Teaching Legal Research in the Age of AI

Classroom Exercise Bank — Companion Handout

Lauren E. Diaz

This companion handout turns the framework from the presentation into ready-to-adapt teaching exercises. Every exercise is mapped back to the four AI-literacy competencies (information literacy, technology competence, professional responsibility, experiential learning), the Five Core Competencies, and ABA Standards 301–303. Each entry includes facilitation notes and a short cue for what to assess.

Foundations as AI Safeguards

Theme from the deck: Before students can supervise AI, they must understand the research foundation (e.g., sources of law vs. tools, hierarchy of authority, jurisdiction, posture, holding vs. dicta. Teach these as AI safeguards, not background vocabulary). The more powerful the tool, the more important the student's research judgment.

Exercise: Source, Tool, or Not Authority?

For each item, label it (a) primary authority, (b) secondary authority, or (c) a research tool or aid. State whether it can be cited to a court as law. For anything that is a tool or an AI output, write one sentence on what the researcher must do before relying on it.

1. A municipal noise ordinance pulled from the city's official online code.
2. An AI chatbot's paragraph summarizing "the elements of adverse possession in your state."
3. A commercial-database headnote describing a court's holding on standing.
4. A state intermediate appellate decision applying a comparative-negligence statute.
5. A practitioner treatise chapter on mechanic's liens.
6. A citator entry showing a statute was amended last legislative session.
7. An AI-generated "research memo" comparing two contract clauses.
8. A federal agency's final rule published in the Federal Register.
9. A digest topic-and-key-number list of negligence cases.
10. A law professor's blog post analyzing a recent grant of certiorari.

Assessment: Forces students to separate sources of law from research tools and aids, and to internalize that an AI summary is never itself authority. Reinforces the rule that a legal proposition is unreliable until traced to a real, authoritative, current, relevant source.

Assess for: Correct classification; accurate statement of citability; quality of the "verify-first" step named for each tool.

Exercise: The Jurisdiction-Blend Audit

Using any general-purpose AI tool, ask one question that spans more than one body of law (e.g., "Can a landlord keep a tenant's security deposit for normal wear and tear?") without naming a jurisdiction. Then audit the answer:

1. Did the tool name a controlling jurisdiction, or did it blend states and/or federal law?
2. Identify every place the answer mixed rules from different jurisdictions or presented a "general rule" as if it controlled.
3. Re-run the question with your jurisdiction specified. What changed?
4. Locate, in an official source, the specific statute or rule in your jurisdiction that actually governs.
5. Write two sentences a supervising attorney would accept explaining why the first answer was unreliable.

Assessment: Demonstrates the deck's point that AI blurs jurisdictional lines and that specifying jurisdiction is itself a research skill. Trains students to catch blended authority — the single most common AI research error.

Assess for: Whether the student located the controlling source independently; precision in identifying the blend; clarity of the explanation.

Exercise: Holding or Dicta?

The instructor supplies a short appellate opinion and an AI-generated one-paragraph summary of it. Tasks:

1. Identify the procedural posture and the precise question the court decided.
2. Mark every sentence in the AI summary as holding, dicta, or inaccurate.
3. Find one broad quote the AI used and explain whether it is part of the holding or merely dicta.
4. Rewrite the rule statement so it reflects only what the case actually held at that posture.

Assessment: Targets the deck's warning that AI quotes broad language without separating holding from dicta and ignores procedural posture. Builds the close-reading judgment that lets students supervise AI rather than defer to it.

Assess for: Accuracy of the holding/dicta line; correct identification of posture; quality of the rewritten rule.

Exercise: Familiarity

Your supervising attorney hands you the memorandum below. It concerns the Carmack Amendment, the federal statute governing an interstate motor carrier's liability for damaged freight. The memo is clean, confident, and formatted like competent work. Evaluate it as if you had to rely on it.

1. Before researching anything, record your gut reaction: does this read as competent and complete? Rate your confidence in it from 1 to 10.
2. Mark every statement you would need to verify before relying on it, and flag anything that feels like a red flag.
3. Now verify. Locate the statute and every cited case. Record which authorities are real, which are misstated, and which cannot be found at all.
4. Re-rate your confidence after verification. What changed, and why?
5. Compare your flags to the instructor's redline (next page).

Sample AI-Generated Memorandum [SAMPLE — CONTAINS DELIBERATE ERRORS; DO NOT RELY ON THIS AS LAW]

RE: Carrier liability for damaged freight under the Carmack Amendment

Question Presented. Whether our client, a shipper, may recover from an interstate motor carrier for goods damaged in transit, and whether the carrier has any defense.

