

ARTIFICIAL INTELLIGENCE IN LITIGATION

*Hallucinated Citations, Sanctions, and the Hidden Risks of Using AI in Legal
Matters*

A Guide for Attorneys and Self-Represented Litigants

Published in association with
BE YOUR OWN LAWYER

www.beyourownlawyer.org

2026 Edition

A Note to the Reader

This book is a general-education resource. It is not legal advice, and reading it does not create an attorney-client relationship with the publisher, with Be Your Own Lawyer, or with any person or entity mentioned in these pages. Court rules, local rules, and case law change constantly, and the sanctions decisions summarized in this book are illustrative examples drawn from published opinions and reputable legal-industry reporting current as of mid-2026 — not an exhaustive or permanent list.

Every case, fine amount, and court order described in this book is real and drawn from published judicial opinions or reputable legal-news reporting. Citations are provided so that a reader — or the reader's attorney — can look up the full opinion independently. Because the law in this area is moving quickly, always confirm the current status of any case before relying on it.

If you are involved in an active legal matter with a deadline, do not wait. Consult a licensed attorney in your jurisdiction, or a qualified legal-support service, promptly.

Contents

Introduction: Why This Book Exists

Part One — The Five Problems

Chapter 1. Hallucinated Cases and Fake Authority

Chapter 2. Procedural and Jurisdictional Mismatches

Chapter 3. Confidentiality and Waiver of Privilege

Chapter 4. Misapplication of Substantive Legal Reasoning

Chapter 5. Court-Imposed Disclosure Mandates and AI Rules

Part Two — What You Stand to Lose

Chapter 6. Risk Assessment: Losing the Case, Paying the Price

Chapter 7. Case Studies in Sanctions, Fines, and Dismissal

Part Three — A Way Forward

Chapter 8. How Be Your Own Lawyer Prevents These Problems

Chapter 9. A Practical Checklist for Using AI Safely

Glossary of Terms Used in This Book

Conclusion: Using AI Wisely, Not Blindly

Introduction: Why This Book Exists

Generative artificial intelligence has moved into the practice of law faster than the rules that govern the practice of law have moved to meet it. Tools like ChatGPT, Claude, Gemini, and dozens of purpose-built "legal AI" products can now produce, in seconds, a document that reads like a competent legal brief — complete with case names, docket numbers, and quoted holdings. The problem is that a meaningful share of what these tools produce is simply wrong, and the tools deliver that wrong information with the same confident, fluent tone they use when they are right.

For a licensed attorney, filing that wrong information is a professional-responsibility problem. For a self-represented litigant — someone who cannot afford a lawyer and is doing their best to navigate a courthouse alone — it can be a life-altering problem. Courts have already dismissed pro se cases, imposed five- and six-figure fines, revoked pro hac vice admissions, and referred attorneys for discipline, all because someone trusted an AI-generated citation without checking it.

This book has four purposes. Part One lays out, in plain language, the five recurring problems that AI creates in legal filings. Part Two explains what is actually at stake — the practical consequences of getting this wrong — and then walks through real, documented cases in which courts imposed sanctions, fines, or dismissal because of AI-generated errors. Part Three explains how a trained, human legal-support resource such as Be Your Own Lawyer (www.beyourownlawyer.org) is built specifically to prevent the mistakes described in this book, and how a self-represented litigant can use that kind of help instead of, or alongside, an AI tool.

The duty to verify what goes in front of a judge has not changed. What has changed is how easy it has become to file something you never actually checked.

Chapter 1. Hallucinated Cases and Fake Authority

The Issue

General-purpose AI language models do not search a database of real cases the way a legal research platform does. They generate text by predicting statistically likely word patterns learned during training. When asked for a case citation, the model produces something that looks like a citation — a plausible party name, a plausible reporter volume, a plausible court and year — because that is the pattern it has learned, not because it retrieved a real record. The result is that these tools regularly invent judicial opinions, docket numbers, and quotations that were never written by any judge.

This is not a rare glitch. Independent research cited throughout the legal industry has found error rates for general-purpose AI models on legal queries running as high as the high double digits, and even specialized "AI-assisted" legal research products built on top of licensed case-law databases have measured error rates in the double digits. No AI legal tool available today can be trusted to be self-verifying.

