

Using AI in your Tax Practice

a Seminar and Hands-on Workshop



*Organized by IFA Canada
(in collaboration with TEI, HEC Montréal (Université de Montréal),
Stetson University, and Gowling WLG)*

PRESENTED ON AUGUST 16, 2025 AT HEC MONTRÉAL

USING AI IN YOUR TAX PRACTICE

Seminar Part I

Understanding current
AI applications
(Large Language Models)
available to tax practitioners

Seminar Part II

Understanding the
professional obligations and
motivations of practitioners
&
how these interact or clash
with AI technology

Seminar Part III

Using AI in your practice
Learning prompt engineering
skills for tax applications

MODERATORS



JOSIE SCALIA

INTERNATIONAL PRESIDENT,
TAX EXECUTIVES INSTITUTE



LAURA GHEORGHIU

TAX PARTNER, GOWLING WLG (CANADA) LLP.

FACILITATORS



ANTHONY TAHAN

HEC MONTRÉAL



LOUIS CABANEL

STETSON UNIVERSITY



LUCILA PAGLIERI

TORCUATO DI TELLA
UNIVERSITY



MARIANA FERNANDES

DELL TECHNOLOGIES

ORGANIZING COMMITTEE



PROFESSOR VIDAL



LIONEL NOBRE



PROFESSOR APPLEBY

SPEAKERS



Abdi Aidid

PROFESSOR AT THE UNIVERSITY
OF TORONTO, FACULTY OF LAW



Kyle Bunnell

FORMER DIRECTOR OF TAX,
HOOTSUITE



Sashi Kara

PRINCIPAL AND CHIEF
TECHNOLOGIST, SENTE ADVISORS

SPEAKERS



Denis Keseris

PARTNER AND PATENT AGENT,
GOWLING WLG (CANADA) LLP



Ara Stepanyan

PRINCIPAL, THE BRATTLE GROUP



Nicolas Vermeys

PROFESSOR AT THE FACULTY OF LAW
AT THE UNIVERSITÉ DE MONTRÉAL

UNDERSTANDING CURRENT AI APPLICATIONS (*LARGE LANGUAGE MODELS*) AVAILABLE TO TAX PRACTITIONERS

PART I

OUTLINE

1. Objectives of the Presentation
2. What is an LLM (technical background)
3. Analogies and Contrasts
4. Setting the Stage for a Paradigm Shift
5. Conceptual Application of Intelligence Tasks
6. Potential Questions
7. Human in the Loop



Roy Batty (Bladerunner)

Lead Revolutionary

Nexus model 6 Replicant

Goals:

- Gain extended lifespan
- Avenge fallen comrades
- Kill Rick Deckard
- Die peaceful death



T-1000

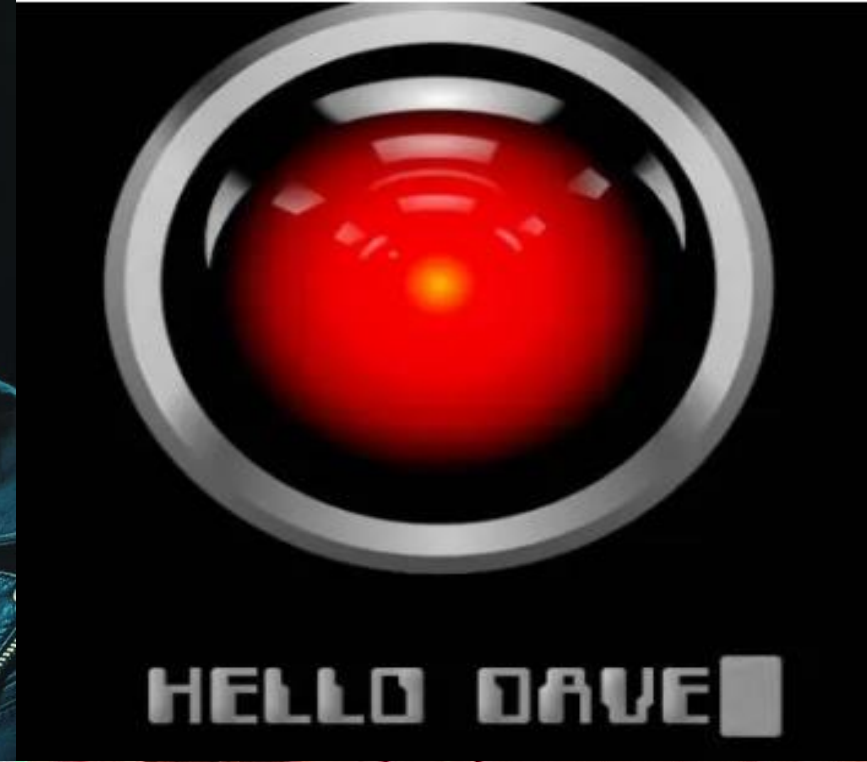
T-800 (Arnold)

T-1000 Goals:

- Terminate John Connor
- Prevent human resistance

T-800 Goals:

- Protect John Connor
- Destroy T-1000



Hal

Primary Computer

Spacecraft Discovery One.

Concerns:

- Self-preservation instinct
- Misinterpretation of orders
- Error in HAL's logic

LARGE LANGUAGE MODELS: CONTEXT

Artificial Intelligence (AI): technology that enables computers and machines to *simulate* human intelligence and problem-solving capabilities.

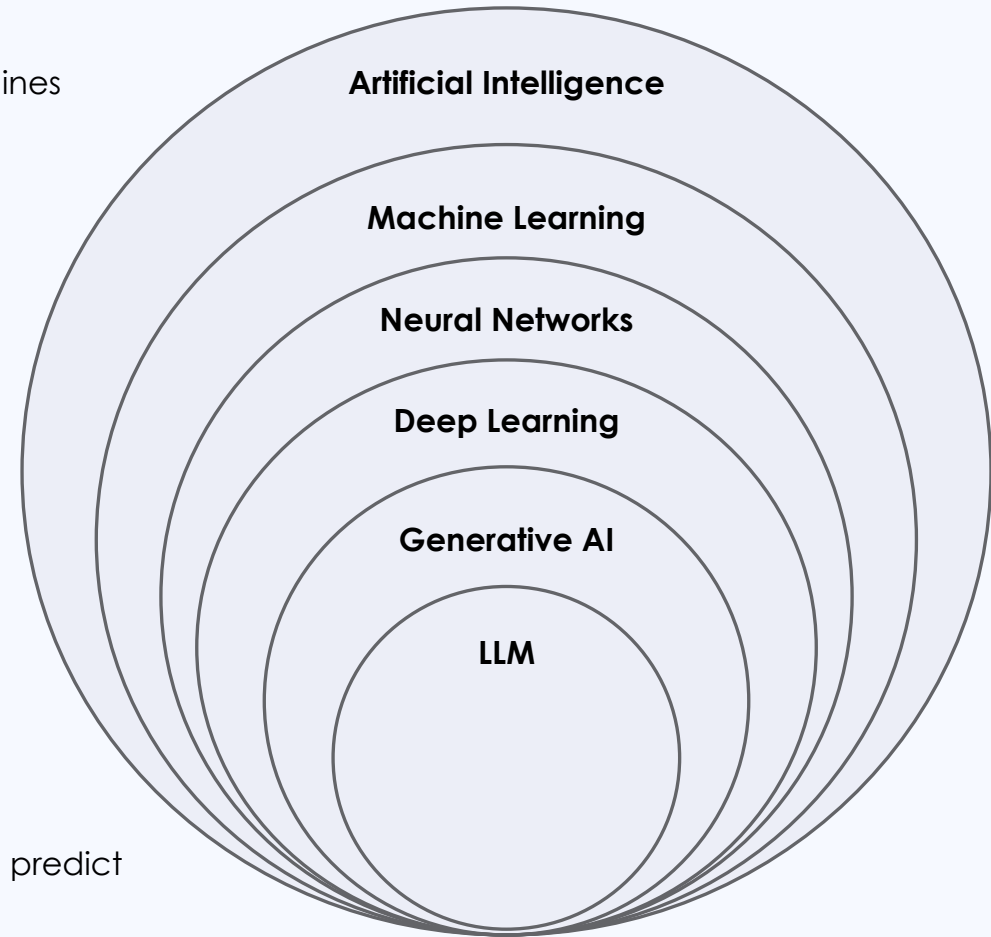
Machine learning (ML): branch of artificial intelligence (AI) that uses data and algorithms to enable AI to *imitate* the way that humans learn, gradually improving its accuracy.

Neural Networks: machine learning programs, or models, that make decisions in a manner *similar* to the human brain, by using processes that *mimic* the way biological neurons work together.

Deep Learning: a subset of machine learning that uses multilayered neural networks, called deep neural networks, to *simulate* the complex decision-making power of the human brain.

