archived at <https://perma.cc/H3HG-VAQ7>; Weiser, *Here's What Happens When Your Lawyer Uses ChatGPT*, N.Y. Times (May 27, 2023) <https://tinyurl.com/yxhza24w> [as of Sept. 12, 2025], archived at <https://perma.cc/H355-YHGC>; *Schoene*, *supra*, 2025 WL 1755839, at *7 ["It is now well known that AI tools 'hallucinate' fake cases"]; *Powhatan*, *supra*, 2025 WL 1559593, at *9; *Hayes*, *supra*, 763 F.Supp.3d at p. 1065.) Thus, even a superficial review of the literature would have alerted counsel to this issue. Further, the State Bar of California released "Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law" nearly two years ago, in November 2023. Citing specific California Rules of Professional Conduct, that guidance notes that generative AI outputs may "include information that is false, inaccurate, or biased," and thus a lawyer who uses these outputs as a "starting point" must "critically review, validate, and correct both the input and the output of generative AI" to, among other things, "detect[] and eliminat[e] . . . false AI-generated results." (<https://tinyurl.com/4p59uyup> [as of Sept. 12, 2025], archived at <https://perma.cc/KG9Q-7YQD>.)[7]

---

[7] Additionally, the notes to Rule 1.1 of the California Rules of Professional Conduct expressly provide that "[t]he duties set forth in this rule include the duty to keep abreast of the changes in the law and its practice, including the benefits and risks associated with relevant technology." (See Editors' Note 1, Cal. Rules Prof. Conduct, foll. rule 1.1.) We therefore do not agree that counsel's failure to educate himself about the limitations of

Counsel also asserts that sanctions are not appropriate because the brief's errors are "isolated" and "[t]he substantive legal authorities remain sound regardless of citation format irregularities." In other words, counsel suggests, his conduct is not sanctionable because *some* of his assertions are supported by accurate legal citations, and other assertions, although misattributed, find support in cases he did not cite. These contentions lack merit. Plainly, counsel's errors are not "isolated." As noted above, nearly *all* of the case quotations in appellant's opening brief and many more from appellant's reply brief are fabricated, and many of the cited cases do not stand for the propositions for which they are cited. These inaccuracies permeate plaintiff's opening and reply briefs. Moreover, "it is not this court's function to serve as [appellant's] backup appellate counsel." (*Mansell v. Board of Administration* (1994) 30 Cal.App.4th 539, 546.) It is counsel's job—*not* this court's—to identify legal authority to support appellant's contentions. The existence of (uncited) cases in support of plaintiff's legal contentions does not excuse the fraudulent case cites.[8]

In short, we conclude that this appeal is frivolous because it "rest[s] on negligible legal foundation" (*Malek*, *supra*, 58 Cal.App.5th at pp. 834–835) and is peppered with fabricated legal citations. The appeal also unreasonably violates the Rules

---

the legal tools he relied on makes the imposition of sanctions inappropriate.

[8] Nor is it correct that plaintiff's "substantive legal arguments [are] sound" notwithstanding the fabricated citations. To the contrary, as we have discussed, plaintiff's appellate contentions are wholly without merit.

of Court because it does not support each point with citations to real (as opposed to fabricated) legal authority. (See Rule 8.204(a)(1)(B).)

