AI, LEGAL RESEARCH, AND PROFESSIONAL JUDGMENT

Legal Slop: What Happens When Legal Words Become Free?

Vibe Lawyering, AI Hallucinations, and Legal Judgment After the Drafting Barrier Falls

By Brian Edward Walters | www.TexasTowingLawyer.com

ABSTRACT

Generative artificial intelligence has collapsed the cost of producing legal-looking text. That is an access breakthrough, but it also creates a new professional and institutional risk: legal slop. This article defines legal slop as legal text that has the surface form of legal analysis without verified authority, contextual legal judgment, or reliable source ownership. Using recent sanctions and warning cases involving fabricated citations, false quotations, AI-drafted filings, pro se litigants, attorney verification failures, and contaminated orders, the article argues that the legal profession is moving from an era of scarce expression to an era of scarce verification and judgment. The problem is not only fake cases. It is the persistence of legal-looking error in adversarial systems, public dockets, legal databases, contracts, and copied forms. The legal skill that matters most in the age of AI is not merely producing plausible legal text, but knowing whether the words are true, useful, lawful, and capable of producing a remedy the institution can actually provide. Legal words are now free. Legal judgment is not.

Legal words are now free. Legal judgment is not.

WHAT THIS ARTICLE COVERS

- How vibe coding maps onto vibe lawyering.
- Why legal slop is worse than ordinary drafting error.
- What recent sanctions and warning cases already show.
- Why legal research now requires source forensics and contextual judgment.

I. Introduction: Vibe Lawyering And The End Of Expensive Legal Words

Silicon Valley has a name for software built by prompt and hope: vibe coding. A user tells an artificial intelligence system what the software should do, the system produces code, and the user keeps prompting until the product appears to work. At its best, the practice can make experienced developers faster. At its worst, it produces what software engineers have begun calling vibe slop: code that looks productive in the moment but carries defects, security risks, technical debt, and maintenance problems that the user did not understand well enough to see.[1]

Law has entered the same moment.

For most people, both code and law have always been foreign languages. They use familiar words and symbols, but the words do not mean what ordinary readers assume they mean. In code, meaning depends on syntax, dependencies, runtime behavior, security assumptions, tests, and maintainability. In law, meaning depends on jurisdiction, procedural posture, authority hierarchy, statutory text, pinpoint support, later history, evidentiary posture, remedy, and forum. Both systems punish false fluency.

Generative artificial intelligence has changed the price of that fluency. It has made code newly accessible to people who could not previously write it. It has done the same with law. Anyone with an internet connection and a free account can now generate a petition, motion, response, declaration, appellate brief, proposed order, demand letter, settlement agreement, operating agreement, or table of authorities. The old drafting barrier has fallen.

Legal words are now free. Legal judgment is not.

That distinction is the center of this article. Artificial intelligence has not made legal expertise obsolete. It has made legal imitation cheap. The legal system has entered a new era of "vibe lawyering": legal work generated from instinct, grievance, desired outcome, or plain-English instruction before anyone verifies whether the law, facts, remedies, citations, or transactional assumptions are true. Vibe lawyering is the method. Legal slop is the residue.

Legal slop is legal text that has the surface form of legal analysis without the discipline of verified authority, contextual legal judgment, or reliable verification. It can exist as hidden land mines inside otherwise good-appearing legal documents. It includes fake cases, fake quotations, real cases cited for false propositions, invented procedural histories, misquoted statutes, proposed orders with nonexistent authority, and polished legal documents that conceal missing terms or unasked questions. In litigation, legal slop often lies about authority. In transactional work, legal slop often lies about completeness.

The comparison to coding is useful, but it is not symmetrical. Code slop usually meets a machine. Legal slop meets an institution. More precisely, legal slop meets an adversarial institution. A defective legal filing is not merely a private error in a private repository. It assigns work to a court, a clerk, opposing counsel, and another party. It may force an opposing party to pay a lawyer to prove that the case does not exist. It may ask a judge to rely on a false premise. It may enter a docket that is designed to preserve, not quarantine, what lawyers file.

The court system is not a debugger. Litigation is a brutal way to debug language.

That is why vibe lawyering carries a different kind of risk than vibe coding. Code slop can create technical debt. Vibe lawyering can create legal debt: sanctions exposure, false statements to courts, wasted judicial time, corrupted orders, procedural failures, distorted legal research, and archival pollution that may persist long after the mistake is discovered. Code slop usually has to run. Legal slop only has to remain.

The early case law already shows the pattern. Lawyers have filed briefs containing nonexistent cases and fabricated quotations. Self-represented litigants have filed legal-looking briefs filled with citations that appear ordinary but do not exist. Courts have vacated orders that relied on bogus authority. Bankruptcy courts have sanctioned lawyers for hallucinated case law. Texas appellate courts have begun warning both lawyers and pro se litigants that hallucinated citations are not mere formatting errors. The most visible cases have produced sanctions, warnings, referrals, and published opinions. But those sanctions cases are not the universe. They are only the visible edge of the universe.

