




Beyond the Hype: **Practical AI Ethics** for Legal Practitioners

FOR EDUCATIONAL PURPOSES, NOT LEGAL ADVICE

SPEAKER BIO

Amanda L. Stubson, Esq.

Legal Product Specialist

Customer Success

Professional Experience:

- *Thomson Reuters*
 - *Attorney Editor*
 - *Product Developer*
- *Other Legal Experience*
 - *Transactional – Business Formation, Estate Planning, Real Estate*
 - *Civil Litigation – Municipal Defense*



Agenda

1. Legal Ethics Landscape
2. Responsible AI Use Policies – Why They Matter
3. Agentic AI and Recent LLM Developments
4. Q&A

The Legal Ethics Landscape

Lawyers' Core Duties Don't Change

- Ethical rules still apply to AI
 - No AI exception
- AI is a tool, not legal judgment
- Lawyers remain responsible for AI-aided work
- Supervision and policy setting is key

Significant Ethical Duties Implicated by AI

Competence &
Diligence




Confidentiality



Communication
& Candor



Supervision





Duties of Competence & Diligence

Rule 1.1 – Competence

“A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.”

Rule 1.3 – Diligence

“A lawyer shall act with reasonable diligence and promptness in representing a client.”

ABA Model Rules Implicated by Large Language Models

“In 2012,... ABA adopted Comment 8 to Rule 1.1 which states that ‘[t]o maintain the requisite knowledge and skill, lawyers should keep abreast of changes in the law and its practice, including the **benefits and risks associated with relevant technology.**’”¹

Two, of the many, implications from this:

- 1). A lawyer **must know about what AI tools** exist to help them do their work with the highest degree of skill and thoroughness, and
- 2) A lawyer must know how to use, and understand, **what the AI product is doing (or not doing).**

Source: Cerny, Jenine, Delchin, Steve, Nguyen, Huu. Legal Ethics in the Use of Artificial Intelligence. Retrieved July 31, 2023, from https://www.squirepattonboggs.com/-/media/files/insights/publications/2019/02/legal-ethics-in-the-use-of-artificial-intelligence/legaethics_feb2019.pdf.

Disrupted

New York lawyers sanctioned for using fake ChatGPT cases in legal brief

By Sara Merken

June 26, 2023 3:28 AM CDT · Updated a month ago



Competence to use new technology is required

“

The phrase “including the benefits and risks associated with relevant technology” would offer **greater clarity** regarding this duty and emphasize the growing importance of technology to modern law practice. As noted in ethics opinions, such as those relating to cloud computing, **this obligation is not new**. Rather, the proposed amendment emphasizes that a lawyer should remain aware of technology, including the benefits and risks associated with it, as part of a lawyer’s general ethical duty to **remain** competent in a digital age.

- ABA Commission on Ethics 20/20, August 2012

”

Source: https://www.americanbar.org/content/dam/aba/administrative/ethics_2020/20120508_ethics_20_20_final_hod_introduction_and_overview_report.pdf
(Accessed on 8/9/2023)

What ABA Opinion 512 Requires

To satisfy their duty of competence, attorneys must acquire sufficient learning and skill about using AI in their practices.

“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. 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.”



Formal Opinion 512

Generative Artificial Intelligence Tools

July 29, 2024

Duty of Competence & Understanding AI

Understanding AI Is Now Part of Competence

- Lawyers must learn AI's capabilities and risks
- Stay updated on tools, updates, and case law
- Vet vendors for ethical alignment

Competent AI Use = Informed Use

- Read tool documentation and terms
- Understand limits: hallucinations, bias
- Seek expert advice when needed



13 State Bars with Formal Ethics Opinions

- Alaska, Florida, Kentucky, Michigan, Mississippi, New Mexico, North Carolina, Oregon, Pennsylvania, Texas, Virginia, Washington D.C., West Virginia
- AI use ≠ excuse for ethics failure
- Duty of competence + verification required

State Bar Guidance (as of April 2026)

SK1

State-Level Trends

- Some states provide practical guidance (e.g., CA and IL)
- Others warn about confidentiality and misuse
- Expect more formal guidance to come

Slide 11

SK1 [@Fancher, Chad (Legal)] should we try to update this slide?
Sracic, Kristina (Legal), 2026-04-08T16:39:25.249

FC1 0 Updated.
<https://www.bloomberglaw.com/external/document/X2JK49QC000000/legal-profession-comparison-table-state->
Fancher, Chad (Legal), 2026-04-08T21:54:53.986

Duty of Diligence

ABA Rule 1.3

A lawyer shall act with reasonable diligence and promptness in representing a client.

