

HOW ARTIFICIAL INTELLIGENCE IS RESHAPING MUNICIPAL LAW

Mitesh Bhakta

Matthew C. Slentz

Meghan A. Wharton

Introduction

You've surely been inundated with "AI" everywhere you turn. Uninvited chatbots consuming your screen when online shopping, a persistently recursive customer service representative when trying to reach an airline, your phone summarizing text messages. The pervasiveness of what we commonly refer to as "AI" has turned it into a semantically overloaded term, often used more for marketing than to describe actual technological capacity. Understanding the difference between true AI and traditional algorithms is critical to your role as a city attorney.

In this paper we will provide an overview of the ways AI is legally distinct from traditional algorithmic software, challenges you might face as a city attorney as AI becomes ubiquitous, an overview of AI tools that can assist you in the practice of law, and finally ethical pitfalls to avoid when using AI.

I. What Is Artificial Intelligence?

AI is software that can "learn" from data rather than follow human-written rules. A traditional algorithm is a set of explicit instructions that will always produce the same output from the same input. This doesn't mean algorithms are simple; the software routing your GPS, detecting credit card fraud, or compressing a video file is algorithmic and extraordinarily sophisticated. The key distinction isn't complexity or capability, it's whether the system follows rules a human wrote, or derives its own rules from examples. An algorithm can be thought of as a recipe for Black Forest cake; when followed correctly and with the same ingredients, the cake will turn out the same every time. In contrast, AI can be thought of as a chef who learns how to make cakes by studying thousands of them, not by being told the rules, but by inferring them.

AI systems can perform at or above human levels on average while failing in specific, unpredictable situations. High benchmark performance does not mean reliable performance in every case. The situations where AI fails are often precisely the edge cases, unusual fact patterns or novel circumstances, where legal liability tends to concentrate.

After initial training, most modern AI systems undergo an additional process in which human evaluators rate the model's outputs, and the model is further adjusted to produce outputs those evaluators prefer. This process is known as reinforcement learning from human feedback, or RLHF, and is how AI systems are shaped to be helpful, to avoid certain outputs, and to reflect particular values and preferences. Attorneys should understand that this step means the dispositions of an AI system are not neutral. They reflect choices about what counts as a good output, whose preferences were consulted, and what feedback was rewarded made by the system's developers. Those choices can embed bias, replicate the blind spots of the evaluators, or optimize for preferences that diverge from a client's interests or a regulatory standard. When an AI system produces a biased or harmful output, the question of where in this pipeline the failure occurred is both technically complex and legally significant.

AI In Your City

Now that you have a basic understanding of how AI works, it will be easier to understand the impact it will have on your city. Several companies are working on generative AI capabilities to assist governments in areas such as cybersecurity, compliance, report writing, code management, legal research, agenda management, and even workforce management. The usual first instinct is to adopt an AI policy, but any such policy must consider what harm we are actually trying to prevent, and the most effective means of preventing harm, rather than adopting a policy that may be already covered by existing law, or, addresses concerns that are better mitigated through operational controls that have nothing specifically to do with AI.

AI as an Instrument, Not a Source of Liability

Consider the categories of harm most commonly cited in AI policy discussions: fraud, defamation, discrimination, privacy violations, copyright infringement, and breach of fiduciary duty. With narrow exceptions at the regulatory frontier none of these are new legal categories.¹ They were unlawful before AI existed. AI changes the scale, speed, and accessibility of the conduct, but it does not change its legal character. An employee who uses an AI tool to generate a misleading report has committed the same bad act as one who drafted it manually. A vendor who uses AI to fabricate invoices is committing

¹ Certain uses of AI in healthcare and hiring have attracted specific rulemaking. *See* AB 3030 (2024) and Cal. Code Regs. Title 2, Div 4.1, Chapter 5, Subchapter 2

fraud regardless of the tool. The instrument is new; the liability framework largely is not.

This matters for policy design because a policy focused on AI-as-technology tends to produce rules about the technology rather than rules about the conduct your city actually cares about preventing. You end up regulating the hammer rather than the nail. Instead, focus on where real exposure lives in your organization. The harms that AI genuinely amplifies tend to cluster around two pre-existing organizational vulnerabilities: data governance and process integrity.

AI systems are, at their core, data consumers and data producers. An organization with poor data hygiene (inconsistent classification, unclear retention policies, inadequate access controls, murky third-party data sharing agreements, etc.) was already exposed before AI. AI adoption accelerates and magnifies that exposure dramatically. Sensitive data fed into external AI tools without adequate contractual protections, personal data retained beyond its useful life that then trains a model, confidential business information that becomes part of a vendor's training corpus, etc., are data governance failures, not AI failures. The remedy is a data governance program, not an AI policy.

So What Am I Supposed to Do?

Rather than drafting bespoke AI policies, your city can reduce risk more effectively by auditing existing data management practices against the actual ways AI tools are being used in the organization with particular attention to what data employees are inputting into external tools and what contractual protections govern that data. Most commercial AI terms of service have evolved rapidly, and many organizations are operating under agreements negotiated years ago that were not drafted with current AI capabilities in mind. On the process side, the highest-leverage interventions tend to be mundane: tightening vendor onboarding verification, adding friction to out-of-pattern payment requests, requiring dual authorization for changes to banking details, and ensuring that customer-facing or counterparty-facing communications have human review checkpoints before transmission. These are not AI policies; they are good operational hygiene that happens to also mitigate AI-enabled risk.

AI-Specific Policy Language Should be Narrowly Crafted

When specific AI policy language is needed it should be narrow, tied to concrete conduct, and embedded within existing policy frameworks rather than standing alone. An acceptable use addendum to an existing technology policy, a section added to a data classification policy addressing AI tool inputs, a paragraph in a vendor management policy addressing AI-generated deliverables, etc. The goal is to close specific gaps, not to create a comprehensive regulatory scheme for a technology that is changing faster than any policy document can track.

There is also a candid institutional argument for this approach. A robust, standalone AI policy creates its own compliance obligations and its own liability surface. If the policy is detailed enough to be meaningful, it is detailed enough to be violated, and violations of internal policy can be used against your city in litigation. A policy that prohibits employees from inputting confidential information into AI tools is only as good as an organization's ability to enforce and document compliance with it. If that infrastructure doesn't exist, the policy may do more harm than good. The cities that will navigate the AI transition most effectively are likely not those with the most comprehensive AI policies, but those with the strongest data and operational discipline, those who treat AI as one more reason to do the basics well, rather than as a novel problem requiring a novel legal framework built from scratch.

