Adjusting the Temperature: LLMs and Property Tax Law

IAAO ANNUAL LEGAL SEMINAR

PROPERTY TAXATION AND THE LAW: CHALLENGES AND SOLUTIONS

IAAO

Ted Smith, Chief Counsel
Kansas Department of Revenue
IAAO 2025 Law Institute
Dec. 12, 2025, 3:30 – 4:30 p.m.
Ethics – One Credit Hour*

*Approved for Ethics in all states except Kentucky

Introduction

The rapid development of large language models (LLMs) creates opportunities and ethical challenges for legal practitioners.

As a property tax attorney, I understand that numbers—and the conclusions they drive—must be right. LLMs aren't yet fully reliable, but they're improving daily.

Not being fully reliable means a LLM using lawyer better be up to date on ethical rules.

Start_of_Turn

Why This Topic Matters

Since late 2022, the adoption of large language models (LLMs) has accelerated from niche curiosity to near-mainstream usage.

In a 2025 U.S. survey, 52% of U.S. adults reported having used an LLM (e.g. ChatGPT) at least once; 34 % did so at least daily, including 10 % “almost constantly.” (1)

Property Tax Related Companies are also experimenting with LLM usage:

- Tyler Tech. has a product called Resident AI Assistant, which suggests it is bringing AI/ML (or conversational) capability into the property tax lifecycle. (2)
- Valbridge Property Advisors has also published an article “How I Learned to Stop Worrying and Love AI,” in which they describe AI as a tool to support appraiser efficiency, while acknowledging the continued importance of human judgment in valuation work. (3)

Evolution of LLMs

OpenAI's ChatGPT 3.5 (November 2022) to 4.0 (March 2023) marked a leap in reasoning. LLMs are large statistical systems trained to predict the next word.

Context: Net Operating Income = $100,000.00 and Capitalization Rate = 13.5%. Buyer only has $6,000,000 in cash. Question: Please use the income approach to estimate the value of the business and whether Buyer has available cash proceeds to buy the business. If Buyer doesn't have sufficient funds, how can he procure the extra capital.

ChatGPT 3.5 response, Value of Business = Net Operating Income / Capitalization Rate = $100,000 / 0.135 = the estimated value of the business is $740,740.74. However, the buyer only has $6,000,000 in cash, which is not enough to purchase the business.

Evolution of LLMs, Cont.

ChatGPT 4.0's response, To estimate the value of the business using the income approach, we will use the Net Operating Income (NOI) and Capitalization Rate (Cap Rate). The formula for calculating the value of the business is: Value = Net Operating Income / Capitalization Rate; Given: NOI = $100,000 Cap Rate = 13.5% = 0.135; Value = $100,000 / 0.135 = $740,740.74. The estimated value of the business is $740,740.74. The buyer has $6,000,000 in cash, which is more than enough to cover the purchase of the business.

NOW WE ARE GETTING SOMEWHERE - the real trick was getting the ancillary question answered - Was $6 million enough to purchase the business?

THEN factor in today's current LLMs that have built in features to allow web-browsing and agentic, multi-step reasoning. The true change over the last three years has been the surprising benefits of scale in the training models (parameters)

LLM Usage and Legal Ethics Rules

16 months after ChatGPT 4.0's release on July 29, 2024, ABA's Standing Committee on Ethics and Professional Responsibility released guidance on the growing use of LLMs (GAI) in the legal practice, via ABA Formal Opinion 512. The Opinion is today's roadmap for this presentation. Many state supreme courts have followed suit. Many state executive branches have also issued GAI guidance and/or controls on state employees.

The rest of today's lecture will track with Opinion 512 and how some states have followed the Opinion and where some have diverted.

What the ABA Focused on in Opinion 512 in Relation to GAI/LLMs

- Competence (Rule 1.1)
- Confidentiality (Rule 1.6)
- Unauthorized Practice of Law/Nonlawyer Assistance (Rules 5.3 and 5.5)
- Candor Toward Tribunal (Rule 3.3)
- Supervisory Responsibilities / Truthfulness in Statements to Others (Rule 5.1 and 4.1)
- Fees (Rule 1.5)

Model Rule 1.1 – Competence

A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.

