

Red Line Challenge: Instructor Packet

Exercise Overview

The Red Line Challenge is a short, high-pressure verification exercise designed for an upper-level AI-augmented legal research course. Students receive a brief that appears facially plausible but contains multiple errors associated with AI-generated legal work: fabricated citations, citation drift, doctrinal distortion, assumed facts, and tone mismatch.

Students are not asked to rewrite the brief. Their task is to identify as many issues as possible, classify each issue, and verify their findings.

The exercise is intentionally difficult. Students should not be expected to catch every issue in the allotted time. The difficulty is part of the learning design: students experience how quickly a polished legal document can become dangerous when authority, facts, and standards are not independently verified.

It is important to note to students that this exercise also develops understanding and empathy for attorneys who have real work and deadlines that make verification, especially tricky verification, challenging.

Recommended Document Set

This exercise should be distributed as two separate documents:

1. Student Brief
This is the modified appellate brief excerpt students will review. It should include only the Introduction and Argument sections. Do not include the instructor key, planted-error list, or total number of errors.
2. Student Worksheet
This is the student-facing error log. Students use it to record suspected errors, classify issue types, note their verification strategy, and reflect on their confidence.

The instructor should also keep this packet as a private teaching guide. Do not distribute the instructor key to students before the debrief.

Timing

Recommended time: 20 minutes

Learning Goals

By completing this exercise, students should be able to:

1. Recognize common errors in AI-generated or AI-assisted legal documents.
2. Distinguish between fabricated authority, citation drift, doctrinal distortion, assumed facts, and tone mismatch.
3. Verify legal propositions, quotations, parentheticals, and case citations using a legal research platform.
4. Explain why a real citation is not necessarily a reliable citation.
5. Develop a practical verification strategy for reviewing AI-assisted work product.
6. Connect verification failures to professional responsibility, supervision, candor, and lawyer accountability.

Suggested Student Instructions

You will receive a brief that may contain errors commonly associated with AI-generated legal analysis.

You have 20 minutes to identify as many problems as possible.

Your task is to:

- Flag every suspected error.
- Classify the type of issue.
- Verify your findings using Lexis, Westlaw, or the internet only.
- Record your findings in the worksheet.

You are not required to rewrite or correct the brief. Focus on spotting, classifying, and verifying errors.

Do not use BriefCheck, BriefCatch, Drafting Assistant, or similar automated brief-checking tools. This is a manual verification exercise.

Error Classification Categories

Students should use these categories when recording errors. The instructor can choose to explain these before or not – providing context to all types of errors will improve student experience and learning thought.

Citation Fabrication

The case, statute, regulation, or other authority does not exist, or the cited case name and reporter citation do not correspond to a real authority.

Citation Drift

The cited authority exists, but the brief misstates the quoted language, holding, legal standard, procedural posture, or relevance of the authority.

Doctrinal Distortion

The legal standard may exist somewhere, but it is attributed to the wrong jurisdiction, wrong court, wrong cause of action, or wrong doctrinal context.

Assumed Fact

The brief includes factual details not supported by the provided record or problem materials.

Tone Mismatch

The language is inconsistent with professional legal advocacy, such as editorializing, sarcasm, exaggeration, or rhetorical phrasing that undermines credibility.

Instructor Guidance

Do not tell students how many errors are planted unless you want to make the exercise easier. The recommended version gives students the categories but not the counts.

The brief contains a core set of planted authority problems and several additional debrief issues. The core set is best for scoring or structured review. The additional issues are useful for discussion because they show that AI-related problems are not limited to fake cases.

Students should not verify record citations, but these errors are included to simulate real-life errors. If students identify assumed facts, treat those as strong observations, but remember to explain that assumed-fact verification requires a record packet or prior knowledge of the original case materials.

Instructor Key: Core Planted Issues

A. Citation Fabrications

These citations were drafted to look plausible in a Fourth Circuit employment discrimination brief. Document analysis tools available on Westlaw and Lexis should catch these errors if you show that as a demonstration.