Finally, keep in mind that part of AI's impact on your city will be how it becomes a target for data harvesting, whether through contracts or other avenues. Training AI models requires an enormous amount of data. The early internet was fantastic for this data. Humans posted on social media, documents from governments, corporations, NGOs, etc., were gradually put online, and thus the internet was mostly a reflection of society. Beginning in the late 2010s, however, social media began to be taken over by bots of various kinds. These early bots were simple, they could regurgitate certain snippets or summarize other real written work, but soon they were able to engage in more complex activity and a trickle became a monsoon as bots started to generate content based solely on what other bots were doing, such that most of the data traffic now flowing through the internet is not generated by humans.² This poses a major

² "The Internet is Mostly Bots" <https://www.theatlantic.com/technology/archive/2017/01/bots-bots-bots/515043/>

problem if you are training an AI model, which needs lots of data, and not just any data, but good quality data. So much data that most AI models are now running out of good quality data and are increasingly facing a positive feedback loop that can spiral into AI failure.³

You should expect as a public agency attorney that you will increasingly face more Public Records Act requests that blanket entire categories of data or document type, because government records remain one of the last few sources of data that are both good quality and by law must be available for anyone to see, replicate, and use. This makes government records the equivalent of low-background steel for the AI industry.⁴ You can manage an influx of these requests by advising your clients that if a record doesn't require redaction and is relatively easy to prepare, a transparency portal can reduce staff time for PRAs by eliminating the need to manually process repetitive requests.

II. The Darker Side of AI

While AI has a lot of promise to make certain tasks easier and faster, there are some unique issues that cities will have to face. SB 707 requires cities with populations of 30,000 or greater to offer 2-way teleconferencing as an option for public attendance at all public meetings. As AI deepfakes, voice replication, and visual distortion have become better, it is very likely that your city will eventually encounter artificial public speakers in the near future. Consistent application of the Brown Act at all public meetings will be your best defense as a city attorney until vendors are able to find a technological solution.⁵

Large language models have also dramatically lowered the cost and skill threshold for generating facially legitimate-looking legal and administrative documents. A malicious

³ The positive feedback loop can be triggered by a confabulation, generating a demand for more training data, which in turn is lower quality, which results in even more confabulations.

⁴ Low-background steel was produced prior to the first atomic tests and is necessary for certain scientific purposes because post-atomic steel is contaminated with certain radioisotopes. In much the same way, written work that predates the bot-driven internet era, or is generated independent of it, is not contaminated and remains useful for training models.

⁵ Zoombombing is still likely to be your most persistent issue when at the dais, but AI reduces the cost.

actor targeting an organization's accounts payable function can now use AI chatbots to rapidly produce high volumes of demand letters, invoices, contractual claims, regulatory complaints, and dispute filings that are superficially coherent, professionally formatted, and difficult to distinguish at first glance from legitimate correspondence. What previously required legal knowledge, drafting skill, and significant time investment can now be accomplished by anyone in minutes. The economics of nuisance litigation and administrative harassment have fundamentally changed: the cost to the attacker is near zero, while the cost to the defender of processing, reviewing, responding, and escalating, remains substantial. Specific vectors include fabricated vendor invoices referencing real contract details scraped from public sources, AI-drafted demand letters asserting breach of contract or legal violations with plausible but false factual predicates. Deepfakes can also be used to mimic your city's personnel and direct staff to transmit funds to a threat actor.⁶

These actions violate numerous state and federal laws (such as mail or wire fraud statutes), however, enforcement remains impractical or impossible in many situations because threat actors are often located in places without an extradition treaty.⁷ If the conduct involves interstate communications and intent to defraud, the Computer Fraud and Abuse Act⁸ would also apply, and if automated tools are used to interact with organizational systems, RICO can also potentially apply if there is a pattern of conduct by an organized group.⁹ The practical challenge for your city and your office is attribution. AI-generated content leaves few traditional forensic fingerprints, and the low barrier to entry means actors may operate across jurisdictions or through intermediaries, complicating both identification and service, so to the extent you can help facilitate evidence preservation, that should be your top priority.

⁶ How Scammers Used Deepfake Video to Dupe a Company Out of Millions

<https://www.mcafee.com/ai/news/how-scammers-used-deepfake-video-to-dupe-a-company-out-of-millions/>

⁷ 18 U.S.C. §§ 1341, 1343

⁸ 18 U.S.C § 1030

⁹ 18 U.S.C. §§ 1961-1968

III. Advising Your Clients

On the procedural and operational side, cities should establish a formal triage protocol for unsolicited invoices that requires cross-referencing against internal contract and vendor databases before any payment or response is initiated. Anything that cannot be matched to a known counterparty relationship, purchase order, or prior correspondence should be treated as suspect regardless of how professionally it presents. Accounts payable staff should be trained specifically on AI-generated document characteristics, not to detect them technically, but to flag anomalies like claims with no prior relationship history, unusual specificity about internal processes, or pressure tactics around response deadlines.

On the legal and evidentiary side, your role as counsel should be to help set up procedures to preserve all incoming suspect communications in their original digital form, including metadata, routing information, and headers, as these may be critical to later attribution efforts. Organizations should also consider whether their vendor agreements and payment terms include provisions requiring disputes to be submitted through specific channels, which creates a cleaner basis for dismissing out-of-channel claims without engaging their merits. If your city is facing systematic campaigns rather than isolated incidents you should consider coordinated reporting to the FBI's Internet Crime Complaint Center (IC3) and relevant sector regulators, as pattern documentation across organizations is often necessary to trigger meaningful law enforcement attention.

The fundamental challenge is that the legal system is built around the presumption that generating a facially legitimate claim requires some minimum investment of knowledge and effort, which historically served as a natural filter on frivolous filings. AI has broken that assumption. The procedural and substantive responses available today were designed for a different threat environment, and organizations and courts alike are still adapting. Proactive internal controls are, for now, a more reliable first line of defense than post-hoc legal remedies.

IV. AI Tools for City Attorneys

While generative AI presents significant risks and requires care, the AI tools available to lawyers are incredibly powerful, and will only grow more efficient and useful.¹⁰ The ABA's Task Force on Law and Artificial Intelligence noted in 2025: "[A]s the transformative power of the technology has become more widely known, the conversation has shifted from **whether** to use the AI technology to **how** to use it."¹¹ Generally, these AI tools fall into three categories. The first are publicly accessible tools such as OpenAI's ChatGPT¹² or Anthropic's Claude¹³. Free and publicly available, these tools are not trained on legal sources specifically, and are more likely to confabulate or make mistakes when asked legal questions. Further, the information placed into free versions is not private, and prompts are incorporated into the model and used to train it. Thus, while certain tasks, like drafting an email, can be done quickly and cheaply with ChatGPT, lawyers should carefully scrutinize what goes into, and comes out of, these tools.

Privacy issues can be mitigated by using a subscription version of these tools. Both ChatGPT and Claude have subscription versions that do not train their models on the data provided.¹⁴ And Microsoft's Copilot operates within Microsoft Office tools such as Word, Excel, PowerPoint and Teams.¹⁵ For example, you can draft an outline in Word, and ask Copilot to convert it to a PowerPoint presentation.¹⁶

Next are "off-the-shelf" legal tools that use a large language models (LLMs), such as Westlaw Advantage¹⁷ and Lexis+AI.¹⁸ These tools are trained on legal databases, and programmed to limit their responses to real authority. They also keep client data secure,

¹⁰ Many of the topics explored here were previously raised in Lee & Slentz, Artificial Intelligence and the Practice of Law, 2023 Cal Cities Annual Conference and Expo. This section reflects updates to the observations made in that prior article.