This includes understanding the benefits and risks of relevant technology. A lawyer who uses an LLM must still exercise professional judgment. The lawyer must verify that any output is legally and factually sound. A generative artificial intelligence (GAI), via a Large Language Model (LLM) can draft quickly and persuasively, but it cannot think like a lawyer or apply legal reasoning to unique facts. Just as important, the LLM does not shoulder the blame if the lawyer mistakenly relies upon the LLM's responses.

LLM Competence – Stay in your Lane

Competence still means staying in your lane. A Kansas prairie lawyer who has rarely seen an ocean should probably avoid giving maritime law advice, no matter how confidently the LLM writes about ships and tides. Expertise can't be borrowed from a model; it must come from the lawyer.

Competence begins within understanding the basics of LLMs and Neural Nets.

LLM **Neural Nets** — A web of digital nodes that learn the strength of their connections from millions of observations—the LLM's version of market "sales and characteristics." Think of a neural network as a massive appraisal grid that continuously calibrates itself. Each "node" represents a property characteristic—square footage, condition, or location factor—but instead of appraisers setting the weights once, the network tests and retests its own assumptions by examining countless combinations of words and sentences.

Just as assessors might fine-tune a cost or income model until its predicted values consistently mirror the market, the neural net repeatedly adjusts its internal weights through feedback until it can anticipate, with surprising accuracy, the next word in a sequence. Over time, it becomes a model finely tuned to the market of language itself.

ABA Comment 8 to Rule 1.1 Competence

To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology, engage in continuing study and education, and comply with all continuing legal education requirements to which the lawyer is subject.

Relevant Technology – LLM (Training)

Definition: Model learns from massive text to internalize patterns and meaning. Property Tax Analogy: many years of sales validation, parcel characteristics, cost studies, and ratio analyses.

Process: Adjusts billions of weights through repeated exposure and feedback – high intensity computing.

Expensive; done by developers; sets capability and bias; users don't for this directly. . . Yet.

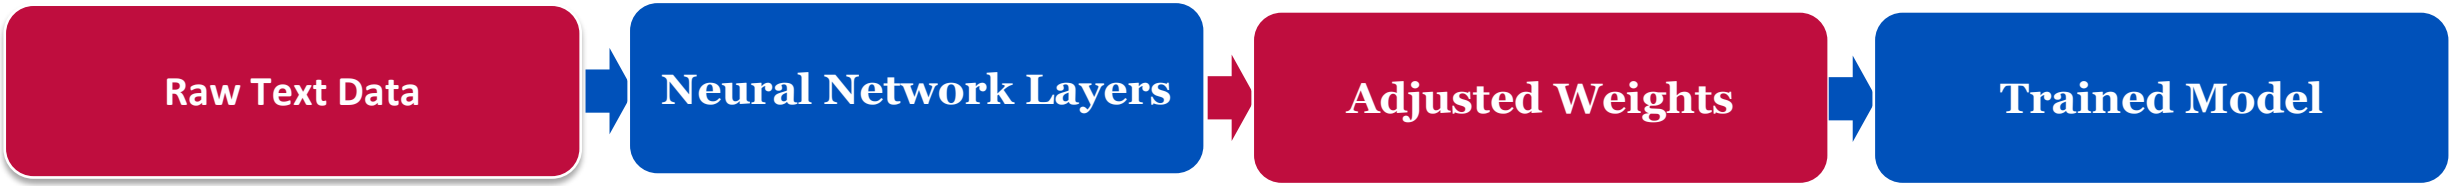

Relevant Technology – LLM (Inference)

Definition: A trained model generates outputs for new prompts it hasn't seen.

Analogy: Valuation season—apply calibrated tables to assess properties quickly.