The deeper problem is measurement. We can count the hallucinations that courts catch and write about. We cannot yet count the legal slop that sits undetected in pleadings, motions, proposed orders, trial records, settlement files, contracts, docket entries, and legal databases. That hidden mass matters because the law is a memory system. Once a document is filed, archived, indexed, downloaded, quoted, or copied, its errors can outlive the case in which they were made.

This article argues that legal writing in the age of artificial intelligence must shift from expression to contextual legal judgment. Verification is part of that judgment, but it is not the whole of it. The essential legal skill of the next decade will not be the ability to produce plausible legal text. The machines can already do that. The essential skill will be understanding what the words are supposed to accomplish in the broader context of the client's desired outcome, the facts that matter, the court or institution that must act, the remedies actually available in the jurisdiction, and the future events the document must survive. Source ownership remains indispensable: the lawyer must know whether the words are true, whether the authority supports the sentence, and whether the remedy exists. But that has always been true. What has changed is the volume and polish of false or incomplete legal text. AI makes falsity easier to generate, easier to disguise, and easier to copy.

II. Law And Code As Technical Languages

The analogy between law and code is not a slogan. It works because both are technical languages that govern real-world consequences through formal systems.

Code is not simply words typed into a computer. It is a set of instructions that must operate inside a larger architecture. A short function may depend on libraries, permissions, inputs, outputs, memory, security assumptions, and the behavior of other systems. Code that looks clean on the screen may fail in production because the user did not understand the environment in which it would run.

Law is similar. A legal sentence may look like ordinary English. But its meaning depends on the system in which it operates. "The court has jurisdiction" means one thing in a justice court, another in a bankruptcy adversary proceeding, another in a state court appeal, and another in a federal criminal case. "The case supports the proposition" depends not only on the case name, but on court, date, procedural posture, holding, dicta, later history, controlling law, and the exact sentence being supported.

Artificial intelligence can imitate both languages with remarkable fluency. It can produce code that looks like code and legal prose that looks like legal analysis. But fluency is not verification. A model can write a plausible appellate standard of review without knowing whether the cited case actually exists. It can write a polished contract without knowing the business contingency that will matter three years later. It can write a proposed order

without knowing that the authority embedded in it is invented.

Software engineers have begun to recognize this distinction in the debate over vibe coding. The danger is not that a tool helped produce code. The danger is that the human has stopped understanding, testing, or owning the system. The same distinction should govern law. The danger is not that an AI system helped draft a motion, brief, contract, or memo. The danger is that the lawyer, litigant, or business user treats the output as law before understanding what it must accomplish and verifying that it can do so.

The mistake is easy to make because legal writing is an artifact profession. Lawyers are judged by documents. Courts receive documents. Clients buy documents. The document has always been the visible thing. But the document is not the lawyering. Lawyering is the judgment behind the document: what outcome the client is trying to achieve, what forum or institution can act, what remedy actually exists, what source controls, what fact changes the answer, what risk matters, what to leave out, and what happens if the document is tested.

AI collapses the price of the artifact. It does not collapse the price of judgment.

III. From Vibe Coding To Vibe Lawyering

"Vibe lawyering" is not a claim that every AI-assisted legal task is defective. That would be wrong and unhelpful. AI can summarize, format, brainstorm, compare drafts, generate checklists, identify missing issues, and accelerate routine work when a competent human remains responsible for the result.

The useful line is not AI versus human. The useful line is judged and verified work versus unexamined output.

Vibe coding becomes dangerous when software is generated faster than it can be understood. Vibe lawyering becomes dangerous when legal consequences are generated faster than they can be understood, situated, and verified. In both settings, the product may look good. The code may compile. The website may load. The motion may sound lawyerly. The contract may look comprehensive. The question is what happens when the system is stressed.

In software, stress may come from users, tests, attackers, integrations, or maintainers. In law, stress comes from institutions: courts, opposing parties, regulators, trustees, buyers, lenders, heirs, judges, and time. A legal document is not merely read. It is used against facts, parties, rules, remedies, institutional limits, client objectives, and incentives that may not have appeared in the prompt.

The risk is higher in law because the test environment is often someone's life, liberty, property, custody, business, or money. A defective legal filing is not a broken prototype. It is a demand on an adversary and a tribunal. A defective contract is not a failed demo. It may allocate risk between real parties for years.

This is the first major difference between vibe coding and vibe lawyering: law meets an institution.

The second difference is that law meets an adversary. The opposing lawyer's job is to verify, challenge, distinguish, and defeat. In a healthy system, that adversarial pressure helps expose weak claims and false authority. But relying on litigation to catch legal slop externalizes the cost. The lawyer or pro se litigant who files a fake citation has assigned research work to the other side. The party who signs an AI-generated contract with missing protections has assigned future litigation work to whoever has to repair the failure.