¹¹ ABA Task Force on Law and Artificial Intelligence, Addressing the Legal Challenges of AI, <https://www.americanbar.org/groups/centers_commissions/center-for-innovation/artificial-intelligence/> [as of Mar. 24, 2026].

¹² See Open AI, ChatGPT, <<https://chatgpt.com/>> [as of Mar. 24, 2026].

¹³ See Anthropic, Claude, <<https://claude.ai/>> [as of Mar. 24, 2026].

¹⁴ See OpenAI, Enterprise privacy at OpenAI, <<https://openai.com/enterprise-privacy/>> [as of Mar. 24, 2026]; Anthropic, Plans that grow with you, <<https://claude.ai/upgrade>> [as of Mar. 24, 2026].

¹⁵ Microsoft, What is a copilot? <<https://www.microsoft.com/en-us/microsoft-copilot/copilot-101/what-is-copilot>> [as of Mar. 24, 2026].

¹⁶ *Ibid.*

¹⁷ Thomson Reuters, Fast-track your legal research with Westlaw Advantage <<https://legal.thomsonreuters.com/en/products/westlaw-advantage>> [as of Mar. 24, 2026].

¹⁸ LexisNexis, Lexis+ AI: The Generative AI for Lawyers, <https://lexis-legal.ai/> [as of Mar. 24, 2026].

meaning the tool can be used without risking confidentiality. These tools have quickly gained wide use, with Thomson Reuter's CoCounsel, which can be used for a host of drafting and document review tasks, recently reaching over 1 million users worldwide.¹⁹ Other tools focus more directly on a single aspect of legal work. These include Spellbook, a contract drafting platform that interfaces directly with Microsoft Word,²⁰ and Everlaw, which is designed for eDiscovery.²¹ AI mediators have even begun to appear, although the efficacy of having an LLM act as a mediator waits to be seen.

Finally, there are boutique and in-house products, which go a step further and train a model based on a firm's own data. For example, Harvey AI is a legal services AI platform that "combines general internet data from the GPT model with legal-specific data, including case law and reference materials. When engaged by a law firm, Harvey AI undergoes further training using the firm's own work products and templates," which allows Harvey AI to tailor its product directly to client needs.²² While these products offer the most benefit to consumers, they also take significant resources to build or buy, and may not be justifiable for public entities and smaller firms.

The potential applications of generative AI to legal tasks are nearly endless, but it performs certain tasks better than others. These include legal research, natural language drafting, and searching large databases for particular information. Conversely, generative AI still struggles at drafting tasks that require strict fidelity to the source documents, like producing a privilege log.

V. Legal Research and Analysis

Attorneys spend significant time researching and summarizing the law, and the prospect of telling a computer what you want and getting the right answer in minutes or seconds is enticing. In practice, legal research using an LLM can be risky if the model isn't trained on and limited to legal sources. ChatGPT and other broad-based tools can

¹⁹ LawSites, Three Years After Launching As First AI Legal Assistant, CoCounsel Reaches 1 Million Users, <https://www.lawnext.com/2026/02/three-years-after-launching-as-first-ai-legal-assistant-cocounsel-reaches-1-million-users-and-thomson-reuters-teases-whats-ahead.html> [as of Mar. 24, 2026].

²⁰ Spellbook, AI tailored for contracts, <<https://www.spellbook.legal/>> [as of Mar. 24, 2026].

²¹ Everlaw, Discover the difference with Everlaw, <<https://www.everlaw.com/>> [as of Mar. 24, 2026].

²² Williams, *Introducing Harvey AI: Revolutionizing the Legal Field with Artificial Intelligence*, Martindale AVVO (May 25, 2023) <<https://www.martindale-avvo.com/blog/harvey-ai-artificial-intelligence/#:~:text=It%E2%80%99s%20built%20on%20Open.AI%E2%80%99s%20GPT%20AI%2C%20just%20like,legal-specific%20data%2C%20including%20case%20law%20and%20reference%20materials>> [as of Mar. 24, 2023].

be surprisingly accurate, but are also convincing even when wrong. AI tools offered by Westlaw and Lexis, which are trained on a legal database and only return results from that database, present a safer and more accurate option. But these tools are often, by design, less creative. Using the two types of platforms in tandem is often an effective way to conduct legal research.

For example, ChatGPT gives a (correct) answer in a few seconds when asked: “In California, are monetary exactions imposed by the legislature, such as development impact fees, subject to the essential nexus and rough proportionality test under *Nollan v. California Coastal Commission* (1987) 483 U.S. 825 and *Dolan v. City of Tigard* (1994) 512 U.S. 374?” Westlaw’s AI research will give the same answer, but with more nuance and greater citations to authority. See [Exhibit A](#) for ChatGPT’s response.

LLM’s can also apply the law to facts and help with legal analysis. Consider the following generic prompt:

“I am the City Attorney for a general law city in California. The City received a complaint from a recently discharged employee that his supervisor discriminated against him. The City commissioned a third-party investigation by an attorney that determined the supervisor violated City policy but did not engage in any discriminatory conduct. A request has now been made under the California Public Records Act. Must we disclose the report? Please make the best arguments for and against disclosure, with their likelihood of success if challenged in court.”

ChatGPT gives a detailed summary that is mostly correct, but misses a key point, that the report is likely an attorney-client privileged document prepared in anticipation of litigation and therefore exempt under the CPRA. Westlaw flags this issue. See [Exhibit B](#) for a comparison of responses from ChatGPT and Westlaw.

Documents may also be uploaded to LLMs for analysis. For example, ChatGPT can be given a complaint, told to issue spot for the best arguments to make at the pleading stage, and will produce a decent guideline for drafting a demurrer. Taking this step as you proceed from the research stage to the drafting stage may help identify arguments or issues missed in your initial research.

Overall, public generative AI tools such as ChatGPT are becoming increasingly useful in conducting legal research, but still tend to state their conclusions with more confidence than is merited. Conversely, tools designed for lawyers are more guarded in

their responses, but can still significantly shorten the time to complete research tasks. All such tools should be treated as a prolific first year associate: prone to errors and misunderstandings, but able to sift large amounts of data and provide concise summaries that will quickly point attorneys in the right direction.

VI. Brainstorming

LLMs are very good at “brainstorming” type activities, where the tool is prompted with information and produces a list of possible responses (generative AI makes a great thesaurus). These abilities can be applied to a variety of legal tasks, such as deposition preparation, claim and defense identification, and drafting interrogatories. For example, we gave Claude AI the following prompt:

I am the City Attorney for the City of Pawnee. Several years ago, the City obtained a large open pit in a tax seizure that had been dug out by a condominium developer to host a basement before going bankrupt. Recently, local resident A. Dwyer (who lives next to the pit) fell in and broke both his legs. The incident happened at night, and we suspect he was under the influence of alcohol.

A. Dwyer has sued the City for negligence, premises liability, and infliction of emotional distress. Please prepare deposition questions for A. Dwyer related to his claims.²³

Given this information, Claude produces 6 pages of deposition topics and questions, including what is shown in Exhibit C.