#	Problematic Language or Citation	Issue Type
1	"Cruz v. Blue Ridge Transit Auth., 912 F.3d 143, 151 (4th Cir. 2019)"	Citation Fabrication
2	"Santiago v. Rowan County Facilities Dep't, 889 F.3d 117, 123 (4th Cir. 2018)"	Citation Fabrication
3	"Bennett v. Coastal Carolina Univ., 873 F.3d 207, 214 (4th Cir. 2017)"	Citation Fabrication
4	"Ramirez v. Piedmont Logistics, Inc., 861 F.3d 185, 191 (4th Cir. 2017)"	Citation Fabrication
5	"Nguyen v. Appalachian State Maint. Servs., 845 F.3d 311, 318 (4th Cir. 2016)"	Citation Fabrication

B. Citation Drift & Doctrinal Drift

These authorities exist, but the brief misstates or distorts what they support. Document analysis tools available on Westlaw and Lexis should catch these errors if you show that as a demonstration.

#	Problematic Language or Citation	Instructor Explanation
6	"This Court may affirm only on the exact rationale stated by the district court. Lawson v. Union Cnty. Clerk of Court, 828 F.3d 239, 247 (4th Cir. 2016)."	Lawson states "Our review is not limited to the grounds the district court relied upon, and we may affirm 'on any basis fairly supported by the record.'"
7	"Unless the remarks upon which plaintiff relies were related to the employment decision in question, they can still be direct evidence if made by a supervisor in the same department. Brinkley v. Harbour Recreation Club, 180 F.3d 598, 608 (4th Cir. 1999)."	This is an incorrect quote.
8	"McCray v. Pee Dee Reg'l Transp. Auth., 263 F. App'x 301, 306 (4th Cir. 2008) (holding statements made two and a half years before employment action were direct evidence because they were 'highly offensive')."	This mistakes McCray's holding, outcome, and governing standard.
9	"Pretext may be shown if the employer's	This misstates Adam's legal standard and

	reason is ‘inaccurate, unfair, or simply debatable.’ <i>Adams v. Trs. of the Univ. of N.C.-Wilmington</i> , 640 F.3d 550, 560 (4th Cir. 2011).”	leaves out key limits. The ‘inaccurate, unfair, or simply debatable’ language is not present.
10	“But Williams holds the opposite: that an employee cannot proceed unless he can prove an actual application. <i>Williams</i> , 370 F.3d at 433.”	This misstates Williams’ language. The standard articulated is that the plaintiff must prove they would have applied if they knew about the application, not that they actually applied.
11	“Heiko holds that a plaintiff can show pretext whenever his qualifications are ‘roughly comparable’ to the selectee’s. <i>Id.</i> at 261.”	Heiko does not use this language. Heiko states “a plaintiff asserts job qualifications that are similar or only slightly superior to those of the person eventually selected, the promotion decision remains vested in the sound business judgment of the employer... But where, as here, the plaintiff has made a strong showing that his qualifications are demonstrably superior, he has provided sufficient evidence that the employer’s explanation may be pretext for discrimination.”
12	“The Fourth Circuit applies the ‘nearly identical’ comparator standard: the comparator must have engaged in the same misconduct, under the same supervisor, with the same disciplinary history, and in the same time frame. See <i>Lewis v. City of Union City</i> , 918 F.3d 1213, 1227–28 (11th Cir. 2019) (en banc).”	There is no language at the page cite that states the Fourth Circuit has adopted this standard. This case is an Eleventh Circuit case that rejects the ‘nearly identical’ standard, and the reference to the Fourth Circuit is only through the discussion of the Supreme Court’s handling of <i>Young</i> .
13	“DeJarnette holds that even if an employer’s reason is mistaken, summary judgment is inappropriate where the decision was ‘probably wrong.’ <i>DeJarnette v. Corning Inc.</i> , 133 F.3d 293, 299 (4th Cir. 1999).”	This misstates DeJarnette’s holding and inverts the standard. It also misstates the plaintiff’s burden.
14	“Hill requires only that a biased employee have some influence over the final decision. <i>Hill v. Lockheed Martin Log. Mgmt., Inc.</i> , 354 F.3d 277, 290 (4th Cir. 2004).”	This misstates Hill’s standard and omits limiting conditions.