That is why the coding metaphor must be handled carefully. Law does have tests. They are motions, objections, sanctions, appeals, defaults, bankruptcies, enforcement disputes, and malpractice claims. Those are not cheap tests. They are litigation.

IV. Legal Slop Defined

Legal slop is legal text that has the surface form of legal analysis without the discipline of verified authority, contextual legal judgment, or reliable verification. It may appear as an obvious fake citation. More dangerously, it may appear as a hidden land mine in an otherwise good-appearing legal document.

The most obvious form is the fake case. In Mata v. Avianca, Inc., lawyers submitted nonexistent cases generated by ChatGPT in an opposition to a motion to dismiss. One fictional case, Varghese v. China Southern Airlines Co., even came with a citation that led to a different real case and listed a real Fifth Circuit judge as a panel member in an opinion that did not exist.[2] The court sanctioned the lawyers under Rule 11 and its inherent authority.[3]

The next form is the fake citation that looks ordinary. In Kruse v. Karlen, a self-represented appellant filed a brief in which only two of twenty-four cited cases were genuine. The court noted that some false citations used plausible party names and reporter formats, while others used real names for propositions the cases did not support.[4] This is the pro se version of free legal words: the filing looks like an appellate brief until someone does the work of opening the cases.

The more dangerous form may be the real case with a fake quotation. In Noland v. Land of the Free, L.P., the California Court of Appeal explained that nearly all legal quotations in the opening brief, and many in the reply brief, were fabricated. Most of the cited cases existed. The quotations did not.[5] This matters because citation-existence checks alone will not catch this category. A database can confirm that the case exists and still leave the false proposition undetected unless the user reads the case.

Another form is order contamination. In Shahid v. Esaam, the Georgia Court of Appeals vacated and remanded after a trial-court order relied on non-existent case law. The appellee's brief added more bogus citations, including citations used to support a request for attorney's fees on appeal.[6] That case shows why legal slop is not confined to party filings. It can move from a lawyer's draft into a judicial order and then into the appellate record.

There is also specialized-court slop. In In re Kheir, a bankruptcy court in the Southern District of Texas sanctioned an attorney who filed hallucinated caselaw in a response to a motion to dismiss. The court treated citation verification as ordinary professional work and rejected carelessness, good faith, and ignorance as excuses for a Bankruptcy Rule 9011 violation.[7]

And there is Texas warning slop. In Campbell v. Campbell, the Third Court of Appeals warned a pro se litigant whose filings appeared to include nonexistent cases and nonexistent quotations.[8] In In re Miller, a First Court of Appeals concurrence warned that filing fictitious or misleading citations, whether AI-generated or not checked by a human, is a serious breach of candor that may lead to corrective action, including striking briefs or reporting counsel to the State Bar.[9]

These examples show that "legal slop" is not one defect. It is a family of defects:

1. A fake case.
2. A real case name with a fake citation.
3. A real case cited for a false proposition.
4. A fake quotation from a real case.
5. A fake procedural history.
6. A proposed order that imports bad authority.

7. A permanent filing that preserves false legal material.
8. A polished transactional document that omits the issue the user did not know to ask about.

That final category is important. Legal slop is not only a litigation problem.

V. The Pro Se Drafting Barrier Has Collapsed

For decades, one barrier to self-representation was not merely knowledge of the law. It was production of the artifact. A person might know that something went wrong but not know how to turn that grievance into a petition, motion, response, affidavit, brief, or proposed order. The courthouse required legal words in legal formats. Those words were expensive.

Generative AI changes that variable. It can produce a motion. It can produce an appellate brief. It can produce a table of authorities. It can turn a layperson's narrative into a legal-looking filing. That is an access breakthrough, and it should not be dismissed.

But access to legal words is not access to verified law.

Nor is access to free legal words access to a lower legal standard. Texas courts have put the rule plainly. In Ho v. University of Texas at Arlington, the Amarillo Court of Appeals stated that self-represented litigants are "held to the same standards as licensed attorneys" and must comply with applicable procedural law.[10] The reason is fairness to the opposing party: a different rule would give self-represented litigants an unfair advantage over represented litigants.[11] In Fox v. Wardy, the El Paso Court of Appeals applied that rule to reject a pro se litigant's argument that he should receive a less strict standard because he was not an attorney.[12] Liberal construction, where it applies, is a reading principle. It is not immunity from deadlines, record rules, evidence rules, briefing rules, sanctions rules, or the duty not to file false legal authority.

Kruse shows the problem in appellate form. The brief contained the architecture of legal argument but not the source discipline of legal argument. The court found that twenty-two of twenty-four citations were inaccurate and that even the two genuine citations did not support the propositions asserted.[13] Al-Hamim v. Star Hearthstone, LLC shows the same problem with a different judicial response. The Colorado Court of Appeals identified eight fake citations in a self-represented litigant's opening brief, required the litigant to respond, and accepted his admission that he had relied on AI and failed to inspect the brief.[14] The court declined monetary sanctions in that case, but warned that future filings containing hallucinations may result in sanctions.[15]

Ferris v. Amazon.com Services, LLC shows why the institutional cost matters. There, a pro se plaintiff filed a complaint containing a fictional Eastman Kodak citation, then responded to a motion to dismiss with six more fake cases. The court described the basic imbalance: one party can create a fake legal brief quickly, while the opposing party and the court must parse through the case names, citations, and legal propositions to determine what, if anything, is true.[16]

That imbalance is the litigation version of the AI access problem. The system should want self-represented litigants to produce clear, organized filings. It should not want them to file legal fiction that forces courts and opponents to spend time disproving nonexistent law.