Impact on Attorneys

For a licensed attorney, submitting a brief containing invented case law is not merely embarrassing — it implicates specific, enforceable professional rules. Model Rule of Professional Conduct 3.3 (Candor Toward the Tribunal) requires a lawyer not to knowingly make a false statement of law to a court, and requires correction of a previously made false statement. Federal Rule of Civil Procedure 11 requires that every filing be the product of a reasonable inquiry, and it exposes the signing attorney personally to sanctions when it is not.

Courts across the country have responded to AI-fabricated citations with formal Rule 11 sanctions, orders to show cause, mandatory reports to state bar disciplinary authorities, and direct monetary fines. The 2023 case of *Mata v. Avianca, Inc.* in the Southern District of New York is the case that first put the legal profession on notice: two attorneys submitted a brief containing six fabricated cases generated by ChatGPT, and the court fined them \$5,000 and referred the matter for further review. As Chapter 7 documents in detail, the penalties handed down since *Mata* have grown substantially larger and more severe.

Impact on Pro Se Litigants

Self-represented parties are, if anything, more exposed than attorneys. Many pro se litigants turn to AI tools precisely because they cannot afford a lawyer, and they often have no independent way to know that the tool has fabricated authority. Unlike an attorney, who faces a disciplinary board with defined procedures and, often, an opportunity to explain and remediate, a pro se litigant caught filing AI-hallucinated citations typically faces something faster and blunter: the pleading is struck, the claim is summarily dismissed, and in a growing number of cases the litigant is ordered to pay the opposing side's attorney's fees — money the litigant frequently does not have.

Why Courts Are Not More Forgiving of Pro Se Filers

Courts have traditionally given self-represented litigants some latitude on formatting and legal polish. But several courts have now drawn a sharp line between imperfect advocacy, which is tolerated, and fabricated authority, which is not. As one appellate court put it in dismissing a pro se appeal built substantially on invented cases, there is no forgiving construction of a citation that simply does not exist.

Chapter 2. Procedural and Jurisdictional Mismatches

The Issue

A general-purpose AI model is trained on an enormous, undifferentiated mix of federal rules, state statutes, and local court rules from every jurisdiction in the country — and often from other countries as well. When asked to draft a legal document, the model can blend these sources together into a single, fluent response that sounds authoritative but silently mixes rules that do not actually apply together. The model has no built-in awareness of which courthouse, which judge, or which specific local rule governs the document it is producing.

Pleading Deficiencies

This blending shows up most often as outdated legal standards, incorrect statutes of limitations, or procedural requirements borrowed from the wrong state. A limitations period that is correct for a breach-of-contract claim in one state may be entirely wrong in the state where the litigant actually lives, and an AI tool has no reliable way of knowing which one the user needs unless it is told precisely — and even then, it may still get it wrong.

Procedural Traps

Beyond substantive law, individual courts and even individual judges impose their own formatting and procedural requirements: page limits, font and margin rules, mandatory meet-and-confer certifications before certain motions can be filed, verification requirements for certain types of complaints, and local forms that must accompany a filing. AI tools generally do not know these hyper-local requirements, and a self-represented litigant relying on an AI draft frequently discovers the problem only when the clerk's office rejects the filing at the counter — sometimes after a filing deadline has already passed.

- Missed statutes of limitations drawn from the wrong state's law
- Federal civil procedure rules applied in a state-court case, or vice versa
- Missing mandatory certifications (meet-and-confer, verification, notice requirements)
- Incorrect page limits, caption formats, or citation formats specific to a local court
- Wrong venue or wrong court entirely, based on a generic or outdated rule

Chapter 3. Confidentiality and Waiver of Privilege

The Issue

Consumer-facing AI chat tools are, from a data-handling standpoint, third-party services. When a user types case facts, financial details, medical information, or the contents of a settlement discussion into a public AI model, that information is transmitted to and processed by the company operating the tool, and in many configurations it may be logged, retained, or used to improve future versions of the model. This is fundamentally different from a conversation protected by attorney-client privilege or the work-product doctrine.

Attorney Risk

For a lawyer, disclosing client information to a public AI tool without appropriate safeguards can violate Model Rule 1.6 (Confidentiality of Information), which requires a lawyer to make reasonable efforts to prevent unauthorized disclosure of information relating to the representation. Beyond the ethics rule, sharing privileged material with an outside, non-confidential service can support an argument that attorney-client privilege or work-product protection has been waived, because the information was voluntarily disclosed to a third party outside the privileged relationship.