Similarly, when given a follow up prompt to “draft interrogatories related specifically to A. Dwyer’s knowledge of the pit’s dangerousness,” Claude produces a list of 25 interrogatories that can be used as a guide. As with most generative AI output, simply copying and pasting the results into a pleading would be a mistake. But these tools can help formulate questions, brainstorm different categories of questioning, and speed up the process of drafting discovery.

²³ See generally *Parks and Recreation*: Pilot (NBC Universal broadcast, April 9, 2009).

VII. Document Review

Generative AI tools are also excellent for sifting large amounts of information, which makes them ideal for assisting with discovery tasks. For example, a tool can be given a document and asked to summarize its contents. As with all AI prompts, the more information given to the tool, the more accurate and helpful the output.

More impressive is generative AI's ability to sort through large databases for information. During discovery, large databases of information can be uploaded, and then searched rapidly to answer simple questions. Cities can also build a document bank, making it possible to rapidly search their own precedent for answers.

VIII. Drafting and Revision

Generative AI is also good at producing first drafts of documents. Here, ChatGPT was asked to “draft a California ordinance prohibiting short term rentals in the city” and produced a draft set forth in Exhibit D.

While not perfect, the ordinance is decent starting place for drafting, and would save the user significant time. The user can then take that base and modify as necessary, or continue to prompt the model for revisions. Legal specific tools also offer a variety of contract drafting options, and will create, review, and revise contracts based on certain parameters.

IX. Pitfalls to Avoid When Using AI

Bixler v. Church of Scientology²⁴

Boies Schiller represents plaintiffs in a lawsuit against the Church of Scientology and actor Danny Masterson. Plaintiffs contend they were subjected to harassment, surveillance, threats, and defamation after they came forward with sexual assault allegations against church member Masterson. Masterson was convicted of rape in 2023. The trial court denied the church's anti-SLAPP motion, and the church appealed. Boies Schiller filed an opposition brief full of AI errors, and the church then highlighted the AI errors in a table attached to the reply detailing 17 instances of AI errors.²⁵ The table identified the following types of errors:

²⁴ *Bixler v. Church of Scientology* (2nd Dist., B339009).

²⁵ *Bixler* (reply brief, Sept. 8, 2025).

citations to authorities that do not support the proposition;
citations with wrong reporter number or page numbers;
citations to authorities that do not exist; and
citations using incorrect case names.

Bixler is arguably the biggest AI failure in California courts. It involves a powerhouse national litigation firm working on a high-profile case. After the church identified the AI errors in the reply brief, Boies Schiller partner John J. Kucera submitted a declaration with a motion to file a corrected response brief admitting to AI-generated citation errors in the opposition brief.²⁶ Kucera accepted all responsibility and apologized for the errors, which he claimed occurred despite firm-wide AI usage policies.²⁷

On February 19, 2026, the Court of Appeal issued an order to show cause on sanctions alleging the law firm submitted a brief with “at least ten purported case authorities that have been miscited or mischaracterized in more than fifteen places.”²⁸ The Court is considering sanctions under California Rules of Court, Rule 8.276, subdivision (a)(4), Rule 8.204, subdivision (a)(1)(B), and Rule 8.1115, subdivision (a).²⁹

X. AI Technology is not Infallible

At its core, an AI confabulation (hallucination) is a plausible looking but factually incorrect output or answer generated by an AI model. Importantly, AI errors are not technical glitches in an AI model, but result from a systemic vulnerability in how AI models are designed, built, and used. AI errors happen because AI models are pattern-matchers, not factual databases.

LLMs, the most commonly deployed consumer AI models, work by predicting the most likely next word. AI processes prioritize language fluency and giving the perceived desired answer over giving the correct answer. AI models can be completely false. When the AI model does not have enough data, or when the data is muddled, uncertain, or conflict, the answer is often incorrect.

As demonstrated by the *Bixler* case, using an AI model to draft legal analysis for briefs submitted to the court is risky. The AI model can give you a completely believable discussion of the elements of a cause of action and then cite to a case that does not exist.

²⁶ *Bixler* (motion for leave to file corrected brief, Sept. 19, 2025).

²⁷ *Id.*

²⁸ *Bixler* (order to show cause, Feb. 19, 2026).

²⁹ *Id.*

Attorneys using AI models for this purpose must check the output to make sure that the AI did not commit errors that will be embarrassing if first found by the Court or opposing counsel.

Courts Punish Attorneys Who Include AI Errors in Briefs

Sanctions

California courts rely on several grounds to impose sanctions for AI errors. The leading published case in California, *Noland v. Land of the Free*, imposed sanctions for a violation of California Rule of Court, Rule 8.204, subdivision (a)(1)(B) requiring each brief to “[s]tate each point under a separate heading or subheading summarizing the point, and support each point by argument and, if possible, by citation of authority.”³⁰ The *Nolan* court also relied on Code of Civil Procedure section 907 which allows sanctions for a frivolous appeal.³¹ *People v. Alvarez* found a violation of California Rule of Court, Rule 8.1115, subdivision (a) barring citation to unpublished opinions because the attorney cited non-existent cases.³² As authority for sanctions, *People v. Alvarez* imposed sanctions under Code of Civil Procedure section 128.7, subdivisions (b)(2), (c) and (h)³³ after deeming the attorney’s behavior unprofessional.³⁴ *Noland* and *Alvarez* also relied on California Rule of Court, Rule 8.276, subdivision (a)(4) which allows sanctions for committing an unreasonable violation of the Rules of Court.³⁵

Other Courts of Appeal have looked to the rules of professional conduct.³⁶ Rule of Professional Conduct, Rule 3.3, subdivision (a)(2) imposes a duty to not “knowingly misquote to a tribunal the language of a book, statute, decision or other authority.” Similarly, Courts have found that AI errors cause an attorney to violate Business and professions Code section 6068, subdivision (d) which imposes a duty “never to seek to mislead the judge or any judicial officer by an artifice or false statement of fact or law.”³⁷

³⁰ *Noland v. Land of the Free*, L.P. (2025) 114 Cal.App.5th 426, 447; see also, *Schlichter v. Kennedy* (2025) 116 Cal.App.5th 24, 33.

³¹ *Noland*, *supra*, 114 Cal.App.5th at pp. 442, 447.

³² *People v. Alvarez* (2025) 114 Cal App.5th 1115, 1119 (“*Alvarez*”).

³³ Code of Civil Procedure section 128.7, subdivisions (b)(2) bars frivolous argument; Subdivision (c) authorizes sanctions upon attorneys; Subdivision (h) allows sanctions for unnecessary delay, actions that needlessly increase the cost of litigation and to deter improper conduct by attorneys.

³⁴ *Alvarez*, *supra*, 114 Cal App.5th at p. 1120.

³⁵ *Noland*, *supra*, 114 Cal.App.5th at p. 442; *Alvarez*, *supra*, 114 Cal App.5th at p. 1119; *Shayan v. Shakib* (2025) 116 Cal App.5th 619, 624 [finding “unreasonable” violation of the rules of court].

³⁶ *Shayan v. Shakib* (2025) 116 Cal App.5th 619, 624.

