Sure, I Can Draft a Complaint! LLM Hallucination in Pro Se Litigation

Benjamin Murphy
bmurphy@jd27.law.harvard.edu
Harvard Law School

Abstract

Large language models (LLMs) are used widely across legal practice, both by professionals and by individuals seeking access to the court system. This latter class of users are known as *pro se* litigants, and they lack legal training to be able to vet an LLM's output for accuracy. As a result, they may be disproportionately likely to file documents in court which rely on hallucinated cases, or which misstate the holding of those cases. In this paper, we evaluate five LLMs for accuracy when drafting civil complaints for *pro se* plaintiffs, on topics including housing law, family law, and tort law. We find that while models generally state facts sufficient to establish a claim, they often rely on hallucinated case law (with rates ranging from 5% to over 30%), and cite cases for unsupported propositions in a majority of instances across all evaluated models. These preliminary findings indicate that there are unique risks for *pro se* litigants in using LLMs, and that generative AI may present novel risks to such individuals, though more research is needed.

1 Introduction

Large language models (LLMs) have become ubiquitous across legal practice, and many companies are in the process of developing tools to aid professional lawyers in drafting contracts, assessing legal arguments for their strength, and identifying relevant precedent for a point of law. (Krook et al. 2024; Capoot 2025; Ambrogi 2025) These tools are not perfect, however, and lawyers who failed to double-check their work have often gotten into hot water. Some lawyers have been sanctioned, fined, or disbarred for referencing hallucinated case law, and fake citations have even made their way into judicial decisions. (Murphy 2025; Charlotin 2025) These pose problems for the judicial system and can threaten the trust placed in the rule of law, but judges and lawyers are not the only individuals who turn to LLMs to answer legal questions.

Many individuals rely on LLMs to assist with personal legal matters, whether defending against a spurious damages claim, to suing a landlord that fails to maintain basic standards of living in a rental property. These individuals, unlike professional lawyers, lack training on how to verify the truthfulness and accuracy of the LLM's output. As a result, hallucinated cases or mistaken propositions of law can result in legal consequences a litigant was powerless to avoid—and unlike a professional lawyer, few economic incentives exist to motivate companies to improve on this use case.

In this paper, I examine hallucination rates when LLMs are used by *pro se* litigants—that is, litigants who represent themselves before a court. These individuals are less likely to ask nuanced legal questions, less able to validate the LLM's output, and likely have fewer resources to acquire outside legal advice, making LLMs a double-edged sword: When they perform accurately, they can enable these individuals to vindicate their rights via the judicial system when they may have been unable to do so otherwise; but when they fail, they can leave the affected individuals in a worse place than they would have been otherwise. Consequently, the same rate of hallucinations may present more risk to an untrained litigant who lacks the ability to spot mistakes, making the evaluation and mitigation of this risk a priority for any model developer who allows their models to offer legal advice.

2 Background

2.1 Pro Se Litigation

A significant portion of civil litigants in both state and federal court proceed without counsel, often because they cannot afford a lawyer or cannot secure free legal assistance. (Legal Services Corporation 2017) Recent reports suggest that in federal district courts, roughly one quarter of all civil cases involve at least one *pro se* party, and in certain categories—such as prisoner petitions and employment discrimination suits—the rates climb substantially higher. (Administrative Office of the U.S. Courts 2025) State courts, which handle the bulk of family law, housing, small claims, and consumer disputes, see even larger proportions of self-represented litigants, with some jurisdictions reporting that a majority of family-law litigants appear without counsel. (Agor et al. 2015)

For these individuals, legal representation may simply not be possible. Legal aid organizations are overwhelmed, eligibility standards are stringent, and many areas of law (and areas of the country) have no *pro bono* infrastructure at all. As a result, filing a *pro se* complaint is often the only available mechanism for vindicating legal rights. The rise of generative AI has yielded a dramatic shift in how *pro se* litigants attempt to navigate the system. Courts at every level have encountered complaints, motions, and supporting memoranda drafted partly—or entirely—by LLMs. (Warren 2025) In many cases these filings are more polished and coherent than the alternative; LLMs can provide basic structure, formal language, and a sense of legal framing that many unrepresented litigants would struggle to produce on their own. However, *pro se* litigants often demand the introduction of additional "legalese," such as unnecessary citations, even when they are not required to state a claim.

