


TEACHING LEGAL RESEARCH IN THE AGE OF AI

A Framework for Building AI-Literate, Practice-Ready Graduates

Lauren E. Diaz



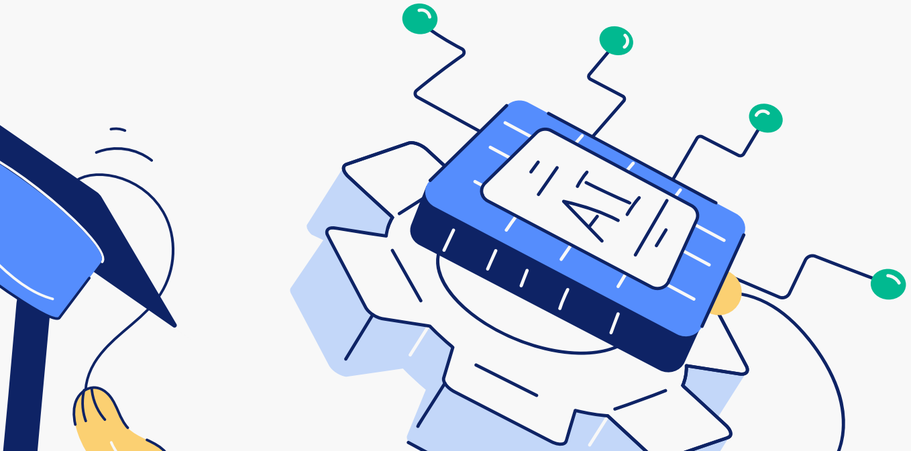


HOW DO WE TEACH STUDENTS TO USE AI IN LEGAL RESEARCH WITHOUT LOSING ACCURACY, AUTHORITY, ETHICS, OR PROFESSIONAL JUDGMENT?

TAKE WHAT YOU NEED

Scaling from one exercise to a full semester

THE CORE TEACHING PREMISE

AI-assisted legal research builds on traditional legal research.
It does not replace it.

Traditional legal research skills gives students the map.

AI gives them the speed.

Without the map, the speed will not get students where they need to go.

THE PROBLEM WE'RE SOLVING

Key Statistics:

- 73% of legal professionals now use AI for research tasks
- AI tools hallucinate at rates of 17-33% on legal questions
- 75% of supervising attorneys expect new hires to cite-check AI materials
- Only 40% of 3L students feel proficient in AI-assisted research
- Multiple documented cases of lawyers filing hallucinated citations with courts

Bottom Line: Students are graduating into a profession that expects AI competence but punishes AI incompetence. Law schools must close this gap. Teaching students how to conduct AI-assisted legal research requires more than tool training, it requires **AI Literacy**.

THE PROBLEM WE'RE SOLVING

- AI output looks like careful legal analysis (e.g., case names, citations, rule statements, policy arguments)
- It may be well-organized, fluent, and competent. And it may be wrong.
- Major capability jumps every few months, faster than any prior technology shift
- **Existing Instruction**: too general, cautionary, or tool-specific
- **Gap**: practical, step-by-step guidance on using AI responsibly in the research process

WHAT IS AI LITERACY AND HOW TO TEACH IT

AI Literacy is not a seperate technical skill. It is the intersection of four competencies:

- Information literacy
- Technology competence
- Ethical and professional responsibility
- Experiential Learning

Students need to understand how AI systems generate outputs, know AI's capabilities and limitations, and identify the ethical and professional responsibility concerns that arise when AI is used in legal research.

AI Literacy: It is the ability to understand, use, evaluate, and responsibly interact with AI systems. It includes knowing, at a practical level, not a theoretical one, what an AI tool is designed to do, what information it can actually access, how it produces outputs, and what its limitations are. In the context of legal research, AI literacy is the ability to use AI tools to support the research process while independently verifying every source, authority, jurisdictional claim, and legal proposition before relying on the output. It means knowing how to prompt an AI tool effectively, how to spot the difference between an output grounded in real authority and one that merely sounds grounded, and how to catch hallucinations or omissions before they make it into a work product.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Information Literacy: “Information literacy is the ability to identify what information is needed, understand how information is organized, identify the best sources of information for a given need, locate these sources, evaluate the sources critically, and share that information.”