Pro se legal slop therefore calls for blunt risk clarity. Many self-represented litigants lack access to Westlaw, Lexis, legal training, and the habit of checking procedural posture or authority hierarchy. A free AI account may be the only legal drafting tool they can reach. But that does not change what happens when the document is filed. A fake case, fabricated quotation, or misstatement of law pollutes the record in the same way whether it came from a licensed attorney, a self-represented party, or a chatbot prompt. Legal slop pollutes equally, regardless of

who produces it.

VI. The Signature Duty And Attorney Verification Failure

The institutional harm does not change with the filer's status. The signature-duty framework already exists.

Texas Rule of Civil Procedure 13 treats an attorney's or party's signature as a certificate that the signer has read the filing and, after reasonable inquiry, that the paper is not groundless and brought in bad faith or for harassment. Chapter 10 of the Texas Civil Practice and Remedies Code similarly ties the signing of a pleading or motion to certifications about improper purpose, existing law, factual support, and proper denials, and authorizes sanctions for violations. Federal Rule of Civil Procedure 11 uses the same basic structure: by presenting a paper to the court, whether by signing, filing, submitting, or later advocating it, an attorney or unrepresented party certifies that the filing is proper and that the legal contentions are warranted after a reasonable inquiry.[17]

AI does not change that architecture. It changes the speed, volume, and polish of material that can be filed without satisfying it.

For lawyers, the problem becomes professional verification failure. Lawyers have tools, training, duties, and licenses. They know, or must know, that a case citation is not a decoration. It is a representation to a tribunal that the authority exists and supports the proposition for which it is offered. When lawyers file legal slop, they do not merely misunderstand a new technology. They launder unverified output through a law license.

Fletcher v. Experian Information Solutions, Inc. is an important current example because the Fifth Circuit placed AI hallucinations inside ordinary lawyer-discipline doctrine. The court noted that the first high-profile federal incident was Mata, but observed that AI hallucinations have become an even greater problem despite news coverage, CLE presentations, scholarship, and judicial warnings.[18] It also explained that the problem is becoming more subtle. In the Fifth Circuit's words, hallucinations now often appear as "false quotes or statements of law attributed to real cases," not merely as obvious fake cases.[19]

The Fifth Circuit's answer was not to invent a new AI-only doctrine. It held that although modern generative AI is new, the ordinary sanctions framework remains capable of addressing the misconduct.[20] The court disciplined the appellate lawyer because she failed to check her own brief, repeatedly misrepresented the law, and was not forthcoming about the source of the errors.[21]

In re Kheir reaches the same principle in bankruptcy practice. The court stated that confirming cited caselaw is basic and routine for a practicing attorney and that carelessness, good faith, and ignorance do not excuse noncompliant filings.[22] It then sanctioned the lawyer with fee reimbursement, continuing education, client notice, and a compliance certificate.[23]

Noland adds an important human point: the client may not know that the lawyer has filed fabricated law. The Court of Appeal addressed the merits of the client's appeal but sanctioned counsel, observing that counsel had used several AI tools and failed to read the cases the tools cited.[24] The case therefore separates the tool from the professional duty. AI did not sign the brief. The lawyer did.

This should be the rule of the profession: if your name is on the filing, you own the source.

VII. Vibe Lawyering In The Transactional Context

The cases collected so far mostly involve litigation. That is where fake citations become visible. But vibe lawyering is broader than citations, and the transactional context may carry an even longer half-life.

In litigation, the defect often announces itself as authority: the case is fake, the quotation is fake, or the cited case does not support the point. In transactional work, the defect may be silence. A contract can look professional and still fail because the AI user did not know what issue needed to be included in the prompt.

This is the unknown-unknowns problem. A user asks for an operating agreement. The model generates provisions for management, capital contributions, allocations, and transfers. The document looks complete. But perhaps it omits deadlock mechanics, buy-sell rights, tax distributions, manager removal, lender consent, drag-along language, or what happens when one member dies, divorces, files bankruptcy, loses a license, or becomes a competitor.

A user asks for an asset purchase agreement. The model produces definitions, purchase price, closing, representations, and covenants. But perhaps it misses lien releases, excluded liabilities, successor-liability allocation, customer consents, licenses, permits, transition services, sales-tax clearance, employee issues, or indemnity procedures that matter in that industry.