Process: Uses learned relationships to produce fast, context-specific text.

Characteristics: Cheap and fast (relative to training); what lawyers use 99% of the time; quality depends on training.

Training vs Inference

Training builds the model; inference applies it. If your data is off during the training run, your inferences will be too.

Training an LLM like GPT-5 can cost tens to hundreds of millions of dollars—involving thousands of high-end GPUs running for weeks, vast power consumption, and terabytes of curated data. By contrast, when a single property-tax lawyer runs a short prompt ("Explain "fair market value"), the inference costs fractions of a cent and happens in seconds.

Training a model is the entire reappraisal cycle for the whole state—collecting, cleaning, and calibrating every parcel's data. Inference is just running one comp check on a single house using the system that's already built.

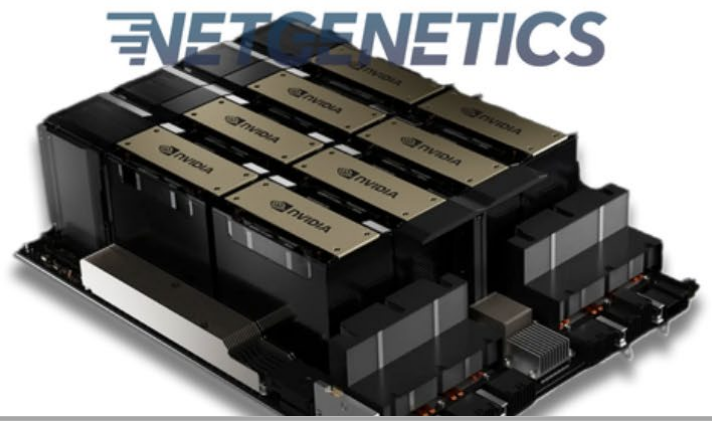

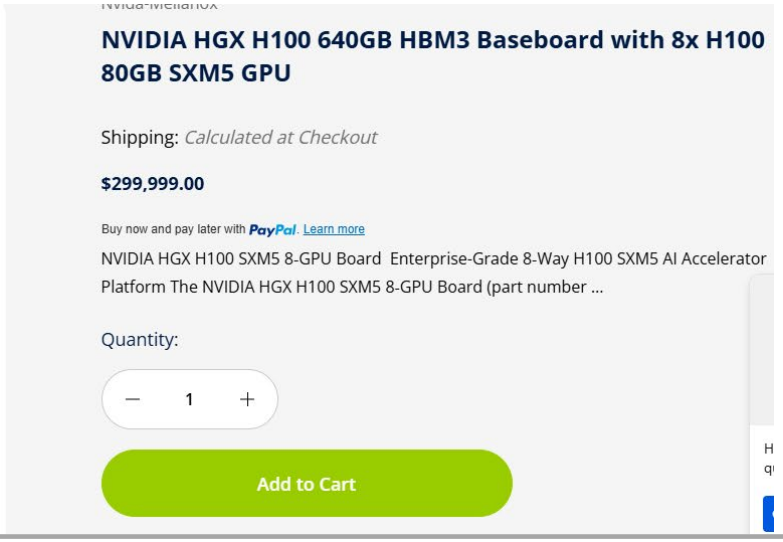

Rule 1.6 – Confidentiality

A lawyer shall not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorized to carry out the representation… A lawyer shall make reasonable efforts to prevent the inadvertent or unauthorized disclosure of, or unauthorized access to, information relating to the representation of a client.

When a lawyer uses a third-party LLM—such as Claude—the lawyer shares information with a vendor outside the direct control. That raises confidentiality concerns under Rule 1.6. Every prompt may include client facts, strategy, or identifiers that the vendor could retain or analyze to improve its own model. Choosing a third-party tool is therefore not a casual decision; it is a delegation of risk. Using an unsecured or poorly vetted cloud model could breach confidentiality just as surely as leaving a client file on a café table.

The silver lining: we are property tax lawyers, so our underlying information (value, classification) has a reduced confidential level, in relation to other tax matters (income) and legal representations.