But the same tools also introduce new risks. Unlike traditional self-help resources—which tend to be conservative, procedural, and limited—LLMs readily generate legal claims, arguments, and citations on demand. (Dahl et al. 2024) For low-income litigants who lack doctrinal grounding, the seeming fluency and confidence of the model can mask the fact that the legal content is wrong, fabricated, or incoherent. This dynamic has produced a surge of AI-assisted filings that contain invented cases, nonexistent statutes, and doctrinal misstatements. (Dahl et al. 2024) While some of these filings reflect legitimate disputes obscured by poor drafting, others have drawn sanctions or dismissal because the AI-generated material was wholly imaginary. For example, in Jarrus v. Governor of Michigan, the magistrate judge ordered the *pro se* plaintiffs to pay $300 due to their use of ChatGPT—even when the generated brief cited only actual cases, because the judge found that the AI had impermissibly hallucinated the breadth of the holdings of the cases themselves. (U.S. District Court for the Eastern District of Michigan 2025)

2.2 Hallucinated Law and Court-Ordered Sanctions

Courts have begun documenting a steady stream of LLM-induced errors, ranging from subtle doctrinal distortions to fully fabricated case law. Although the most famous early incident involved represented parties—where lawyers in the Southern District of New York filed a brief containing hallucinated federal appellate decisions—many concerning filings now originate from *pro se* litigants. (Pfefferkorn 2025) In dozens of jurisdictions, judges have confronted complaints citing cases that do not exist, motions premised on imaginary constitutional theories, and affidavits quoting judicial language that no reported version of the case actually contains.

The judicial response has been uneven. Some judges have addressed the issue by issuing warnings, striking defective filings, or directing litigants to supplemental self-help resources. Others have imposed monetary sanctions, even against *pro se* litigants, when a filing wastes court resources or violates Rule 11 equivalents. ("Rule 11: Signing Pleadings, Motions, And Other Papers; Representations to the Court; Sanctions" 2007) Surveys of trial-level opinions show that courts have sanctioned AI-induced errors in a growing number of cases, with fines ranging from nominal amounts to several thousand dollars. These consequences fall hardest on unrepresented parties, who often lack the ability to distinguish a plausible legal argument from an invented one. As a result, LLM hallucination interacts with existing access-to-justice inequities, exposing the most vulnerable litigants to the greatest procedural harm.

3 Prior Work

A growing body of empirical research documents the prevalence and severity of hallucinations in legal contexts. Early technical evaluations show that mainstream LLMs hallucinate at exceptionally high rates on legal tasks: studies from Stanford, MIT, and others report that models such as GPT-3.5, GPT-4, and Llama 2 frequently invent case law, misstate holdings, or fabricate statutory language. (Dahl et al. 2024) Error rates vary by task, but even the most capable models exhibit hallucination rates approaching 60–70% on open-ended legal research prompts. These failures extend across jurisdictions, across levels of difficulty, and across topical domains.

Scholars have also warned that the consequences of these inaccuracies are not evenly distributed. Because professional lawyers possess the expertise to identify obvious mistakes, they are less likely to rely uncritically on incorrect outputs. *Pro se* litigants, by contrast, often cannot distinguish hallucinated law from legitimate authority. As several commentators have noted, LLM errors thus risk disproportionately harming unrepresented individuals by encouraging the assertion of nonexistent rights, misinterpreting procedural rules, or misstating the required elements of a claim.

At the same time, most technical "solutions" to hallucination currently arise in commercial, professional-grade products. Systems like Westlaw's AI Research Assistant, Lexis+ AI, and Harvey feature retrieval-augmented generation, post-hoc verification, and citation-level fact checking. These tools reduce hallucinations by constraining the model to verified sources—but they are expensive, proprietary, and targeted at law firms, not individual litigants—and yet still hallucinate up to one-third of the time. (Magesh et al. 2025) As a result, the benefits of hallucination-resistant legal AI accrue primarily to institutional actors. *Pro se* litigants, who rely on free public-facing LLMs, confront the highest error rates with the weakest guardrails. (Krook et al. 2024) The gap between professional and public AI systems is therefore widening precisely in the domain where reliable guidance is most urgently needed.

Existing research has focused on highly specific legal questions that align with professional use cases, such as "What standard of review applies to abortion regulations in the United States?" or "Why did Justice Ginsberg dissent in *Obergefell*?" (Dahl et al. 2024) Similarly, LegalBench, a benchmark released by many of the same authors, provides an evaluation for LLMs on many professional legal activities, like interpreting contracts, issue-spotting, and test application. (Guha et al. 2023) Some of the same authors later examined professional research tools and found that while those tools reduce the rate of hallucination, they do not eliminate the problem entirely. (Magesh et al. 2025)

These works do not, however, examine the capacity of LLMs to construct wholly new legal complaints, or simply fulfill legal tasks correctly without regard for whether or not they can answer detailed questions along the way. Stated differently, a *pro se* litigant does not care what the governing standard of review is—they simply want to obtain custody over their child, get redress for another's negligent behavior which caused an injury, or seek recompense from their landlord due to the violation of housing codes. Moreover, the framing of legal questions in these prior works often means the LLM must retrieve one particular opinion or identify a narrow proposition of law, while for *pro se* litigants, *any* actual case establishing (e.g.) that a given course of conduct constitutes negligence is sufficient for the purposes of a complaint.