A user asks for a settlement agreement. The model produces a release and confidentiality clause. But perhaps it misses Medicare, Medicaid, ERISA, workers' compensation, hospital liens, tax treatment, confidentiality carveouts, non-disparagement limits, enforcement venue, indemnity, attorneys' fees, or whether the release reaches unknown claims, affiliates, officers, employees, insurers, and successors.

None of those failures looks like a hallucinated citation. They look like a contract that was not good enough.

This is why legal slop in transactional work may be more dangerous than legal slop in litigation in half-life terms. A fake case in a motion can be caught when opposing counsel opens the source. A defective contract may sit undisturbed for years. It may not be tested until default, bankruptcy, death, divorce, acquisition, lender diligence, regulatory review, termination, casualty, or litigation. When the stress event arrives, the failure will not be labeled "AI hallucination." It will simply be an ambiguity, an omitted term, an unenforceable clause, or a lost remedy. The legal slop was there all along, not as an obvious defect, but as a hidden land mine in a document that looked good enough to sign.

In litigation, legal slop lies about authority. In transactional work, legal slop lies about completeness.

This does not mean AI cannot help draft contracts. It can. It may become very good at routine drafting. But law is too broad and context-dependent for a user to assume that a general model knows every factual, legal, business, tax, regulatory, insurance, industry, and enforcement consequence that should control the draft. The most dangerous contract term may be the term no one knew to ask for.

VIII. False Statements To Courts

Once legal slop is filed with a tribunal, it changes character. It is no longer merely an internal research failure. It is a representation to an institution.

United States v. Hayes gives the point its sharpest form. The case involved a criminal defense lawyer who submitted a fictitious case and quotation, then continued to characterize the problem as an inadvertent citation error after the government and the court raised concerns. The court rejected that framing and imposed sanctions.[25]

Hayes is important because it connects hallucinated authority to the duty of candor. The court invoked rules of professional conduct prohibiting false statements of fact or law to a tribunal and prohibiting knowing misquotation of authority.[26] It then stated the principle that should become central to AI-era legal writing: "Citing nonexistent case law or misrepresenting the holdings of a case is making a false statement to a court."[27]

That formulation matters. Lawyers may be tempted to describe AI hallucinations as mistakes, glitches, or drafting problems. Some will be. But once the defective authority is filed, the court and the adversary are entitled to treat the citation as a representation. The lawyer is not merely saying "this sounds right." The lawyer is saying: this law exists, this source says this, and this court may rely on it.

Hayes also reinforces the adversarial cost. The court quoted Mata for the harms caused by fake opinions: wasted time for the opposing party, wasted judicial resources, deprivation of authentic arguments for the client, harm to the reputations of judges and courts falsely invoked, harm to parties attributed with fictional conduct, and cynicism about the legal profession and the judicial system.[28]

The lesson is simple. A hallucinated citation is not harmless because a lawyer subjectively hoped it was true. Legal writing submitted to a court carries institutional weight. A false citation puts that weight behind something that does not exist.

IX. The Court System Is Not A Debugger

The adversary system is a truth-testing device, but it is not a free one.

When a lawyer files a false citation, opposing counsel must spend time opening the case, searching for the quotation, checking parallel citations, looking for similar names, determining whether the case is nonexistent or merely miscited, and deciding whether to raise the issue. A client pays for that time. A judge or clerk may repeat the same work. If the bad authority enters an order, the cost may move to an appeal.

Ferris captures the imbalance well. The court observed that AI can produce realistic legal fiction much faster than courts and opposing parties can verify it.[29] That is not merely inconvenience. It is a cost-shifting mechanism. The person who files the fake authority has externalized verification costs to everyone else.

That externalization is more serious than a technical bug because litigation is coercive. A party served with a motion must respond or risk losing rights. A party facing a hallucinated argument may need to pay to disprove a source that never existed. A judge in a crowded docket may face an order that appears supported until someone catches the defect.

This is why the "test suite" metaphor must be used carefully. In law, the test is often high-stakes litigation. The bug report may be a motion for sanctions. The regression test may be an appeal. The failed deployment may be a voided order, an unenforceable agreement, a lost right, or a fee award.

The adversary system can catch legal slop. But it catches it by making someone pay to find it.

X. How Legal Slop Enters The Legal Record

Legal slop becomes dangerous because law is archival.

The path is not complicated. An AI system generates legal-looking text. A pro se litigant, lawyer, staff member, or business user files or circulates it. The document enters a docket, transaction file, proposed order, or archive.

Opposing counsel may or may not check it. The court may or may not catch it. The document may be quoted, adopted, sanctioned, indexed, or copied. Later readers encounter the artifact as part of the legal environment.

Shahid shows this chain in its most troubling form. The court did not merely identify bogus authority in a party's brief. It addressed bogus authority in a trial-court order and vacated the order. The opinion noted that the order appeared to have been prepared by counsel, that counsel had cited the fictitious cases in the trial court, and that more fake cases appeared in the appellate brief.[30]

That is archival contamination. The defect did not remain in the drafting tool. It did not remain in a lawyer's notes. It entered a court order, then an appeal, then a published appellate opinion explaining the problem. The correction is now public. But so is the contaminant. The false cases have been preserved because the system must preserve the record of the error.