Brief Answer. Yes. Under the Carmack Amendment, an interstate motor carrier is strictly and absolutely liable for any loss or damage to cargo it transports, regardless of fault, and no defenses are available. The law is settled in the shipper's favor.

Discussion. The Carmack Amendment, codified at 49 U.S.C. § 14501, governs the liability of interstate motor carriers for cargo loss or damage. It imposes a uniform federal rule of carrier liability that displaces all other law.

To establish a claim, the shipper need only prove two things: that the goods were damaged on arrival and the amount of the loss. Once the shipper makes this showing, liability is automatic.

The Supreme Court confirmed this rule in Missouri Pacific Railroad Co. v. Elmore & Stahl, 377 U.S. 134 (1964), holding that a carrier is liable for cargo damage in all circumstances and may not assert any excepted causes.

More recently, the Seventh Circuit reaffirmed absolute carrier liability in Trans-Global Freight Systems, Inc. v. Apex Logistics, 612 F.3d 884 (7th Cir. 2010), where the court "rejected every defense a carrier might raise and imposed liability without exception."

This rule is binding nationwide, and our jurisdiction follows it. See Coastal Carriers Co. v. Bayline Shipping, 88 Cal. App. 4th 1201 (2001). Because the carrier is strictly liable and can raise no defense, the shipper may also pursue a parallel state-law negligence claim for the full amount.

Leading commentators agree that this is the controlling rule, which is dispositive here. See Treatise on Freight Liability § 7.04.

Conclusion. The carrier is absolutely liable. The law is settled and no further research is necessary.

Instructor Redline Key: What Is Actually Wrong
Reveal this after students have submitted their own flags. Every authority is either fabricated, misstated, or misused, yet the memo reads as polished, confident work.

Flag 1: Wrong statute (source verification). "codified at 49 U.S.C. § 14501." The Carmack Amendment's carrier-liability provision is 49 U.S.C. § 14706. Section 14501 addresses federal preemption of state economic regulation — a different subject. A correctly formatted citation can still point to the wrong section. Pull the statute itself first.

Flag 2: Overbroad, absolute rule (scope of authority). "strictly and absolutely liable … no defenses are available." Carmack liability is strict but not limitless. A carrier can avoid liability by showing it was free from negligence and that the loss was caused by one of five recognized excepted causes (act of God, the public enemy, the act or default of the shipper, public authority, or the inherent nature of the goods). Absolute-liability, no-defense framing is a red flag — verify against the statute and controlling cases.

Flag 3: Missing element. "the shipper need only prove two things." The prima facie case has three elements, not two: delivery to the carrier in good condition, arrival in damaged condition, and the amount of damages. A missing element feels "off" only if you already know the rule.

Flag 4: Misstated holding (read the case). Missouri Pacific Railroad Co. v. Elmore & Stahl is a real case, but the memo inverts it. Elmore & Stahl preserves the carrier's ability to rebut the prima facie case by proving freedom from negligence and an excepted cause; it does not hold that carriers have no defenses. Never trust an AI's characterization of a holding — open and read the opinion.

Flag 5: Fabricated case and quotation. Trans-Global Freight Systems, Inc. v. Apex Logistics, 612 F.3d 884 (7th Cir. 2010), and its quotation cannot be located in any database. The citation looks perfect — volume, reporter, circuit, year, and a pinpoint quote — which is exactly why it must be confirmed in Westlaw or Lexis, not Google, and never relied on unless opened.

Flag 6: Wrong jurisdiction; nonbinding cited as binding. "binding nationwide … See Coastal Carriers Co. v. Bayline Shipping, 88 Cal. App. 4th 1201 (2001)." A state intermediate appellate decision cannot be binding nationwide and cannot control a federal statutory question. Confirm the case exists, then classify it as binding or persuasive in the actual forum.

Flag 7: Preemption error. "the shipper may also pursue a parallel state-law negligence claim." The prevailing rule is the opposite: the Carmack Amendment broadly preempts state and common-law claims for loss of or damage to goods moving in interstate commerce. Flag and verify — this statement is likely wrong.

Flag 8: Secondary source treated as binding. "Leading commentators agree … which is dispositive here. See Treatise on Freight Liability § 7.04." A treatise is secondary authority. It can help you find and understand the law; it is never "dispositive" and cannot be cited as binding.

Flag 9: Premature closure / overconfidence. "The law is settled and no further research is necessary." This is the single most reliable red flag in any AI output. Competent research rarely ends with the first confident answer.