- That definition has long grounded legal research pedagogy, and the instructional approach remains the same even with the introduction of generative AI in legal research.
- Students must learn to evaluate and verify AI outputs.
- Rather than treating AI output as authoritative, students should be training to interrogate them by asking familiar research questions: Is the information accurate? Does it address the specific legal issue? Is it grounded in the correct juridiction and court level? Does it reflect bias or omit important perspectives? Is the analysis complete, or does it exclude key authority?
- Prompting instruction also fits naturally within an information literacy framework. Long before generative AI, instruction included legal research strategies such as Boolean, connectors, field restrictions, and platform specific syntax. AI tools require that same kind of literacy. A prompt is the user’s research query, and the quality of the output depends heavily on how well that query is framed, scoped, and constrained, including by jurisdiction, authority type, time frame, and the specific task requested.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Information Literacy

- Just because AI tools are widely available does not automatically make users expert researchers.
- Without instruction on how to formulate strong prompts and refine them in response to results, students may receive answers that sound plausible but are incomplete, misdirected, or legally unreliable.
- Teaching students how to prompt effectively therefore functions as modern query formulation and is a core component of AI literacy as information literacy.
- Teaching students how to prompt well, and then critically evaluate and verify what they receive, reinforces that AI literacy is information literacy applied to new research systems.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Technology Competence: ABA Model Rule 1.1 requires lawyers to provide competent representation, including the legal knowledge, skill, thoroughness, and preparation reasonably necessary for the matter. Comment 8 clarifies that competence includes staying current with changes in law and practice, including understanding the benefits and risks of relevant technology through continuing study.

- For students to truly assess the benefits and risks of AI, they need a foundational understanding of AI.
- Technology competence in AI-assisted legal research begins with a foundational understanding of how AI systems work and why they fail.
- Students do not need to become computer scientists, but they should be able to use core AI vocabulary accurately and understand how the major concepts relate to one another.
- Students need a basic understanding of what artificial intelligence is as the umbrella term, how machine learning and deep learning fit within it, how LLMs operate as a type of deep learning system, and how GAI uses those models to produce new text and other outputs.
- They should also have a practical grasp of how these systems are trained and generate responses, what data and design choices shape their outputs, and what common limitations and risks follow from that process such as hallucinations. This conceptual map helps students recognize what the tool is doing, what it cannot do on its own, and why the risks in AI-assisted legal research are structurally different from the risks in traditional database research.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Technology Competence:

- Technology competence also requires tool differentiation and selection judgment.
- AI tools vary substantially by design and risk profile, including public versus proprietary systems, tools embedded in legal research platforms versus general-purpose chatbots, and increasingly agentic systems that can execute multi-step tasks. Students need to be able to distinguish between general-use AI, general-use AI with research, specific-use or domain-specialized AI, proprietary AI, and open-source AI. Some tools are mostly generating text, some are combining generation with retrieval, and others are constrained by curated datasets, each with different implications for accuracy, transparency, confidentiality, and verification.
- Students should learn to match the tool to the task and to the sensitivity of the information involved.
- They should also be prepared to compare performance across platforms, benchmark outputs against authoritative sources, and track improvements over time as tools evolve. This comparative approach reinforces that competence is not tied to one product.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Ethical and Professional Responsibility: students must understand how integrating AI into legal workflows can affect their ethical and professional duties (e.g., competence, diligence, communication, fees, confidentiality, candor, misconduct)

- Students must learn to identify and manage confidentiality risks, understand what AI use may be permissible in practice, recognize when disclosure or client consent may be required, and internalize that verification duties cannot be outsourced to a tool.
- Confidentiality requires deliberate control over what information is entered into an AI system, where that information may go, and how it might be stored, shared, or used for model improvement depending on the platform and settings.
- Students should understand that pasting client facts, draft pleadings, or sensitive strategy into public or consumer-facing AI tools can create serious confidentiality problems, particularly when the platform's data use policies are unclear or when prompts and outputs may be retained.
- Students therefore need practical guidance on tool selection and safe-use norms, including when to avoid entering sensitive information entirely, when to anonymize or sanitize facts, and when internal, enterprise, or vendor-approved tools with stronger safeguards may be appropriate. Students also need to get in the habit of checking each tool's terms of service, privacy policy, and data sharing or retention settings because these rules can differ significantly across platforms and can change over time.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Ethical and Professional Responsibility:

- In addition to AI Bias (i.e., systematic errors or unfair outcomes that can arise in artificial intelligence due to biased training data or algorithmic flaws), students should learn to identify and counter automation bias, the tendency to over trust a tool because it appears confident, fluent, or efficient.
- Ethical and professional responsibility literacy must be framed as ongoing.
- AI-related court rules, standing orders, local filing requirements, bar guidance, and GAI capabilities continue to evolve rapidly. Students should learn how to stay current in their future jurisdictions by monitoring relevant court guidance, state bar opinions, and ABA materials, and by adapting their workflows as tools change.
- Students must also understand how courts are beginning to treat AI use through the lens of confidentiality, privilege, and work product, including why the outcome can turn on how the tool was used.
- Because this area is evolving quickly, students must stay current on how courts are treating AI-related privilege, work product, and confidentiality questions. Treating AI literacy as continuing professional responsibility reinforces that competence is not a one-time skill acquisition, but a continuing obligation. By equipping students with confidentiality judgment, disclosure awareness, verification discipline, and automation-bias resistance, law librarians help ensure that AI-assisted legal research and drafting is not only effective, but professionally responsible in modern practice.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Experiential Learning: To become AI literate, the students need to use the products.