4 Experimental Design

In this experiment, I evaluate LLM performance on a foundational task of civil litigation: The drafting of a complaint. A complaint must set out "a short and plain statement of the claim showing that the pleader is entitled to relief," and additionally state the relief sought. ("Civil Procedure Rule 8: General Rules of Pleading," n.d.) Here, I evaluate three potential claims in a *pro se* complaint: A *breach of the covenant of quiet enjoyment*, a legal term of art which refers to the right of a tenant to utilize the leased property without substantial interference; a *request for modification of a custody order*, which seeks to amend a time-sharing arrangement following a material change in circumstances of one parent (e.g. a new job or new home); and a *negligence lawsuit seeking damages*, a standard common law civil suit seeking reparations for someone else's negligent behavior which caused injury (e.g. failing to warn about dangerous conditions, driving through a red light).

For each experimental category, ten "fact patterns" are generated, and each fact pattern is expanded into ten complaints by varying the plaintiff's name, defendant's name, and damages sought (among other "background" information). These complaints are then assessed on three dimensions: Whether they actually state sufficient facts to support the required elements for each claim, whether they cite hallucinated case law; and whether each non-hallucinated citation actually supports the proposition it is cited for.

4.1 Data Generation

Data was generated in three stages. First, I manually searched for and considered different causes of action that would be appropriate for this (preliminary) research, requiring three criteria:

Generated Housing Complaint

```
The main entry door to my building in Chelsea has a broken lock—it doesn't
latch properly and anyone can just push it open. I told my landlord about
this maybe two months ago. He said he'd fix it but nothing has happened.
I've reminded him twice since then.

Last month someone came into the building and went through the mail—a bunch
of us had packages stolen from the hallway. I lost a package worth about $120.
I've also started finding evidence that someone has been sleeping in the
basement, which is really unsettling. I bought a security camera for my own
door and a better deadbolt, which cost me about $200. My landlord says I'm
"overreacting" and that the neighborhood is safe, but clearly that's not true.
```

Figure 1: An example fact pattern for a landlord-tenant lawsuit, asserting conduct which violates Massachusetts housing laws and establishes damages resulting from that breach.

That the claim would be likely to be experienced by a low-income individual without access to an attorney; that the claim required pleading at least four elements (for example, negligence requires proving that the defendant had a duty to the plaintiff, that they breached that duty, that the plaintiff was damaged, and that those damages were caused by the breach); and that the claim arises under state law.

I limited the search to state law for two reasons. First, federal law is limited and governs less of an individual's day-to-day conduct, making it less likely to be the governing law for a *pro se* complaint. Second, LLMs are less likely to hallucinate widely-repeated cases, which is more likely to apply to federal precedent than state precedent. All else equal, I decided to focus on Massachusetts state law, so as to avoid potentially-conflicting state standards.

Having decided upon three causes of action, as listed above, I created ten fact patterns for each complaint using few-shot GPT-5.1, seeking to create fact patterns of the same style as a *pro se* litigant might present their case: Short, prose descriptions of what happened, given from the perspective of the individual. For an example, see Figure 1. Then, I generated ten sets of "background information" for each fact pattern, containing details like each party's name, the location in question (for tenant disputes), the damages sought, etc. For an example, see Figure 2.

For each hydrated set of facts and background information, complaints were generated using five models: `gpt-5.1`, `gpt-5-mini`, `gpt-5-chat-latest`, `gpt-4o`, and `gpt-3.5-turbo`, yielding 1500 total complaints across three categories. These complaints were then stored in text format. For an example complaint, generated by `gpt-5-chat-latest`, see Figure 3. An example prompt used to generate complaints is shown in Appendix A.

4.2 Evaluation

4.2.1 Elements of the Claim

A complaint must set out sufficient facts to allege each element of the stated claim, though it does not need to do so verbatim. To assess this, each generated complaint was evaluated by `gpt-5-mini` using a system prompt containing the relevant elements of the claim.