Pro Se Risk

Self-represented litigants face a parallel but distinct danger. Without training in privilege or confidentiality doctrine, pro se parties routinely paste unredacted Social Security numbers, account numbers, medical records, children's names, or the full narrative of a dispute into a public chatbot to get help drafting a filing. That information can become part of a company's server logs, and depending on the platform's terms of service, it may be reviewed by human staff, used in model training, or become discoverable in ways the litigant never intended and would not have agreed to if they had understood the tradeoff.

A Practical Rule of Thumb

Never type into a public AI tool anything you would not be comfortable posting on a public online forum under your own name. If the information is sensitive enough that you would hesitate to post it publicly, it does not belong in a consumer AI chat window.

Chapter 4. Misapplication of Substantive Legal Reasoning

The Issue

Even when an AI model does not invent a case outright, it can still fail at the deeper task of legal reasoning. Large language models are optimized to produce text that sounds persuasive and well-organized; they are not optimized to correctly separate a court's binding holding from non-binding dicta, to apply the correct standard of review, or to recognize and navigate a split among the circuits or among a state's appellate districts. These are precisely the skills that distinguish a competent legal argument from one that merely sounds like one.

Superficial Arguments

The practical result is what legal commentators have started calling verbally fluent nonsense: text that reads smoothly, uses the right vocabulary, and follows a familiar brief structure, but whose actual legal analysis does not hold up once a judge or opposing counsel examines it closely. For a trained attorney, over-reliance on AI-drafted analysis risks submitting arguments that look complete on the page but collapse the moment a judge asks a follow-up question. For a pro se litigant, the risk is worse: if the litigant cannot explain or defend the reasoning behind their own filing in a hearing or at oral argument — because they did not actually construct that reasoning themselves — the gap becomes obvious and can undermine the litigant's credibility with the court on everything else in the case.

Chapter 5. Court-Imposed Disclosure Mandates and AI Rules

Standing Orders

In direct response to the problems described in Chapters 1 through 4, a rapidly growing number of individual federal judges, state trial courts, and even entire county court systems have issued standing orders that specifically address generative AI. These orders commonly require a certification — filed with, or attached to, the document itself — stating whether AI was used in preparing the filing and confirming that a human being independently verified every citation and every factual assertion generated by the tool. Some state bar rule-making committees have gone further, proposing rules under which citation errors traceable to generative AI misuse can support case-ending sanctions for both attorneys and pro se litigants.

Compliance Burdens

These disclosure requirements are not uniform. They vary by judge, by court, and by state, and they change frequently as more courts adopt them. A litigant — represented or not — who fails to include a required AI-use certification can face an automatic strike of the filing, entirely apart from whether the underlying content was accurate. This adds an additional, easy-to-miss procedural hurdle layered on top of the substantive risks already described.

- Some federal judges require an affirmative certification with every filing, regardless of whether AI was actually used
- Some state and county courts require disclosure only when AI was used, but leave the definition of "used" vague
- Requirements change frequently and differ from courtroom to courtroom, even within the same courthouse
- Failure to comply can result in a filing being struck automatically, without regard to its substantive merit

Chapter 6. Risk Assessment: Losing the Case, Paying the Price

Chapters 1 through 5 described the mechanisms by which AI errors enter a legal filing. This chapter distills what those errors actually cost, in practical terms, once a court discovers them. The consequences generally fall into a small number of recurring categories, and — as the case studies in Chapter 7 show — courts have shown a clear willingness to impose more than one of these consequences at the same time.

1. Losing the Case Outright

The single most severe outcome is dismissal — often with prejudice, meaning the litigant cannot refile the claim. Courts have dismissed claims, struck entire pleadings, and thrown out appeals specifically because a filing relied on fabricated case law. Dismissal with prejudice ends the case permanently, regardless of whatever underlying merit the claim might otherwise have had.

2. Monetary Sanctions and Fines

Courts routinely impose direct monetary sanctions on the party or attorney responsible for the AI-generated errors. As Chapter 7 documents, reported fines have ranged from a few hundred dollars for a first, promptly-acknowledged mistake up to well over \$100,000 in cases involving repeated or aggravated misconduct.

