Advanced Legal Research with AI Tools Syllabus

Fall 2026

Professor Lauren E. Diaz

ld1468@nova.edu

1. **Course Overview**

This course introduces students to the use of generative artificial intelligence in legal research and practice. Emphasis is placed on integrating AI with traditional research methods, evaluating accuracy, preventing errors such as "hallucinated" citations, and ensuring ethical use. Through readings, discussions, and hands-on assignments, students will gain practical experience comparing AI-generated results with conventional sources and will develop the skills to responsibly incorporate AI into legal research workflows.

2. **Learning Outcomes**

By the end of this course, students will achieve the following learning outcomes:

1. *Research legal issues thoroughly and efficiently.*

2. *Use technology to meet ethical duties of the legal profession (e.g. to address duties of confidentiality for all communications, to fulfill filing and other judicial obligations, and to keep abreast of technologies that affect accuracy of information provided to clients).*

3. *Demonstrate creativity and innovation when providing legal services.*

4. *Demonstrate strategic planning skills.*

5. *Demonstrate self-directed learning practices for life-long learning*

In addition to the faculty-approved learning outcomes, students will achieve the following:

- Analyze facts and develop efficient research plans.
- Locate and critically assess sources of information for credibility, relevance, limitations, and bias.
- Integrate generative AI-enhanced techniques into legal research and drafting workflows.
- Track, document, and organize research findings.
- Use research, writing, and case management tools proficiently while maintaining confidentiality.
- Understand ethical and professional responsibility duties when using AI in legal research.

- Recognize risks, limits, and benefits of generative AI and large language models.
- Critically assess AI outputs for accuracy, completeness, and reliability.
- Gain practical experience using a range of AI systems.
- Develop effective prompt strategies for legal research tasks.
- Use AI to enhance doctrinal research across cases, statutes, regulations, and secondary sources.
- Validate AI outputs with traditional research methods and correct errors.
- Communicate research with proper citations and transparent notes on AI use.
- Explore and evaluate the future role of AI in legal research and practice.
- Apply the verification-first method: Plan, Prompt, Verify, Document to AI-assisted research tasks, and maintain a documented research trail.
- Understand the AI Stack.

3. **Required Materials**

Lauren E. Diaz, THE AI LEGAL RESEARCHER: LEGAL RESEARCH IN THE AGE OF ARTIFICIAL INTELLIGENCE (CAROLINA ACADEMIC PRESS)

Thomas G. Martin, *Generative AI and the Delivery of Legal Services: A Law Student Textbook and Workbook for Understanding and Implementing AI in Law*

Barbara J. Busharis, Anne E. Mullins, & Suzanne E. Rowe, *Florida Legal Research 5th Edition*

4. **Artificial Intelligence Use Policy**

ABA Formal Opinion 512 recognizes that AI competency is essential for legal practice and requires lawyers to understand these tools to make informed professional judgments. Aligned with this guidance, students in this course will use generative AI tools for research, text generation, critique, and document analysis. Students must be prepared to discuss when AI use is and is not appropriate, the limitations of generative AI, which tools best suit specific tasks, ethical limitations, and identify and assess any AI-generated content.

Disclosure and verification obligations are also evolving at the court level. Students must check the current rules, standing orders, and administrative orders of any court before filing, and must independently verify every citation regardless of how AI was used.

Throughout this course, there will be assignments where you are required to use AI tools and others where the use of AI is prohibited. I will indicate in the instructions when AI use is required or prohibited. Students are expected to read all posted readings. Using AI to summarize the reading will be treated the same as if students failed to read. Effective use of AI is a crucial skill that you must develop for the modern practice of law. Part of that skill is ensuring the

accuracy of all AI outputs and maintaining intellectual ownership over your work product. Ultimately, in practice and in this class, you are responsible for the accuracy of all submitted assignments.