- AI literacy cannot be learned passively.
- Students need practice and repetition (e.g.., hands-on work, short simulations, guided exercises that require verification).
- In order for students to truly learn the capabilities and limitations of AI, students must use it, test it, challenge, break it, and verify it.
- Students must use these tools on real research problems, test the outputs, and verify results using traditional legal research methods.
- Iterative practice with feedback turns students into active practitioners rather than passive recipients of information.

WHAT IS AI LITERACY AND HOW TO TEACH IT

Key Teaching Points

- AI literacy is not a one-time lesson. It is an ongoing professional habit.
- Distinguish AI literacy (transferable evaluation skills) from tool fluency (knowing specific platforms)

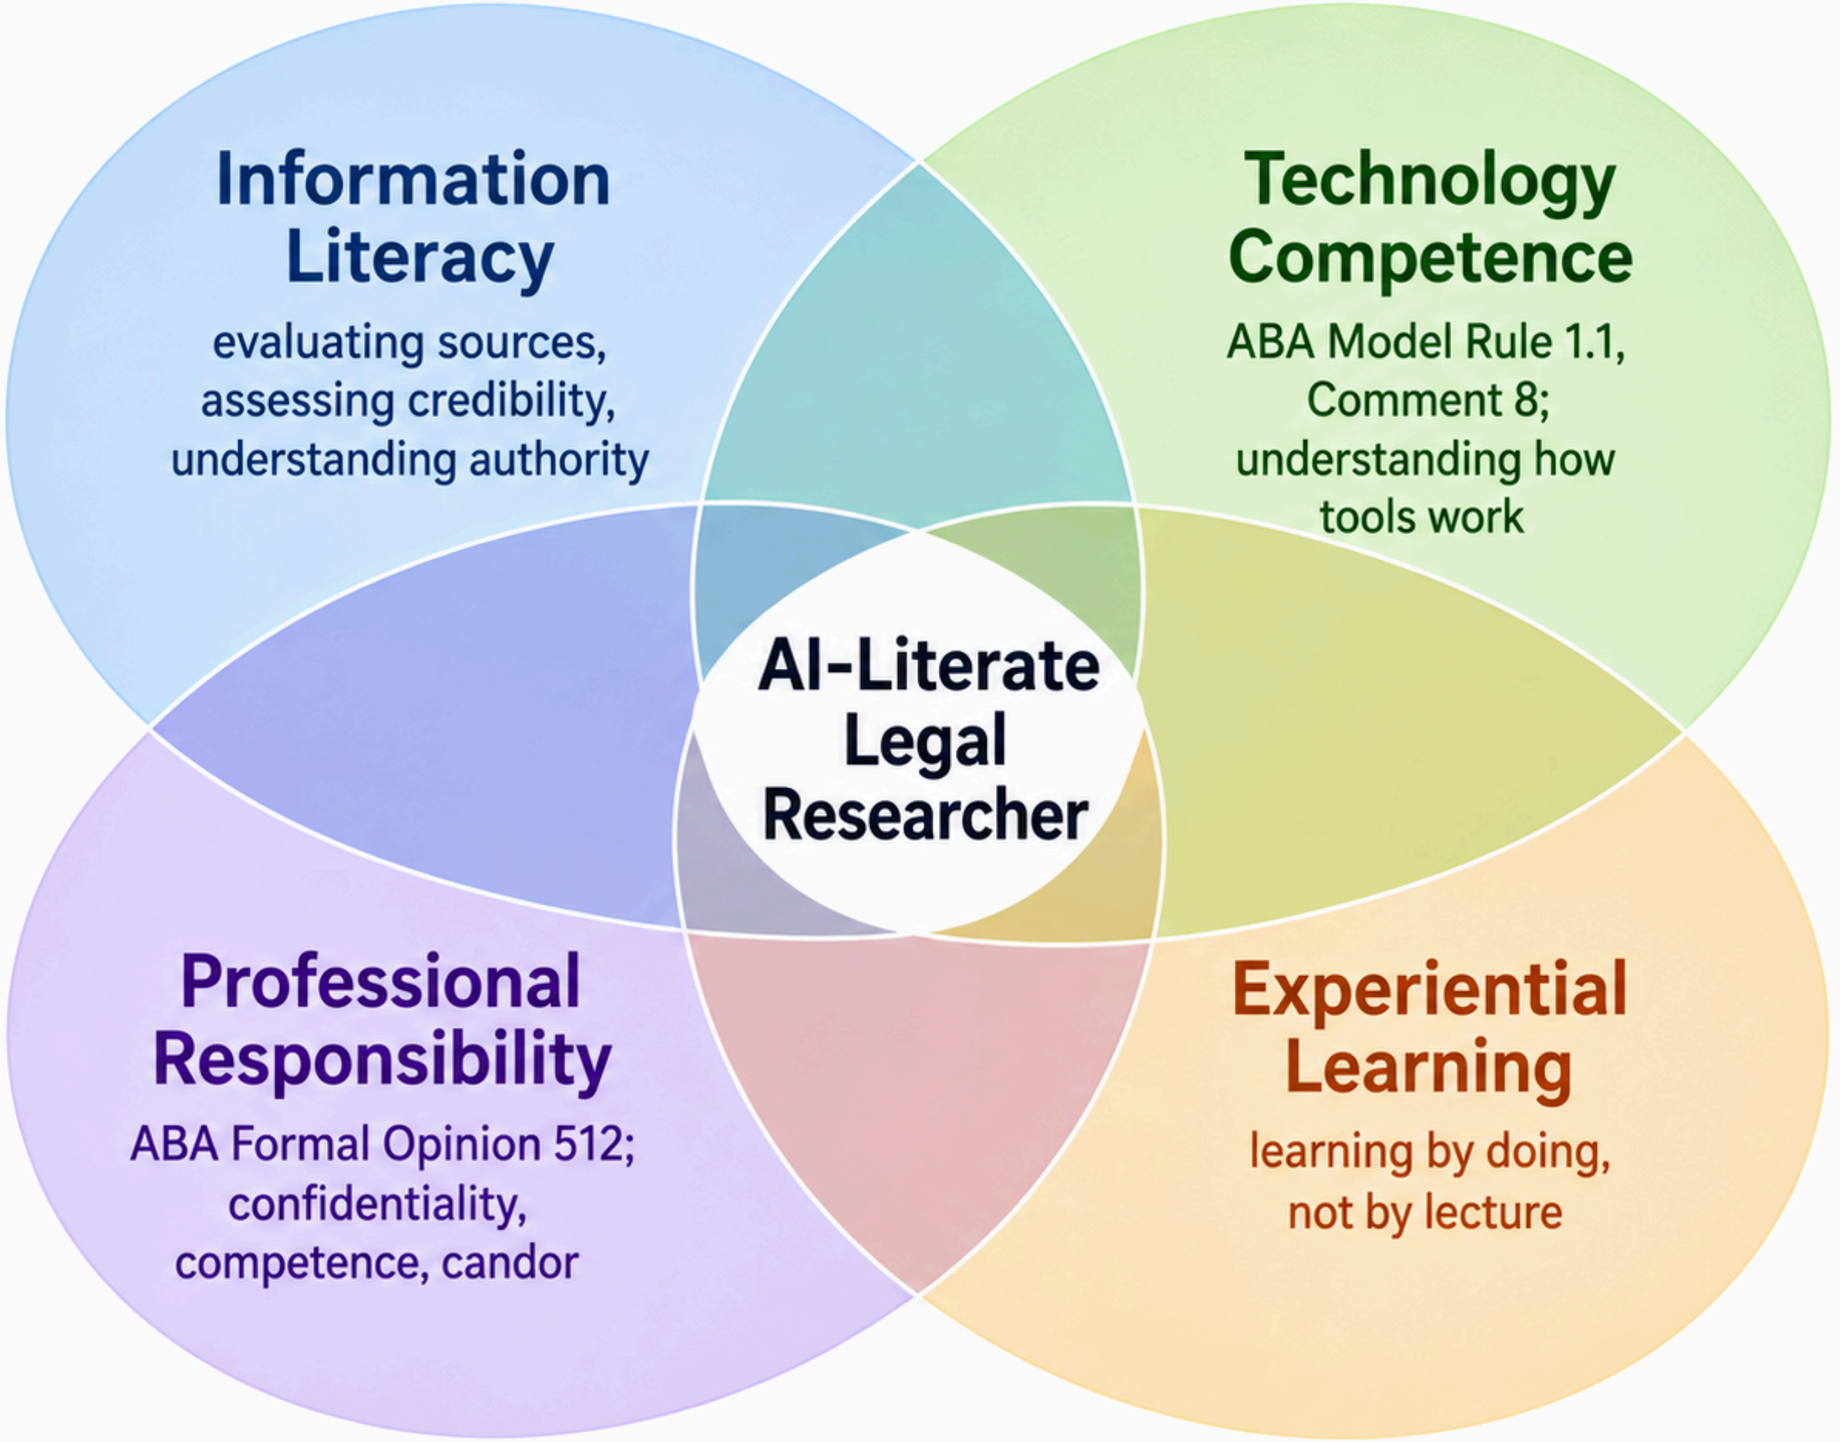
Information Literacy
evaluating sources, assessing credibility, understanding authority
Technology Competence
ABA Model Rule 1.1, Comment 8; understanding how tools work
AI-Literate Legal Researcher
Professional Responsibility
ABA Formal Opinion 512; confidentiality, competence, candor
Experiential Learning
learning by doing, not by lecture