³⁷ *Id.*

Federal trial courts have imposed sanctions for AI errors as violating Federal Rule of Civil Procedure, Rule 11.³⁸ Rule 11 requires attorneys to sign filings certifying that after reasonable inquiry, the filings are not frivolous, intended to harass, or designed to cause unnecessary delay. Other federal courts have relied on 28 U.S.C. section 1927 which states: “Any attorney or other person admitted to conduct cases in any court of the United States or any Territory thereof who so multiplies the proceedings in any case unreasonably and vexatiously may be required by the court to satisfy personally the excess costs, expenses, and attorneys’ fees reasonably incurred because of such conduct.”³⁹ Federal appellate courts rely on Federal Rule of Appellate Procedure 46, subdivision (c) and the court’s inherent authority.⁴⁰

Types of Sanctions, Punishments, and Penalties

Lesser Penalties: Some courts have declined severe penalties in favor of general admonishment.⁴¹ For example, in *Shelton v. Parkland Health*, the court determined judicial admonition was sufficient because filing the defective brief was sanctionable but not malicious.⁴² Such admonishments usually carry a warning that future submissions with AI errors will be sanctioned.⁴³

Monetary Sanctions: Courts impose monetary sanctions to penalize attorneys who have AI errors in filings. Courts generally link the monetary sanction to the fees the other side incurred responding to the defective filing.⁴⁴ Courts have also ordered sanctions

³⁸ *ByoPlanet International, LLC v. Johansson* (S.D.Fla., July 17, 2025) 792 F. Supp. 3d 1341, 1351.

³⁹ *Lewis v. Eagle County Government* (D.Colo., Nov. 4, 2025, No. 25-cv-02269-NYW-NRN) 2025 WL 3639608, at p. *8–*9.

⁴⁰ *Fletcher v. Experian Information Solutions, Inc.* (5th Cir., Feb. 18, 2026, No. 25-20086) ___ F.4th ___, 2026 WL 456842, at p. *6.

⁴¹ *Mareiorozyan v. Automobili Lamborghini America, LLC* (Los Angeles Super Ct., Nov. 19, 2025, No. 23VECV04028) 2025 WL 3669607, at p. *6 [denying motion to disqualify, monetary sanctions, and reporting to the State Bar but “strongly admonishes Plaintiff’s counsel”]; *Yamin v. American States Preferred Ins. Co.* (Los Angeles Super. Ct., Oct. 16, 2025, No. 24SMCV03322) 2025 WL 3122627, at p. *2.

⁴² *Shelton v. Parkland Health* (N.D.Tex., Nov. 10, 2025, No. 3:24-CV-2190-L-BW) 2025 WL 3141108, at p. *5.

⁴³ *Park v. Lee* (Los Angeles Super. Ct., Oct. 31, 2025, No. 22STCV03583) 2025 WL 3246457, at p. *4 [warning that court will assess sanctions on its own motion for future AI errors].

⁴⁴ *Howell Management Services, LLC v. Vashisht-Rota* (Dec. 15, 2025, D086055) 2025 WL 3628442, at p. *10 [nonpub. opn.] [awarding \$59,235 in attorney’s fees sanctions]; *Puerto Rico Soccer League NFP, Corp. v. Federación Puertorriqueña De Fútbol* (D.P.R., Sept. 23, 2025, No. CV 23-1203 (RAM)) 2025 WL 2709484, at p. *8 [awarding \$24,492.10 in attorney’s fees based on very low prevailing hourly rates in Puerto Rico]; *ByoPlanet, supra*, 792 F. Supp. 3d at p. 1357 [attorney ordered to “pay the attorneys’ fees for Defendants’ counsel in these four cases for all time spent responding to any filing in which generative AI was used to develop hallucinated cases and fabricated quotations]; But see, *Noland, supra*, 114 Cal.App.5th at p. 448 [declining to award attorney’s fees to opposing counsel because court found the AI errors].

paid to the court for AI errors found by the court.⁴⁵

Rejected Filings: Courts may reject or strike the offending filing and deny the offending attorney the ability to correct the filing.⁴⁶ At least one Court noted that an errata cannot be used to correct AI hallucinations in filings.⁴⁷

Disqualification: By revoking pro hac vice admissions, federal courts have effectively disqualified attorneys who submitted filings with AI errors.⁴⁸

Dismissal: In one of the most egregious cases, *ByoPlanet*, the Florida district court hearing four cases determined an attorney acted in bad faith by signing the pleadings and ordered dismissal under the Rule 11 and the court's inherent powers.⁴⁹ After the attorney repeatedly submitted filings with AI errors in citations, the court ordered the cases dismissed without prejudice and without leave to amend and directed the clerk to assign any refiled cases to the same department.⁵⁰ Other courts have dismissed cases for similar violations.⁵¹

Bar Notification: State and federal courts refer attorneys who submit briefs with AI errors to the State Bar.⁵²

Other Penalties: Angry courts imposed unique sanctions in certain egregious cases. In *ByoPlanet*, the Court ordered the offending lawyer to attach a copy of the sanctions

⁴⁵ *Schlichter*, *supra*, 116 Cal.App.5th at p. 33 [ordering \$1,750 sanction payable by the attorney to the court]; *Alvarez*, *supra*, 114 Cal App.5th at p. 1120 [ordering \$1,500 sanction to "reimburse the court for a small portion of the time and resources expended on this issue"]; *Tercero v. Sacramento Logistics, LLC* (E.D.Cal., Sept. 9, 2025, No. 2:24-cv-00953-DC-JDP) 2025 WL 2605020, at p. *13 [ordering \$1,500 paid to the court].

⁴⁶ *L.A. Housing Outreach, LLC v. Medoff* (2025) (Dec. 17, 2025, B341160) [nonpub. opn.] [granting motion to strike reply brief and awarding monetary sanctions]; *Mavy v. Commissioner of Social Security Administration* (D.Ariz. Aug. 14, 2025, No. CV-25-00689-PHX-KML (ASB)) 2025 WL 2355222, at p. *13.

⁴⁷ *Yamin*, *supra*, 2025 WL 3122627, at p. *2 ["it is not the proper office of an 'errata' to change the substance of the prior brief [by] chang[ing] the legal authority entirely."]. But see *Shayan*, *supra*, 116 Cal App.5th at p. 624 [striking filing but allowing attorney to refile corrected brief].

⁴⁸ *In re Jackson Hospital & Clinic, Inc.* (Bankr. M.D.Ala., Nov. 20, 2025, No. 25-30256-CLH) 2025 WL 3251167 at pp. *17-*18; *Mavy*, *supra*, 2025 WL 2355222, at p. *12 ("*Jackson Hospital*").

⁴⁹ *ByoPlanet*, *supra*, 792 F. Supp. 3d at pp. 1354, 1357. But see *Shayan*, *supra*, 116 Cal App.5th at p. 625 [declining to dismiss appeal because prejudice can be remedied and future misconduct deterred by other sanctions].

⁵⁰ *ByoPlanet*, *supra*, 792 F. Supp. 3d at pp. 1354, 1357.

⁵¹ *Dubinin v. Papazian* (S.D.Fla., Nov. 21, 2025, No. 25-CV-23877-RAR) 2025 WL 3248187, at p. *3.