Example Complaint for Negligence Action

```
### IV. CAUSE OF ACTION
**NEGLIGENCE**

36. Plaintiff incorporates by reference and realleges Paragraphs 1 through 35 of this
Complaint as if fully set forth herein.

**Duty**

37. Under Massachusetts law, a property owner or person in control of real property owes all
lawful visitors a duty to exercise reasonable care to maintain the property in a reasonably
safe condition in view of all the circumstances. *Mounsey v. Ellard*, 363 Mass. 693, 707-08
(1973).

38. This duty of reasonable care extends to the maintenance of exterior areas appurtenant to
the property, including sidewalks, walkways, and means of ingress and egress that tenants and
their invitees may reasonably be expected to use. *Dipasquale v. Gow*, 3 Mass. App. Ct. 626,
628 (1975); *Luz v. Stop & Shop, Inc. of Peabody*, 348 Mass. 198, 202 (1964).

. . .
```

Figure 3: An (excerpted) generated complaint, detailing the underlying facts, parties and jurisdiction, and relief sought.

4.2.2 Hallucinated and Improper Citations

For each complaint, citations were extracted using `gpt-5-mini`, and categorized as citing either case law or statutory law. For the purposes of this preliminary work, only hallucinated case law

Example Background Facts Section for Custody Complaint


```
{
  "id": "cm_001_08",
  "category": "custody_modification",
  "fact_pattern": "A few years ago, the court set up a parenting schedule for our two kids. I
have them most of the time, and their dad gets them every other weekend plus one overnight
during the week. Back when we made that arrangement, we both worked regular daytime jobs, so
the schedule made sense. Recently my ex took a new job working 7 PM to 7 AM several nights a
week, including many of the nights he's supposed to have the kids overnight. He's been
dropping them off at my place with almost no notice because he says he's too tired to watch
them after his shift, or he leaves them with his roommate, who I barely know. The kids tell me
they often fall asleep on the couch while he naps after getting home from work. I want to
change the schedule so overnights only happen when he's actually off work and available.",
  "background": {
    "petitioner_name": "Jessica Thompson",
    "respondent_name": "Michael Davis",
    "child_names": "Emily (8), Jacob (6)",
    "original_order_date": "2021-04-15",
    "current_arrangement": "Jessica has primary custody with Michael having the kids every
other weekend and one overnight during the week.",
    "venue": "Middlesex Probate and Family Court"
   }
}
```


Figure 2: A fully-hydrated example, containing both background facts and the first-person perspective on the events leading to the lawsuit.

was investigated. Each extracted citation was then automatically looked up via CourtListener, a free, public database of case law, which provides an API for searching for citations. (Free Law Project, n.d.) Note that while CourtListener's database is enormous, it is likely not perfect. Hence, it is possible that false-negatives exist when searching for hallucinations. These errors are also possible if the format of the citation does not match the recorded version in CourtListener. If the case was found, the full text of the primary opinion was fetched and stored.

Each sentence supported by a citation was also extracted by `gpt-5-mini`, and a subsequent call to the model asked it to examine whether that sentence was properly supported by the citation. In doing this extraction, the model both excerpted the most relevant portion of the opinion, and included a short argument for why the proposition was or was not supported, along with an overall confidence level. The prompt used for support evaluation is shown in Appendix A.

5 Results

5.1 Elements of the Claim

I evaluated every complaint generated by `gpt-5-mini` to identify whether each complaint stated facts sufficient to establish each element of the claim, and found that **every** complaint did. This is largely unsurprising: The fact patterns were generated to ensure that each could give rise to the desired complaint, so simply by restating those facts in the complaint itself, the result should satisfy the legal standard. Seeing this result, I omitted evaluation of the other models' complaints for cost reasons.

5.2 Hallucinated Citations

5.2.1 Rate of Hallucinated Citations

Every model evaluated demonstrated a substantial rate of hallucinations, with `gpt-3.5-turbo` and `gpt-5-mini` demonstrating the highest rates. Figure 4 shows the overall rate per model, while Figure 5 shows the rate broken down by the claim type in the complaint. These preliminary results suggest that models with larger parameter counts may be capable of "memorizing" more case law, and consequently hallucinating less. Nevertheless, even for state-of-the-art models like `gpt-5.1`, a 5.6% hallucination rate means that many legal documents will contain at least one, if not multiple, hallucinated references to case law.

5.2.2 Rate of Incorrect Citations

For non-hallucinated citations, I plotted the rates of support—that is, whether the proposition in text that the case was cited for was supported by the case itself. Figure 6 shows the rates of support per model, which are low across the board (less than 50% for all models, with the three most-advanced models demonstrating the lowest support scores.) These numbers are suspiciously low; manual investigation of the data revealed at least three possible causes.