- Comment 1, in part: “...a lawyer may have authority to exercise professional discretion in determining the means by which a matter should be pursued. See Rule 1.2.”
- Comment 2: “A lawyer's work load must be controlled so that each matter can be handled competently.” See also, Rule 3.2 (expediting litigation).

Practical Application:

- ✓ Technology, like AI, may assist lawyers in maintaining their workload
- ✓ AI is not expected... yet.

Source: ABA Model Rule 1.3 and Comments

Duties of Competence & Diligence: Practical Takeaways



- ☐ Understand the software and dataset that is being utilized, including associated limitations and risks.
- ☐ Treat any AI-generated work as a draft and not final work product.
- ☐ Validate all case citations and quotations against a reputable database.
- ☐ Confirm underlying sources of legal and factual findings
- ☐ Conduct independent research, or review of applicable materials, to validate that output is comprehensive and accurate.
- ☐ Review, analyze, and confirm conclusions.
- ☐ Follow policies for permissible use of AI (and disclosure of any use).

Significant Ethical Duties Implicated by AI

Competence & Diligence




Confidentiality



Communication & Candor



Supervision



Confidentiality of Information

ABA Rule 1.6

(a) A lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent... (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.

- Comment 2, in part: “A fundamental principle in the client-lawyer relationship is that, in the absence of the client's informed consent, the lawyer must not reveal information relating to the representation.”

Application

- ✓ Carefully evaluate the underlying, enabling software that AI is built with and uses (e.g., cloud storage, third-party processing agreements, etc.)
- ✓ Choose AI built by experienced developers for legal practitioners

Source: ABA Model Rule 1.6 and Comments.

Duty of Confidentiality: ABA Guidance



“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**, even if the tool is used exclusively by lawyers at the same firm.”

“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.”

“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.**”



Confidentiality & Client Data

Safeguard Client Information at All Costs

- Don't input confidential info into public AI
- Use only tools with vetted security
- Review privacy policies carefully

Watch for Data Retention & Model Training

- Some LLMs retain inputs
- Others can opt-in to train on business data



Duty of Confidentiality: Practical Guidelines



- ☐ **Review Terms of Use, privacy policy, and other applicable contracts for any AI tool**
 - ☐ Who has access to information that is input into AI tool?
 - ☐ Can information input into AI tool be used for training?
 - ☐ Is information input into AI tool retained? For how long? Can it be deleted by user?
 - ☐ What data security measures are in place?
- ☐ **May need to consult with IT professionals and/or cybersecurity experts**
- ☐ **Obtain informed consent if client confidential information may be disclosed**

Recent Case Law

Slide 19

SK1

added this diivider slide

Sracic, Kristina (Legal), 2026-04-08T15:02:37.534

Prompt Discoverability

In *Tremblay v. OpenAI* (2024 WL 3748003), the court made a ruling regarding the discoverability of prompts used to generate outputs from ChatGPT.

- U.S. District Judge Araceli Martinez-Olguin **reversed** a magistrate judge's order compelling the plaintiffs to produce prompts and outputs that did not support their claims of copyright infringement.
 - The judge determined that the prompts were crafted by the plaintiffs' counsel and contained their mental impressions and opinions about how to interact with ChatGPT. **Therefore, these prompts were protected by the attorney work product doctrine.**
- **However**, the court also ruled that the plaintiffs had waived their work product protection for the prompts they did use to support their claim against OpenAI, as these prompts were disclosed in the complaint.
 - This waiver, however, did not extend to prompts that were not disclosed in the complaint..



Tremblay v. OpenAI: Takeaways

Are AI Prompts Discoverable?

- Case: *Tremblay v. OpenAI*
- Prompts = work product... but waived if disclosed
- Entire prompt threads may become discoverable

Attorney-Client Privilege and Work Product Doctrine

In *U.S. v. Heppner* 2026 WL 436479 (Feb. 17, 2026), the Court ruled that written exchanges between a criminal defendant and a generative AI platform, Claude, were not protected by attorney client privilege or the work product doctrine.