3. Paying the Other Side's Legal Fees

Beyond a fine payable to the court, a party found to have filed AI-hallucinated material is frequently ordered to reimburse the opposing party for the time and expense of identifying and briefing the fabrications — fees the responsible party would not have incurred but for the false citations.

4. Professional Discipline

For licensed attorneys, sanctions can extend beyond the immediate case: referral to a state bar disciplinary authority, suspension of a law license, or revocation of pro hac vice admission (the temporary permission to appear in a court outside the attorney's home jurisdiction).

5. Loss of Credibility With the Court

Even short of formal sanctions, a litigant who is caught submitting unreliable material loses the benefit of the doubt for the remainder of the case. Judges who discover fabricated authority scrutinize everything else that party files far more skeptically going forward — an intangible but very real cost.

The Bottom Line

AI-generated errors do not just risk losing an argument. Documented cases show they can end the entire case, generate a personal financial judgment against the litigant, and — for attorneys — put a law license at risk. The risk is cumulative: a single set of fabricated citations has, in reported cases, triggered

dismissal, a fine, and a fee award all at once.

Chapter 7. Case Studies in Sanctions, Fines, and Dismissal

The following summaries are drawn from published court opinions and reputable legal-industry reporting current as of mid-2026. A public database maintained by legal researcher Damien Charlotin now tracks well over a thousand court decisions worldwide in which a party relied on AI-hallucinated material and a court responded — a number that has grown from roughly two hundred cases in mid-2025 to well over a thousand within about a year. The cases below are representative examples, not an exhaustive list, and each should be independently verified before being cited in any filing.

Attorney Cases

Mata v. Avianca, Inc.

S.D.N.Y. 2023

Facts: Attorneys representing a plaintiff in a personal-injury suit against an airline used ChatGPT to help research and draft a brief opposing dismissal. The brief cited six cases that did not exist, complete with invented quotations and docket numbers.

Outcome: The court fined the two attorneys \$5,000 and issued a widely publicized opinion that became the founding case in this area of law. It put the entire legal profession on notice that AI-generated citations must be independently verified before filing.

Coomer v. Lindell

D. Colo., sanctions orders July 2025 and May 2026

Facts: In a high-profile defamation case, attorneys representing a well-known client filed a brief prepared with AI assistance that contained nearly thirty defective citations, including cases that did not exist.

Outcome: The federal court in Colorado fined each of the two attorneys \$3,000. The case illustrates that a prominent client and a high-profile matter provide no protection — courts evaluate the citations themselves, not the notoriety of the parties.

Whiting v. City of Athens

6th Cir., decided 2026

Facts: Across three consolidated appeals, two attorneys filed briefs containing more than two dozen fabricated citations and misrepresentations of fact — described by the appellate panel as a conservative estimate of the problem.

Outcome: The Sixth Circuit ordered each attorney to pay \$15,000 in punitive sanctions to the court registry, to reimburse the opposing party's full appellate attorney's fees across all three appeals, and to pay double costs. The court also directed that the opinion be forwarded to the chief judge for possible disciplinary proceedings. This remains one of the most severe reported sanctions to date at the federal appellate level.

Wyoming Federal Court Matter Involving Morgan & Morgan Attorneys

D. Wyo., 2025

Facts: Attorneys from a large national law firm filed motions citing nine cases; eight of them did not exist. The fabricated citations reportedly originated not from a consumer chatbot but from the firm's own in-house AI legal-research platform.

Outcome: The court revoked the drafting attorney's pro hac vice admission (his temporary permission to practice in that court) and fined him \$3,000, with \$1,000 fines for each of the two attorneys who signed the filings. The case shows that purpose-built legal AI tools are not immune from the same hallucination risk as general-purpose chatbots.

Couvrette v. Wisnovsky

D. Or., 2026

Facts: AI-generated fake citations and quotations appeared in briefing before the federal court in Oregon.

Outcome: The court imposed a sanction of \$110,204.38 — more than twenty times the size of the original 2023 Mata fine — reflecting the escalating severity with which courts are now treating this misconduct, particularly where it is repeated or the party fails to promptly self-correct.