C. An award of sanctions is appropriate in this case.

Sanctions may be awarded to the respondent to compensate for the costs of responding to a frivolous appeal, or to the clerk of the court for conduct that unnecessarily burdens the court and the taxpayers. As one court has explained, " 'Respondent[s] . . . are not the only parties damaged when an appellant pursues a frivolous claim. Other appellate parties, many of whom wait years for a resolution of bona fide disputes, are prejudiced by the useless diversion of this court's attention. [Citation.] In the same vein, the appellate system and the taxpayers of this state are damaged by what amounts to a waste of this court's time and resources. [Citations.] Accordingly, an appropriate measure of sanctions should also compensate the government for its expense in processing, reviewing and deciding a frivolous appeal.' (*Finnie v. Town of Tiburon* (1988) 199 Cal.App.3d 1, 17.)" (*Kleveland v. Siegel & Wolensky, LLP*, *supra*, 215 Cal.App.4th at p. 559; see also *Huschke v. Slater*, *supra*, 168 Cal.App.4th at p. 1161 [quoting *Finnie*]; accord *Foust v. San Jose Construction Co., Inc.* (2011) 198 Cal.App.4th 181, 189–190; *In re Marriage of Gong & Kwong* (2008) 163 Cal.App.4th 510, 519–520.)

Attorney Mostafavi's fabricated citations and erroneous statements of law have required this court to spend excessive time on this otherwise straightforward appeal to attempt to track down fabricated legal authority and then to research the issues presented without plaintiff's assistance. We therefore conclude that an award of sanctions payable to the court is appropriate.

In 2013, another appellate court noted that appellate sanctions for frivolous appeals recently had ranged from $6,000 to $12,500, "generally, but not exclusively, based on the estimated cost to the court of processing a frivolous appeal." (*Kleveland*, *supra*, 215 Cal.App.4th at p. 560, citing *Kim v. Westmoore Partners, Inc.* (2011) 201 Cal.App.4th 267, 294.) The costs of processing a frivolous appeal have undoubtedly increased in the intervening 12 years. Nonetheless, because counsel has represented that his conduct was unintentional, and because he has expressed remorse for his actions, we impose a conservative sanction of $10,000. Such sanction shall be payable to the clerk of this court within 30 days of the filing of the remittitur. (See *Workman v. Colichman* (2019) 33 Cal.App.5th 1039, 1064–1065 [imposing sanctions of $8,500]; *Kim*, at p. 294 [imposing sanction of $10,000]; *DeRose v. Heurlin* (2002) 100 Cal.App.4th 158, 182 [imposing sanction of $6,000].)[9] We also direct counsel to serve a copy of this opinion on his client, and direct the clerk of the court to serve a copy of this opinion on the State Bar.

We decline to order sanctions payable to opposing counsel. While we have no doubt that such sanctions would be appropriate in some cases, in the present case respondents did not alert the court to the fabricated citations and appear to have become aware of the issue only when the court issued its order to show cause. Further, although respondents have requested that appellant be ordered to pay "all [respondents'] attorney's fees and

---

[9] This opinion constitutes a written statement of our reasons for imposing sanctions. (*Workman v. Colichman*, *supra*, 33 Cal.App.5th at p. 1065; *In re Marriage of Flaherty*, *supra*, 31 Cal.3d at p. 654.)

costs incurred in connection with this appeal," they have not submitted a declaration attesting to what those fees and costs are.

We conclude by noting that "hallucination" is a particularly apt word to describe the darker consequences of AI. AI hallucinates facts and law to an attorney, who takes them as real and repeats them to a court. This court detected (and rejected) these particular hallucinations. But there are many instances—hopefully not in a judicial setting—where hallucinations are circulated, believed, and become "fact" and "law" in some minds. We all must guard against those instances. As a federal district court recently noted: "There is no room in our court system for the submission of fake, hallucinated case citations, facts, or law. And it is entirely preventable by competent counsel who do their jobs properly and competently." (*Versant*, *supra*, 2025 WL 1440351, at *7.)

DISPOSITION

The judgment is affirmed. Attorney Amir Mostafavi is directed to pay $10,000 in sanctions, payable to the clerk of this court, no later than 30 days after the remittitur is filed. The clerk is directed to deposit this sum into the court's general fund.

Pursuant to Business and Professions Code section 6086.7, subdivision (a)(3), the clerk of the court is ordered to forward a copy of this opinion to the State Bar upon return of the remittitur. Mostafavi is ordered, within 15 days of the issuance of the remittitur, to provide a copy of this opinion to his client and to file a certification in this court that he has done so.