Fletcher underscores the same preservation problem in a different way. Westlaw's editor note on several of these cases warns that the questionable AI-linked citations are reproduced because they are part of the official record, but Westlaw does not provide links for them.[31] That is exactly the problem: legal archives preserve the artifact even when they cannot validate the embedded law.

Legal archives were designed to preserve law. They were not designed to quarantine legal-looking debris.

Nor are formal legal databases the only reservoirs from which bad legal material can be drawn. Lawyers, nonlawyers, and pro se litigants routinely search the open web for sample petitions, motions, briefs, contracts, and appellate filings. A document does not need to be a published opinion to become influential. A filed motion may appear in a docket search. A brief may be uploaded to a website. A proposed order may be copied from a public record. A contract form may circulate without anyone knowing whether it was drafted for a different jurisdiction, a different transaction, or a different risk profile.

These are slop troughs: collections of legal-looking text that can be copied, repurposed, summarized, scraped, or used as examples without anyone verifying whether the embedded law is good. The problem existed before AI. Lawyers have always borrowed forms, copied language, and reused filings. AI accelerates the practice and obscures the source. A litigant may ask a model to draft from examples found online. A lawyer may paste a public filing into a prompt. A model may have been trained on or retrieved from legal documents that themselves contained hallucinations, obsolete law, bad quotations, or context-specific provisions.

That creates a feedback loop. Legal slop enters a filing. The filing becomes public. A later user copies it. A model summarizes it. A new filing reproduces the error in more polished form. The repetition makes the proposition look more ordinary, not more true. Without a verification system that separates good law from legal-looking debris, the contamination can compound faster than the profession can audit it.

XI. Legal Slop Has A Longer Half-Life Than Code Slop

Code slop often decays in use. It breaks a build, fails a test, triggers a security alert, causes an outage, frustrates a user, or becomes unmaintainable. That feedback loop is imperfect, but it exists.

Legal slop can hibernate.

A hallucinated citation can sit in a filed motion that never produces a written order. A fake quotation can remain in a brief if the case settles. A proposed order can be signed before anyone checks every case. A docket entry can preserve the filing indefinitely. A researcher can later find the document without knowing whether the proposition was real, rejected, quoted, sanctioned, or merely repeated.

This persistence is stronger than the persistence of most bad code because legal records are built for permanence. A court system cannot simply delete a filed document because it contains bad law. Westlaw and Lexis cannot pretend the filing or opinion never happened. The record must preserve both the infection and the immune response.

The same is true, sometimes more quietly, in transactions. A defective AI-generated contract may never be litigated. It may sit in a company's files, govern the parties' relationship, be assigned to a buyer, be reviewed by a lender, or become relevant only years later. The defect may not be visible until the missing term matters. At that point, nobody will call it a hallucination. They will call it ambiguity, breach, waiver, indemnity failure, uninsured loss, tax exposure, governance deadlock, or malpractice.

> *Code slop usually has to run. Legal slop only has to remain.*

XII. The Westlaw And Lexis Blind Spot

This article should not overstate the problem. Westlaw and Lexis are powerful tools. They can help determine whether a case exists, whether it has been cited, whether it has negative treatment, and whether authority appears in reported legal materials. No serious legal researcher should work without those capabilities.

But AI-era legal slop raises a different question. The question is not only: does this case exist? The question is: did this legal proposition enter the ecosystem through unverified AI output?

That is a provenance question. It is not the same as KeyCite or Shepard's.

Noland illustrates why. If the cited cases exist but the quotations are fabricated, a search result confirming the existence of the cases does not solve the problem. The researcher must open the case and check the sentence. Fletcher makes the same point in broader terms: the modern hallucination problem includes fake propositions attached to real cases, not just obviously fabricated authorities.[32]

The legal research future will therefore require more than finding cases. It will require forensic reading. The researcher must ask where the proposition came from, whether the case actually supports it, whether the court adopted or rejected it, whether the quotation is real, and whether repetition is being mistaken for reliability.

Legal databases can preserve and index the legal record. But preservation is not provenance. A database may tell the researcher that a document exists. It may not tell the researcher whether the document is contaminated.

The open web makes the blind spot larger. Westlaw and Lexis at least organize legal material around citations, treatments, courts, and source hierarchy. Google does not. A public docket filing, a sample brief, a lawyer blog, a scraped PDF, an advocacy packet, a trial-court order, and a law review article may all appear as search results. To a user seeking language, each may look useful. To a model seeking patterns, each may look like legal text. But legal text is not legal authority.

The next legal research problem is therefore not only whether professional databases can flag hallucinated cases. It is whether the legal information ecosystem can prevent unverified legal artifacts from becoming raw material for the next generation of legal artifacts.

XIII. The Visible Edge Of An Unmeasured Problem

The known AI hallucination cases are serious, but they do not measure the problem.