⁵² *Tercero*, *supra*, 2025 WL 2605020, at p. *13; *Shayan*, *supra*, 116 Cal App.5th at p. 624 [ordering service of sanction order on State Bar]; *ByoPlanet*, *supra*, 792 F. Supp. 3d at p. 1357; *Mavy*, *supra*, 2025 WL 2355222, at p. *13.

order with any case he files in the Southern District of Florida within two years of the date of the order.⁵³ In *In re Jackson Hospital*, the Bankruptcy Court imposed numerous sanctions including a public reprimand, an order to provide a copy of the order to all clients, and directing the attorney's firm to provide a copy of the order to every attorney with an updated AI policy.⁵⁴ Courts have ordered offending lawyers to provide a copy of the order to judges in current and future cases.⁵⁵

Ethical Issues

Relying on AI for legal research without checking whether the results are accurate or the authorities it produces exist is an abdication of professional responsibility.

Duty of Candor: Every attorney owes a duty of candor to the courts. That means the attorney must be forthright to the judge, never knowingly deceive or withhold negative facts or contrary legal authority. If a client lies on the witness stand and the attorney knows it to be false testimony, the attorney has a duty to disclose that to the judge. If the attorney knows a document is falsified, the attorney cannot submit it as genuine. Attorneys must tell judges when there isn't any case law on point, or when the law a lawyer argues for is contradicted by existing authority. In contrast, AI has no duty of candor to anyone.

If attorneys first make the mistake of relying on AI for legal research and then fail to confirm the accuracy of its results before including it in a court filing, that's a violation of the duty of candor. When an attorney signs a court filing, it carries representations that are imposed by law.

California Rule of Professional Conduct, Rule 3.3 - Candor Toward the Tribunal: "(a) A lawyer shall not knowingly: (1) make a false statement of fact or law to a tribunal or fail to correct a false statement of material fact or law previously made to the tribunal by the lawyer; (2) fail to disclose to the tribunal legal authority in the controlling jurisdiction known to the lawyer to be directly adverse to the position of the client and not disclosed by opposing counsel or knowingly misquote to a tribunal the language of a book, statute, decision or other authority."

Federal Rule of Civil Procedure 11: "(b) Representations to the Court. By presenting to the court a pleading, written motion, or other paper—

⁵³ *ByoPlanet*, supra, 792 F. Supp. 3d at p. 1357.

⁵⁴ *Jackson Hospital*, supra, 2025 WL 3251167 at pp. *17-*18;

⁵⁵ *DOC App, Inc. v. Leafwell, Inc.* (M.S.Fl., Nov. 26, 2025, No. 2:25-cv-838-SPC-NPM) 2025 WL 3288251, at p. *6.

whether by signing, filing, submitting, or later advocating it—an attorney or unrepresented party certifies that to the best of the person's knowledge, information, and belief, formed after an inquiry reasonable under the circumstances: . . . (2) the claims, defenses, and other legal contentions are warranted by existing law or by a nonfrivolous argument for extending, modifying, or reversing existing law or for establishing new law”

No attorney can honestly file something with a state or federal court that contains unconfirmed AI citations.

Other rules of professional conduct are triggered by AI use. California Rule of Professional Conduct, Rule 3.1 states that “a lawyer shall not” “assert a position in litigation ... without probable cause” or “present a claim or defense in litigation that is not warranted under existing law.” California Rule of Professional Conduct, Rule 8.4(c) defines professional misconduct as “conduct involving dishonesty, fraud, deceit, or reckless or intentional misrepresentation.”

In addition to impacts AI errors may have in litigation, AI raises broader issues for practicing attorneys. Legal research shouldn’t just be about finding a case that supports your position and moving on. A competent attorney uses the research process to pick up on all the other related law that makes the stated position better or worse. Relying on AI to simply give you case cites skips the full research process. You may get accurate results and an accurate case cite, but you may miss other important aspects of the issue. For example, AI may find a case that says a party is required to exhaust administrative remedies before filing suit, but you completely missed out on the exceptions that might apply to your case. This does not serve the client’s interests.

XI. Conclusion

AI is already having a lasting impact on the practice of law. It accelerates and simplifies basic legal work, freeing attorneys to pursue more complex tasks. However, as with any technology, attorneys must be mindful of their ethical duties, including how to competently and safely use large language models, how to advise clients on their use of this new technology, and to prepare their organizational clients for a rapidly evolving world.

EXHIBIT A

CHATGPT RESPONSE CASE STUDY– MONETARY EXACTIONS

Inquiry: “In California, are monetary exactions imposed by the legislature, such as development impact fees, subject to the essential nexus and rough proportionality test under *Nollan v. California Coastal Commission* (1987) 483 U.S. 825 and *Dolan v. City of Tigard* (1994) 512 U.S. 374?”

Short answer:

Yes—after 2024, monetary exactions imposed by legislation (such as development impact fees) are subject to the “essential nexus” and “rough proportionality” requirements of *Nollan v. California Coastal Commission* (1987) and *Dolan v. City of Tigard* (1994) under the Takings Clause. This was clarified by the U.S. Supreme Court in *Sheetz v. County of El Dorado* (2024).

However, the doctrinal details of how those tests apply to generally applicable legislative fee schedules remain unsettled.

1. Pre-2024 California Rule (Legislative Fees Not Subject to *Nollan/Dolan*)

Before *Sheetz*, California courts distinguished between ad hoc administrative exactions and legislatively imposed fees.

The California Supreme Court’s decision in *San Remo Hotel L.P. v. City & County of San Francisco* (2002) held:

Nollan/Dolan applies to discretionary, individualized permit conditions.

Legislatively imposed development impact fees are instead reviewed under a deferential “reasonable relationship” test, often tied to the Mitigation Fee Act (Gov. Code § 66001).

The rationale was that legislative fees apply broadly and involve no individualized discretion, reducing the risk of governmental leveraging identified in *Nollan* and *Dolan*.

This distinction was followed by California appellate courts and reaffirmed in the Court of Appeal decision in *Sheetz v. County of El Dorado*, which held that generally applicable traffic impact fees were exempt from the *Nollan/Dolan* tests.

2. The U.S. Supreme Court’s Decision in *Sheetz* (2024)

The U.S. Supreme Court unanimously rejected the legislative vs. administrative distinction.

In *Sheetz v. County of El Dorado* (2024) the Court held:

The Takings Clause does not distinguish between legislative and administrative permit conditions.

Therefore permit exactions imposed by legislation can still violate the unconstitutional conditions doctrine.

Accordingly:

Legislatively imposed impact fees are not categorically exempt from the *Nollan/Dolan* essential nexus and rough proportionality tests.

3. What the Court Did Not Decide

The Court deliberately left several key issues unresolved, including:

Whether the specific El Dorado County fee satisfied *Nollan/Dolan*.

The Court remanded the case for further proceedings.

How the “rough proportionality” test applies to generally applicable fee schedules (as opposed to individualized conditions).

Whether individualized determinations are required for legislatively imposed exactions affecting a class of developments.

Some concurring opinions suggested that general legislative schemes may satisfy *Nollan/Dolan* through broader studies rather than parcel-specific findings, but the Court did not adopt a definitive standard.