Students may responsibly incorporate AI into their research and writing process. Responsible use of AI entails (1) verifying all information and sources gained through AI-assisted research for their validity, (2) using generative AI as a drafting tool to assist with your writing, not as a tool to do your writing for you, and (3) retaining cognitive control over your assignments by engaging in critical analysis and legal reasoning and not allowing AI to think for you.

To ensure responsible use of AI on their assignments, students are required to disclose all uses of AI. Disclosure requirements are not meant to discourage the use of AI. Using AI responsibly will likely improve the overall quality of your final work product.

This course's AI use policy is consistent with and supplements both the Nova Southeastern University Shepard Broad College of Law AI Use Policy and the Nova Southeastern University Artificial Intelligence Governance Policy. Students must adhere to all institutional and college-wide standards for responsible and ethical AI use, including maintaining data privacy, avoiding the disclosure of confidential or nonpublic information, and ensuring compliance with all applicable university, state, and federal laws.

5. **Assignments**

The following provides a brief description of the various types of assignments students will be expected to complete during this course. As with the course calendar itself (below), the specific assignments are subject to modification by the instructor, upon prior notice and instruction. Non-AI assignments must be completed without using generative AI. They must be done solely by the student engaging in analysis without the aid of any generative AI. AI assignments require the use of generative AI, and the specific AI system identified in the assignment.

Readings (Non-AI)

Students must read all assigned materials designated in the course calendar as "required reading." Suggested readings, which are optional, are also noted in the course calendar and are recommended for students wanting to supplement the required materials.

Class Exercises/Participation (5%)

This includes participation in class discussions, completion of in-class assignments, and completion of practice problems. In-class assignments will be graded on a credit/no credit basis. Students who put in a good-faith effort to complete the assignments will receive full credit for them.

Daily "AI Current Event" Elevator Pitch (Attendance and 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.

Skill Assessments (5%)

Students will be assigned various skill assessment tasks as part of their homework. These tasks may be provided via CALI, Lexis, or Westlaw. Students must be registered for all three as a prerequisite to completing the course. Please reach out to Professor Diaz for instructions on how to register if you have not previously done so.

Knowledge Checks (5%)

Students are required to complete a small set of questions before and after designated topics.

Oral Presentation of Research Results (Non-AI) (5%)

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 (5%)

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 (5%)

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 (5%)

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 (5%)

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 (20%)

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.

AI Evaluation (Professional Research Tool Review) Final Assignment (5%)

Students will evaluate generative AI research platforms (e.g., Lexis+ AI, Westlaw Precision AI, CoCounsel, ChatGPT Enterprise) for 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?). Structure your review using the AI Tool Evaluation Framework in *The AI Legal Researcher*. 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 (35%)

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 may use the Research Trail Template to structure the strategy log and the Verification Checklist in *The AI Legal Researcher* to document validation 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.

Weekly Modules

Students first rebuild the traditional research judgment that makes AI supervision possible (Weeks 1–3), then learn how AI systems actually work and where they fail (Weeks 4–6), master the verification-first method at the center of the book—Plan, Prompt, Verify, Document (Weeks 7–9), learn to choose tools and apply the method across real research contexts (Weeks 10–12), and close with agentic systems, the future of the field, and synthesis (Weeks 13–14).

Week 1 – Why the AI Legal Research Still Starts With Traditional Legal Research Skills

Readings: The AI Legal Researcher Introduction & Chapter 1; Florida Legal Research (pgs. 1-30)

Class Discussion: What legal research actually is—issue identification, authority evaluation, updating, synthesis—not merely "finding cases." How AI changes the pace and polish of research but not the hierarchy of authority or the duty of accuracy. Three framing ideas from Chapter 1: AI amplifies traditional research errors rather than inventing new ones; AI exposes research deficiency, it does not create it; and the "disappearing apprentice" problem—why foundational skills matter more, not less, in an AI environment. Research judgment as the durable skill that transfers across every tool.