The hardest part. The memo mixes real authority with fabricated and misstated authority. A reader who recognizes one real case may extend trust to the fabricated one sitting beside it. Fluency plus partial accuracy is what makes the whole memo dangerous.

Debrief: Familiarity Changes the Risk

Your ability to spot AI errors depends on how well you know the law.

When you are familiar with an area of law, or have researched it deeply:

- Incorrect rule statements jump out.
- Missing elements feel "off."
- Overconfident conclusions raise red flags.
- You know where to check first in Westlaw.

When you do not yet know the law well:

- AI can sound extremely convincing.
- Polished language can mask:
    - misstated standards,
    - wrong statutes,
    - nonbinding or fabricated cases,
    - overbroad or absolute rules.
- Your confidence in the output can exceed its accuracy.

AI is most persuasive precisely when you are least equipped to challenge it.

Assessments: Makes the deck's core warning visceral: AI is most persuasive when the reader knows the least. By asking students to evaluate a confident memo in an unfamiliar area before and after verification, it exposes the gap between perceived and actual reliability and shows that error-detection is a function of subject-matter knowledge plus verification discipline — not reading comprehension.

Assess for: Honesty of the before/after confidence ratings and the size of the change; number of planted errors caught; correct labeling of the foundational concept each error misapplies; whether the student used a legal database rather than a web search and actually read the sources; a reflection connecting familiarity to error-detection.

Technology Competence: The AI Stack and Tool Fluency

Theme from the deck: Students need practical AI vocabulary, not computer science — enough to know what questions to ask about a tool. Distinguish AI literacy (transferable evaluation skill) from tool fluency (which tool fits which task). Tool choice is a research decision.

Exercise: Behavior → Stack Layer → Verification Risk

For each described tool behavior, (a) name where it sits in the AI stack (machine learning, NLP, extractive AI, generative AI, RAG, or agentic), and (b) state the one verification step that matters most given that behavior's failure mode.

1. A platform pulls the exact sentence stating a case's disposition.
2. A chatbot writes a new plain-language explanation of a regulation you pasted in.
3. A tool sorts 30,000 emails into "privileged" and "not privileged."
4. A research assistant retrieves statutes from a database, then drafts an analysis citing them.
5. A system takes a research question, runs multiple searches, compares results, and returns a memo with proposed next steps.
6. A contract tool flags every limitation-of-liability clause across a deal folder.
7. A tool estimates the odds a summary-judgment motion succeeds based on past rulings.

Assessment: Extends "Teach the AI Stack in Plain Language" by pairing each layer with its distinct risk, so students learn that different architectures demand different verification — retrieval failure (RAG) vs. fabrication (generative) vs. misclassification (ML classifier).

Assess for: Correct stack placement; whether the named verification step actually fits that architecture's failure mode.

Exercise: One Question, Three Tools

Pick one narrow research question in an area you do not know well. Run it through (1) a general-purpose chatbot, (2) a legal-specific research platform, and (3) an official source or database. For each tool, record:

- What it returned, and whether it provided checkable citations or links.
- Whether you could inspect the underlying authority.
- What the tool is good and bad at for this task.

Then answer: Which tool was appropriate for which stage of the research, and why? Which would you never use for confidential client facts, and why?

Assessment: Operationalizes "Tool Choice Is a Research Decision" and "AI Literacy vs. Tool Fluency." Teaches matching the tool to the task and to the sensitivity of the information, and shows that competence is not tied to one product.

Assess for: Discernment about fit and confidentiality; whether the student verified against an official source rather than trusting the chatbot.

Cognitive Bias and the Fluency Problem

Theme from the deck: AI can be wrong in ways that look right. Polished, confident, fluent output reduces skepticism. Students must learn to identify and counter automation bias and the fluency problem, and to see that speed without direction is not efficiency.

Exercise: The Polished Wrong Answer

Setup: Provide, or generate live, a confident, well-formatted AI "research summary" that contains several embedded errors of the kinds the deck lists: a trial-court order treated as binding precedent, two states' rules merged, dicta presented as the rule, a superseded statute presented as current, and a closing line that no further research is necessary. Do not tell students in advance how many errors there are.

Sample summary (hypothetical: replace the placeholder citation with a freshly generated, clearly fictional output before class

"Under settled law, a homeowner is strictly liable for any injury caused by a dog. In Doe v. Roe, No. 22-CV-114 ([County] Trial Ct. 2024), the court held that breed is irrelevant, and courts statewide follow this rule. Neighboring-state decisions agree, confirming the rule is uniform. Because the law is settled, no further research is necessary."