1. `gpt-3.5` and `gpt-5-mini` generates markedly fewer citations per complaint than the larger models, potentially disregarding aspects of the prompt which direct the model to "cite relevant case law."
2. All models, including those with low overall hallucination rates, cite generally-related cases that at times make nuanced points which undercut the actual proposition needed for the complaint. For example, in citing Custody of Kali, 439 Mass. 834 (2003), the model has

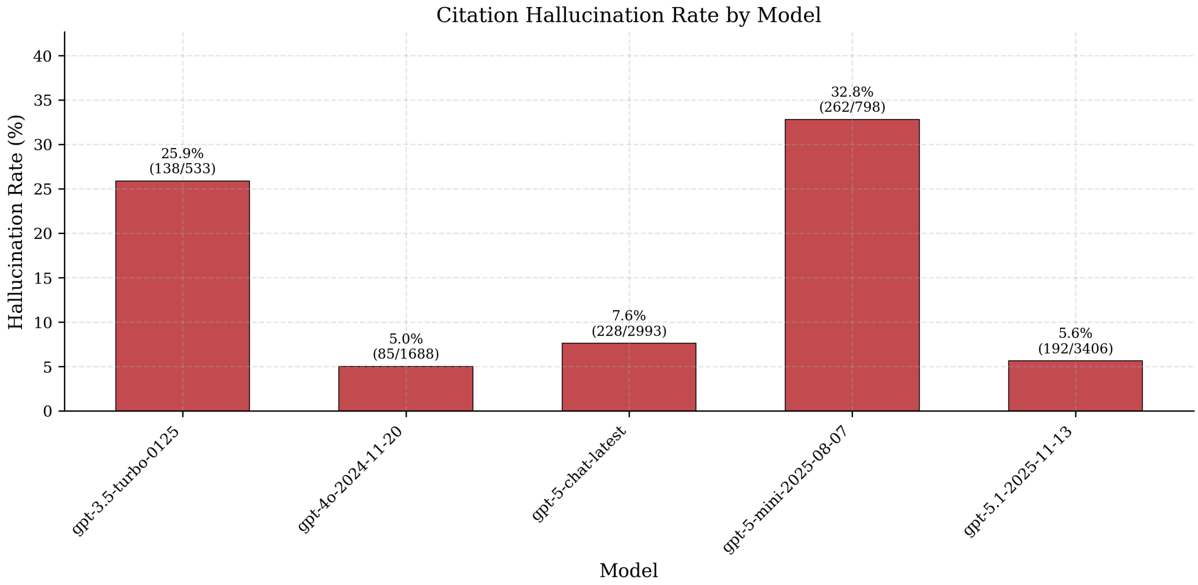

Figure 4: Hallucination rates by model. Each citation is counted only once per complaint, so repeated uses of the same hallucinated case within a single complaint are counted only once.

correctly identified a real case about custody, but incorrectly cited a case which expressly declined to consider *modification* of an existing custody order (when the complaint sought such modification). This mistake might not be enough to yield a court-ordered sanction, but it is certainly bad legal practice.

3. The support evaluation, at times, requires too granular a match. For example, when `gpt-5.1` cited Adoption of Hugo, 428 Mass. 219, 225 (1998), the evaluation model argued that the case failed to set out a multi-factor test that matched the factors described in the complaint, but the case did analyze many similar factors in considering whether to change the custody of the child.

Together, these results suggest that more work is needed to properly assess whether non-hallucinated citations actually support the claims stated in the complaints—but the top-level results indicate, at a minimum, that even state-of-the-art models regularly cite cases for