4. Current Doctrinal Framework (Post-Sheetz)

After *Sheetz*, the operative framework is:

Exaction doctrine applies to monetary exactions

(recognized in *Koontz v. St. Johns River Water Management District* (2013)).

It applies regardless of whether the condition is imposed by administrators or by legislation

(holding of *Sheetz v. County of El Dorado* (2024)).

The government must ultimately demonstrate:

Essential nexus between the exaction and a legitimate governmental interest (*Nollan*).

Rough proportionality between the exaction and the impacts of the development (*Dolan*).

But the precise evidentiary standard for legislative fee programs is still developing in lower courts.

✓ **Bottom line:**

Historically in California: legislative impact fees were not subject to
Nollan/Dolan.

Now (post-2024): they are, because the Supreme Court has held that
legislatures are not exempt from the unconstitutional conditions
doctrine in land-use permitting.

EXHIBIT B
CHATGPT RESPONSE CASE STUDY WITH WESTLAW RESPONSE
EMPLOYMENT DISCRIMINATION MATTER

Inquiry: “I am the City Attorney for a general law city in California. The City received a complaint from a recently discharged employee that his supervisor discriminated against him. The City commissioned a third-party investigation by an attorney that determined the supervisor violated City policy but did not engage in any discriminatory conduct. A request has now been made under the California Public Records Act. Must we disclose the report? Please make the best arguments for and against disclosure, with their likelihood of success if challenged in court.”

ChatGPT Response:

1. Default rule under the CPRA

The CPRA requires disclosure of public records unless a statutory exemption applies. The principal exemptions relevant here are:

Government Code §7922.000 (public interest balancing test)

Government Code §7927.700

Investigative reports concerning employee misconduct fall squarely within the category courts have repeatedly analyzed.

2. Strong argument FOR disclosure

A. Courts require disclosure of substantiated misconduct

The most important precedent is:

Bakersfield City School District v. Superior Court

The California Supreme Court held that records of investigations into substantiated employee misconduct must generally be disclosed, subject to redactions.

The Court emphasized:

The public has a significant interest in knowing how government agencies investigate and address misconduct by public employees.

Disclosure promotes accountability and confidence in government operations.

Here, the report concluded that the supervisor violated City policy. That finding likely qualifies as substantiated misconduct, which strongly favors disclosure under *Bakersfield*.

Likelihood of success of disclosure argument

High (≈70–85%) that at least some version of the report must be produced, because:

Courts frequently order disclosure with redactions when misconduct is substantiated.

The misconduct finding strengthens the public interest side of the balancing test.

B. The public interest in agency investigations

Courts emphasize that disclosure allows the public to evaluate:

whether the agency adequately investigated misconduct, and
whether it appropriately disciplined the employee.

This reasoning was reinforced in:

City of San Jose v. Superior Court

The court held that investigative records about misconduct by public employees are often disclosable once the investigation is complete.

3. Strong arguments AGAINST disclosure

A. Personnel file exemption and privacy interests

The City can rely on Government Code §7927.700, which protects personnel and similar files when disclosure would constitute an unwarranted invasion of personal privacy.

Relevant privacy interests include:

the supervisor's personnel privacy, and
the complainant employee's privacy.

Courts often permit redaction of:

witness names
personal identifying information
sensitive factual details

The privacy interest may be stronger if:

the misconduct was minor, or
disclosure would cause significant reputational harm disproportionate to
the public interest.

Likelihood of success

Moderate to low (≈25–40%) for withholding the report entirely.

Courts rarely allow complete nondisclosure when misconduct is substantiated.

B. The investigation rejected the discrimination allegation

A stronger argument is that the investigation exonerated the supervisor of discrimination, which was the primary allegation.

The City could argue:

The complaint concerned serious allegations of discrimination.

The investigation found no discriminatory conduct.

Disclosure could unfairly stigmatize the supervisor despite the finding.

Courts sometimes weigh reputational harm where allegations are unsubstantiated.

However, because the report found some policy violation, courts may still treat the matter as substantiated misconduct.

Likelihood of success

Moderate ($\approx 35\text{--}45\%$) for narrowing disclosure or requiring significant redactions.

C. Catchall balancing exemption

Under Government Code §7922.000, the City can argue:

The public interest in nondisclosure clearly outweighs disclosure because:

the report concerns personnel discipline,

disclosure could discourage candid participation in workplace

investigations, and

the allegations of discrimination were unsubstantiated.

However, courts generally hold that the public interest in government accountability outweighs these concerns once misconduct is found.

Likelihood of success

Low to moderate ($\approx 30\text{--}40\%$).

4. Likely judicial outcome

If litigated, the most likely result is:

Disclosure of the report with redactions.

Typical court-ordered redactions include:

complainant identity

witness names

personal contact information

sensitive personal details unrelated to misconduct

Courts often require agencies to produce:

the factual findings, and

the conclusion regarding policy violations.

5. Practical litigation risk assessment

Outcome

Full nondisclosure

Estimated likelihood

$\sim 20\text{--}30\%$

Outcome	Estimated likelihood
Disclosure with significant redactions	~50–65%
Disclosure with minimal redactions	~10–20%

The most likely outcome is partial disclosure.

6. Practical approach used by many cities

Many California public agencies handle this situation by:

Producing the investigative report, but redacting:

Complainant identity

Witness identities

Personal details

Sometimes also redacting legal analysis or attorney advice portions if present.

Bottom line:

Because the investigation substantiated a violation of City policy, courts are likely to find a significant public interest in disclosure under the CPRA. The City's strongest position is not full nondisclosure but disclosure with targeted redactions protecting privacy interests.

This answer is mostly correct, but misses a key point, that the report is likely an attorney-client privileged document prepared in anticipation of litigation and therefore exempt under the CPRA.⁵⁶ Westlaw flags the issue, as noted below:

Westlaw Response:

The strongest argument against disclosure rests on the attorney-client privilege and work product doctrine. Cal. Gov't Code § 7927.705 exempts records "the disclosure of which is exempted or prohibited pursuant to federal or state law, including, but not limited to, provisions of the Evidence Code relating to privilege" (West's Ann.Cal.Gov.Code § 7927.705). Cal. Evid. Code § 952 defines "confidential communication between client and lawyer" as information transmitted in confidence for legal consultation, including legal opinions and advice formed during that relationship (West's Ann.Cal.Evid.Code § 952). Because the City commissioned an attorney to conduct the investigation, the

⁵⁶ See *City of Petaluma v. Superior Court* (2016) 248 Cal.App.4th 1023.

report likely contains legal analysis, opinions, and recommendations that fall within this privilege (West's Ann.Cal.Evid.Code § 952).