Fletcher cited Damien Charlotin's AI hallucination database and reported that, as of February 18, 2026, the database identified 239 United States cases involving lawyer hallucinations.[33] That number is useful because it shows that the problem is not anecdotal. But the database itself is also useful because of what it does not claim. It tracks cases where a court or tribunal addresses AI hallucinations. It does not measure all fake citations or all AI use in filings.

That limitation is not a weakness. It is the point.

We know about Mata because a court sanctioned lawyers. We know about Kruse because an appellate court published an opinion. We know about Shahid because fake law entered a trial-court order and the appellate court vacated it. We know about In re Kheir because a bankruptcy court issued a published sanctions opinion. But there is no reliable count of hallucinated citations that were never checked, never challenged, never outcome-determinative, never published, or never preserved in a searchable system.

The sanctions cases are not the universe. They are the visible edge of the universe.

That is why the legal system should not wait for a crisis visible enough to count. By the time legal slop is visible in a reported opinion, the cost has already been imposed.

The unmeasured universe also includes documents that were never designed to be authority but are nevertheless used as drafting material. Those documents may never be cited in a sanctions order. They may never be indexed as cases. They may never be Shepardized or KeyCited because they are not authorities in the first place. But they can still shape legal arguments, pleadings, contracts, and AI outputs. That is why counting hallucinated citations in opinions, while useful, cannot measure the real contamination risk.

XIV. From Legal Research To Legal Forensics

Legal research has become legal forensics.

The traditional research questions remain necessary:

1. Does the authority exist?
2. What does it say?
3. Is it still good law?
4. Is it binding or persuasive?
5. What is the correct citation?

AI adds another layer:

1. Did the source actually support the sentence?
2. Is the quoted language in the source?
3. Did a party, court, or AI tool introduce the proposition?
4. Did the court adopt the proposition or merely describe a defective filing?
5. Is the authority being cited as law or as an example of hallucination?

6. Has repetition made the proposition look more reliable than it is?
7. Is the document being used as authority, example, form, or mere language?
8. Did the language come from a verified source or from a public filing copied without context?

That work is not glamorous. It is opening sources, reading holdings, checking pinpoints, verifying procedural posture, and refusing to cite what one has not confirmed. It is also the work that separates legal judgment from legal text.

But source forensics are not the whole professional skill. A verified sentence can still be useless, misleading, or harmful if it does not serve the client's objective, fit the forum, connect to an available remedy, or survive the factual stress the document is supposed to manage. AI makes this distinction more important because it can produce true sentences without understanding whether those sentences matter.

The profession should teach and reward that work. Law schools, courts, firms, and legal technology vendors should stop treating source verification as clerical. This work is and always has been central to the larger exercise of legal judgment.

XV. Source Ownership And Contextual Judgment As The New Legal Duty

The professional standard can be stated plainly in two related rules:

If your name is on the filing, you own the source.

If your name is on the work, you own the legal judgment.

The first rule is not new. Every case, quotation, parenthetical, procedural-history statement, statute, rule, and legal proposition had to be verified before AI existed. What AI changes is the risk environment. The lawyer is now more likely to encounter false, incomplete, or contextless legal text that looks polished enough to trust. It does not matter whether the first draft came from a lawyer, associate, paralegal, client, contract attorney, AI system, search platform, or public filing. The signer owns the source.

That is not merely an ethical slogan. It is built into the rules that govern signed filings.

The second rule is the deeper one. Legal work is not sound merely because the citations exist. The lawyer must own the fit between the words and the client's goal, the facts, the available remedies, the procedural vehicle, the jurisdiction, and the institution being asked to act. A brief that accurately quotes a real case may still ask for a remedy the court cannot give. A contract that uses elegant language may still omit the term that matters at default, bankruptcy, death, divorce, acquisition, or enforcement.

For litigators, that means AI-assisted briefs should have source binders. The binder should tie each proposition to the authority, pinpoint, quote check, status check, jurisdictional relevance, requested remedy, and the forum's power to grant it. For transactional lawyers, ownership means more than checking citations. It means owning the factual assumptions and stress conditions: who the parties are, what can go wrong, what law applies, what remedies matter, what consents are required, what defaults look like, and what future event the document must survive.

For courts, the response should begin from equal ownership of filed work. A hallucinated case does not become less contaminating because it was filed by a self-represented litigant. A fabricated quotation does not impose less work on the adversary because the filer lacked a lawyer. Courts may reasonably provide warnings, self-help materials, and plain-language notices before filing so unrepresented parties understand that AI can invent law. But

those tools should be framed as risk warnings, not as post-filing leeway. Once legal slop enters the docket, the relevant facts are the same: the filing contains false legal material, the opponent must respond, the court must spend resources, and the record has been polluted. The baseline rule should be the same for everyone: if your name is on the filing, you own the source.