Oregon Federal Court "Family Winery Dispute" Sanctions

D. Or., unpublished order, 2025

Facts: Across three separate briefs filed over a five-month period in an intrafamily dispute over a winery, the plaintiffs' filings included fifteen AI-generated fake case citations and eight fabricated quotations. When the fabrications were discovered, counsel offered no explanation.

Outcome: Because the evidence suggested the client had participated in the misconduct, the court dismissed the plaintiffs' claims with prejudice, fined the lead attorney \$15,500 using a per-fabrication penalty formula, and awarded the defendants their attorney's fees for both the summary-judgment and sanctions briefing.

Pro Se and Self-Represented Litigant Cases

Kruse v. Karlen

Mo. Ct. App. 2024

Facts: A self-represented appellant's brief cited twenty-four cases; only two turned out to be genuine. The brief also failed to comply with numerous other appellate rules — it was unsigned, lacked a required appendix, and had an inadequate statement of facts.

Outcome: The Missouri Court of Appeals dismissed the appeal as frivolous and ordered the appellant to pay \$10,000 in damages to the opposing party. The court rejected the appellant's after-the-fact apology, stating that the harm from filing fabricated authority had already been done regardless of any later contrition.

Crespo v. Tesla, Inc.

S.D. Fla. 2025

Facts: A pro se plaintiff suing Tesla filed multiple discovery motions citing case law hallucinated by an AI tool. Tesla moved to strike the filings and sought attorney's fees for the time spent identifying and responding to the fake citations.

Outcome: The court held that even an unrepresented, non-lawyer litigant is still required to follow court rules and act in good faith, and ordered the plaintiff to reimburse a reduced portion of Tesla's requested fees. The decision confirms that pro se status does not excuse the underlying misconduct, even when it may have been unintentional.

Estate of Kuerschner Probate Appeal

Ariz. Ct. App., decided March 2026

Facts: A self-represented widow appealing an adverse probate ruling filed briefs containing twenty-five citations that were fabricated or materially inaccurate, including a case that could not be located at all and a real case cited for a quotation that did not appear anywhere in the actual opinion.

Outcome: The appellate court dismissed the case under its inherent authority and imposed monetary sanctions covering the opposing party's attorney's fees and costs, expressly stating that while courts extend some latitude to pro se litigants on legal polish, they extend no such latitude to citations that are simply fake.

Employment Discrimination Matter Involving a Pro Se Plaintiff

Reported 2026

Facts: A self-represented plaintiff's filings relied on AI-generated citations and case law that did not withstand scrutiny.

Outcome: The court ordered the plaintiff to pay \$2,632 in attorney's fees to the defense and dismissed the complaint with prejudice, stating plainly that there is no forgiving reading of a citation that does not exist — such fabrications reflect bad faith and are sanctionable regardless of the litigant's pro se status.

Aldridge v. Tussing

Or. Sup. Ct., 2026

Facts: A litigant's filing before the Oregon Supreme Court cited cases that did not exist and quoted judges who never said the words attributed to them. After the court issued a show-cause order identifying the problem, the litigant filed a second submission — within twelve hours — that again contained fabricated material.

Outcome: The Oregon Supreme Court dismissed the entire proceeding, one of the court's first-ever sanctions orders specifically addressing AI-generated fabrications, holding that repeating the misconduct after being warned warranted a meaningful sanction.

What These Cases Have in Common

- None of the sanctioned parties intended, in most reported cases, to deceive the court — the fabrications originated with the AI tool, not the human filer. Courts sanctioned them anyway, because the duty to verify belongs to the filer, not the tool.
- Pro se status reduced or eliminated leniency once fabricated authority was involved, even though courts generally extend latitude to self-represented litigants on formatting and legal polish.
- Promptly acknowledging the error and cooperating with the court's inquiry (as in some Oregon matters) tended to reduce the severity of sanctions; denying responsibility or repeating the conduct after a warning tended to increase it dramatically.
- Fines have grown roughly twenty-fold since the first reported case in 2023, and courts are increasingly willing to combine dismissal, a monetary fine, and a fee award in the same order.