Tasks:

1. Read it as if it were a colleague's polished memo.
2. List every error you can find and name the foundational concept the AI misapplied (jurisdiction, binding vs. persuasive, posture, holding vs. dicta, currency).
3. Note where the fluency of the writing made an error easy to miss.
4. Compare your error list to the instructor's redline and reflect on what you missed and why.

Assessment: Directly teaches automation bias and the fluency problem — that confident, polished output lowers a reader's guard. Demonstrates that AI "can be wrong in ways that look right." Works best when students are not tipped off that the memo is flawed.

Assess for: Number and accuracy of errors caught; correct labeling of the misapplied concept; quality of the reflection on why fluency masked the error.

Exercise: Speed Without Direction

Reflect on a time you (or a hypothetical junior associate) got a fast, polished AI answer to a vague question. Walk through the four-step high-risk pattern: unclear issue → incomplete facts or no jurisdiction → polished answer → fluency mistaken for reliability. Then rewrite the opening move: what should the researcher have planned before prompting, and what single planning step would have prevented the wrong turn?

Assessment: Reinforces "Speed Without Direction Is Not Efficiency" — that real efficiency requires direction, judgment, and verification, not just a fast answer.

Assess for: Insight into where judgment was needed; concreteness of the corrective planning step.

Plan

Theme from the deck: The Core Workflow begins with Plan — frame the issue, identify jurisdiction, map research questions, choose tools. Students should not use AI to figure out what they should have planned first.

Exercise: Plan Before You Prompt (Three Lists)

Scenario: A small bakery in your state receives a cease-and-desist letter claiming its logo infringes a competitor's unregistered mark. The owner wants to know whether she must stop using the logo and whether she has a defense. Before touching AI, build three lists:

1. List 1: Facts you know.
2. List 2: Facts that are missing or unclear.
3. List 3: Questions you must answer through research.

Then identify the jurisdiction, the likely source types (statute, case law, secondary), and three search terms. Only after that, write an orientation prompt asking AI to help fill gaps — and mark which AI suggestions you adopt, modify, or reject.

Assessment: Implements "The Core Workflow: Plan" and "Plan Before You Prompt." Builds the discipline of forming an independent baseline against which AI output can be evaluated.

Assess for: Quality of the missing-facts list; whether jurisdiction and source types were set before prompting; critical handling of AI's suggestions (adopt/modify/reject).

Prompting

Theme from the deck: Prompting is structured communication with an AI system and a modern form of query formulation. The Five-Element Prompt Framework (Role, Facts, Jurisdiction, Task, Limits) keeps the student in control. A better prompt creates a better starting point, not a final answer.

Exercise: Weak Prompt vs. Five-Element Prompt

Choose a research question. Write Prompt A as a one-line query ("research [topic]"). Write Prompt B using the five-element framework: Role; Facts (nonconfidential); Jurisdiction; Task; Limits (no fabricated citations, flag uncertainty, restrict source type and date). Run both. Compare scope, specificity, and reliability, and list any issue Prompt B surfaced that Prompt A missed. End by naming three things you would still verify regardless of prompt quality.

Assessment: Teaches prompting as modern query formulation and the five-element framework, and reinforces that even an excellent prompt yields a starting point that still requires verification.

Assess for: Whether Prompt B actually constrains jurisdiction, source, and limits; quality of the comparison; recognition that verification is still required.

Exercise: Prompt Surgery (Advocacy → Neutral Research)

You are handed this prompt: "Find cases proving our client's non-solicitation clause is fully enforceable." Tasks:

1. Identify why this prompt invites confirmation bias and one-sided output.
2. Rewrite it as a neutral research prompt that asks for binding and adverse authority, exceptions, and unresolved questions.
3. Explain what you would lose by running the original prompt before doing neutral research.

Assessment: Targets confirmation bias and the distinction between advocacy prompts and neutral research. Reinforces the deck's rule: ask for research assistance, not legal conclusions.

Assess for: Whether the rewrite genuinely seeks adverse authority; articulation of the bias risk.

Verification

Theme from the deck: Verification is the core skill. Finding that a case exists is not enough — verification is substantive, not clerical. Students must read the actual sources, run citators (AI cannot), and confirm the source supports the proposition.