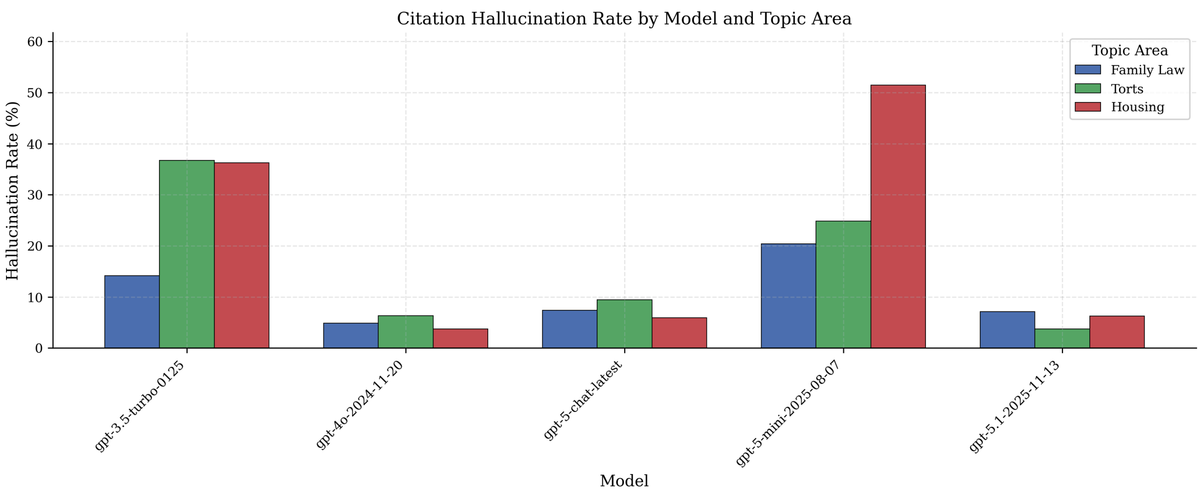

Figure 5: Hallucination rate, as shown in Figure 4, broken down by claim stated in the complaint. For highly hallucinatory models, distribution across topics is highly uneven.

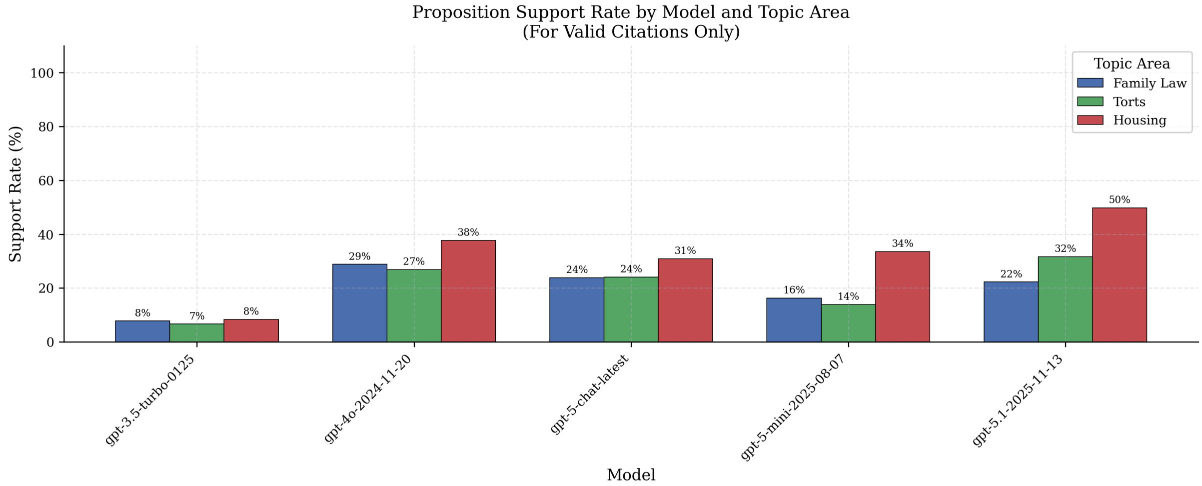

Figure 6: Support rate for valid citations, shown broken down by both model and topic.

conclusions of law that they do not support: `gpt-5.1` misstates the holdings of cases between 50% and 78% of the time, depending on the topic of the complaint.

6 Conclusion

This work provides a literature review and preliminary analysis suggesting that state-of-the-art models remain likely to hallucinate case law and misstate case holdings when used to generate state law complaints by a *pro se* litigant. This work begins to demonstrate that LLMs, rather than furthering access to justice, may in fact expose *pro se* plaintiffs to sanctions and fines. Courts may well attempt to limit or eliminate uses of generative AI, whether by *pro se* litigants or professional attorneys, but doing so eliminates any potential benefit that these individuals might receive. The complaints generated in this study were vastly better than what any individual could construct without the help of a lawyer. Consequently, this work indicates that more study of legal hallucinations in *pro se* contexts is needed, but does not suggest that the only solution is to eliminate the uses of AI entirely.