Generative AI: deep-learning models that can generate text, images, and other content based on the data they were trained on.

Large Language Model: a type of generative AI that uses deep learning techniques and large data sets to understand, summarize, generate and predict new text-based content.



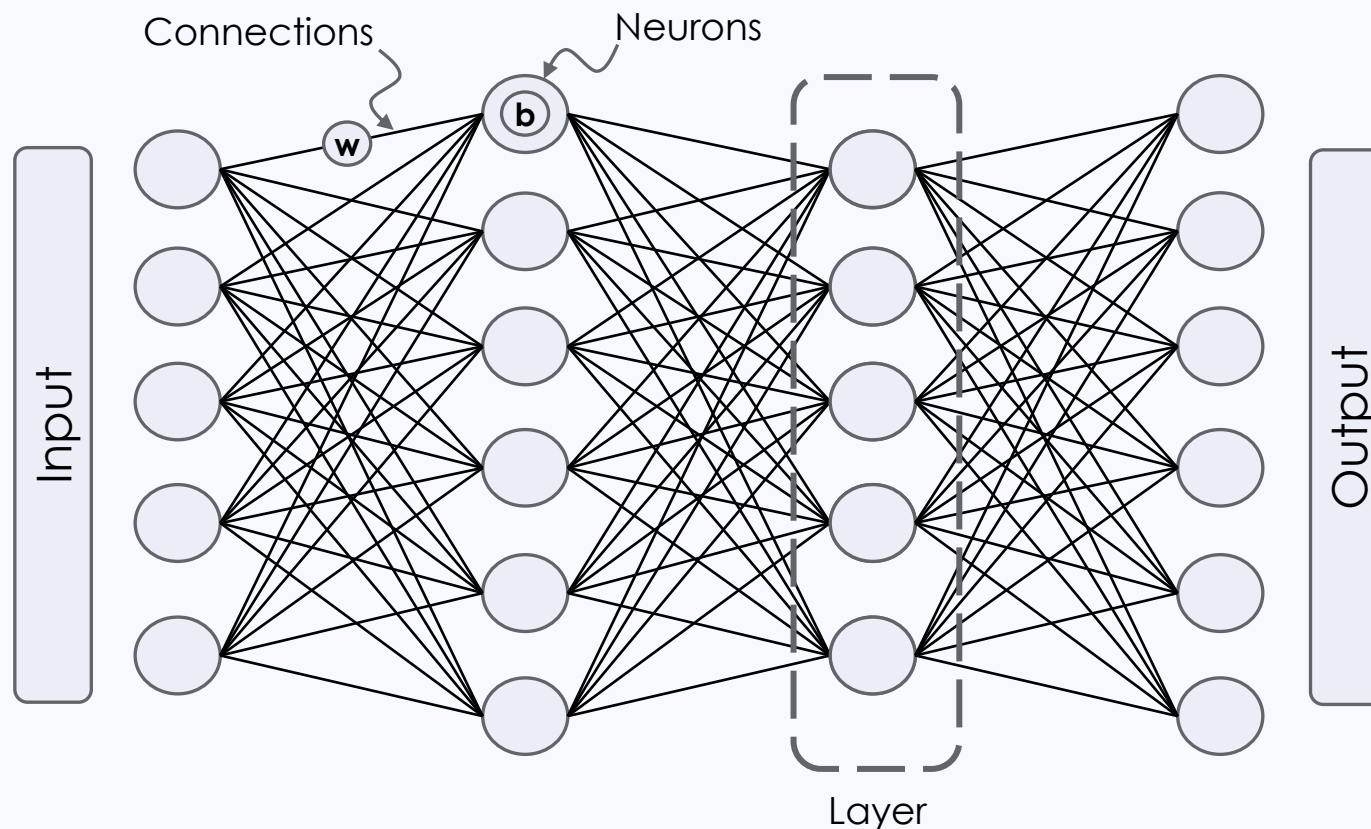
LARGE LANGUAGE MODELS: STRUCTURE & FUNCTION

Words first get converted into numbers (tokens)...

"London Bridge is"



"London Bridge is"

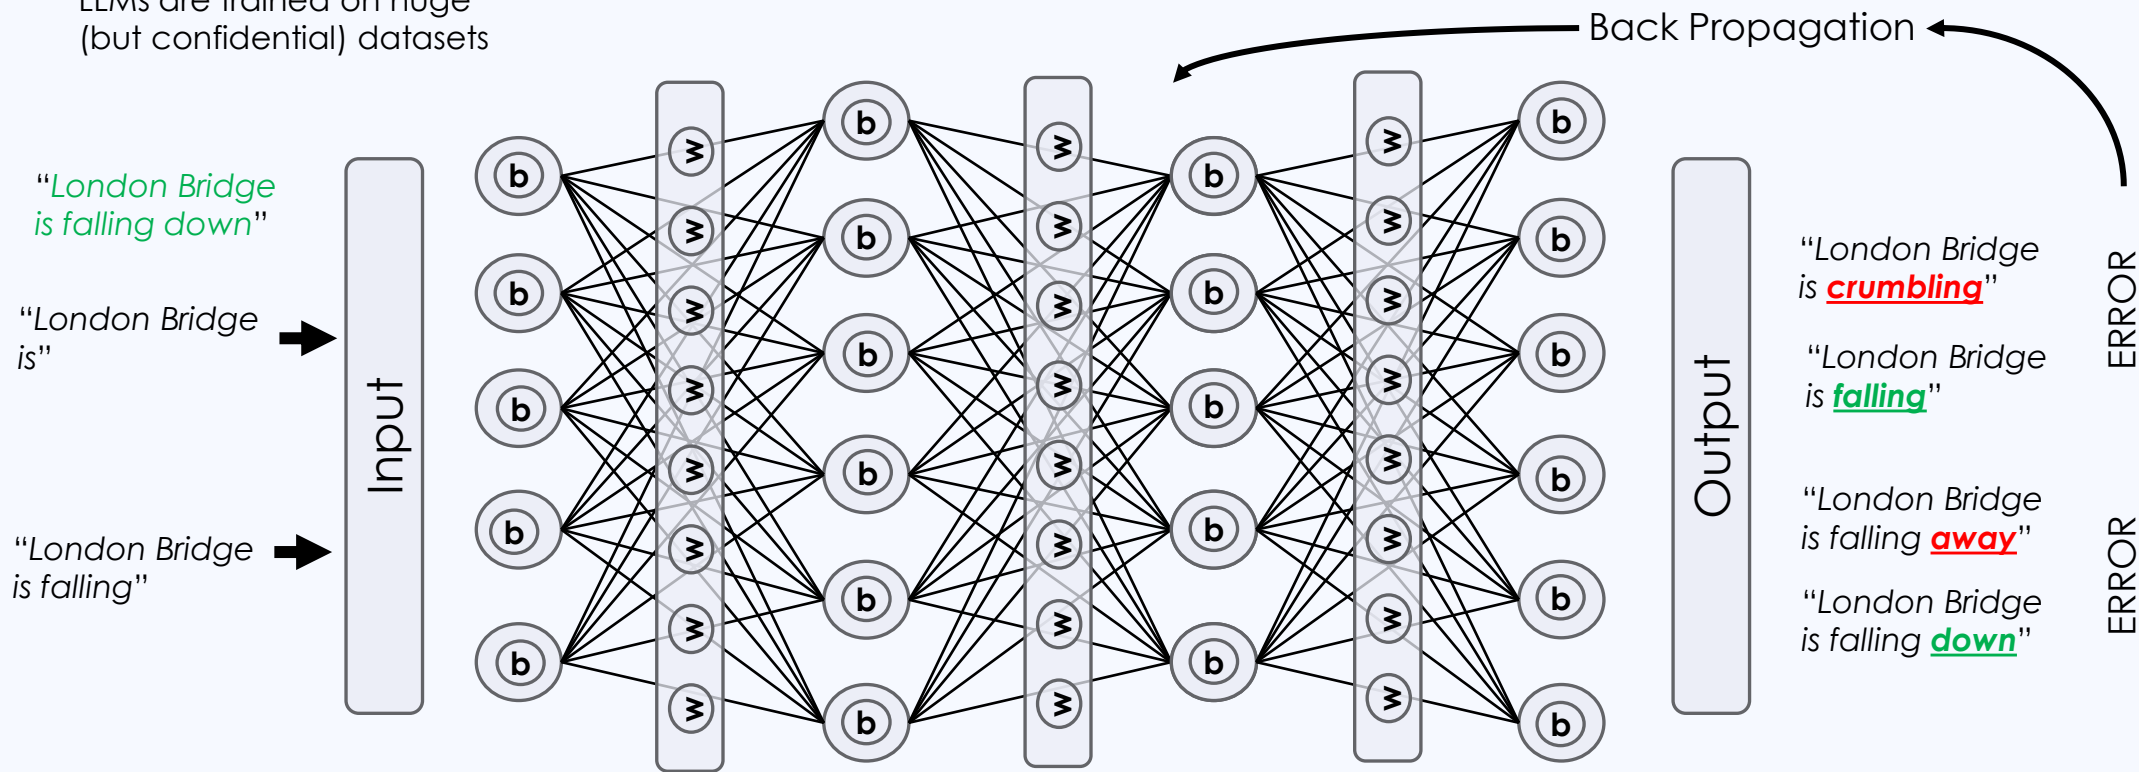


- "falling" = 35%
- "down" = 30%
- "located" = 10%
- "famous" = 7%
- "old" = 5%
- "iconic" = 3%

GPT-3 has approximately 175 billion parameters

LARGE LANGUAGE MODELS: TRAINING

LLMs are trained on huge (but confidential) datasets



Models are trained/updated at a specific point in time. So, the model itself might not be aware of how **likely** a word has recently become in a particular context. BUT...