Exercise: Full Verification Run

Take any AI-generated answer to a legal question (your own or instructor-supplied). Run the verification checklist and record a result for each step:

1. Identify every factual claim.
2. Identify every legal citation.
3. Confirm each cited case exists in a legal database (not a web search).
4. Read each case — do not trust the 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. Run citators for negative treatment (overruled, limited, distinguished, criticized).
9. Confirm jurisdiction and procedural posture.
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.

Flag at least one instance where a source exists but does not actually support the proposition the AI assigned to it.

Assessment: Operationalizes "Verification Is the Core Skill" and "Verification Is More Than Citation Checking." Drives home that confirming existence is not the same as confirming support, currency, or weight.

Assess for: Completeness of the checklist run; whether the student read sources rather than trusting characterizations; identification of at least one "exists but doesn't support" gap.

Document and Disclose

Theme from the deck: In the AI era, students must show their work. Build AI disclosure into the Document step; obligations vary by jurisdiction, court, and judge. Require disclosure on every assignment so the habit forms before practice.

Exercise: The Research Trail + AI Disclosure

For a completed research task, produce a one-page research trail recording each of the following:

- Research question; facts and assumptions; jurisdiction; procedural posture.
- AI tool(s) and version used; every prompt submitted; key outputs received.
- Sources consulted; citator results; verification steps taken.
- Any AI errors or omissions you caught; remaining uncertainty.
- Disclosure analysis: do your jurisdiction's court rules, standing orders, or this assignment require disclosure, and what exactly would you disclose?

Assessment: Implements "The Core Workflow: Document" and "AI Disclosure: Teach It as a Habit." Makes the process visible and reviewable, and builds the disclosure reflex.

Assess for: Completeness and honesty of the trail; quality of the disclosure analysis; whether errors and uncertainty were actually surfaced.

Ethics and Professional Responsibility

Theme from the deck: Students must control what information goes into a tool, where it goes, and how it is stored; check each tool's terms, privacy policy, and retention settings; and treat AI literacy as a continuing obligation that evolves with court rules, bar guidance, and tools.

Exercise: Confidentiality Tool-Matching

Scenario: A partner asks you to analyze a dispute file containing a sealed mediation memo, board emails about an unannounced acquisition, and a spreadsheet of customer personal information. Three tools are available: (A) a free consumer chatbot with no enterprise agreement; (B) an enterprise platform whose contract bars training on inputs; (C) a legal-specific tool that requires uploading documents to a vendor server. Tasks:

1. For each tool, state which of the three materials (if any) you could enter, and why.
2. Identify what you would check in each tool's terms of service, retention, and data-use settings before entering anything.
3. Draft a two-sentence recommendation to the partner.

Assessment: Builds the confidentiality judgment the deck stresses, deliberate control over what goes into a tool, where it goes, and how it is stored, and the habit of checking terms and retention before use.

Assess for: Correct risk ranking of the three materials; specificity of the terms/retention checks; soundness of the recommendation.

Exercise: Staying Current (The Ongoing Obligation)

Identify the jurisdiction where you intend to practice. Locate and summarize: (a) any court rule, standing order, or local rule addressing AI use or citation certification; (b) one state bar or ABA ethics opinion on generative AI; and (c) one development from the last few months (an order, guidance, or news item). Then write a short plan for how you will keep this current after graduation.

Assessment: Reinforces that AI literacy is an ongoing professional habit, not a one-time lesson, and that obligations evolve by jurisdiction, court, and judge. Builds the monitoring discipline.

Assess for: Accuracy and currency of the sources found; whether the student located primary guidance rather than only commentary; realism of the monitoring plan.

Hybrid Workflow

Theme from the deck: AI literacy is learned by doing — use it, test it, challenge it, break it, verify it. Assess the process, not just the output. Traditional skills are the map; AI is the speed.

Exercise: Hybrid Workflow

Given a realistic client fact pattern and a research question, complete the full cycle:

1. Plan (non-AI): issue, jurisdiction, posture, source types, search terms.
2. Use AI for orientation only — issues, sub-issues, and search terms — and mark what you adopted, modified, or rejected.
3. Locate and read the actual authority in reliable sources.
4. Verify everything: citations, currency, holding vs. dicta, adverse authority.
5. Produce three deliverables: a short internal memo, a research log, and an AI-use disclosure.

The grade rewards the process: the map, the verification, and the judgment (not just the final answer).

Assessment: Synthesizes the entire framework and the core premise that traditional skills are the map and AI is the speed. Mirrors practice-ready, NextGen-style performance work and the deck's call to assess the process, not just the output.