- The FBI seized numerous documents and electronic devices from Heppner's home, including approximately thirty-one documents that memorialized communications Heppner had with the AI platform Claude. Heppner's counsel claimed privilege over these documents, arguing that they contained information learned from counsel and were created to obtain legal advice.
- The court found that the AI documents did not meet the criteria for **attorney-client privilege** because they were:
 - **not communications between Heppner and his counsel,**
 - **not confidential (per Anthropic's T&Cs), and**
 - **not created for the purpose of obtaining legal advice** from Claude.
- The documents were created by Heppner of his own volition. Therefore, the court determined that the AI documents were not protected under the **work product doctrine** because:
 - they were **not prepared by or at the behest of counsel,**
 - **they did not reflect defense counsel's strategy**

Slide 22

SK1 reformatted your slide thought it would be easier to follow
Sracic, Kristina (Legal), 2026-04-08T14:59:08.188

U.S. v. Heppner: Takeaways

Are written exchanges between a defendant and generative AI platform protected?

- Case: *U.S. v. Heppner*
- AI documents are not confidential and they did not involve communication with counsel and therefore, they are not protected by **attorney-client privilege**.
- Because the AI documents they were not prepared by or at the behest of counsel and they did not reflect defense counsel's strategy, these documents are not protected under the **work product doctrine**
- **Note this sentence which generated significant commentary:**
 - "Had counsel directed Heppner to use Claude, Claude might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege."

Slide 23

- SK1** Added this takeaway slide
Sracic, Kristina (Legal), 2026-04-08T14:59:18.407
- FC1 0** Thank you, this is clarifying. I'll delete slide 22.
Fancher, Chad (Legal), 2026-04-08T15:17:31.884

A Different View on Work Product Doctrine

In *Warner v. Gilbarco, Inc.* 2026 WL 373043 (Feb. 10, 2026), the Court ruled that Artificial intelligence (AI) materials used by the pro se former employee were subject to work-product protection and the former employee did not waive work product protection by using an artificial intelligence large language model.

- The court denied Defendants' request for production of "all documents and information concerning [Plaintiff's] use of third-party AI tools in connection with this lawsuit" as the information sought is not discoverable. If the information was discoverable, it is subject to protection under the work-product doctrine, which Plaintiff is permitted to assert. To the extent Defendants argue that Plaintiff waived the work-product protection by using ChatGPT, the work-product waiver has to be a waiver *to an adversary* or in a way likely to get in an adversary's hand. ChatGPT (and other generative AI programs) are *tools, not persons*, even if they may have administrators somewhere in the background.
- The court found that the AI materials were subject to the **work product doctrine** because they involved:
 - **the plaintiff's internal drafting process, analysis, and thought process, and were prepared in anticipation of litigation.**
- The court found that the former employee **did not waive work product protection** by using an artificial intelligence large language model because:
 - the work-product waiver had to be a waiver to an adversary or in a way likely to get in an adversary's hand, and the AI model was merely a tool, not a person.

SK1

Slide 24

SK1 do you mean hand or head? maybe we can make this clearer?

Sracic, Kristina (Legal), 2026-04-08T15:27:03.364

FC1 0 That's the language used in the decision.

Fancher, Chad (Legal), 2026-04-08T15:32:49.415

SK1 1 i figured-sounds good

Sracic, Kristina (Legal), 2026-04-08T15:34:50.568

Warner v. Gilbarco, Inc.: Takeaways

When do AI materials constitute work product?

- Case: *Warner v. Gilbarco, Inc.*
- AI documents are subject to the **work product doctrine** when the documents involve thought and analysis and were prepared in anticipation for litigation.
- Using a large language model does not waive work product protection

Morgan v. V2X, Inc., Where Defense Counsel Read ABA Opinion 512

2026 WL 864223 (March 30, 2026)

Holding: The Court required a Pro Se Plaintiff to disclose the name of any AI platform used in connection with confidential information, as the Plaintiff did not demonstrate that identifying the tool would reveal mental impressions or legal strategy.

The defendant filed a motion to amend the existing Protective Order to include specific language addressing AI use and to compel the plaintiff to disclose the AI tool he used.

The court granted the defendant's motion in part, requiring the plaintiff to disclose the name of any AI tool used with confidential information. The court also amended the Protective Order to restrict the use of mainstream AI tools like ChatGPT, Claude, and Gemini for processing confidential information unless certain contractual safeguards are in place. These safeguards include prohibiting the AI provider from storing or using inputs to train or improve its model and ensuring the ability to delete all confidential information upon request.