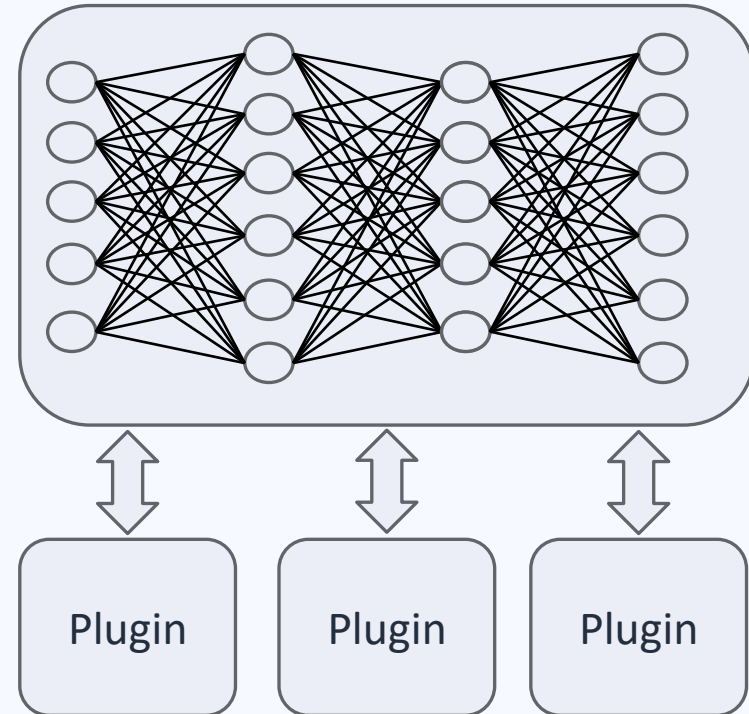
Fine Tuning (e.g., for tasks such as sentiment analysis, text classification, or conversational agents)

LARGE LANGUAGE MODELS: ENHANCEMENTS

LLM Plugins are external tools that expand an LLM's capabilities by allowing it to perform specific tasks or access external data sources that are not part of the model's inherent functions.

Examples include:

- 1) **API Integration** (weather, currency exchange, news)
- 2) **Database Query** (SQL database, knowledge base)
- 3) **Task Automation** (Lists, emails, etc.)
- 4) **Specialized Data Retrieval** (Medical, legal, financial, **tax**)
- 5) **Content Generation** (Code generation, image generation)
- 6) **Communications** (Slack, Teams, Discord)
- 7) **E-Commerce** (Amazon, e-bay)
- 8) **Computation** (Wolfram)



LARGE LANGUAGE MODELS: ENHANCEMENTS

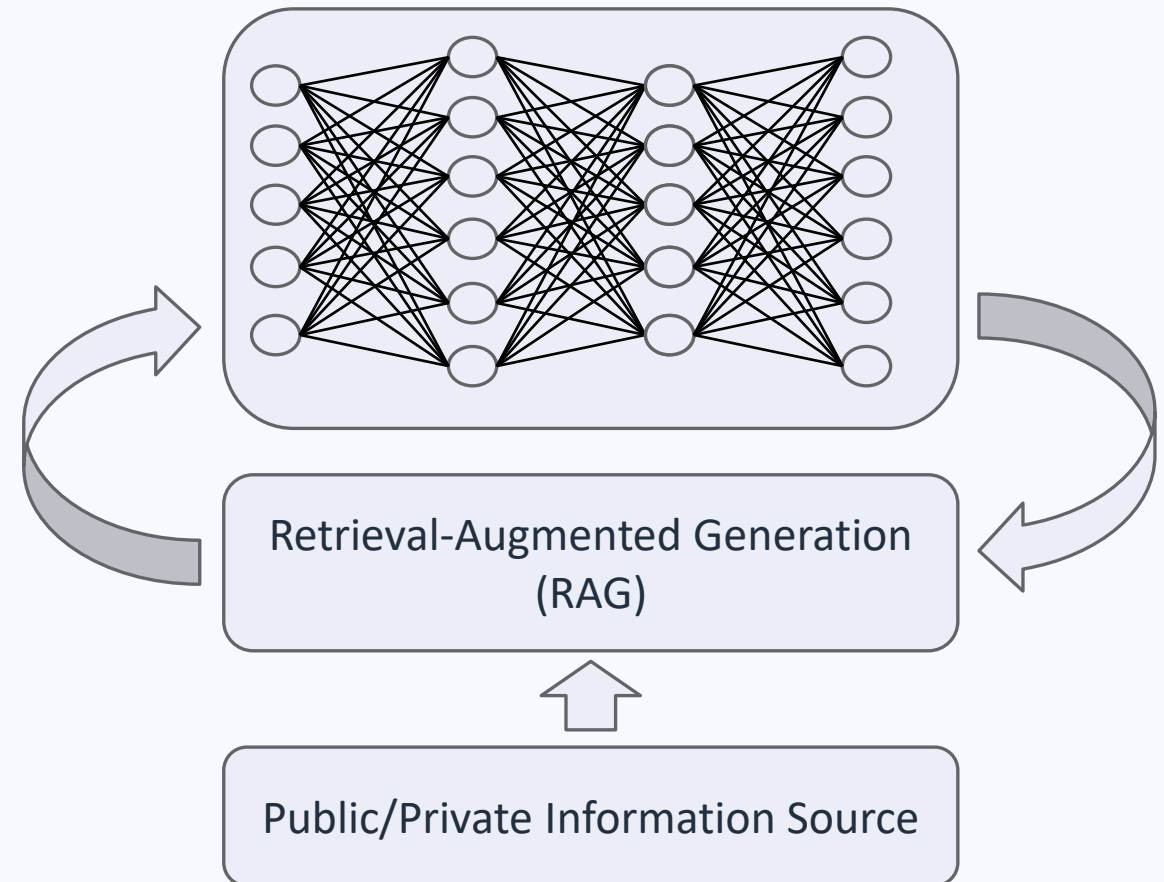
Retrieval-Augmented Generation (RAG) is a framework that combines the power of retrieval-based methods with generation-based models.

Examples include:

Open-Domain Question Answering: Instead of relying solely on pre-trained knowledge, the system retrieves relevant documents from sources like Wikipedia, then generates a precise answer based on the most pertinent information.

Conversational Agents: A customer service chatbot powered by RAG can provide accurate and up-to-date responses to customer inquiries by retrieving information from a company's knowledge base, FAQs, or other documentation.

Tax Law Research for Professionals: A RAG system designed for tax professionals could retrieve and summarize complex tax laws, regulations, and case law from vast legal databases.



LARGE LANGUAGE MODELS: ANALOGIES & CONTRASTS

Large Language Models:

- Are versatile **tools** with multiple functions, particularly when used in combination with other enhancements such as plugins and RAGs.
- Are next-word prediction machines: **probabilistic**, as opposed to **deterministic** (as in the case of traditional software).
- Cannot be timid: beware of anthropomorphizing LLMs, as they are designed to generate authoritative-sounding statements based on probabilistic processes (e.g., the use of the term “hallucination” instead of “mistake”).
- **Can** be safe from a data-privacy perspective: Instances of LLMs can be set up in secure/closed environments (e.g., Microsoft Azure).
- Not intelligent (IMHO)...



...tools, not thinkers: LLMs are powerful **tools** that can assist with various tasks, but they do not “think” like humans.

...no understanding: LLMs generate responses based on patterns and probabilities learned during training, without any comprehension or awareness of what they are saying.

...no consciousness: LLMs do not have consciousness, emotions, or self-awareness (as far as I can tell, anyway).

...no learning (after training): Once an LLM is trained, it does not learn or adapt in real time. It does not gain new knowledge unless it is explicitly retrained with new data (or is enhanced with other tools).