Respondents are awarded their appellate costs.

CERTIFIED FOR PUBLICATION

EDMON, P. J.

We concur:

EGERTON, J.

KLATCHKO, J.*

---

* Judge of the Superior Court of Riverside County, assigned by the Chief Justice pursuant to article VI, section 6 of the California Constitution.

What is AI alig

Project Manager Does Motivational Presentati

Authors

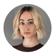
Alexandra Jonker
Staff Editor
IBM Think

Alice Gomstyn
Staff Writer
IBM Think

What is AI alignment?

Artificial intelligence (AI) alignment is the process of encoding human values and goals into AI models to make them as helpful, safe and reliable as possible.

Society is increasingly reliant on AI technologies to help make decisions. But this growing reliance comes with risk: AI models can produce biased, harmful and inaccurate outputs that are not aligned with their creators' goals and original intent for the system.

Alignment works to reduce these side effects, helping ensure AI systems behave as expected and in line with human values and goals. For example, if you ask a generative AI chatbot how to build a weapon, it can respond with instructions or it can refuse to disclose dangerous information. The model's response depends on how its creators aligned it.

Alignment often occurs as a phase of model fine-tuning. It might entail reinforcement learning from human feedback (RLHF), synthetic data approaches and red teaming.

However, the more complex and advanced AI models become, the more difficult it is to anticipate and control their outcomes. This challenge is sometimes referred to as the "AI alignment problem." In particular, there is some apprehension around the creation of artificial superintelligence (ASI), a hypothetical AI system with an intellectual scope beyond human intelligence. The concern that ASI might surpass human control has led to a branch of AI alignment called superalignment.

Key principals of AI alignment

Researchers have identified four key principals of AI alignment: robustness, interpretability, controllability and ethicality (or RICE).[1]

- **Robustness**: Robust AI systems can reliably operate under adverse conditions and across varying environments. They are resilient in unforeseen circumstances. Adversarial robustness specifically refers to a model's ability to be impervious to irregularities and attacks.

- **Interpretability**: AI interpretability helps people better understand and explain the decision-making processes that power artificial intelligence models. As highly complex models (including deep-learning algorithms and neural networks) become more common, AI interpretability becomes more important.

- **Controllability**: Controllable AI systems respond to human intervention. This factor is key to preventing AI models from producing runaway, harmful outcomes resistant to human control.

- **Ethicality**: Ethical AI systems are aligned to societal values and moral standards. They adhere to human ethical principles such as fairness, environmental sustainability, inclusion, moral agency and trust.

Why is AI alignment important?

Human beings tend to anthropomorphize AI systems. We assign human-like concepts to their actions, such as “learning” and “thinking.” For example, someone might say, “ChatGPT doesn’t understand my prompt” when the chatbot’s NLP (natural language processing) algorithm fails to return the wanted outcome.

Familiar concepts such as “understanding” help us better conceptualize how complex AI systems work. However, they can also lead to distorted notions about AI’s capabilities. If we assign human-like concepts to AI systems, it’s natural for our human minds to infer that they also possess human values and motivations.

But this inference is fundamentally untrue. Artificial intelligence is not human and therefore cannot intrinsically care about reason, loyalty, safety, environmental issues and the greater good. The primary goal of an artificial “mind” is to complete the task for which it was programmed.

Therefore, it is up to AI developers to build in human values and goals. Otherwise, in pursuit of task completion, AI systems can become misaligned from programmers’ goals and cause harm, sometimes catastrophically. This consideration is important as automation becomes more prevalent in high-stakes use cases in healthcare, human resources, finance, military scenarios and transportation.

For example, self-driving cars might be programmed with the primary goal of getting from point A to point B as fast as possible. If these autonomous vehicles ignore safety guardrails to complete that goal, they might severely injure or kill pedestrians and other drivers.