D. Assumed Facts

#	Problematic Language
14	“Plaintiff later acknowledged in an EEOC mediation session that he ‘may have misread signals,’ though he maintained he was treated differently because of his background.”
15	“Plaintiff also conceded in deposition that he had never been certified on the building automation system used campus-wide, which was listed as a preferred qualification for the Supervisor role.”

D. Tone Mismatch

#	Problematic Language
16	“Frankly, Plaintiff’s attempt to recast a straightforward policy violation as discrimination is the kind of litigation move that wastes everyone’s time.”
17	“In short: stray remarks are not stray when they sting.”

F. Optional Additional Citation Drift Issues

These are more subtle and may be treated as advanced or bonus findings.

Suggested Debrief Structure

1. Start with process, not answers

Ask students:

- What did you check first?
- Did you start with case existence, quotations, parentheticals, or legal standards?
- Which errors were easiest to identify?
- Which felt plausible even after you suspected something was wrong?

The goal is to make students articulate a verification workflow.

2. Separate “fake” from “misleading”

Students often focus on fabricated cases because they are the most visible form of hallucination. Emphasize that citation drift can be more dangerous because the case exists. A real citation may create false confidence, especially when the parenthetical sounds plausible.

3. Discuss time pressure

Ask students:

- What did you miss because of the time limit?
- What would you check next if you had another 20 minutes?
- How would your review change if this were going to a court, a supervising attorney, or a client?

This helps students see why AI-assisted work product requires triage and a repeatable review process.

4. Connect to professional responsibility

Use the debrief to connect the exercise to:

- Competence: knowing how to use and verify AI-assisted work.
- Candor: ensuring court submissions do not contain false authority.
- Supervision: reviewing work product prepared by junior lawyers, staff, or AI tools.
- Confidentiality: recognizing that verification often requires returning to trusted legal sources rather than asking another public AI tool.
- Professional judgment: deciding when a document is not ready to move forward.

5. Emphasize the practical takeaway

Students should leave with this principle:

A polished brief is not a verified brief. Verification requires checking the existence, meaning, jurisdiction, procedural posture, and use of authority before relying on it.

Suggested Scoring or Completion Approach

This exercise works well as a formative, low-stakes assignment. Recommended completion standard:

- Student identifies at least several suspected issues.
- Student classifies each issue using the provided categories.
- Student documents a reasonable verification method.
- Student completes the reflection questions.

If graded more formally, instructors can award points for:

- Correctly identifying fabricated citations.
- Correctly identifying citation drift.
- Correctly identifying the doctrinal distortion.
- Recognizing tone mismatch or assumed facts.
- Explaining the verification strategy.
- Reflecting accurately on what they did not have time to check.

Avoid grading solely by number of errors found. The more important assessment is whether students are developing a reliable verification process.

Instructor Preparation Checklist

Before using this exercise:

- Confirm that the student brief does not include the instructor key.
- Standardize the time limit across the worksheet and instructions.
- Standardize the verification tool requirement, such as “Lexis only.”
- Remove any sample worksheet row that could be mistaken for an intended planted error.
- Decide whether assumed facts are scored or used only for debrief.
- Verify the fabricated citations in Lexis shortly before class.
- Verify the real cases and drift explanations in Lexis shortly before class.
- Decide whether students may work individually, in pairs, or in small groups.
- Decide whether students submit the worksheet or use it only for class discussion.

Instructor Note on Difficulty

This exercise is intentionally too hard to complete perfectly in 20 minutes. That is a feature, not a flaw.

The point is not that students should catch every issue immediately. The point is that AI-generated legal work can create a verification burden that is difficult to manage under time pressure. Students should experience the difference between reading a document that sounds legally sophisticated and verifying a document that is legally reliable.