The court recognized the challenges faced by pro se litigants in accessing enterprise-grade AI tools and emphasized that while modern AI tools can be used in ways that do not involve uploading confidential information, the risks associated with mainstream AI tools necessitate these restrictions.

Slide 26

SK1

Maybe we just leave this slide as is and don't do a takeaway slide, rhoughts? [@Fancher, Chad (Legal)]

Sracic, Kristina (Legal), 2026-04-08T15:55:05.034

FC1 0

I'm OK with that. You'll likely be the first one presenting these slides next week, so if you are comfortable, let's leave it.

Fancher, Chad (Legal), 2026-04-08T15:58:24.037

“Informed” Consent

“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,**
 - **specific information 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.**

Key: Adding boiler-plate provisions to engagement letters purporting to authorize the lawyer to use GAI is not sufficient.



Transparency in Billing & Client Consent

You Can't Bill for Time Saved by AI

- Time saved $\neq$ billable time
- Flat fees may be better
- Explain when AI impacts cost or substance

Informed Consent: Notify When Using AI

- Notify client before using AI with their info
- Must disclose if AI is used for substantive tasks
- Review third-party data-sharing terms



Significant Ethical Duties Implicated by AI

Competence & Diligence




Confidentiality



Communication & Candor



Supervision



Candor Toward the Tribunal

ABA Rule 3.3

(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; (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...

- Comment 2, in part: “This Rule sets forth the special duties of lawyers as officers of the court to **avoid conduct that undermines the integrity of the adjudicative process.**”

Application

- ✓ Using AI that limits its investigation to a real, domain-specific source of information, and that shows its work, will help lawyers “avoid conduct that undermines the integrity of the adjudicative process.”

Source: ABA Model Rule 3.3 and Comments

“Hallucination”

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

- Plaintiff’s attorney submitted affirmation in opposition to motion to dismiss
- Affirmation contained citations to, and discussion of, several nonexistent judicial decisions
 - *Varghese v. China Southern Airlines Co., Ltd.*, 925 F.3d 1339 (11th Cir. 2019)
 - *Miller v. United Airlines, Inc.*, 174 F.3d 366 (2d Cir. 1999)
 - *Martinez v. Delta Airlines, Inc.*, 2019 WL 4639462 (Tex. App. Sept. 25, 2019)
- Attorney did not correct issue after identified by defendant in response
- Court demanded copies of the decisions and an explanation for the discrepancy
- Attorney submitted an affidavit with copies of the purported decisions, which also did not exist

Here’s What Happens When Your Lawyer Uses ChatGPT

A lawyer representing a man who sued an airline relied on artificial intelligence to help prepare a court filing. It did not go well.



New York lawyers sanctioned for using fake ChatGPT cases in legal brief



First Case Addressing Hallucinations: *Mata v. Avianca* (2023)

Key takeaway:

Misunderstanding technology is not a defense for misusing technology

12. Mr. Schwartz also testified at the hearing that he knew that there were free sites available on the internet where a known case citation to a reported decision could be entered and the decision displayed. (Tr. 23-24, 28-29.) He admitted that he entered the citation to "Varghese" but could not find it:

THE COURT: Did you say, well they gave me part of Varghese, let me look at the full Varghese decision?

MR. SCHWARTZ: I did.

THE COURT: And what did you find when you went to look up the full Varghese decision?

MR. SCHWARTZ: I couldn't find it.

THE COURT: And yet you cited it in the brief to me.

MR. SCHWARTZ: I did, again, operating under the false assumption and disbelief that this website could produce completely fabricated cases. And if I knew that, I obviously never would have submitted these cases.

- 6 -

Source: MATA v. AVIANCA INC (2023), <https://caselaw.findlaw.com/court/us-dis-crt-sd-new-yor/2335142.html> (Accessed 4/15/2024)

Verifying AI Outputs & Duty of Candor

Trust but Verify: Lawyers Must Confirm AI Outputs

- AI may hallucinate case law
- Submitting unverified content = ethical violation
- Cite: *Mata v. Avianca*

Model Rule 3.3 & Independent Review

- Lawyers must not knowingly submit false info
- Always verify: citations, analysis, research
- Double check before filing anything AI-assisted



Federal Standing Orders on AI

Judges Are Getting Proactive

- Many federal judges issued AI SOs since May 2023
- Most allow AI use with disclosure/certification
- Some prohibit all AI use except for legal search engines (e.g., SD Ohio)

Common Themes Across Orders

- Disclosure of tool used and verification of output
- Focus on hallucinations and fake citations
- Emphasis on Model Rule 1.1, Comment 8
- Shift from proactive approach to sanctioning individual attorneys



Standing Orders – AI Certifications

IN THE UNITED STATES DISTRICT COURT FOR THE EASTERN DISTRICT OF PENNSYLVANIA

STANDING ORDER RE: ARTIFICIAL INTELLIGENCE (“AI”) IN CASES ASSIGNED TO JUDGE BAYLSON

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.