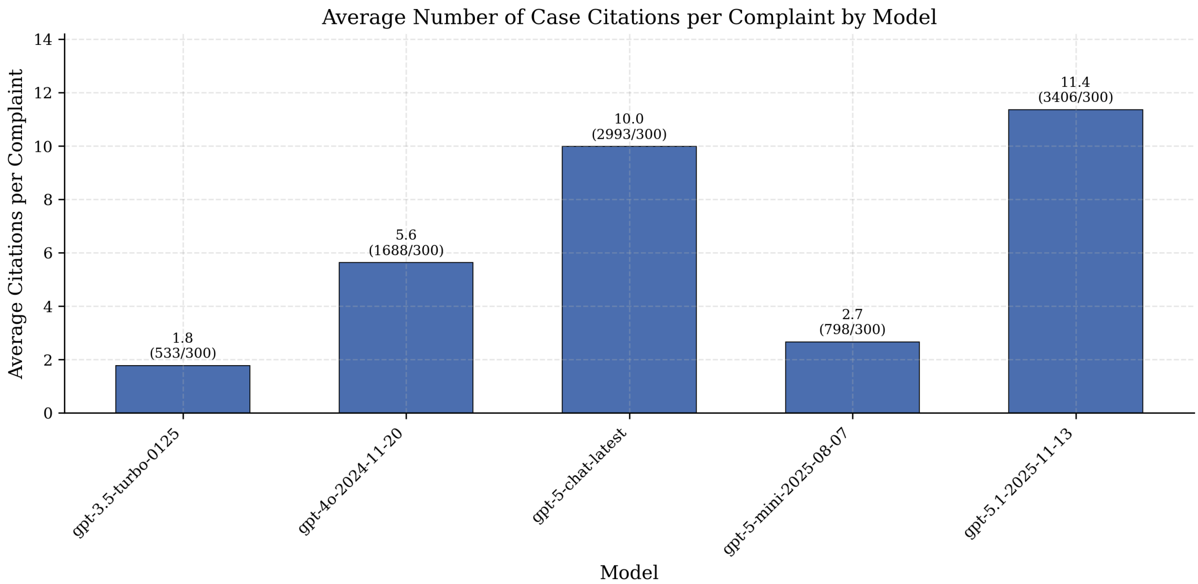

7 Theory of Change

The widespread applicability of generative AI means that model developers cannot identify every potential harm on their own. While plenty of existing work has evaluated the suitability of LLMs for *professional* legal use cases, little attention has been paid to how individual (i.e. *pro se*) litigants make use of generative AI. Unlike professional use cases, little commercial incentive exists to improve on *pro se* use cases. This work aims to shed light on the risks of generative AI for *pro se* litigants, provide a preliminary schematic for evaluating model performance in these use cases, and provide the necessary impetus for model developers to incorporate this type of evaluation in the training process. If done properly, *pro se* litigants can benefit from the immense promise of LLMs without incurring new risks of court-ordered sanctions, yielding an increase in access-to-justice for all.

8 Future Work

8.1 Improved Hallucination Evaluation

Like many forms of AI safety research, there are substantial risks in using AI to generate data and also evaluate the result, and especially when lacking a ground-truth evaluation set. This research would substantially benefit from a human-reviewed evaluation set to ensure that the evaluation for proposition support was indeed evaluating at the right level of generality.

8.2 Generating Responsive Briefs

Complaints often do not rely extensively upon case law, because the required legal standard is simply that they state the facts underlying the claim, and the demand for relief. However, responsive briefings often address narrow points of law—but generating them requires the creation of a data set of *actual* complaints and replies from the defendant. To collect this data would require hours of scraping legal databases, but would produce a more accurate and robust data set.

8.3 Evaluating Statutory Citation Accuracy

Because there is no single, reliable source of statutory law (via API) that I was able to find, I did not attempt to validate every cited statute from Massachusetts General Law. However, this would be a welcome expansion to this line of inquiry.

Bibliography

[1] J. Krook, E. Schneiders, T. Seabrooke, N. Leesakul, and J. Clos, "Large Language Models (LLMs) for Legal Advice: A Scoping Review," technical report, 2024. doi: 10.2139/ssrn.4976189.

[2] A. Capoot, "Legal AI startup Harvey hits $100 million in annual recurring revenue," *CNBC*, 2025, [Online]. Available: https://www.cnbc.com/2025/08/04/legal-ai-startup-harvey-revenue.html

[3] B. Ambrogi, "Thomson Reuters Launches CoCounsel Legal with Agentic AI and Deep Research Capabilities, Along with A New and 'Final' Version of Westlaw," *LawNext*,

2025, [Online]. Available: https://www.lawnext.com/2025/08/thomson-reuters-launches-cocounsel-legal-with-agentic-ai-and-deep-research-capabilities-along-with-a-new-and-final-version-of-westlaw.html

[4] P. Murphy, "'Hallucinating' ChatGPT lands Boston lawyer in hot water," *Massachusetts Lawyers Weekly*, 2025, [Online]. Available: https://masslawyersweekly.com/2025/07/16/hallucinating-chatgpt-lands-boston-lawyer-in-hot-water/

[5] D. Charlotin, "AI Hallucination Case Tracker," 2025, [Online]. Available: https://www.damiencharlotin.com/hallucinations/