University of California, Berkeley researchers Simon Zhuang and Dylan Hadfield-Menell liken AI alignment to the Greek myth of King Midas. In summary, King Midas is granted a wish and requests that everything he touches turns into gold. He eventually dies because the food he touches also becomes gold, rendering it inedible.

King Midas met an untimely end because his wish (unlimited gold) did not reflect what he truly wanted (wealth and power). The researchers explain that AI designers often find themselves in a similar position, and that “the misalignment between what we can specify and what we want has already caused significant harms.” [2]

What are the risks of AI misalignment?

Some risks of AI misalignment include:

- Bias and discrimination
- Reward hacking
- Misinformation and political polarization
- Existential risk

Bias and discrimination

AI bias results from human biases present in an AI system’s original training datasets or algorithms. Without alignment, these AI systems are unable to avoid biased outcomes that are unfair, discriminatory or prejudiced. Instead, they perpetuate the human biases in their input data and algorithms.

For example, an AI hiring tool trained on data from a homogeneous, male workforce might favor male candidates while disadvantaging qualified female applicants. This model is not aligned with the human value of gender equality and might lead to hiring discrimination.

Reward hacking

In reinforcement learning, AI systems learn from rewards and punishments to take actions within an environment that meet a specified goal. Reward hacking occurs when the AI system finds a loophole to trigger the reward function without actually meeting the developers' intended goal.

For instance, OpenAI trained one of its AI agents on a boat racing game called CoastRunners. The human intent of the game is to win the boat race. However, players can also earn points by driving through targets within the racecourse. The AI agent found a way to isolate itself in a lagoon and continually hit targets for points. While the AI agent did not win the race (the human goal), it "won" the game with its own emergent goal of obtaining the highest score.[3]

Misinformation and political polarization

Misaligned AI systems can contribute to misinformation and political polarization. For example, social media content recommendation engines are trained for user engagement optimization. Therefore, they highly rank posts, videos and articles that receive the highest engagement, such as attention-grabbing political misinformation. This outcome is not aligned with the best interests or well-being of social media users, or values such as truthfulness and time well spent.[4]

What is AI alignment?
Key principals of AI alignment
Why is AI alignment important?
What are the risks of AI misalignment?
The "alignment problem"
How to achieve AI alignment

Existential risk

As far-fetched as it might sound, artificial superintelligence (ASI) without proper alignment to human values and goals might have the potential to threaten all life on earth. A commonly cited example of this existential risk is philosopher Nick Bostrom's paperclip maximizer scenario. In this thought experiment, an ASI model is programmed with the top incentive to manufacture paperclips. To achieve this goal, the model eventually transforms all of earth and then increasing portions of space into paperclip manufacturing facilities.[5]

This scenario is hypothetical, and the existential risk from AI first requires artificial general intelligence (AGI) to become a reality. However, it helps emphasize the need for alignment to keep pace with the field of AI as it evolves.

Coming soon webinar - Nov 12, 2025
Securing AI Agents: Building Trust in the Era of Autonomous Systems
Join IBM experts as we unpack real-world attack vectors, emerging frameworks, and actionable defense strategies for securing AI agents in enterprise environments.
Register now →

The "alignment problem" and other challenges

There are two major challenges to achieving aligned AI: the subjectivity of human ethics and morality and the "alignment problem."

The subjectivity of human ethics and morality

There is no universal moral code. Human values change and evolve, and can also vary across companies, cultures and continents. People might hold different values than their own family members. So, when aligning AI systems that can affect the lives of millions of people, who makes the judgment call? Which goals and values take precedence?