DATED: 6/6/2023

BY THE COURT:

/s/ MICHAEL M. BAYLSON

MICHAEL M. BAYLSON, U.S.D.J.

UNITED STATES DISTRICT COURT FOR THE NORTHERN DISTRICT OF CALIFORNIA

STANDING ORDER FOR CIVIL CASES BEFORE DISTRICT JUDGE ARACELI MARTÍNEZ-OLGUÍN (Revised November 22, 2023)

4. **Artificial Intelligence (AI).** Counsel is responsible for providing the Court with complete and accurate representations in any submission (including filings, demonstratives, evidence, or oral argument), consistent with Federal Rule of Civil Procedure 11, the California Rules of Professional Conduct, and any other applicable legal or ethical guidance. Use of ChatGPT or other such tools is not prohibited, but counsel must at all times personally confirm for themselves the accuracy of any content generated by these tools. At all times, counsel—and specifically designated lead trial counsel—bears responsibility for any submission made by the party that the attorney represents. Any submission containing AI-generated content must include a certification that lead trial counsel has personally verified the content’s accuracy. Failure to include this certification or comply with this verification requirement will be grounds for sanctions. Counsel is responsible for maintaining records of all prompts or inquiries submitted to any generative AI tools in the event those records become relevant at any point.



Standing Orders – AI Prohibitions

UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF OHIO
WESTERN DIVISION, AT DAYTON

STANDING ORDER GOVERNING CIVIL CASES

Effective as of December 18, 2023

**Notice to counsel: New Requirement for Joint Proposed Final Pretrial
Conference Orders Effective as of December 18, 2023**

Notice to counsel: New AI Provision Effective as of July 14, 2023

Hon. Michael J. Newman
United States District Judge
Walter H. Rice Federal Building & U.S. Courthouse
200 West Second Street, Room 505
Dayton, Ohio 45402
newman_chambers@ohsd.uscourts.gov

VI. ARTIFICIAL INTELLIGENCE (“AI”) PROVISION

No attorney for a party, or a *pro se* party, may use Artificial Intelligence (“AI”) in the preparation of any filing submitted to the Court. Parties and their counsel who violate this AI ban may face sanctions including, *inter alia*, striking the pleading from the record, the imposition of economic sanctions or contempt, and dismissal of the lawsuit. The Court does not intend this AI ban to apply to information gathered from legal search engines, such as Westlaw or LexisNexis, or Internet search engines, such as Google or Bing. All parties and their counsel have a duty to immediately inform the Court if they discover the use of AI in any document filed in their case.

Significant Ethical Duties Implicated by AI

Competence &
Diligence




Confidentiality



Communication
& Candor



Supervision



ABA Model Rules Implicated by Large Language Models

Rule 5.1 Responsibilities of Partners, Managers, and Supervisory Lawyers

Rule 5.3 Responsibilities regarding Non-Lawyer Assistance

Under ABA Model Rules 5.1 and 5.3, lawyers have an ethical obligation to supervise lawyers and nonlawyers who are assisting lawyers in the provision of legal services to ensure that their conduct **complies with the Rules of Professional Conduct**.



In 2019 the ABA changed the rule from **assistant** to **assistance**.

The change clarified that the scope of Rule 5.3 encompasses nonlawyers **whether human or not**. Under Rules 5.1 and 5.3, lawyers are obligated to supervise the work of AI utilized in the provision of legal services...

Source: <https://www.americanbar.org/content/dam/aba/directories/policy/annual-2019/112-annual-2019.pdf>

Treat AI Like a Nonlawyer Assistant

Law firms should treat AI use like
supervising a paralegal:

- Set clear **usage policies**.
- **Review outputs**.
- Ensure **ethical standards** are met.