Week 2 – The Traditional Legal Research Baseline: Jurisdiction, Authority, and the Research Plan

Readings: The AI Legal Researcher Chapter 1; Florida Legal Research Chapter 4, 6, and 9

Class Discussion: Building a research plan before touching a search bar. Primary vs. secondary authority; binding vs. persuasive; holding vs. dicta; "good law," citators, and updating. The two concepts AI handles worst—jurisdiction and procedural posture—and why a researcher who cannot catch those errors cannot supervise AI that makes them. Legal research as an iterative process, not a single query.

Week 3 – Oral Presentation of Research Results Assignment (Non-AI)

Week 4 – AI Literacy and the AI Stack

Readings: The AI Legal Researcher Chapter 2; Generative AI and the Delivery of Legal Services Introduction & Chapters 1, 2, 3 & 7; ABA Model Rule 1.1, Comment 8 (technological competence).

Class Discussion: AI literacy as an extension of information literacy and a professional habit, not a one-time lesson. The AI stack in plain language: machine learning, deep learning, NLP, large language models, generative vs. extractive AI, and agentic AI. Why an LLM is a next-token predictor rather than a reasoning lawyer, and why fluent output is not reliable output. Grounding and retrieval-augmented generation (RAG): the difference between a tool that searches live legal authority and one that answers from training patterns alone. Hallucination and confabulation; non-determinism (the same prompt, different answers); context windows; and what happens to confidential information you paste into a tool.

Week 5 – Cognitive Bias, Human Error, and AI Risk

Readings: The AI Legal Researcher Chapter 3

Discussion: Human-side risk: automation bias, confirmation bias, anchoring, availability, the framing effect, premature closure, and "cognitive surrender" to fluent output. System-side risk: training-data and coverage bias, majority-view bias, sycophancy, ranking and recency bias. The core AI risks in legal research—hallucination, citation fabrication vs. distortion, false confidence, omitted and misclassified authority, retrieval failure, and outdated law. "Familiarity changes the risk": why AI is most persuasive precisely when you know the area least, and why unfamiliar questions demand more verification, not less. Bias-interruption questions as a workflow habit.

Week 6 – Ethics and Professional Responsibility

Readings: The AI Legal Researcher Chapter 4*;* Generative AI and the Delivery of Legal Services Chapters 4 & 8; ABA Formal Opinion 512; selections on court AI rules.

In-Class Discussion: The Model Rules most implicated by AI-assisted research—competence (1.1), diligence (1.3), communication (1.4), fees (1.5), candor to the tribunal (3.3), and

misconduct (8.4)—and how ABA Formal Opinion 512 applies them to generative AI. Confidentiality and privilege: every prompt is a potential disclosure, and why client-identifying information should never go into a consumer or unapproved tool. The duty to disclose AI use, and how court requirements are evolving—including the Florida Supreme Court's May 2026 amendment to Rule of General Practice and Judicial Administration 2.515(d)(2) and its preemption of circuit-level AI orders. Evaluating tools beyond vendor marketing, including independent findings that leading legal AI tools still hallucinate at meaningful rates.

Week 7 – The Core Method: Plan, Prompt, Verify, Document

Readings: The AI Legal Researcher Chapter 5

Class Discussion: Mapping the four steps onto the traditional research process and identifying where AI fits at each stage. Why planning—not prompting—comes first, and how a plan creates the baseline that makes verification meaningful. The non-negotiables: verify before relying, document before moving on.

Week 8 – Prompting for Legal Research

Reading: The AI Legal Researcher Chapter 6 (Prompting for Legal Research); Generative AI and the Delivery of Legal Services Chapter 6

Class Discussion: The five-element prompt for legal research—Role, Facts, Jurisdiction, Task, Limits—and why omitting jurisdiction or procedural posture produces confident, incomplete answers. Prompts that invite hallucination; iterative prompting (never stop at the first output); the prompt log as part of the research trail. Building a customized AI research assistant, context engineering, and prompt techniques that reduce confirmation and sycophancy bias. The Boolean analogy: shortcuts in prompting fail the same way shortcuts in search always have.