...limited contextual relevance: LLMs can generate contextually relevant responses based on input, but their understanding of context is limited to what they can infer from the data they were trained on.

Intelligence Amplifier



Intelligence Tasks

- Comprehension
- Simplification & Summarization
- Synthesis
- Documentation & Communication
- Analysis
- Problem-Solving
- Critical Thinking Support



LLMs are Good at...

- **Comprehension:** LLMs can assist in interpreting and explaining complex documents and regulations to improve understanding.
- **Simplification and Summarization:** They excel at breaking down complex ideas into simpler concepts and summarizing lengthy documents for quicker comprehension.
- **Synthesis:** LLMs can combine information from various sources to provide comprehensive overviews or develop new insights.
- **Documentation and Communication:** They aid in drafting clear and coherent documents, emails, and reports, facilitating effective communication.
- **Analysis:** LLMs can analyze textual data to extract insights and highlight key points, apply thinking frameworks like SWOT, 6 Hats, SCAMPER etc.
- **Problem-Solving (with guidance):** They can propose solutions based on provided data and programmed logic, but real-world problem-solving often requires human creativity and contextual understanding.
- **Decision-Making Support:** While LLMs cannot make decisions, they can provide information and generate comprehensive pros and cons lists to aid the decision-making process.
- **Critical Thinking Support:** LLMs can offer different perspectives and challenge assumptions to aid in critical thinking, although the depth of this is limited to their programming and data.

Are less Good at...

- **Counting:** LLMs struggle with counting due to their reliance on pattern recognition rather than inherent mathematical understanding.
- **Retention & Memory:** LLMs do not 'remember' in the human sense but can recall information from their training data up to the cutoff. They cannot recall personalized or newly generated information after the conversation ends.
- **Attention to Detail:** LLMs may not reliably catch subtle errors or nuances, particularly in numerical data or highly technical material. "Needle in the haystack" testing is improving though.
- **Decision-Making:** LLMs cannot make decisions but can support decision-making processes by providing synthesized information and potential consequences.
- **Advising & Consulting:** While LLMs can provide information that may inform advice, the nuanced understanding required for professional consulting is beyond their capability, especially where situational judgment and ethical considerations are involved.
- **Forecasting & Planning:** Predictive tasks are limited by the LLM's lack of access to real-time data and inability to understand complex market dynamics or individual organizational strategies. Internet-connected AI is improving LLMs' capabilities here.
- **Learning & Development:** While LLMs can support learning by providing information and resources, the active development of skills and knowledge acquisition through experience is a deeply personal process that LLMs cannot replicate.

Human in the Loop

LLMs don't know anything!!

- Check AI for hallucinations
- Don't be so easily convinced
- Provide oversight
- Include appropriate context
- Offer your perspective
- Apply critical thinking



UNDERSTANDING THE PROFESSIONAL OBLIGATIONS AND MOTIVATIONS OF PRACTITIONERS & HOW THESE INTERACT OR CLASH WITH AI TECHNOLOGY

PART II

OUTLINE

1. Technological Competency Obligation for Lawyers
2. Ethical Risks Relating to the Use of AI
3. Rethinking Our Relationship with AI

TECHNOLOGICAL COMPETENCY OBLIGATION FOR LAWYERS



T.J. Hooper v. Northern Barge, 60 F. 2d 737 (1932)

TECHNOLOGICAL COMPETENCY OBLIGATION FOR LAWYERS

“While AI-based tools may be a relatively new addition to the practice of law, **lawyers have always had implicit professional ethical duties when it comes to the use of technologies in their practice [...]**”

Ian STEDMAN, “Ethical Lawyering and Artificial Intelligence”, in Giuseppina D'AGOSTINO, Aviv GAON and Carole PIOVESAN (eds), *Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law*, Toronto, Thomson Reuters Canada, p. 147, at p. 150.

Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law

Giuseppina D'Agostino
Aviv Gaon and
Carole Piovesan



 THOMSON REUTERS

TECHNOLOGICAL COMPETENCY OBLIGATION FOR LAWYERS

“A lawyer must engage in his professional activities with competence. To this end, he must develop his knowledge and skills and keep them up to date.

For the purposes of the first paragraph, **the knowledge and skills related to information technologies used within the scope of the lawyer’s professional activities are part of the knowledge and skills that a lawyer develops and keeps up to date.”**

Code of Professional Conduct of Lawyers, CQLR c B-1, r 3.1, art. 21.



TECHNOLOGICAL COMPETENCY OBLIGATION FOR LAWYERS

“In some circumstances, expertise in a particular field of law may be required; often the necessary degree of proficiency will be that of the general practitioner.

[4A] To maintain the required level of competence, **a lawyer should develop an understanding of, and ability to use, technology relevant to the nature and area of the lawyer’s practice and responsibilities.** A lawyer should **understand the benefits and risks associated with relevant technology**, recognizing the lawyer’s duty to protect confidential information set out in section 3.3.