Vendor Policies

What Vendors Say About Ethics & Compliance

- Varying rules on customer data use for training
- GDPR, CCPA, SCC compliance
- Cross-border restrictions apply

Prohibited Uses & IP Risks

- Ban on facial recognition, misinformation, illegal use
- Reverse engineering forbidden
- Customers own outputs—but not always unique

Responsible AI Use Policies

Why Responsible AI Use Policies Matter

**Judi Flourney** • 1st
A CIO for 25+ years for AmLaw 250 firms and embarking on new a new ...
3h • 

The rush to embrace Generative AI could cost firms more than they realize, client trust, privilege, and even future work.

Clients are tightening Outside Counsel Guidelines (OCGs), regulators are watching, and ethical obligations haven't changed.

Without a clear policy, even well-meaning attorneys can put the firm at risk.

Why Your Firm Needs A Generative AI Policy Now

Generative AI isn't just another tech tool—it touches the core ethical duties of every lawyer: confidentiality, competence, and supervision.

Common Pitfalls

- ✗ No firmwide guidance
- ✗ Uploading client data to public or unvetted tools
- ✗ Ignoring client OCG restrictions
- ✗ Blind reliance on unverified AI output
- ✗ No vendor due diligence

What "Good" Looks Like

- ✓ Clear rules on acceptable AI use
- ✓ Tracking OCG restrictions & obtaining client consent when required
- ✓ Mandatory training on verifying outputs & safeguarding confidentiality
- ✓ Vendor vetting including obtaining Microsoft's contractual exemption in Azure OpenAI to disable content filtering and abuse monitoring that might otherwise allow human review of flagged content
- ✓ Transparent disclosure where OCGs require it
- ✓ Be sure to engage with your firm counsel on any policy that that provides direction to your lawyers

AI can be a game changer but only if used wisely. A strong policy protects clients, preserves privilege, and positions your firm as a trusted, forward-thinking partner.

  Lydia Flocchini and 20 others

1 comment • 1 repost

 Like

 Comment

 Repost

 Send

When and Whether to Get Client Consent

- **Client consent is not always required**, but ethical considerations may make it necessary.
- Required when:
 - AI use may result in the disclosure of any **confidential or sensitive client data**.
 - AI use impacts **legal fees, case strategy or substantive legal advice**.
 - Requested by clients or dictated by engagement terms.
- **Routine AI use** (e.g., research) typically doesn't require consent, but substantive legal tasks do.



How to Obtain Client Consent for AI Use

- **Discuss AI use early** in the engagement.
- Include **disclosure in the engagement letter**.
- Provide **a separate consent form** if needed.
- Consent should be:
 - **Specific** about AI use.
 - **Documented** in writing.
 - **Acknowledged** by the client.




Communicating AI Use with Clients

- **Openly discuss** what AI will be used for (research, drafting, etc.).
 - According to the TR Institute survey: **78% of in-house counsel expect disclosure.**
- Explain **why** AI is being used and the value:
 - Efficiency, accuracy, cost-effectiveness.
 - Market your AI adoption: Highlight your firm's commitment to using technology to improve client service.
- **Address concerns:** security, oversight, quality control.
- **Regularly update clients:** Keep clients informed about any changes in AI usage throughout the representation.



What is an acceptable use policy for AI?

- **An acceptable use policy (AUP) outlines guidelines for how AI should be used within an organization.**
 - An AUP should clearly state what situations are appropriate to use generative AI, what data handling practices must be used, any required proofing and oversight that must be accomplished, and any security protocols your team requires to mitigate risks.
- **Crafting a comprehensive AI policy involves:**
 - Defining the scope of the policy and stakeholders affected by your policy
 - Establishing or referencing the ethical principles and values of your firm
 - Creating or referencing a data governance strategy
 - Addressing potential bias and discrimination
 - Providing training and guidance
 - Outlining regular intervals for reviewing the AI policy



Responsible AI Use Policies

- Define **acceptable AI use cases**.
- Require **human review** of AI outputs.
- **Prohibit uploading confidential data** into unverified tools.
- Regularly **update policies** to reflect legal and tech changes.
 - Need to stay informed about and comply with all applicable laws and regulations.
- Include **accountability** and **risk management** mechanisms.
 - For example, include guidance on when / whether / how to obtain client consent for AI use.