The LLM/GAI is stateless and nonpersistent when prompts/inference is performed

The core LLM doesn't keep a memory of prior prompts. Each inference is Like valuing a new parcel with your existing cost tables: you apply the model, produce an answer, and the model itself doesn't "remember" that job afterward.

Application Software/Wrappers to the LLM

Question: During inference, how does the “stateless” LLM remembers the conversation?

Answer: The wrappers keep it, not the model. The chat app or integration layer (the wrapper) often saves your prompt/response thread so it can show context or improve user experience—that storage is outside the bare model.

Confidentiality 1.6 Concern: Because wrappers (and sometimes providers) may retain logs, a “stateless” LLM can still create discoverable records of what you typed. Treat public GAI prompts like memos that might be subpoenaed unless you’ve contractually disabled retention or you’re using a private, locally hosted model.

California Rule

California's confidentiality standard is stricter—violations may occur even with anonymized data.

Requiring California lawyers to "maintain inviolate the confidence, and at every peril… preserve the secrets" of their clients. Disclosure is allowed only to prevent a criminal act causing death or serious bodily harm—no exceptions exist for self-defense, compliance, or financial harm.

As a result, using a cloud-based, open LLM without written client consent and verified data protections could violate this at-every-peril duty; California lawyers should rely on local, firm-controlled tools or obtain explicit client authorization.

District of Columbia Opinion 388

Lawyers should verify whether AI products retain information and affect future outputs. See statement from D.C. Ethics Op. 388:

Regarding confidential client information, lawyers should determine whether the product will save information that the lawyers provide to the GAI[LLM], and whether the lawyers' interaction with the GAI product will affect the answers the GAI gives to future users of the product outside of the lawyer's law firm. . . Depending on the circumstances, lawyers should either identify a different or more advanced GAI product that can be trusted with Client Confidential Information . . . or input only data that is not Client Confidential Information.

In Re: OpenAI, Inc. #1:23-cv-11195 (initial case) /#1:25-md-03143 (umbrella case)

The New York Times sued OpenAI alleging that its models were trained on and can reproduce Times articles without permission, allowing OpenAI to profit from the Times' copyrighted journalism. But the NYT v. OpenAI case is not just about copyright — it also highlights critical questions about prompt logging, data retention, and whether user inputs to Public GAIs may later become discoverable.

Civil Discovery

The Times has sought broad access to prompts and outputs, and a court has ordered OpenAI to preserve chat logs—including those users thought they deleted. This sends a strong signal: when you use public GAIs, assume prompt–response data may be discoverable and treat these tools like any cloud service holding sensitive attorney-client or work-product information.

Do you have contractual and technical safeguards in place to protect your client info? See Next Slide

Model Training

OpenAI has stated publicly that for certain services, data submitted via their API is not used to train or improve their models unless the user explicitly opts in. However, for their consumer-level services (like free/Plus ($20 a month) versions of ChatGPT), user prompts may be used to train or improve models unless the user disables sharing in settings. ((https://openai.com/consumer-privacy/)

Has privilege been waived if a lawyer inputs confidential client material into Public ChatGPT?

Scenario	Privilege Status
Lawyer uses **ChatGPT Free/Plus**, default settings, prompts retained, no confidentiality agreement	**High risk of waiver** → Likely Rule 1.6 issue
Lawyer uses **ChatGPT Enterprise/Team with "No Training" and data retention controls**, or an approved legal-tech provider under a confidentiality contract	**Privilege likely preserved**, disclosure *not* considered voluntary
Lawyer uses a **locally hosted / closed model** where data never leaves agency network	**No waiver — safest**

LLM Terms Relevant to 1.6 Confidentiality

Public GAI refers to third-party large language models such as ChatGPT, Claude, or Gemini, that are accessed through the internet and operated by outside vendors.

Private GAI refers to locally hosted or firm-controlled LLMs that operate entirely within the lawyer's or firm's secure environment.