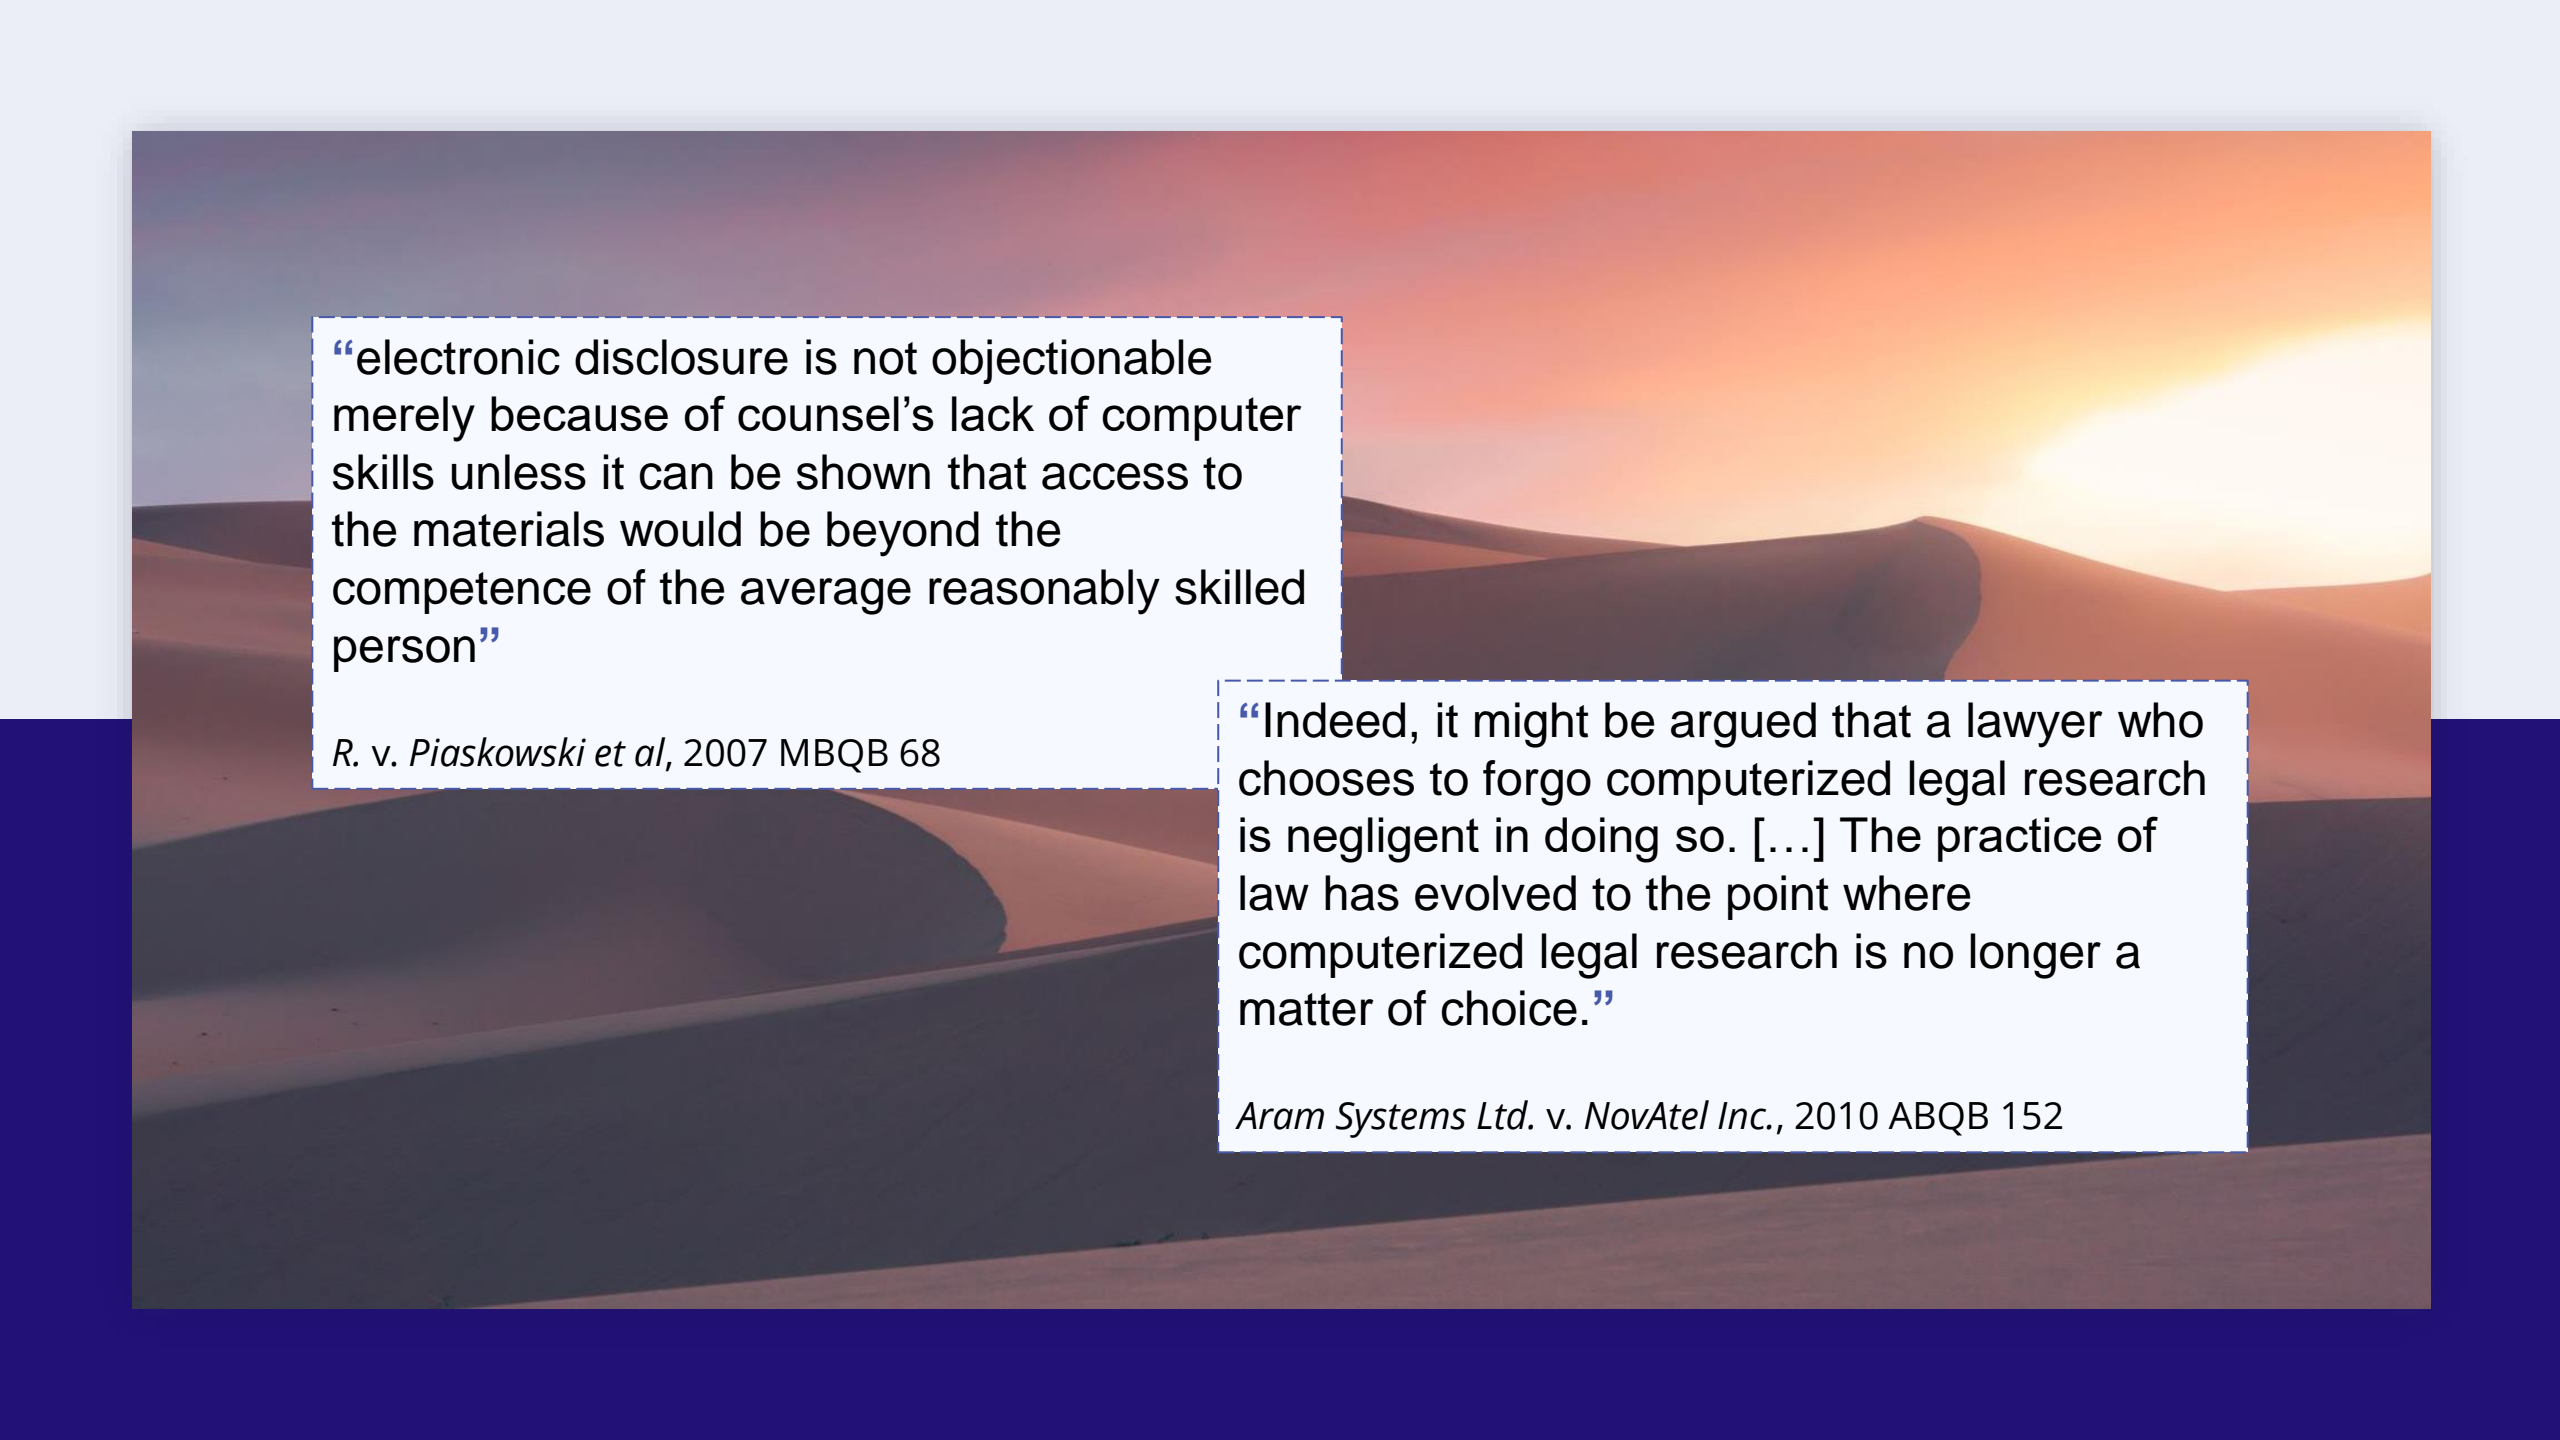
[4B] The required level of technological competence will depend upon whether the use or understanding of technology is necessary to the nature and area of the lawyer’s practice and responsibilities and whether the relevant technology is reasonably available to the lawyer. In determining whether technology is reasonably available, consideration should be given to factors including:

- (a) The lawyer’s or law firm’s practice areas;
- (b) The geographic locations of the lawyer’s or firm’s practice; and
- (c) The requirements of clients.”

LAW SOCIETY OF ONTARIO, *Rules of Professional Conduct*, art. 3.1-2, commentary 4.



**Federation of
Law Societies
of Canada**



“electronic disclosure is not objectionable merely because of counsel’s lack of computer skills unless it can be shown that access to the materials would be beyond the competence of the average reasonably skilled person”

R. v. Piaskowski et al, 2007 MBQB 68

“Indeed, it might be argued that a lawyer who chooses to forgo computerized legal research is negligent in doing so. [...] The practice of law has evolved to the point where computerized legal research is no longer a matter of choice.”