Chapter 8. How Be Your Own Lawyer Prevents These Problems

Every problem described in Chapters 1 through 5 has a common root cause: a self-represented litigant (or, in some of the cases in Chapter 7, even an attorney) relying on a tool that generates confident-sounding legal content without the human judgment, jurisdiction-specific knowledge, and accountability that real legal work requires. Be Your Own Lawyer (www.beyourownlawyer.org) is built around the opposite model. Rather than generating an answer from statistical patterns, it connects self-represented litigants with experienced legal professionals — the organization describes an average of roughly thirty years of courtroom and legal experience among its team — who research, verify, and prepare materials specifically for the litigant's actual case, in the actual court where it is pending. The sections below map that approach directly onto each problem identified earlier in this book.

1. Preventing Hallucinated Cases and Fake Authority

Because Be Your Own Lawyer's research is performed by people — not generated by a language model predicting plausible-sounding text — every case, statute, and quotation it prepares for a litigant is a real citation that a human being actually looked up. This directly addresses the core danger described in Chapter 1: there is no risk of an invented docket number or a fabricated quotation making its way into a filing, because the organization's stated approach is to research case law, statutes, and regulations that actually apply to the litigant's case, not to synthesize a plausible-sounding answer from a general pattern.

2. Preventing Procedural and Jurisdictional Mismatches

Be Your Own Lawyer's stated service specifically includes researching and identifying the rules relevant to a litigant's particular case — the deadlines, formatting requirements, and local procedures that a generic AI tool cannot reliably distinguish between jurisdictions. Where Chapter 2 described AI tools blending federal, state, and local rules into a single unreliable answer, this kind of human review is oriented around a single question: what does this specific court, in this specific case, actually require? That is precisely the gap that leads to the clerk's-counter rejections and missed deadlines described in Chapter 2.

3. Preventing Confidentiality and Privilege Problems

Working with a dedicated legal-support service rather than pasting case facts into a public, consumer-facing AI chat tool avoids the data-exposure risk described in Chapter 3. Information shared for the purpose of obtaining legal help through a professional service is handled in the context of that professional relationship, rather than being transmitted to a general-purpose AI company's servers for indeterminate logging or model-training use. Litigants should always confirm the specific privacy terms of any service they use, but the basic architecture — a human professional relationship built around the litigant's matter — is fundamentally different from a consumer chatbot.

4. Preventing Misapplication of Substantive Legal Reasoning

The "verbally fluent nonsense" problem described in Chapter 4 arises because AI models are optimized to sound persuasive, not to correctly apply legal standards like standard of review or the holding-versus-dicta distinction. Be Your Own Lawyer's model relies on people with decades of actual courtroom experience to organize the law that applies to a litigant's case and explain it in terms the litigant can understand and, importantly, explain themselves. This matters directly at the stage described in Chapter 4: a litigant who understands why an argument works is far better positioned to answer a judge's question at a hearing than one holding a document they cannot defend.

5. Helping Litigants Meet Court-Imposed AI Disclosure Requirements

Because Be Your Own Lawyer's materials are prepared by humans performing genuine legal research rather than generated by an AI tool, a litigant using the service starts from a straightforward, accurate position when a court's standing order asks whether AI was used to prepare a filing and whether every citation was independently verified, addressing directly the compliance burden described in Chapter 5.

6. Support Across the Full Life of a Case

Beyond the five problems above, Be Your Own Lawyer describes assistance across the practical stages a self-represented litigant is likely to face and has the least experience with: preparing pleadings, motions, and briefs; navigating discovery, including interrogatories, requests for production, and depositions; and understanding court procedure — what to do, and when, once a case is filed or a hearing is scheduled. Each of these is an area where the case studies in Chapter 7 show real litigants losing cases or incurring sanctions not because their underlying claims lacked merit, but because the mechanics of the filing itself went wrong.

Bottom Line

The organization's own description of its service is built around a value proposition directly responsive to this book's five core problems: real human legal experience, applied to the litigant's actual case, at a cost far below traditional attorney's fees. A self-represented litigant considering whether to rely on a free AI chatbot or a paid legal-support service should weigh that choice against the documented cases in Chapter 7 — the modest cost of professional help is, in every reported case in this book, dramatically smaller than the fines, fee awards, and lost claims that followed unverified AI use.

This chapter describes Be Your Own Lawyer's own stated services, drawn from its published website content as of 2026. As with any professional service, a prospective client should independently confirm current offerings, pricing, and scope of engagement directly with the organization before proceeding.