FIVE CORE COMPETENCIES

“What do I need to teach my students for them to be AI-literate?”

Students who complete AI-integrated legal research instruction should be able to:

1. Use AI tools to generate research leads and identify relevant sources.
2. Critically evaluate AI-generated content for credibility, accuracy, and completeness.
3. Verify AI outputs against authoritative primary and secondary sources.
4. Explain why AI fails and revise research strategies when it does.
5. Understand ethical constraints on AI use in legal practice.

These map directly to ABA Standards 301, 302, and 303, which already require competent legal analysis, professional responsibility awareness, and practice-ready skills. We are not reinventing accreditation requirements. We’re meeting them through a modern lens.

AI LITERACY CHECKLIST

“What do I need to teach my students for them to be AI-literate?”

1. Traditional Legal Research Skills Are the Foundation
2. AI Literacy Is Information Literacy
3. The Durable Skill Is Research Judgment
4. AI Literacy Is Learned by Doing
5. AI Is Not a Source of Law
6. Distinguish AI-Assisted from AI-Generated
7. Teach the AI Stack in Plain Language
8. Evaluate Tools Beyond Vendor Marketing
9. Cognitive Biases Distort AI-Assisted Research
10. Ethics Apply in Full Force
11. Plan Before You Prompt
12. Prompting Is a Legal Research Skill
13. Verification Is Non-Negotiable
14. Document Everything
15. Hybrid Workflows

THE CORE WORKFLOW: PLAN, PROMPT, VERIFY, DOCUMENT

Step 1 Plan: frame the issue, identify jurisdiction, map research questions, choose tools

Step 2 Prompt: craft effective queries using the Five-Element Prompt Framework

Step 3 Verify: Check every source, claim, quotation, and legal proposition. Confirm holdings, use citators, trace to primary authority

Step 4 Document: record what tools were used, what was relied on, verification steps taken.

Teaching point: This workflow works across tools because it is built on research principles.

- Plan: What exactly am I researching, and what law may control?
- Prompt: How can AI help me build a research path without replacing my judgment?
- Verify: What must I check in authoritative sources before relying on this?
- Document: What did I do, what did I verify, what remains uncertain, and what must I disclose?

AI LITERACY IS LEGAL RESEARCH LITERACY

AI literacy is not separate from legal research literacy. AI literacy is legal research literacy in a new environment.

It is the ability to understand how AI affects:

- How legal questions are framed
- How sources are identified
- How authority is evaluated
- How results are verified
- How uncertainty is documented
- How ethical risks are managed

Key message for students: AI literacy means applying traditional legal research judgment in a new technological environment.

INSTRUCTION BEYOND TOOL DEMOS

Do not make AI instruction only a tool demo.

Students **do not** need:

- A tour of every platform
- A list of every new AI product
- A one-time lecture on hallucinations
- A warning-only approach

Students **need**:

- Transferable workflows
- Verification habits
- Prompting discipline
- Tool discernment
- Ethical judgment
- Practice

THE INSTRUCTOR'S ROLE

“What do I need to teach my students for them to be AI-literate?”

Our job is to teach students to remain in control.

AI can assist.
AI can accelerate.
AI can organize.
AI can suggest possibilities.

But students must still:

- Know what law controls
- Read the law
- Verify the law
- Update the law
- Find adverse authority
- Document the process
- Exercise judgment

START WITH THE RESEARCH FOUNDATION

Before students can supervise AI, they must understand:

- Sources of law
- Primary and secondary authority
- Hierarchy of authority
- Jurisdiction
- Binding vs. Persuasive authority
- Citators and updating
- Procedural posture
- Standard of review
- Elements and material facts
- Published, unpublished, precedential, and nonprecedential opinions

Teaching Tip: Teach these concepts as AI safeguards, not as background vocabulary.