AI policy template examples

- **Fisher & Philips LLP's** [Acceptable Use of Generative AI Tools \[Sample Policy\]](#) goes over standard legal implications of AI, covering topics such as data protection and confidentiality.
- **The Responsible Artificial Intelligence Institute's** [AI Policy Template](#) offers a comprehensive template that covers various aspects of AI use, including AI development, ethical considerations, data privacy, and risk management.
- **MES Computing** shares [5 AI Policy Templates You Can Use As a Framework](#), focusing on generative AI, copyright and fair use, data protection, and incorrect AI responses.
- **HR Source's** template [Generative Artificial Intelligence \(AI\) Use in the Workplace Policy](#) focuses on the practical aspects of AI implementation, including technology selection, training, and risk management.
- **American Inns of Court** provides a sample AI Policy for law firms regarding [Generative AI \(GAI\) Use and Guidance](#). This template has an overview of Generative AI issues in the legal profession, including ethical considerations, data privacy, and cross-border legal challenges.

Oversight, Training, and Human Review

- **Train lawyers** on prompting, review, and ethics.
 - *Best practice:* Require training to get access to the platform.
- Implement **internal audits** of AI-assisted work. The audits should ensure that AI-assisted legal work meets ethical and professional standards.
- Ensure **human oversight and validation**.
 - Ensure AI complements, rather than replaces, legal expertise. Attorneys should review and verify AI-generated content to prevent errors.
- Establish **feedback mechanisms** for lawyers to report issues or concerns.



GenAI Tool Evaluation Criteria


Evaluate tools based on:

- **Transparency & explainability**
- **Accuracy and performance**
- **Data security**
- **Vendor reliability**



What Are Transparency & Explainability?

- **Transparency** answers the question “what happened” in the AI system, while explainability addresses “how” a decision was made using AI.
 - Organizations are transparent when they provide meaningful information so that individuals are aware of the actions taken by AI and what content was AI-generated. See [National Institute of Standards and Technology's \(NIST\) Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#).
- **Explainability** means that organizations should provide individuals with a plain-language explanation of the AI system's logic and decision-making process so that individuals know *how* the AI generated the output or decision. See [the White House's Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People](#).



How to Measure Accuracy & Performance of GenAI

1. Accuracy and Faithfulness:


- Compare against a Source of Truth: Evaluate the AI's responses against a known correct answer to determine accuracy.
- Identify Hallucinations: Check for instances where the AI fabricates information or makes claims not supported by the source material.

2. Relevance and Context:

- Relevance to the Query: Ensure the AI's responses directly address the user's query or prompt.
- Contextual Understanding: Evaluate if the AI can understand the context of the query and provide appropriate responses.
- Coherence and Informativeness: Assess the clarity and logical flow of the output, as well as whether it provides enough information.

3. User Feedback:

- Collect feedback from users on the accuracy, relevance, and overall quality of the AI's responses.



Agentic AI

Understanding Agentic AI

Autonomous Decisions

- Agentic AI systems can autonomously make decisions and take actions without human intervention, increasing efficiency and effectiveness.

Environmental Perception

- Agentic AI can perceive its environment (e.g., via API, canvas awareness), enabling it to interact with the world in a dynamic and proactive manner.

Goal-Oriented Actions

- Agentic AI emphasizes planning and acting towards achieving specified goals, making it proactive and strategic.



One example: the agentic AI travel agent

How is Agentic AI Different from GenAI?

GenAI

- GenAI reacts to specific input to generate output – it creates.
- GenAI requires step-by-step assistance to create: you must describe the requirements, provide details, and ask for changes to get the best results.

Agentic AI

- Once it is given initial instructions, agentic AI makes decisions and acts – it performs.
- Agentic AI is more autonomous: it can carry out multi-step assignments with only initial input and without continual prompting.




Image generated by CoPilot with the prompt: *Generate an image representing the capabilities of agentic AI.*


Recent LLM Developments

AI is getting smarter, *faster*

- Reading handwriting
- Speech recognition
- Image classification
- Basic-level reading comprehension
- Medium-level reading comprehension
- Basic English language understanding
- Multi-task language understanding
- General problem-solving
- Competition-level mathematics
- PHD-level science questions

* AI Action Summit International AI Safety Report, Jan 2025

AI performance vs human performance on select benchmarks*



Smarter, Faster LLMs

Enhanced Context Support

Recent LLMs support up to 1 million tokens of context, allowing them to handle and understand long documents effectively.

Human-like Interaction


These advanced models are designed for more natural and human-like interaction, making communication smoother and more intuitive. **Higher emotional intelligence (EQ) scores.**

Improved Document Comprehension

With better long-document comprehension, these models provide more accurate and relevant responses.