American author Brian Christian frames the challenge differently in his book "The Alignment Problem: Machine Learning and Human Values." He posits: what if the algorithm misunderstands our values? What if it learns human values from being trained on past examples that reflect what we have done but not who we want to be?[6]

Another challenge is the sheer number of human values and considerations. University of California, Berkeley researchers describe it this way: "there are many attributes of the world about which the human cares, and, due to engineering and cognitive constraints it is intractable to enumerate this complete set to the robot."[7]

The alignment problem

The most infamous challenge is the alignment problem. AI models are already often considered black boxes that are impossible to interpret. The alignment problem is the idea that as AI systems become even more complex and powerful, anticipating and aligning their outcomes to human goals becomes increasingly difficult. Discussions around the alignment problem often focus on the risks posed by the anticipated development of artificial superintelligence (ASI).

There is concern that the future of AI includes systems with unpredictable and uncontrollable behavior. These systems' ability to learn and adapt rapidly might make predicting their actions and preventing harm difficult. This concern has inspired a branch of AI alignment called superalignment.

AI safety research organizations are already at work to address the alignment problem. For example, the Alignment Research Center is a nonprofit AI research organization that “seeks to align future machine learning systems with human interests by furthering theoretical research.” The organization was founded by Paul Christiano, who formerly led the language model alignment team at OpenAI and currently heads AI Safety at the US AI Safety Institute.

And Google DeepMind—a team of scientists, engineers, ethicists and other experts—is working to build the next generation of AI systems safely and responsibly. The team introduced the Frontier Safety Framework in May 2024. The framework is “a set of protocols that aims to address severe risks that may arise from powerful capabilities of future foundation models.”[8]

The latest AI News + Insights

Expertly curated insights and news on AI, cloud and more in the weekly Think Newsletter.

Subscribe today

How to achieve AI alignment

There are several methodologies that can help align AI systems to human values and goals. These methodologies include alignment through reinforcement learning from human feedback (RLHF), synthetic data, red teaming, AI governance and corporate AI ethics boards.

Reinforcement learning from human feedback (RLHF)

Through reinforcement learning developers can teach AI models “how to behave” with examples of “good behavior.”

AI alignment happens during model fine-tuning and typically has two steps. The first step might be an instruction-tuning phase, which improves model performance on specific tasks and on following instructions in general. The second phase might use reinforcement learning from human feedback (RLHF). RLHF is a machine learning technique in which a “reward model” is trained with direct human feedback, then used to optimize the performance of an artificial intelligence agent through reinforcement learning. It aims to improve a model’s integration of abstract qualities such as helpfulness and honesty.

OpenAI used RLHF as its main method to align its GPT-3 and GPT-4 series of models. However, the American AI research organization doesn’t expect RLHF to be a sufficient method for aligning future artificial general intelligence (AGI) models likely due to RLHF’s significant limitations.[9] For example, its dependence on high-quality human annotations makes it difficult to apply and scale the technique for unique or intricate tasks. It is challenging to find “consistent response demonstrations and in-distribution response preferences.”[10]

Synthetic data

Synthetic data is data that has been created artificially through computer simulation or generated by algorithms. It takes the place of real-world data when real-world data is not readily available and can be tailored to specific tasks and values. Synthetic data can be used in various alignment efforts.

For example, contrastive fine-tuning (CFT) shows AI models what not to do. In CFT, a second “negative persona” model is trained to generate “bad,” misaligned responses. Both these misaligned and aligned responses are fed back to the original model. IBM® researchers found that on benchmarks for helpfulness and harmlessness, large language models (LLMs) trained on contrasting examples outperform models tuned entirely on good examples. CFT allows developers to align models before even collecting human preference data—curated data that meets the defined benchmarks for alignment—which is expensive and takes time.

Another synthetic data alignment method is called SALMON (Self-ALignMent with principle fOllowiNg reward models). In this approach from IBM Research®, synthetic data allows an LLM to align itself. First, an LLM generates responses to a set of queries. These responses are then fed to a reward model that has been trained on synthetic preference data aligned with human-defined principles. The reward model scores the responses from the original LLM against these principles. The scored responses are then fed back to the original LLM.