AI MAKES FOUNDATIONAL CONCEPTS MORE IMPORTANT

AI can blur distinctions students must know how to catch:

- Treating persuasive authority as binding
- Mixing federal, state, and out-of-state law
- Ignoring procedural posture
- Missing standards of review
- Summarizing secondary sources as if they were law
- Presenting outdated authority as current
- Quoting broad language without identifying holding versus dicta

Takeaway: The more powerful the tool, the more important the student's research judgment becomes.

AI AMPLIFIES TRADITIONAL RESEARCH ERRORS

AI does not create entirely new research mistakes.

It makes old mistakes faster, more polished, and harder to spot.

AI can amplify:

- Wrong jurisdiction
- Overbroad issue framing
- Unsupported propositions
- Failure to update
- Missing adverse authority
- Misread holdings
- Treating dicta as law
- Confusing persuasive and binding authority

Teaching point: AI can be wrong in ways that look right.

LEGAL RESEARCH IS A PROCESS, NOT A SEARCH BAR

Legal research is not simply asking a question and receiving an answer.

It is an iterative process of:

- Understanding the facts
- Identifying legal issues
- Determining jurisdiction
- Selecting source types
- Locating authority
- Reading and evaluating authority
- Updating and verifying
- Searching for adverse authority
- Synthesizing and documenting

Teaching point: AI can support the process, but it cannot replace the process.

LEGAL RESEARCH IS MORE THAN FINDING CASES

Students should see legal research as including:

- Fact investigation and issue identification
- Research planning and strategy
- Primary and secondary authority
- Statutory, regulatory, and administrative research
- Procedural and jurisdictional research
- Adverse authority and weaknesses
- Policy, factual, and interdisciplinary research
- Remedies and practical consequences
- Synthesis, analysis, and documentation

Takeaway: If students define research narrowly, they will use AI narrowly and miss major research tasks.

AI IS NOT A SOURCE OF LAW

AI may:

- Summarize law
- Suggest issues
- Generate search terms
- Draft research plans
- Identify possible authorities
- Produce explanations

But AI is not itself legal authority.

Rule for students: Unless a legal proposition can be traced to a real, authoritative, current, relevant source, it cannot be relied upon.

SOURCES OF LAW VS. RESEARCH TOOLS

Students must distingush between:

Sources of Law

- Constitutions
- States
- Regulations
- Cases
- Court rules
- Administrative decisions
- Executive orders
- Ordinances

Research tools and aids:

- AI summaries
- Headnotes
- Citators
- Digests
- Treatises
- Databases
- Search results

Teaching Point: Research tools help students find and understand the law. They are not substitutes for law.

SPEED WITHOUT DIRECTION IS NOT EFFICIENCY

AI can make a flawed research path efficient.

High-risk pattern:

1. Student starts with an unclear issue.
2. Student gives AI incomplete facts or no jurisdiction.
3. AI produces a polished answer.
4. Student mistakes fluency for reliability.

Takeaway: Efficiency requires direction, judgment, and verification.

THE DISAPPEARING APPRENTICE PROBLEM

Legal research judgment develops through practice, mistakes, feedback, and repetition.

AI creates risk when:

- Students use AI to skip the underlying research process
- New lawyers receive fewer research assignments
- Foundational research training is replaced by polished outputs
- Students never learn to recognize gaps, errors, or unsupported claims

Takeaway: AI cannot replace the training process that builds research judgment.

ORAL PRESENTATIONS OF RESEARCH RESULTS

- Students present their research process and findings in a meeting with an instructor. The emphasis is not only the final answer, but how the student got there. Instructors evaluate the research process, the findings, and the quality of the presentation. Students also complete a self-assessment.
- **Deliverables:**
  - Research several assigned legal issues without relying on generative AI.
  - Present the answer and explain the research path.
  - Explain what sources were used, what was verified, and how the student knew the research was complete.

THE DURABLE SKILL: RESEARCH JUDGMENT

Research judgment is the ability to decide what matters during the research process.

Students need to know:

- Where to start
- When to use secondary sources
- When to move to primary law
- When to broaden or narrow a search
- When to change tools
- When to update
- When to keep researching
- When the answer is sufficiently reliable

Takeaway: The goal is not AI dependence. The goal is supervised, judgment-driven AI use.

THE GOAL IS NOT TOOL MASTERY

Students do not need to learn every AI product.