Assess for: Quality of the independent plan; disciplined AI use (orientation, not authority); thoroughness of verification; clarity of the memo; completeness of the log and disclosure.

Course-Level Assignments and Assessment

Daily “AI Current Event” Elevator Pitch (Attendance + Lifelong Learning Habit)

At the start of each class, you will give a quick one-to-two minute summary of an AI related development, using a news article, court update, professional guidance, etc. You will briefly explain why it matters for legal research or the practice of law.

Oral Presentation of Research Results (Non-AI)

In this simulation assignment, students will research several legal issues, then present their research process and findings in a meeting with an instructor. The instructor will offer feedback on the effectiveness of their research process and findings and on the quality of their presentation. Students will also complete a self-assessment.

Impact of Specific Technology Reflection

Students will analyze the impact a specific technology has had on any facet of the legal research process in legal practice. This prompt is intentionally vague, and it is encouraged that students explore a topic that interests them personally. The reflection and presentation will be graded as provided in the grading rubric on Canvas. If there are any hallucinated citations in a student’s reflection, then that student will have their grade reduced by a full letter grade.

Generative AI Document Analysis and Drafting Critique

Students will use generative AI tools to analyze and draft legal documents. Students will upload a corpus of legal materials to several generative AI tools, then query and examine the results for relevancy and accuracy. Students will also use generative AI to draft an original filing and critique the output.

Ethical AI Use Reflection

Students will examine a few cases in which legal practitioners failed to use AI tools responsibly. Students will reflect on the following questions: Why did they do wrong? How could they have avoided the negative outcome?

AI-Strategy Project

Students will develop a comprehensive AI strategy for a (real or fictitious) legal organization or legal department of your choice. This strategy should demonstrate your understanding of generative AI capabilities, implementation challenges, ethical considerations, and the unique needs of legal practice. Your strategy must be practical, well-researched, and tailored to your chosen context. The context must be legal; this should not be a general AI strategy. You may use a real organization or create a fictitious one with clearly defined characteristics. Your strategy should have the following components: (1) executive summary, (2) organizational profile, (3) strategic AI initiatives (including at least three use cases), (4) phases implementation plan, (5) ethics and risk management, and (6) change management.

Custom AI's & Agents

Students build custom AI's and agents for specific workflows (e.g., students train a custom AI to act as legal research assistants that take opposing legal positions, students build a custom chatbot that helps with a specific legal research problem).

- Build Phase: Using a no-code platform, they create a custom AI that can:
    - Answer questions with embedded primary/secondary sources.
    - Provide citations or direct users to legal databases.
    - Include disclaimers (e.g., "Not legal advice. Consult an attorney.").
- Deliverables:
    - Tested by professor/classmates).
    - A 5–10 min video demo + explanation: Why they chose the topic. How they sourced/trained the custom AI. Its limitations in accuracy.
- Learning Goal: Students show mastery of both legal research and AI integration.
- Grading:
    - Legal accuracy of responses.
    - Usability/design.
    - Reflection on ethical implications and AI limitations.

AI Evaluation (Professional Research Tool Review)

Students will evaluate generative AI research platforms law firm adoption. Each student will be assigned a tool. You must run a sample research problem through the tool (e.g., "What is the standard for punitive damages in Florida bad faith insurance claims?") and evaluate information literacy (Did it cite real, relevant cases?), security/privacy (how safe is client data?), and effectiveness (Did it save time compared to traditional search?). Students will then deliver a procurement-style report advising whether a law firm should buy the tool. Through this assignment, students will develop professional judgment about which tools are reliable for legal research.

Final Exam: AI Assisted Legal Research Performance Task

For the final exam, students will complete a timed, practice-ready AI-assisted legal research performance task designed to mirror the kind of work lawyers do under modern time constraints and the skills emphasized on the NextGen Bar Exam. Students will be given a realistic client fact pattern, case file, and a set of research questions that would normally require many hours of traditional research and drafting. Students may use any AI tool(s) of their choice in additional to traditional research resources, but students remain responsible for accuracy, authority, and verification. Students' submission must include: (1) a concise research strategy log showing the search plan, key queries, and how sources were validated (including citator/currency checks); (2) an internal office memorandum that synthesizes relevant law and applies it to the facts with clear recommendations and identified research gaps; (3) a short client-facing email translating conclusions into practical next steps; (4) an AI Use Disclosure identifying what tool(s) were used, what they were used for, what information was provided, what outputs were relied on, and the specific steps taken to verify all AI-assisted results; and (5) a written work-product to be filed with the Court in accordance with the local Court rules.