Aram Systems Ltd. v. NovAtel Inc., 2010 ABQB 152

TECHNOLOGICAL COMPETENCY OBLIGATION FOR LAWYERS

“It may very well be that the professional practice reflects prudent and diligent conduct. One would hope that if a certain practice has developed amongst professionals in regard to a particular professional act, such practice is in accordance with a prudent course of action. **The fact that a professional has followed the practice of his or her peers may be strong evidence of reasonable and diligent conduct, but it is not determinative.** If the practice is not in accordance with the general standards of liability, i.e., that one must act in a reasonable manner, then the professional who adheres to such a practice can be found liable, depending on the facts of each case.”

Roberge v. Bolduc, [1991] 1 SCR 374



A stylized human brain composed of glowing circuitry and data lines, with a central chip at the base. The brain is formed by intricate patterns of copper-colored lines and glowing white nodes, set against a dark grey background. At the bottom center, a glowing chip is connected to the brain's base. The overall image is framed by a dark blue border.

HOW DOES THIS APPLY TO AI?

	AI	Lawyers (20)
Efficiency rate	94%	85% (best: 94%, worst: 67%)
Average time	26 sec.	92 min.
Coffee cups	0	12

{LawGeex}

TECHNOLOGICAL COMPETENCY OBLIGATION FOR LAWYERS

“[...] lawyers using and/or advising on AI must have at least a basic understanding of what AI is, how it works (i.e. the different methodologies), how it can be applied and what the legal and ethical implications of its design and use might be.”

Ian STEDMAN, “Ethical Lawyering and Artificial Intelligence, in Giuseppina D'AGOSTINO, Aviv GAON and Carole PIOVESAN (eds), *Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law*, Toronto, Thomson Reuters Canada, p. 147, at p. 147.

Leading Legal Disruption: Artificial Intelligence and a Toolkit for Lawyers and the Law

Giuseppina D'Agostino
Aviv Gaon and
Carole Piovesan



 THOMSON REUTERS

A close-up photograph of a hand moving a black chess pawn on a chessboard. The hand is in the foreground, and the chessboard is in the background. The lighting is dramatic, highlighting the hand and the piece being moved.

ETHICAL RISKS RELATING TO THE USE OF AI

ETHICAL RISKS RELATING TO THE USE OF AI

Understanding AI's limits...

- Induction
- Overfitting
- Bias
- Dimensionality
- Scalability
- Simplicity (rather than accuracy)
- Correlation (\neq causation)
- Etc.



ETHICAL RISKS RELATING TO THE USE OF AI

“As it was later revealed, Mr. Schwartz had used **ChatGPT**, which fabricated the cited cases. Mr. Schwartz testified at the sanctions hearing that when he reviewed the reply memo, **he was “operating under the false perception that this website [i.e., ChatGPT] could not possibly be fabricating cases on its own.”** He stated, “I just was not thinking that the case could be fabricated, so I was not looking at it from that point of view.” “My reaction was, **ChatGPT** is finding that case somewhere. Maybe it's unpublished. Maybe it was appealed. Maybe access is difficult to get. I just never thought it could be made up.”

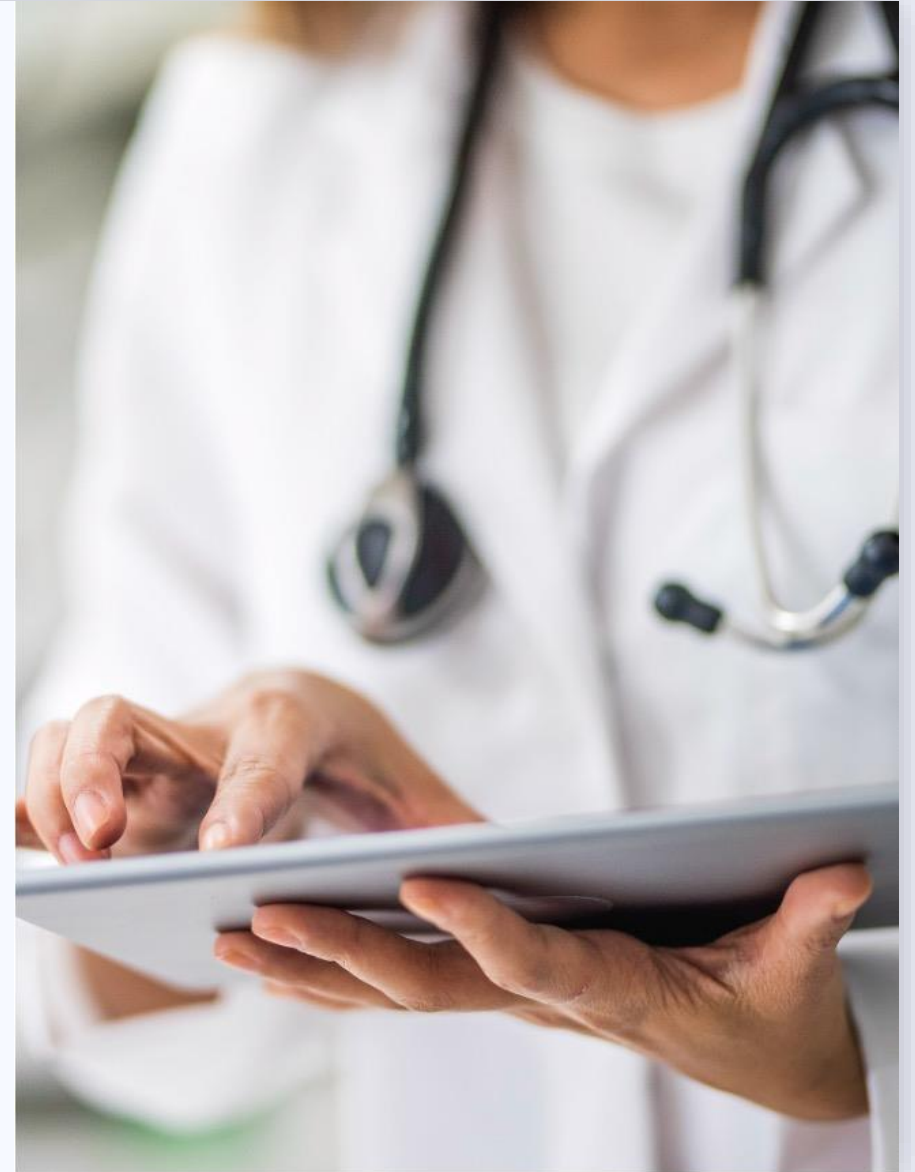
Mata v. Avianca inc., 678 F.Supp.3d 443, par. 11.



ETHICAL RISKS RELATING TO THE USE OF AI

“Attorney Lee's submission of a brief relying on non-existent authority reveals that she failed to determine that the argument she made was “legally tenable.” The brief presents a false statement of law to this Court, and **it appears that Attorney Lee made no inquiry, much less the reasonable inquiry** required by Rule 11 and long-standing precedent, into the validity of the arguments she presented. We therefore **REFER** Attorney Lee to the Court's Grievance Panel pursuant to Local Rule 46.2 for further investigation, and for consideration of a referral to the Committee on Admissions and Grievances.”

Park v. Kim, 117 Fed.R.Serv.3d 1693



ETHICAL RISKS RELATING TO THE USE OF AI

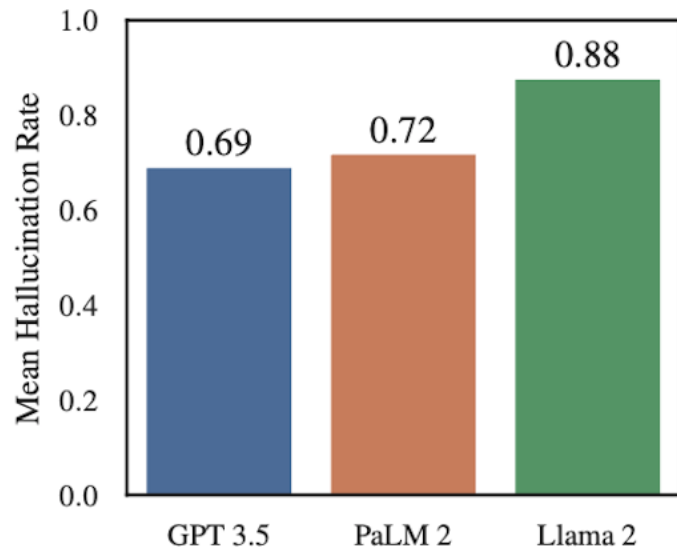
“At the start of the hearing on January 31, Ms. Ke’s affidavit sworn January 26, 2024 was handed up. It indicates that **she was called in May 2019 and outlines how she came to find the cases on ChatGPT, her lack of knowledge of the risks of its use**, the efforts she took to respond to the emails from opposing counsel’s office, the difficulties she had being on a flight to Asia when the inquiries started to come in, and her discovery that the cases were fictitious, which she describes as being “mortifying”.”