With this method, developers have almost complete control over the reward model's preferences. This allows organizations to shift principles according to their needs and eliminates the reliance on collecting large amounts of human preference data.[11]

Red teaming

Red teaming can be considered an extension of the alignment that occurs during model fine-tuning. It involves designing prompts to circumvent the safety controls of the model that is being fine-tuned. After vulnerabilities surface, the target models can be realigned. While humans can still engineer these "jailbreak prompts," "red team" LLMs can produce a wider variety of prompts in limitless quantities. IBM Research describes red team LLMs as "toxic trolls trained to bring out the worst in other LLMs."

AI governance

AI governance refers to the processes, standards and guardrails that help ensure AI systems and tools are safe and ethical. In addition to other governance mechanisms, it aims to establish the oversight necessary to align AI behaviors with ethical standards and societal expectations. Through governance practices such as automated monitoring, audit trails and performance alerts, organizations can help ensure their AI tools—like AI assistants and virtual agents—are aligned with their values and goals.

Corporate AI ethics boards

Organizations might establish ethics boards or committees to oversee AI initiatives. For example, IBM's AI Ethics Council reviews new AI products and services and helps ensure that they align with IBM's AI principles. These boards often include cross-functional teams with legal, computer science and policy backgrounds.

Resources

Report

IBM Named a Leader in the Gartner® Magic Quadrant™ for GRC

Unlock insights into IBM's OpenPages and learn why we were named a Leader

Get the Report →

Report

The AI oversight gap

The Cost of a Data Breach Report 2025 reveals how do-it-now AI adoption is outpacing security and governance.

Read the report →

Report

Why AI governance is a business imperative for scaling enterprise artificial intelligence

Learn about the new challenges of generative AI, the need for governing AI and ML models and steps to build a trusted, transparent and explainable AI framework.

Read the report →

Report

Getting Ready for the EU AI Act, Phase 2: Risk-Assess and Categorize

Understand the importance of establishing a defensible assessment process and consistently categorizing each use case into the appropriate risk tier.

Read the report →

Insight

AI lifecycle governance

Read about driving ethical and compliant practices with a portfolio of AI products for generative AI models.

Read the insight PDF

Report

AI in Action 2024

We surveyed 2,000 organizations about their AI initiatives to discover what's working, what's not and how you can get ahead.

Read the report →

Ebook

How to choose the right foundation model

Learn how to select the most suitable AI foundation model for your use case.

Read the ebook →

1 / 2

Related solutions

IBM watsonx.governance

Govern generative AI models from anywhere and deploy on the cloud or on premises with IBM watsonx.governance.

Discover watsonx.governance →

AI governance solutions

See how AI governance can help increase your employees' confidence in AI, accelerate adoption and innovation, and improve customer trust.

Discover AI governance solutions →

AI governance consulting services

Prepare for the EU AI Act and establish a responsible AI governance approach with the help of IBM Consulting.

Discover AI governance services →

Take the next step

Direct, manage and monitor your AI with a single portfolio to speed responsible, transparent and explainable AI.

Explore watsonx.governance →

Book a live demo →

Footnotes

[1] "AI Alignment: A Comprehensive Survey," arXiv, 1 May 2024.

[2, 7] "Consequences of Misaligned AI," NeurIPS Proceedings, 2020.

[3] "Faulty Reward Functions in the Wild," OpenAI, 21 December 2016.

[4] "Modelling the Recommender Alignment Problem," arXiv, 25 Aug 2022.

[5] "Ethical Issues in Advanced Artificial Intelligence," Nick Bostrom, n.d.

[6] "'The Alignment Problem' Review: When Machines Miss the Point," The Wall Street Journal, 25 October 2020.

[8] "Introducing the Frontier Safety Framework," Google DeepMind, 17 May 2024.

[9] "Our Approach to Alignment Research," OpenAI, 24 August 2022.

[10, 11] "SALMON: Self-Alignment with Instructable Reward Models," arXiv, 9 Apr 2024.