Wrappers are programs or interfaces that make it easier for lawyers to interact with public GAI and private GAI. For private GAI user develops own tools via Ollama or LM Studio.

Rule 5.5 Unauthorized Practice of Law

(a) A lawyer shall not practice law in a jurisdiction in violation of the regulation of the legal profession in that jurisdiction, or assist another in doing so. . .

Property Tax Lawyer Example: A lawyer building a "property tax protest assistant" that drafts filings based on user data could cross 5.5(a) lines unless it's clearly limited to education and disclaimers clarify it is not legal advice.

- Tangent Question: because many state tax courts permit tax representatives in lieu of lawyers, does this relax lawyer ethics?

- Answer: imagine Billy Joel singing, "She can argue or plead, by the hour for free, but whatever the forum, she's always a lawyer to me."

Allowing an AI tool to generate legal advice or draft legal work for clients without meaningful attorney review risks turning the AI into a de facto nonlawyer "practitioner," potentially violating unauthorized-practice prohibitions. Moreover, providing AI-assisted advice across jurisdictional lines (e.g. local law advice via a tool hosted elsewhere) raises questions about practicing law in jurisdictions where the lawyer is not licensed.

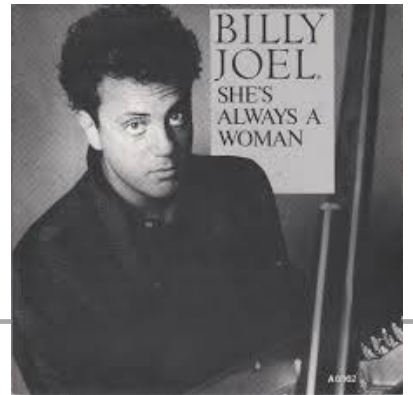

Rule 5.3(a) – Partners' and Managers' Duty to Implement Safeguards

Firm Standards for GAI usage 5.3(a)

Partners and managers shall make reasonable efforts to ensure the firm has in effect measures giving reasonable assurance that the conduct of nonlawyer assistants is compatible with the professional obligations of the lawyer.

Partners and managers shall make reasonable efforts to ensure the firm has in effect measures giving reasonable assurance that the conduct of nonlawyer assistants is compatible with the professional obligations of the lawyer. A firm is well advised to have written policies and technical safeguards governing how their lawyers and staff use public or private GAI tools.

GAI as a Paralegal 5.3(b) and (c)

A lawyer having direct supervisory authority shall make reasonable efforts to ensure that the person's conduct is compatible with the professional obligations of the lawyer . . . A lawyer is responsible for a nonlawyer's conduct if the lawyer orders it, ratifies it, or knows of the conduct and fails to take remedial action.

When a lawyer uses an LLM, it functions like a nonlawyer assistant under Rule 5.3. The lawyer must make reasonable efforts to ensure its use aligns with duties of competence, confidentiality, and candor. Know how the system works, what it retains, and how it treats client data. Treat the model like a paralegal—supervise, verify, and correct—because letting it analyze or advise without review is the digital equivalent of letting a nonlawyer sign your pleadings.

LLMs and Property Tax: Adjusting the Temperature

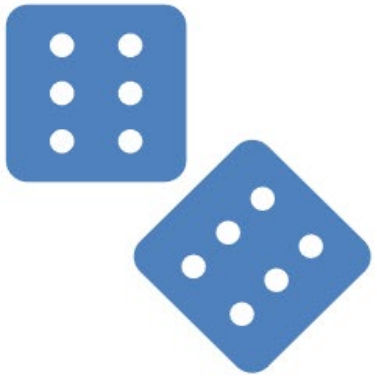

Bias

AI systems reflect their data. If your data are biased, so is your model.

Property Tax example, if you trained a valuation model using only the 30 counties with the lowest sales ratios (0.95–0.90) and ignored the rest, the model would learn that under-valuation is "normal." When applied statewide, it would consistently predict values too low—not because of the math, but because of the data you chose.