They need transferable skills:

- How to evaluate AI tools
- How to decide whether AI is appropriate for a task
- How to prompt with legal research discipline
- How to verify sources and legal propositions
- How to protect confidentiality and privilege
- How to document AI-assisted research

Takeaway: Teach durable research habits, not temporary platform features.

AI LITERACY VS. AI TOOL FLUENCY

AI literacy: Understanding how AI works, evaluating outputs, recognizing limitations, and supervising AI-assisted work.

Tool fluency: Knowing which tool fits which task and why.

Examples:

- General-purpose AI: brainstorming, vocabulary, issue spotting
- Legal-specific AI: source-linked research, citations, jurisdictional filtering
- Enterprise AI: confidential or institutional work with stronger protections
- Official sources and databases: verification, current law, authoritative text

Teaching point: The question is not whether a tool is impressive. The question is whether it is appropriate.

TEACH THE AI STACK IN PLAIN LANGUAGE

Students need practical AI vocabulary, not computer science.

Teach:

- AI as the umbrella term
- Machine learning as pattern recognition from data
- Deep learning as more complex layered pattern recognition
- Natural language processing as AI working with human language
- Large language models as systems that generate text by predicting language patterns
- Generative AI as systems that create new text or content
- Agentic AI as systems that can perform multi-step tasks with more autonomy
- AI is more than a chatbot, it is in e-discovery, search ranking, case management, billing

Teaching point: Students need enough technical literacy to know what questions to ask about the tool.

KEY TECHNICAL CONCEPTS FOR LEGAL RESEARCHERS

Students should understand:

- **Tokens:** pieces of text the system processes
- **Training data:** materials used to build the model's patterns
- **Fine-tuning**: additional training for a specific task or domain
- **Grounding / RAG:** connecting output to external sources
- **Context window**: how much information the tool can consider at once
- **Hallucination:** output that is fabricated, unsupported, or wrong
- **Model vs. tool vs. interface:** the system, the product, and the way users interact with it

Teaching point: Familiarity is not understanding. Students need functional literacy.

AI LITERACY AND AI STACK EXERCISES

Mapping The AI Stack: Below is a list of AI features a legal researcher might encounter. For each, identify where it falls in the AI stack, such as machine learning, NLP, generative AI, extractive AI, RAG, or agentic AI

Questions

1. A legal research platform highlights the most relevant paragraphs of a case based on your search query.
2. A chatbot drafts a three-paragraph summary of a statute you pasted into the prompt.
3. A document review tool classifies 50,000 documents as "responsive" or "non-responsive" during discovery.
4. A tool retrieves three cases from a legal database, then generates an analysis synthesizing them.
5. An AI assistant creates a research plan, searches multiple databases, compares results, and drafts a preliminary memo.
6. A contract analysis tool extracts every indemnification clause from a set of vendor agreements.
7. A tool predicts the likely outcome of a motion to dismiss based on historical data from similar cases.

TOOL CHOICE IS A RESEARCH DECISION

Before choosing a tool, students should ask:

- What task am I trying to complete?
- Does this tool retrieve, summarize, generate, classify, or recommend?
- What sources can it access?
- Are the sources current?
- Does it provide citations or links?
- Can I inspect the underlying authority?
- Does it distinguish binding from persuasive authority?
- Can I verify the output outside the tool?
- What happens to my data?

Takeaway: Tool evaluation is part of legal research competence.

TOOL DISCERNMENT CHECKLIST

1. Appropriateness: Is AI suitable for this task?
2. Confidentiality: Does the prompt include sensitive information?
3. Source Grounding: Is the output tied to identifiable, current authority?
4. Verification: Can every proposition be independently checked?
5. Jurisdiction: Is the tool applying the correct jurisdiction?
6. Disclosure: Do court rules or policies require disclosure?
7. Billing: Is time saved being billed reasonably?

THE CORE WORKFLOW: PLAN

Before prompting, students should identify:

- The assignment
- The client/problem facts
- Missing facts
- Legal issues
- Jurisdiction
- Procedural posture
- Source types needed
- Likely search terms
- Time sensitivity
- Ethical or confidentiality risks

Teaching point: Students should not use AI to figure out what they should have planned first.

THE CORE WORKFLOW: PROMPT

When prompting, students should:

- Use nonconfidential facts
- Specify jurisdiction
- Ask for issues, source types, search terms, and research steps
- Ask AI to flag uncertainty
- Tell AI not to invent citations
- Avoid asking for final legal conclusions too early
- Use follow-up prompts to refine the research path

Teaching point: The prompt should keep the student in control of the research process.

PROMPTING IS LEGAL RESEARCH COMMUNICATION

Prompting is not magic.
Prompting is structured communication with an AI system.