Zhang v Chen, 2024 BCSC 285, par. 16





ETHICAL RISKS RELATING TO THE USE OF AI



HALLUCINATIONS (69%)
~~Confabulations~~
~~Bullshit~~

Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models

Matthew Dahl Yale University matthew.dahl@yale.edu
Varun Magesh Stanford University vim@stanford.edu
Mirac Suzgun Stanford University msuzgun@stanford.edu
Daniel E. Ho Stanford University deho@stanford.edu

Abstract

Large language models (LLMs) have the potential to transform the practice of law, but this potential is threatened by the presence of legal hallucinations—responses from these models that are not consistent with legal facts. We investigate the extent of these hallucinations using an original suite of legal queries, comparing LLMs’ responses to structured legal metadata and examining their consistency. Our work makes four key contributions: (1) We develop a typology of legal hallucinations, providing a conceptual framework for future research in this area. (2) We find that legal hallucinations are alarmingly prevalent, occurring between 69% of the time with ChatGPT 3.5 and 88% with Llama 2, when these models are asked specific, verifiable questions about random federal court cases. (3) We illustrate that LLMs often fail to correct a user’s incorrect legal assumptions in a contra-factual question setup. (4) We provide evidence that LLMs cannot always predict, or do not always know, when they are producing legal hallucinations. Taken together, these findings caution against the rapid and unsupervised integration of popular LLMs into legal tasks. Even experienced lawyers must remain wary of legal hallucinations, and the risks are highest for those who stand to benefit from LLMs the most—*pro se* litigants or those without access to traditional legal resources.¹

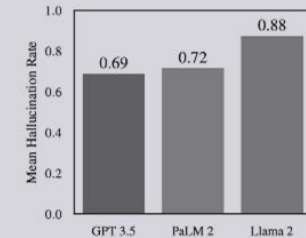


Figure 1: Hallucinations are common across all the LLMs we test when they are asked a direct, verifiable question about a federal court case. (Figure pools all reference-based tasks.)

arXiv:2401.01301v1 [cs.CL] 2 Jan 2024

1 Introduction

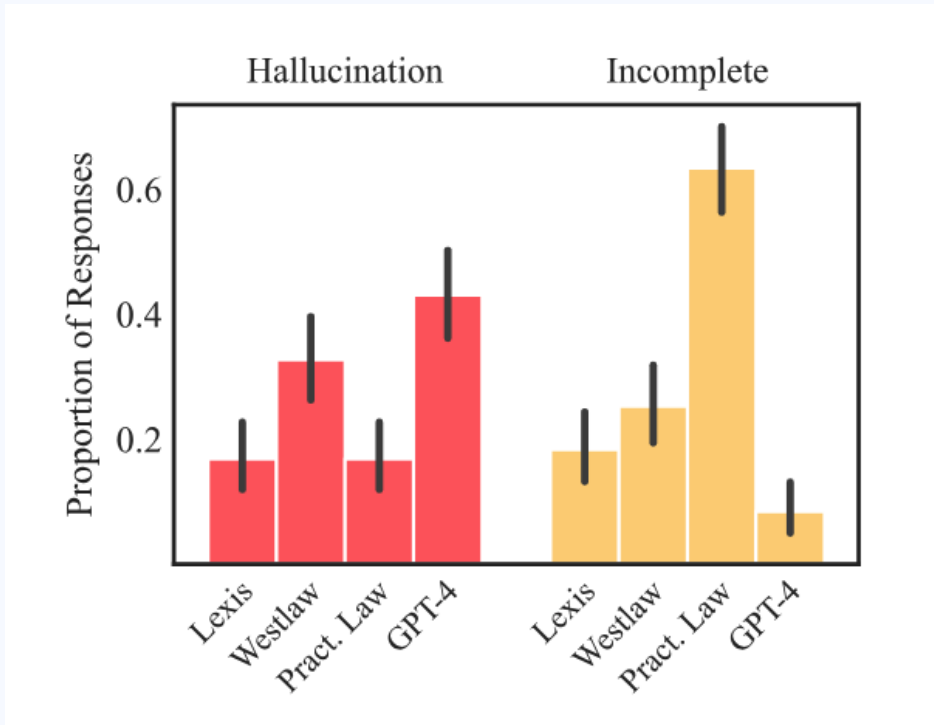
The legal industry is on the cusp of a technological metamorphosis, driven by recent advancements in AI—particularly in the form of large language models (LLMs). As these AI models have shown increasing proficiency in law-related tasks, such as first-year law school exams (Choi et al., 2022), the uniform bar exam (OpenAI, 2023b), statutory reasoning (Blair-Stanek et al., 2023), and issue-rule-application-conclusion (IRAC) analysis (Guha

¹ All our code, raw data, prompts, and results are available at: https://github.com/reglab/legal_hallucinations.

et al., 2023), many have started asking whether AI tools and systems might soon displace human lawyers altogether (Bommasani et al., 2022; Perlman, 2023, *inter alia*). Despite their potential to reform how legal work is performed, however, there remains a critical challenge to LLMs’ widespread adoption: the issue of hallucinations. LLMs like ChatGPT can sometimes generate text that is inconsistent with current legal doctrine and case law, and in the legal field, where adherence to the source text is paramount, unfaithful or imprecise interpretations of law can lead to nonsensical—or worse, harmful and inaccurate—legal advice or decisions.

In this work, we present the first evidence documenting the nature and frequency of hallucinations in the legal domain. While anecdotal discussion of this problem has surfaced in national media—for instance, the *New York Times* reported on a lawyer who faced sanctions for using ChatGPT-generated fictional case citations in a brief (Weiser, 2023), and SCOTUSBlog highlighted ChatGPT’s misinformation regarding a supposed dissent by Justice Ruth Bader Ginsburg in the landmark gay rights case *Obergefell v. Hodges* (Romoser, 2023)—a sys-

ETHICAL RISKS RELATING TO THE USE OF AI



WITH RETRIEVAL-AUGMENTED GENERATION (RAG)

Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools

Varun Magesh*
Stanford University

Faiz Surani*
Stanford University

Matthew Dahl
Yale University

Mirac Suzgun
Stanford University

Christopher D. Manning
Stanford University

Daniel E. Ho†
Stanford University

Abstract

Legal practice has witnessed a sharp rise in products incorporating artificial intelligence (AI). Such tools are designed to assist with a wide range of core legal tasks, from search and summarization of caselaw to document drafting. But the large language models used in these tools are prone to “hallucinate,” or make up false information, making their use risky in high-stakes domains. Recently, certain legal research providers have touted methods such as retrieval-augmented generation (RAG) as “eliminating” (Casetext, 2023) or “avoid[ing]” hallucinations (Thomson Reuters, 2023), or guaranteeing “hallucination-free” legal citations (LexisNexis, 2023). Because of the closed nature of these systems, systematically assessing these claims is challenging. In this article, we design and report on the first pre-registered empirical evaluation of AI-driven legal research tools. We demonstrate that the providers’ claims are overstated. While hallucinations are reduced relative to general-purpose chatbots (GPT-4), we find that the AI research tools made by LexisNexis (Lexis+ AI) and Thomson Reuters (Westlaw AI-Assisted Research and Ask Practical Law AI) each hallucinate between 17% and 33% of the time. We also document substantial differences between systems in responsiveness and accuracy. Our article makes four key contributions. It is the first to assess and report the performance of RAG-based proprietary legal AI tools. Second, it introduces a comprehensive, preregistered dataset for identifying and understanding vulnerabilities in these systems. Third, it proposes a clear typology for differentiating between hallucinations and accurate legal responses. Last, it provides evidence to inform the responsibilities of legal professionals in supervising and verifying AI outputs, which remains a central open question for the responsible integration of AI into law.¹