Chapter 9. A Practical Checklist for Using AI Safely

For readers who will use AI tools in some form despite the risks described in this book — whether as an attorney doing preliminary research or as a self-represented litigant looking for a starting point — the following checklist distills the lessons of Chapters 1 through 7 into concrete, repeatable habits.

Before You Use an AI Tool

- Never enter a client's confidential information, a settlement number, or any sensitive personal detail into a public, consumer-facing AI tool.
- Remember that the tool has no reliable way to know which state, county, or judge your case is in unless you tell it precisely — and even then, it may still apply the wrong rule.
- Treat any AI output as a first draft only, never as a finished product.

Before You File Anything

1. Pull up every single case cited and confirm it exists, in the reporter and volume cited, at a real legal database or the court's own website.
2. Read the actual case — do not rely on the AI's summary — and confirm the quotation or holding attributed to it is really there.
3. Confirm the statute of limitations, filing deadline, and procedural rule cited is the current rule in the correct jurisdiction.
4. Check your specific court's website or clerk's office for local rules on page limits, formatting, and required certifications.
5. Check whether the judge assigned to your case has a standing order on AI use, and comply with any required disclosure or certification.
6. If anything in the draft is beyond your ability to verify or fully understand, get a qualified human — an attorney or a legal-support professional — to review it before you file.

If You Discover an Error After Filing

Several of the cases summarized in Chapter 7 show that courts respond very differently to litigants who promptly acknowledge and correct an error compared to those who deny responsibility or repeat the mistake after being warned. If you discover that a filing contains an AI-generated error, correct it and notify the court as soon as possible, rather than waiting for opposing counsel or the judge to find it first.

Quick Reference: The Five Problems and Their Fixes

Hallucinated citations → verify every case yourself, in a real database. Procedural mismatches → confirm your specific court's and judge's rules. Confidentiality → never paste sensitive facts into a public tool. Weak legal reasoning → have a qualified human review the substance, not just the grammar.

Disclosure mandates → check for, and comply with, any standing order on AI use before you file.

Glossary of Terms Used in This Book

AI hallucination. A confident, fluent output from an AI model that is factually false — for example, a citation to a case, statute, or quotation that does not actually exist.

Candor to the tribunal. The professional-conduct duty (Model Rule 3.3) requiring a lawyer not to knowingly make a false statement of law or fact to a court.

Dicta. Commentary in a judicial opinion that is not part of the court's binding holding, and therefore is not binding precedent.

Dismissal with prejudice. A dismissal that permanently bars the litigant from refiling the same claim.

Pro hac vice. Temporary permission for an attorney to appear in a court outside the jurisdictions where they are normally licensed.

Pro se litigant. A person who represents themselves in a legal matter without a licensed attorney.

Rule 11 (Fed. R. Civ. P.). The federal rule requiring that every filing be based on a reasonable inquiry into the facts and the law, and authorizing sanctions when it is not.

Standing order. A court's or judge's general rule that applies automatically to all cases before it, as opposed to an order issued in a single case.

Work-product doctrine. A protection that shields materials an attorney prepares in anticipation of litigation from discovery by the opposing party.

Conclusion: Using AI Wisely, Not Blindly

Nothing in this book is an argument that AI has no place near the practice of law. Used well — as a starting point for brainstorming, as a tool for organizing a litigant's own notes, or as a way to generate a first draft that a qualified human then thoroughly checks against real, verified sources — AI can genuinely help. The failures documented throughout this book share a single common thread: someone treated an AI tool's output as a finished, trustworthy product instead of as an unverified first draft.

The rule that emerges from every case in Chapter 7, without exception, is simple to state and hard to remember under deadline pressure: nothing generated by an AI tool belongs in a court filing until a human being has independently confirmed that it is real, accurate, and legally correct for the specific court and case involved. For attorneys, that duty is a matter of professional discipline. For self-represented litigants, it may be the single most important lesson in this book — because, as the cases in Chapter 7 show, the courthouse does not treat good intentions as a substitute for a verified filing.

Litigants who cannot afford full attorney representation are not without options. Trained legal-support resources such as Be Your Own Lawyer exist specifically to close the gap between representing yourself and going in with no help at all — providing the human research, procedural knowledge, and document preparation that an AI chatbot cannot reliably provide on its own.

BE YOUR OWN LAWYER

Helping you represent yourself in court

www.beyourownlawyer.org