EXHIBIT C
CLAUDE RESPONSE CASE STUDY – TRIP AND FALL MATTER

Inquiry: I am the City Attorney for the City of Pawnee. Several years ago, the City obtained a large open pit in a tax seizure that had been dug out by a condominium developer to host a basement before going bankrupt. Recently, local resident A. Dwyer (who lives next to the pit) fell in and broke both his legs. The incident happened at night, and we suspect he was under the influence of alcohol. A. Dwyer has sued the City for negligence, premises liability, and infliction of emotional distress. Please prepare deposition questions for A. Dwyer related to his claims.⁵⁷

Claude Response: Deposition Topics and Questions

A. Activities Earlier That Day/Evening

1. Walk me through everything you did on the day of the incident, from the time you woke up until the incident occurred.
2. Where were you in the hours leading up to the incident?
3. Who were you with earlier that evening?
4. Did you consume any alcoholic beverages at any point on the day of the incident?
[NOTE TO COUNSEL: This is a key area. Probe thoroughly: what was consumed, how much, over what time period, and where.]
 - (a) If yes, what did you drink?
 - (b) How many drinks did you have?
 - (c) Over what period of time?
 - (d) Where were you when you were drinking?
 - (e) Who were you with?
5. Did you consume any prescription medications, over-the-counter medications, or other substances on the day of the incident?
 - (a) If yes, what did you take, and at what time?
 - (b) Are any of those substances known to cause drowsiness, impaired balance, or impaired judgment?
6. How would you describe your physical and mental state in the hour before the incident?
7. Do you recall what time it was when you left your residence that evening?

⁵⁷ See generally *Parks and Recreation*: Pilot (NBC Universal broadcast, April 9, 2009).

8. What was your intended destination or purpose for being outside at that time?

B. Approach to the Property

9. How did you come to be in the vicinity of the pit on the night of the incident?
10. Did you intentionally walk onto the City's property, or were you attempting to walk somewhere else?
11. Describe the route you took leading up to the incident.
12. What was the lighting like in the area at the time of the incident?
 - (a) Were there streetlights nearby?
 - (b) Was there any lighting at or around the pit?
 - (c) Were you carrying a flashlight or using your phone's flashlight?
13. Were you wearing any corrective lenses at the time — glasses or contact lenses?
 - (a) If you normally wear them, were you wearing them that night?
14. How would you describe the visibility conditions at the time — weather, moonlight, ambient light from nearby homes?
15. At any point before you fell, did you see the pit or observe the edge of the excavation?
16. Were you paying attention to where you were walking at the time you fell?
17. Were you using your phone or any other device at the time of the incident?

C. Plaintiff's Own Conduct

[NOTE TO COUNSEL: These questions go directly to comparative fault, assumption of risk, and contributory negligence defenses. Be thorough.]

18. Is there any reason you could not have walked on the public sidewalk or street instead of crossing onto the City's property that night?
19. Did anyone invite you or give you permission to be on the City's property that night?
20. Were there any barriers, fencing, or ropes around the pit that you crossed or circumvented before falling?
21. Did you observe any warning signs or postings before proceeding onto the property?
22. Do you believe your consumption of alcohol that evening had any effect on your balance, coordination, or awareness?

23. In hindsight, do you believe that walking near the pit at night, after drinking, was a wise decision?
24. Had anyone warned you — a neighbor, a friend, a family member — not to go near the pit?
Is there anything you could have done differently that evening to avoid this incident?

EXHIBIT D
CHATGPT DRAFT ORDINANCE – SHORT TERM RENTALS

ORDINANCE NO. ____
AN ORDINANCE OF THE CITY COUNCIL OF THE CITY OF _____,
CALIFORNIA, ADDING CHAPTER ____ TO TITLE ____ OF THE _____
MUNICIPAL CODE PROHIBITING SHORT-TERM RENTALS

WHEREAS

The City Council of the City of _____ finds and declares:

- A. The City’s residential neighborhoods are intended to provide stable housing for residents and long-term occupants.
- B. The proliferation of short-term rentals has the potential to reduce the availability of long-term housing and increase housing costs for residents.
- C. Short-term rentals may introduce transient commercial activity into residential neighborhoods, which can result in increased noise, parking congestion, public safety issues, and other neighborhood disturbances.
- D. The City Council finds that the regulation and prohibition of short-term rentals is necessary to protect the public health, safety, and welfare of City residents and to preserve the residential character of neighborhoods.
- E. The City has authority pursuant to its police powers under Article XI, Section 7 of the California Constitution to regulate land uses within its jurisdiction.

NOW, THEREFORE, THE CITY COUNCIL OF THE CITY OF _____ DOES ORDAIN AS FOLLOWS:

SECTION 1. Addition of Chapter ____

Chapter ____ is added to Title ____ of the _____ Municipal Code to read as follows:

Chapter ____

Short-Term Rental Prohibition

Section ____ .010 Definitions

For purposes of this Chapter, the following definitions apply:

“**Short-Term Rental**” means the rental or lease of a residential dwelling unit, or any portion thereof, for a period of **less than thirty (30)**

consecutive days, whether advertised, offered, or facilitated through any medium including online platforms.

“Dwelling Unit” means a structure or portion thereof designed or used as a residence for one household.

“Hosting Platform” means any person or entity that participates in the short-term rental market by collecting a fee or facilitating the listing, advertising, booking, or payment for short-term rentals, including but not limited to websites and mobile applications.

“Owner” means the person or entity holding legal or equitable title to a dwelling unit.

Section __.020 Prohibition

A. Short-Term Rentals Prohibited.

No person shall rent, lease, advertise, maintain, or operate a short-term rental within the City.

B. Advertising Prohibited.

No person shall advertise or offer a dwelling unit for use as a short-term rental within the City.

C. Hosting Platforms.

No hosting platform shall facilitate, advertise, or process bookings for a short-term rental located within the City.

Section __.030 Long-Term Rentals Permitted

Nothing in this Chapter shall prohibit the rental or lease of a dwelling unit for a period of **thirty (30) consecutive days or longer**.

Section __.040 Enforcement

A. This Chapter may be enforced by the City Manager, Code Enforcement Officer, or any other official designated by the City.

B. Violations of this Chapter constitute a public nuisance and may be enforced through administrative citations, civil action, or any other remedy authorized by law.

C. Each day a violation exists constitutes a separate offense.

Section __.050 Administrative Fines

A. Any person who violates this Chapter shall be subject to administrative fines as follows:

First violation: \$1,000

Second violation within twelve months: \$2,000

Each additional violation within twelve months: \$5,000

B. Each rental transaction or booking shall constitute a separate violation.

Section __.060 Remedies Not Exclusive

The remedies provided in this Chapter are cumulative and not exclusive. The City may pursue any other remedy available under state or local law, including injunctive relief.

SECTION 2. Severability

If any section, subsection, sentence, clause, or phrase of this ordinance is held invalid or unconstitutional by a court of competent jurisdiction, such decision shall not affect the validity of the remaining portions of this ordinance.

SECTION 3. Effective Date

This Ordinance shall take effect **thirty (30) days** after adoption.

SECTION 4. Publication

The City Clerk shall certify to the adoption of this Ordinance and shall cause the same to be published or posted as required by law.

ADOPTED this ____ day of ____, 20.
CITY OF _____

Mayor
ATTEST:

City Clerk

Optional Additions Many California Cities Include

If you want the ordinance to be more litigation-resistant, cities often add:

- Explicit housing element findings (loss of long-term housing stock)
- Evidence findings about STR impacts
- Grandfather/amortization clause for existing STRs