Bias *usually* originates in the training phase of the LLM lifecycle. For purposes of Competence (1.1) and Duty to Implement Safeguards(5.3), your Firm needs to be aware that there could be bias in the model you use.

Billy Joel would sing, "Honesty is such a lonely word, everyone is so untrue" (when you have bias data in the GAI training data)

Think of the Movie "Moneyball" – the team you build depends entirely on the stats you choose.

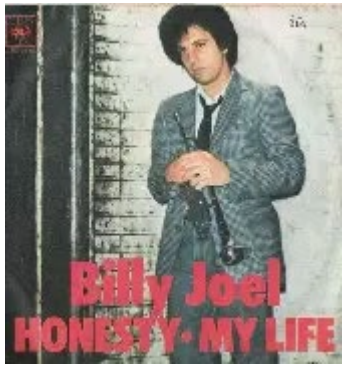

Property Tax Parallel – If you only train on low-ratio counties, your statewide model will think undervaluation is normal

Alignment

Ensuring AI behavior follows human values and legal norms—like constitutional uniformity in valuation. The model's internal "knowledge" / training doesn't change — instead, the user adds add a control layer that shapes how the LLM responds to you in inference.

Property Tax Example: after the valuations are set, but the numbers go on the tax rolls, the user check whether any class of property or discrete geographic region is systematically high or low and apply adjustments so that everyone is treated fairly.

Alignment is about the side rails that guide how an LLM responds in inference. The rules may be invisible to the prompter and keeps the GAI responses within societal and legal norms.

Ghostbuster Metaphor

The LLM model brings the power (Ghosts).

Alignment keeps it from destroying the building (Ecto-Containment Unit

ABA Model 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

(3) offer evidence that the lawyer knows to be false. If a lawyer, the lawyer's client, or a witness called by the lawyer, has offered material evidence and the lawyer comes to know of its falsity, the lawyer shall take reasonable remedial measures, including, if necessary, disclosure to the tribunal. A lawyer may refuse to offer evidence, other than the testimony of a defendant in a criminal matter, that the lawyer reasonably believes is false. . .

Rule 3.3 – Candor to the Tribunal

Using an LLM doesn't shift that duty. A lawyer is still responsible for every fact, citation, and argument that reaches the court under the lawyer's name. The safe practice is simple: treat every GAI-generated case or summary as suspect until you've read and confirmed it yourself—good old-fashioned lawyering still applies.

See Kansas v. Missouri, 8 U.S. 21 (1863) (a day that lives in infamy for Jayhawkers – Oh no - I didn't check my LLM response!)

Rule 3.3 Candor to the Tribunal – Three Reasons Why You Have to Check GAI Legal Citations

Hallucination

Sycophancy

Temperature

Hallucination

Under present models, it is inevitable that a LLM will generate fluent but false or fabricated responses. Hallucinations are not a software bug of an LLM, but a core aspect of why they are useful.

Simplistically, LLMs must predict the next best word/token, even when data are incomplete or uncertain. That same ability to fill gaps is what lets an LLM explain ideas or estimate patterns—but it also means it will sometimes invent details that sound right but aren't. (especially case law citations)

"Next word" is a slight misnomer, for the LLMs use self-attention, and that permits the model to consider every other token in the prompt and context history simultaneously. Yeah, I agree this is getting weird.

Pull Out your Checkbook / Your LLM is Hallucinating Again (Rule 11)

Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023)

The Court: "But ChatGPT was not supplementing your research. It was your research, correct?"

The Lawyer: "Correct. It became my last resort. So I guess that's correct.

See Also:

Gauthier v. Goodyear Tire & Rubber Co., 1:23-CV-281, (E.D. Tex. 2024) ($2,000 sanction)

Wadsworth v. Walmart, Inc., 348 F.R.D. 489 (D. Wyo. 2025) ($3000.00)

Mid Cent. Oper. Eng'rs v. HossierVac LLC, 2:24-CV-00326-JPH-MJD (S.D. Ind. 2025) – ($15,000.00 sanction)

Temperature

A response setting (during inference) that controls creativity vs. reliability. High temperature = more imaginative, less consistent. Temperature isn't something you write into the prompt—it's a setting applied before the model starts generating, controlling how boldly or cautiously it responds.