Week 9 – Verifying AI Output

Readings: The AI Legal Researcher Chapter 7

Class Discussion: The three layers of verification—source (does it exist?), proposition (does it say what the AI claims?), and treatment (is it still good law?)—and why partial verification creates false confidence. The verification echo chamber and the misinformation loop, when AI is used to "check" AI and a fabrication survives unchallenged. Why "I asked the AI to verify it" is not verification, and how to explain that to a supervising attorney. Citation-verification tools and the role of the citator that AI cannot reliably replace.

Week 10 – Choosing and Evaluating AI Research Tools

Readings: The AI Legal Researcher Chapter 8 (Choosing and Evaluating AI Research Tools); Generative AI and the Delivery of Legal Services Chapters 5 & 10

Class Discussion: Categories of AI systems for legal work: general-purpose generative AI assistants, legal-specific AI research platforms (retrieval-augmented and citation-linked to a legal database), agentic legal workflow assistants, and AI document-review/eDiscovery and contract-analysis tools. Enterprise vs. consumer; free vs. paid; open-source and self-hosted options; data hygiene and AI readiness. Search-based vs. generative tools, the hybrid workflow, and a tool-selection framework keyed to the research stage. Evaluating tools through testing and independent research rather than vendor benchmarks.

Week 11 – AI-Assisted Workflows: Planning, Issue-Spotting, and Doctrinal Research

Readings: The AI Legal Researcher Chapter 9

In-Class Discussion: The detailed research-planning workflow: reading facts before AI, separating what you know from what you must find, and using AI to brainstorm issues as hypotheses rather than conclusions. Recognizing issue-framing bias and AI's tendency to over frame (too many tangents) or underframe (missing a dispositive procedural defense, preemption, or jurisdictional bar). Statutory, common-law, regulatory, and constitutional starting points. Multi-jurisdictional surveys: using AI to build a comparison framework while independently verifying each jurisdiction, and watching for blended rules and invented majority/minority splits.

Week 12 – AI-Assisted Workflows: Procedural, Docket, Factual, and Forms (Litigation and Transactions)

Readings: The AI Legal Researcher Chapter 9

In-Class Discussion: Why procedural research is the area where AI fails—local rules, standing orders, and judge-specific requirements change constantly and must be verified against the court's official sources. Docket research: no speculating about filings or histories the tool never retrieved. Factual and background research: fabricated facts about real entities are harder to catch than fabricated citations—ask AI for the research plan, not the facts. Forms and practice materials: the difference between using AI to find authoritative models and letting it generate forms that look polished but omit required elements. Application to transactional drafting and litigation work product, with verification against original records.

Week 13 – Open-Source, Agentic Systems, and the Future of Legal Research

Readings: The AI Legal Researcher Chapter 10

In-Class Discussions: From chat to agents: what makes a system agentic, the capabilities of agentic legal research tools, and why multi-step autonomy compounds error and demands

supervision at multiple points, not just review of the final product. Open-source and self-hosted AI vs. commercial legal AI, and a self-hosted research workflow. “Vibe coding” and lawyer-built legal technology. Emerging issues that frame the profession’s next decade: AI and access to justice, AI and the courts and judges, data breaches and security, prompt injection and data leakage, and the legitimacy questions raised by AI-influenced law. How to keep your AI literacy current as the tools keep changing.

Week 14 – Open-Source AI, Agentic Systems, and the Future of Legal Research

Readings: The AI Legal Researcher Chapter 10

In-Class Discussions: From chat to agents: what makes a system agentic, the capabilities of agentic legal research tools, and why multi-step autonomy compounds error and demands supervision at multiple points, not just review of the final product. Open-source and self-hosted AI vs. commercial legal AI, and a self-hosted research workflow. “Vibe coding” and lawyer-built legal technology. Emerging issues that frame the profession’s next decade: AI and access to justice, AI and the courts and judges, data breaches and security, prompt injection and data leakage, and the legitimacy questions raised by AI-influenced law. How to keep your AI literacy current as the tools keep changing.