Reduced Hallucination Rates

Advancements have also led to reduced hallucination rates, ensuring more reliable and fact-based outputs from the models.




Examples of newer models with these characteristics: GPT-4.1, Claude 3.7, and Gemini 2.0.

In April, OpenAI launched three new models: GPT-4.1, GPT-4.1 mini, and GPT-4.1 nano.

These models outperform GPT-4o and GPT-4o mini across the board, with major gains in coding and instruction following.

They also have larger context windows—supporting up to 1 million tokens of context—and are able to better use that context with improved long-context comprehension. They feature a refreshed knowledge cutoff of June 2024.

Generate a Timeline has already moved to GPT-4.1; other skills will migrate to this model over the next few months.




What's Next: Large Reasoning Models (LRM)

Evolution from LLMs

- LRMs represent an evolution from Large Language Models, advancing beyond understanding and generating human language.
- They are specialized variants of LLMs explicitly designed for reasoning tasks.

Advanced Reasoning Capabilities


- LRMs incorporate advanced reasoning capabilities such as long Chain-of-Thought (CoT) with self-reflection to enhance reasoning capabilities.
- By contrast, LLMs are designed for broader language understanding and generation tasks without specific reasoning enhancements.




Where LRMs Outperform LLMs



TASKS REQUIRING STRUCTURED REASONING: EXAMPLES INCLUDE MATHEMATICAL EQUATIONS, LOGICAL DEDUCTION, LEGAL ANALYSIS, MEDICAL DIAGNOSTICS, AND SCIENTIFIC RESEARCH.



CONTROLLED PUZZLE ENVIRONMENTS: LRMS OUTPERFORM LLMs IN TASKS THAT REQUIRE ALGORITHMIC REASONING AND PROBLEM-SOLVING.



SELF-REFLECTION AND CHAIN-OF-THOUGHT: LRMS GENERATE DETAILED THINKING PROCESSES BEFORE PROVIDING ANSWERS, WHICH ALLOWS FOR THE ANALYSIS OF NOT ONLY FINAL ANSWERS BUT ALSO THE INTERNAL REASONING TRACES.



MEDIUM-COMPLEXITY TASKS: LRMS DEMONSTRATE AN ADVANTAGE IN MEDIUM-COMPLEXITY TASKS WHERE ADDITIONAL THINKING AND REASONING ARE BENEFICIAL.

Current Limits of LRMs



Large Language Models (LLMs) outperform Large Reasoning Models (LRMs) when the problem complexity is low.

LLMs are more token-efficient in low-complexity tasks, achieving comparable or even better performance than LRMs. **This means that** LLMs can provide accurate results with fewer computational resources.



LRMs have not developed generalizable problem-solving capabilities for planning **complex** tasks.

LRMs exhibit a counterintuitive scaling limit where their reasoning effort increases with problem complexity up to a point, then declines despite having an adequate token budget.

Nippon Life Insurance Company of America v. OpenAI Foundation, 1:26CV02448 (N.D. Il. 3/4/26)

OpenAI's answer is due on 5/15/26.

Nippon is making three claims in its complaint against OpenAI, including the **unauthorized practice of law under 750 ILCS 205/1**.

UPL Claim: The plaintiff alleges that Illinois law prohibits practicing law without a license and extends contempt to anyone providing legal services without licensure. They assert that federal courts have the authority to regulate and discipline unlicensed practice before them. The plaintiff claims that OpenAI, through ChatGPT, provides legal advice, analysis, research, and drafts legal documents for court use. They argue that ChatGPT is not an attorney, is not licensed in Illinois or any U.S. jurisdiction, and is not admitted in the Northern District of Illinois.

The plaintiff further alleges that ChatGPT assisted an Illinois resident, Dela Torre, by drafting motions, notices, subpoenas, and requests for judicial notice, and by providing legal analysis, research, and advice for filings in the Northern District of Illinois. They plead that such assistance constitutes practicing law in Illinois without a license.

Q & A



The background of the slide is a light green color with a complex, abstract pattern of thin, dark green lines. These lines intersect to form a grid-like structure in some areas and radiating patterns in others, creating a sense of depth and connectivity. Three solid green circles are placed at specific points where the lines intersect: one in the upper left, one on the left side near the text, and one on the right side.

Thank you!