Good prompting requires:

- Knowing the research task
- Identifying relevant facts
- Specifying jurisdiction
- Naming procedural posture when relevant
- Asking for source types or research paths
- Setting limits
- Requiring uncertainty to be flagged

Teaching point: A better prompt creates a better starting point, not a final answer.

PROMPTING AS A LEGAL RESEARCH SKILL

The Prompt Framework

1. Role: “You are a legal research assistant helping a litigation attorney.”
2. Facts: Provide relevant facts (without disclosing confidential information)
3. Jurisdiction: Always specify. Without it, AI blends jurisdictions.
4. Task: Be specific. “Find cases about negligence” vs. “Identify binding appellate authority in NY addressing commercial landlord duty to protect tenants from third-party criminal acts.”
5. Limits: “Do not fabricate citations. If uncertain, say so. Limit to federal appellate decisions.”

Key Prompting Rules:

- Plan before you prompt. Break complex tasks into separate prompts.
- Ask for research assistance, not legal conclusions. Ask for verification tasks, not verified answers.
- Iterative prompting: never stop at the first output
- Context engineering: brief AI like a colleague, not a search bar

THE CORE WORKFLOW: VERIFY

After receiving AI output, students should:

- Locate all cited sources in reliable databases or official sources
- Read the actual sources
- Check citations and quotations
- Classify authority as binding or persuasive
- Run citators
- Check statutory and regulatory currency
- Search for adverse authority
- Confirm factual and procedural fit

Teaching point: AI output becomes useful only after independent verification.

VERIFICATION IS THE CORE SKILL

Students must verify:

- The source exists
- The citation is correct
- The quotation is accurate
- The source says what AI claims
- The authority is from the correct jurisdiction
- The authority is binding or persuasive
- The law is current
- The procedural posture fits
- The facts are materially similar
- Contrary authority has been checked

Teaching point: Finding that a case exists is not enough.

VERIFICATION IS MORE THAN CITATION CHECKING

Students must ask:

- Does this source actually support the proposition?
- Is the quoted language complete and accurate?
- Is the language holding or dicta?
- Has the case been limited, distinguished, criticized, or overruled?
- Has the statute been amended?
- Are there regulations, rules, or local practices that change the answer?
- What would the other side cite?

Takeaway: Verification is substantive, not clerical.

VERIFICATION CHECKLIST

1. Identify every factual claim
2. Identify every legal citation
3. Check whether cited cases exist (in a legal database, not Google)
4. Read the actual cases — do not trust AI's characterization
5. Verify statutory citations for accuracy and currency
6. Check regulatory citations against official sources
7. Verify quotations word-for-word
8. Check whether law is still good (citators — AI cannot do this)
9. Verify procedural and jurisdictional accuracy
10. Cross-reference with independent research
11. Check dates and temporal accuracy
12. Test the reasoning (does the conclusion follow from the authority?)
13. Document your verification

THE CORE WORKFLOW: DOCUMENT

Students should document:

- Research question
- Facts and assumptions
- Jurisdiction
- Procedural posture
- AI tool used
- Prompts used
- Sources consulted
- Citator results
- Verification steps
- AI errors or omissions
- Remaining uncertainty
- Disclosure analysis, if required

Teaching point: In the AI era, students must show their work.

AI DISCLOSURE: TEACH IT AS A HABIT

- Obligations vary by jurisdiction, court, and judge
- Sources: professional conduct rules, court rules, standing orders, local rules, client instructions
- Build disclosure into the Document step of Plan, Prompt, Verify, Document

Teaching Tips:

- Require AI disclosure on every assignment based on your jurisdiction
- Require students to submit: what tools were used, what they were used for, what verification steps were taken
- This builds the habit before they enter practice

COGNITIVE BIASES AND THE FLUENCY PROBLEM

Biases That Distort AI-Assisted Research:

- Automation Bias
- Confirmation Bias
- Anchoring Bias
- Overconfidence / Premature Closure
- Cognitive Surrender
- The Fluency Problem

Common Human Workflow Errors:

- Vague prompting
- Missing jurisdiction
- No procedural posture
- Copying AI text too quickly
- Skipping citators
- No research trail
- Wrong tool for confidential information

PRACTICAL TAKEAWAYS FOR YOUR CLASSROOM

- Teach foundational legal research concepts with its AI connection: “This is why understanding jurisdiction matters even more now.”
- Adopt Plan, Prompt, Verify, Document as your course’s research method
- Require AI use disclosures on every assignment
- Prioritize transferable skills over specific tools (tools change, skills transfer)
- Assess the process (research process, research log, verification steps) not just the output