For legal databases, the next challenge is provenance. A future research system should not merely identify negative treatment. It should help researchers identify hallucination cases, contaminated orders, and legal propositions that entered the record through unverified AI output. That will be difficult. It may not be fully solvable. But pretending the problem is only fake case names will not be enough.

XVI. Conclusion: Free Words, Legal Judgment

Artificial intelligence has made legal expression cheap. That is a profound change. It may help people understand their rights, organize facts, communicate with courts, and reduce the cost of routine legal work. The answer is not to reject the technology or romanticize the old barriers that kept people out of legal systems.

But the old scarcity of legal words performed one accidental function: it limited the volume of legal-looking material entering courts, archives, and transactions. That limit is gone.

The legal profession must therefore stop treating drafting as the center of legal work. Drafting is becoming easier. Contextual judgment is becoming harder and more important. Verification is part of that judgment, but it is not the whole thing. The lawyer's value is not merely producing words that look like law. It is knowing whether the words are true, lawful, useful, strategically connected to the client's objective, and capable of producing a remedy the institution can actually provide.

Vibe coding can produce technical debt. Vibe lawyering can produce legal debt. The debt may be paid by a court, an opposing party, a client, a buyer, a lender, a trustee, a spouse, a creditor, or a future lawyer trying to understand why a document failed.

> *Legal words are now free. Legal judgment is not.*

Endnotes

1. Christopher Mims, The AI Superstars Who Say a 'Vibe Slop' Crisis Is Coming, Wall St. J. (May 22, 2026, 5:30 AM ET), https://www.wsj.com/tech/ai/vibe-coding-slop-ai-tools-e6a99394.
2. Fletcher v. Experian Info. Sols., Inc., 168 F.4th 231, 233 (5th Cir. 2026); Mata v. Avianca, Inc., 678 F. Supp. 3d 443, 449-53 (S.D.N.Y. 2023).
3. Mata, 678 F. Supp. 3d at 461-66.
4. Kruse v. Karlen, 692 S.W.3d 43, 49-50 (Mo. Ct. App. 2024).
5. Noland v. Land of the Free, L.P., 114 Cal. App. 5th 426, 430-31, 435-36 (2025).
6. Shahid v. Esaam, 376 Ga. App. 145, 145-47, 149-51, 918 S.E.2d 198, 199-200, 202-03 (2025).
7. In re Kheir, 674 B.R. 631, 636-37 (Bankr. S.D. Tex. 2025).
8. Campbell v. Campbell, No. 03-25-00388-CV, 2026 WL 179402, at 2 (Tex. App.—Austin Jan. 22, 2026) (mem. op.).
9. In re Miller, No. 01-26-00319-CV, 2026 WL 1097107, at 1 (Tex. App.—Houston [1st Dist.] Apr. 23, 2026, orig. proceeding) (mem. op.) (Adams, C.J., concurring).
10. Ho v. Univ. of Tex. at Arlington, 984 S.W.2d 672, 679 (Tex. App.—Amarillo 1998, pet. denied).
11. Id.; see also Mansfield State Bank v. Cohn, 573 S.W.2d 181, 184-85 (Tex. 1978).
12. Fox v. Wardy, 234 S.W.3d 30, 33 (Tex. App.—El Paso 2007, pet. dism'd w.o.j.).
13. Kruse, 692 S.W.3d at 49-50.
14. Al-Hamim v. Star Hearthstone, LLC, 2024 COA 128, ¶¶ 25-26, 564 P.3d 1117, 1123.
15. Id. ¶¶ 39-41, 564 P.3d at 1126.
16. Ferris v. Amazon.com Servs., LLC, 778 F. Supp. 3d 879, 880-81 (N.D. Miss. 2025).
17. See Tex. R. Civ. P. 13; Tex. Civ. Prac. & Rem. Code §§ 10.001, 10.004; Fed. R. Civ. P. 11(b)-(c).
18. Fletcher, 168 F.4th at 234-35.
19. Id. at 233.
20. Id. at 239.
21. Id. at 239-40.
22. In re Kheir, 674 B.R. at 636-37.
23. Id. at 647-49.
24. Noland, 114 Cal. App. 5th at 436-37, 441, 443, 445-46, 448-49.
25. United States v. Hayes, 763 F. Supp. 3d 1054, 1064-67 (E.D. Cal. 2025).
26. Id. at 1063-64.
27. Id. at 1067.
28. Id. at 1064 (quoting Mata, 678 F. Supp. 3d at 448-49).
29. Ferris, 778 F. Supp. 3d at 880-81.
30. Shahid, 376 Ga. App. at 145-47, 149-51, 918 S.E.2d at 199-200, 202-03.
31. See, e.g., Fletcher v. Experian Info. Sols., Inc., 168 F.4th 231, 231 (5th Cir. 2026) (Westlaw editor's note); Noland, 114 Cal. App. 5th at 426 (Westlaw editor's note); In re Kheir, 674 B.R. at 631 (Westlaw editor's note).
32. Fletcher, 168 F.4th at 233.
33. Id. at 235.