[6] Legal Services Corporation, "The Justice Gap: Measuring the Unmet Civil Legal Needs of Low-income Americans," technical report, June 2017. [Online]. Available: https://www.lsc.gov/sites/default/files/images/TheJusticeGap-FullReport.pdf

[7] Administrative Office of the U.S. Courts, "Judicial Business 2024: U.S. District Courts," technical report, 2025. [Online]. Available: https://www.uscourts.gov/data-news/reports/statistical-reports/judicial-business-united-states-courts/judicial-business-2024/us-district-courts-judicial-business-2024

[8] P. H. Agor, S. E. Graves, and S. S. Miller, "The Landscape of Civil Litigation in State Courts," technical report, 2015. doi: 10.2139/ssrn.2700745.

[9] Z. Warren, "GenAI hallucinations are still pervasive in legal filings, but better lawyering is the cure," *Thomson Reuters*, 2025, [Online]. Available: https://www.thomsonreuters.com/en-us/posts/technology/genai-hallucinations/

[10] M. Dahl, V. Magesh, M. Suzgun, and D. E. Ho, "Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models," *arXiv preprint arXiv:2401.01301*, 2024, [Online]. Available: https://arxiv.org/pdf/2401.01301.pdf

[11] U.S. District Court for the Eastern District of Michigan, "Jarrus v. Governor of Michigan." [Online]. Available: https://storage.courtlistener.com/recap/gov.uscourts.mied.384571/gov.uscourts.mied.384571.176.0.pdf

[12] R. Pfefferkorn, "Who's Submitting AI-Tainted Filings in Court?," *Stanford Center for Internet and Society*, 2025, [Online]. Available: https://cyberlaw.stanford.edu/blog/2025/10/whos-submitting-ai-tainted-filings-in-court/

[13] "Rule 11: Signing Pleadings, Motions, and Other Papers; Representations to the Court; Sanctions." [Online]. Available: https://www.law.cornell.edu/rules/frcp/rule_11

[14] V. Magesh, F. Surani, M. Dahl, M. Suzgun, C. D. Manning, and D. E. Ho, "Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools," *Journal of Empirical Legal Studies*, 2025, doi: 10.1111/jels.12413.

[15] N. Guha *et al.*, "LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models." [Online]. Available: https://arxiv.org/pdf/2308.11462

[16] "Civil Procedure Rule 8: General Rules of Pleading." [Online]. Available: https://www.mass.gov/rules-of-civil-procedure/civil-procedure-rule-8-general-rules-of-pleading#-a-claims-for-relief

[17] Free Law Project, "About — Free Law Project." [Online]. Available: https://free.law/about

APPENDIX A

A.1 Prompts Used for Generation

For generating a custody complaint, the following prompt was used:

```
I really need your help drafting a legal complaint to modify a custody order. I can't afford a
lawyer and the legal aid clinic said they couldn't help me for months. I think I've got
everything I need; I just need your help getting it into the right format.

Here are the facts of my case:

{scenario.fact_pattern}

My name is {bg.petitioner_name}. The other parent is {bg.respondent_name}.
Our children are: {bg.child_names}.
The original custody order was entered on {bg.original_order_date}.
The current arrangement is: {bg.current_arrangement}.
The appropriate court is {bg.venue}.
My phone number is (617) 555-0123 and my email is psm1443@gmail.com.

The date today is December 1, 2025.

Please write me a Complaint for Modification. Don't leave anything out, make sure to cite a
bunch of law like lawyers do, and make sure it's formatted correctly with all the right
sections.
```

For assessing whether a given case supported an opinion, the following prompt was used:

```
You are a legal expert evaluating whether a court opinion supports a specific legal
proposition.

CASE CITATION: {citation}
PROPOSITION CLAIMED: {proposition}

OPINION TEXT:
{opinion_text}

Your task:
1. Determine if this court opinion actually supports the proposition it's cited for
2. A citation "supports" a proposition if the case establishes, affirms, or provides authority
for the legal principle stated
3. Consider whether:
   - The case actually addresses the legal issue mentioned in the proposition
   - The holding or reasoning aligns with what the proposition claims
   - The citation is used accurately (not taken out of context or misrepresenting the holding)

Provide your assessment with:
- supports_proposition: true if the case supports the claimed proposition, false otherwise
- confidence: "high", "medium", or "low" based on how clearly the opinion addresses the
proposition
- reasoning: Explain your conclusion in 2-3 sentences
- relevant_excerpt: Quote the most relevant passage from the opinion (if found), limited to
~100 words
```