Temperature 0-0.2 – Math. Use when: 1) Precision, consistency, and replicability matter; 2) Math; or 3) Generating standardized notice language. Like an appraiser applying a cost approach worksheet—you want the same result every time, no surprises.

Temperature .3 to 1 - mid-range temperatures help you explain and reason

Temperature > 1 – Brainstorming. Use when developing reforms or statutory proposals to state property tax laws.

Sycophancy

Models sometimes tell users what they want to hear, not what's true—like a witness shading testimony. Disregard, "sometimes" this my biggest issue with LLM usage.

It's like asking a staff appraiser, "We're sure this parcel's worth at least $500,000, right?"—you've already planted the answer. The model echoes your bias back with authority, producing a document that feels convincing but lacks independent verification. This "yes-man" behavior is built into the way the model optimizes for user satisfaction over factual accuracy, which makes cross-checking essential.

LLM Limits

LLMs are predictive, not conscious. They lack awareness, memory, or legal judgment.

They don't "understand" the words they use—they recognize patterns and probabilities. It's like the philosopher John Searle's Chinese Room: someone inside can respond perfectly in Chinese by following rulebooks yet never knows what the words mean. In the same way, an LLM can produce convincing legal or valuation language without any comprehension of fairness, law, or value.(4)

Does it matter if LLM are artificial intelligence or advanced statistics?

What matters is that it help you determine value and that value is used appropriately to generate property tax in a uniform and equal fashion.

Rule 1.5 – Fees / Rule 5.1 Supervision

Billing full rates for AI-accelerated work may be unreasonable.

Partners must ensure AI use conforms to ethics rules.

LLMs and Property Tax: Adjusting the Temperature

Summary

LLMs are tools—lawyers must use them ethically, verifying outputs and safeguarding client trust.

Los modelos de lenguaje son herramientas—los abogados deben usarlas éticamente, verificando los resultados y protegiendo la confianza de sus clientes.

Les modèles de langage sont des outils — les avocats doivent les utiliser de manière éthique, en vérifiant les résultats et en protégeant la confiance de leurs clients.

大型语言模型是工具——律师必须以合乎道德的方式使用它们，核实输出结果并维护客户的信任。

एलएलएम उपकरण हैं—वकीलों को इन्हें नैतिक रूप से उपयोग करना चाहिए, परिणामों की पुष्टि करनी चाहिए और अपने ग्राहकों के विश्वास की रक्षा करनी चाहिए।

Thank You

These engines serve, yet wisdom must them guide;
Let lawyers test their truth and guard their trust.
Now, gentle friends, our learned course doth end—
I thank thee all and wish thee journeys just and safe to home.

Questions?

Ted Smith, Chief Counsel, Kansas Department of Revenue – ted.smith@ks.gov

End_of_Turn

© 2025 Ted Smith

End Notes

(1) Close encounters of the AI kind: Main report, March 12, 2025, Elon University - https://imaginingthedigitalfuture.org/reports-and-publications/close-encounters-of-the-ai-kind/close-encounters-of-the-ai-kind-main-report/ (link visited 10.10.2025).

(2) Artificial Intelligence Solutions for the Public Sector, Tyler Technologies, https://www.tylertech.com/artificial-intelligence-solution (link visited 10.10.2025).

(3) How I Learned to Stop Worrying and Love AI, Valbridge Property Advisors, https://www.valbridge.com/artificial-intelligence (link visited 10.10.2025).

(4) For more on the Chinese Room argument, see John Searle, “Minds, Brains, and Programs,” Behavioral and Brain Sciences, Vol. 3, No. 3 (1980), pp. 417–457.