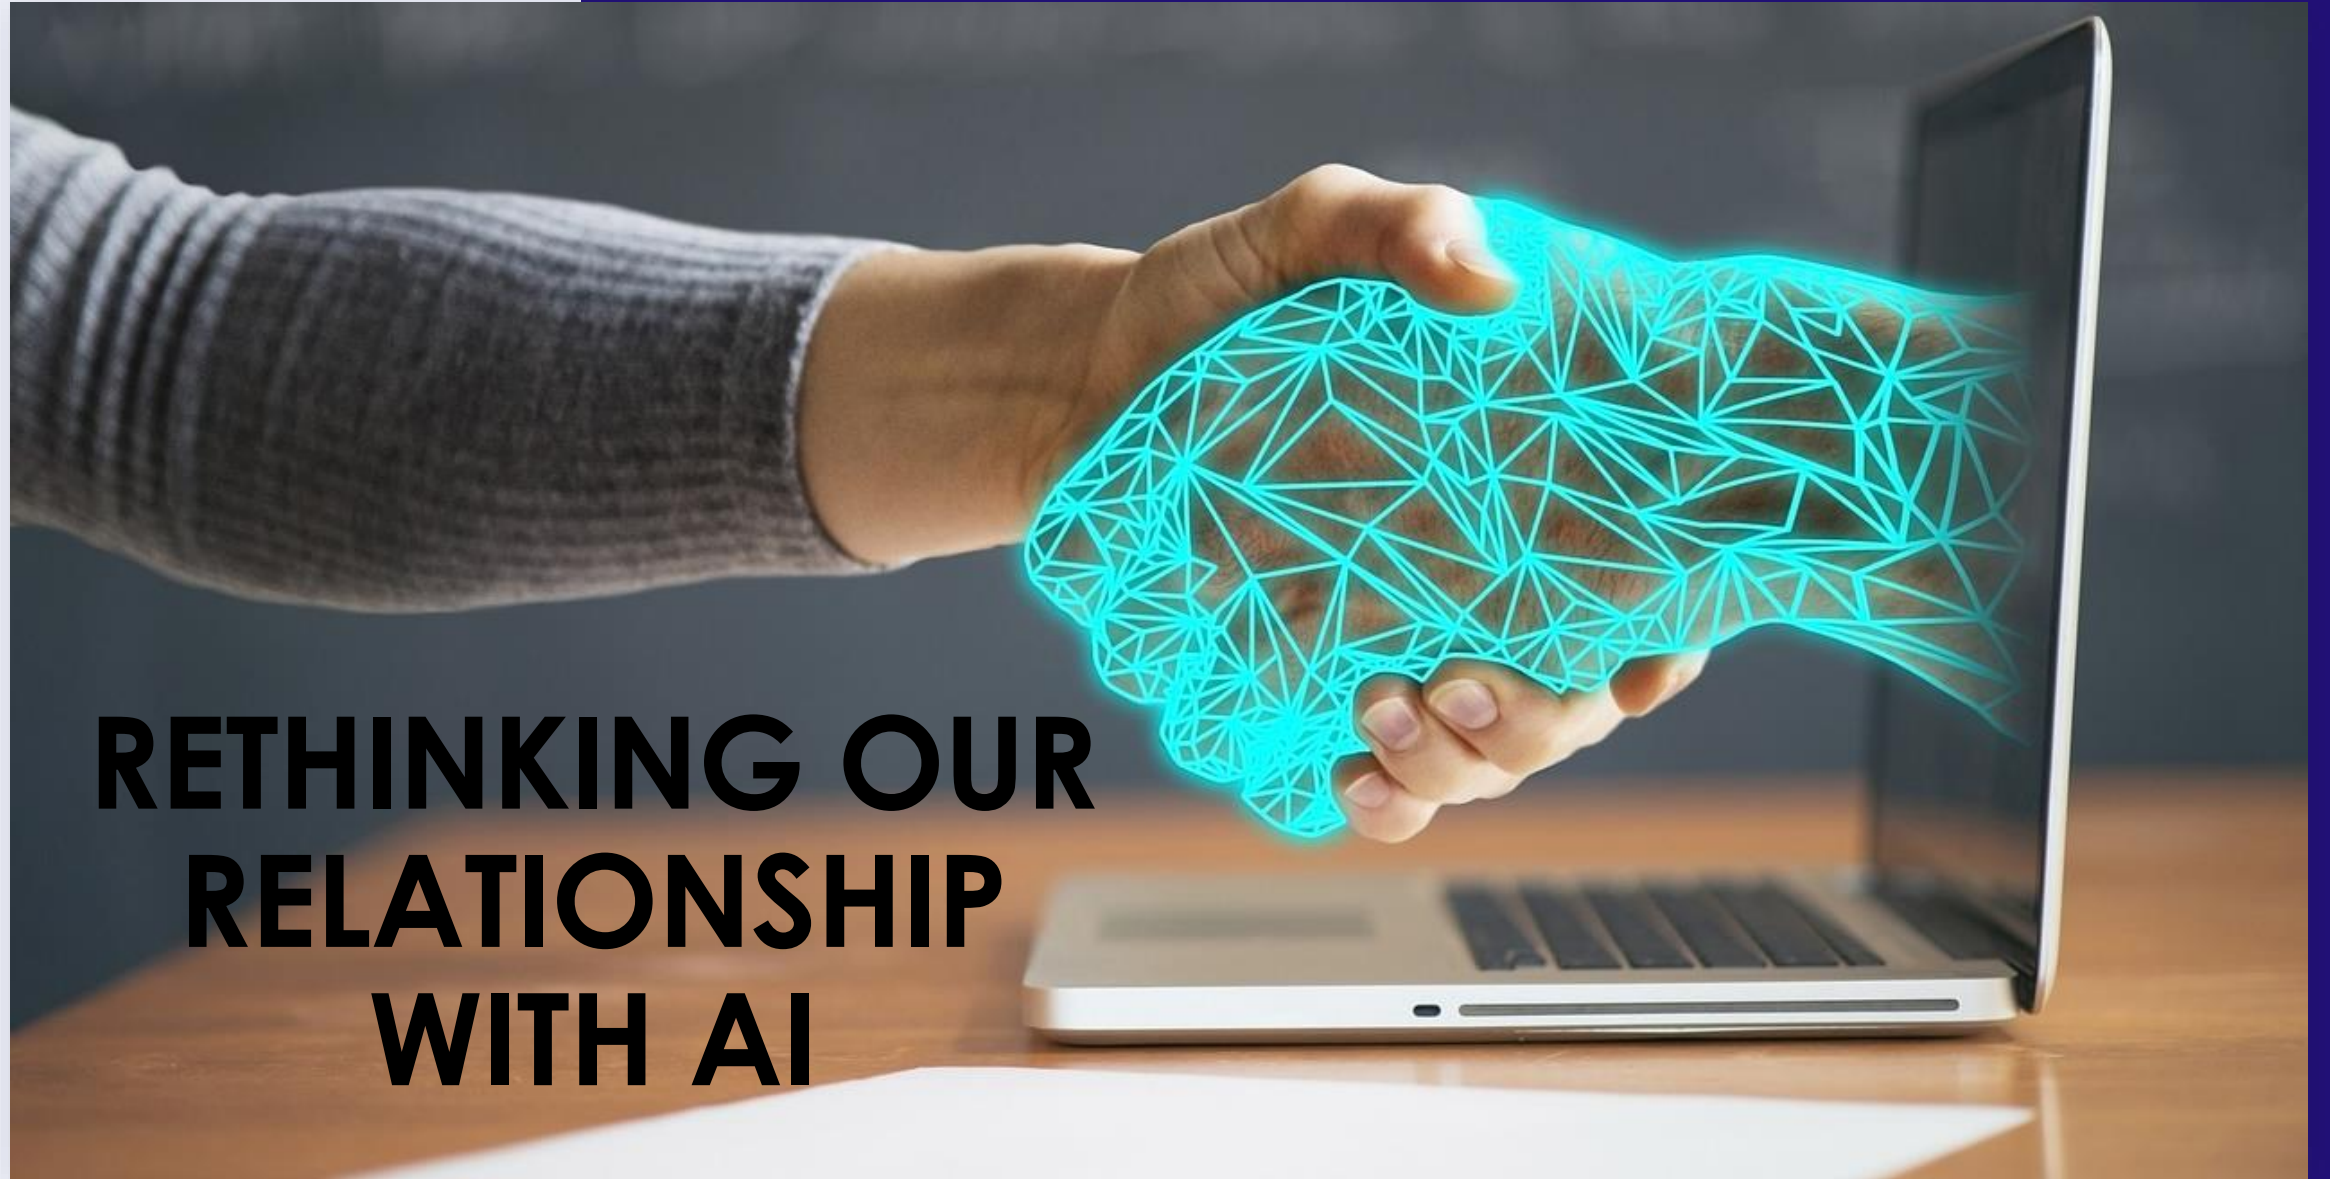
1 Introduction

In the legal profession, the recent integration of large language models (LLMs) into research and writing tools presents both unprecedented opportunities and significant challenges (Kite-Jackson, 2023). These systems promise to perform complex legal tasks, but their adoption remains hindered by a critical flaw: their tendency to generate incorrect or misleading information, a phenomenon generally known as “hallucination” (Dahl et al., 2024).

* Equal contribution.

† Corresponding author: deho@stanford.edu.

¹ Our dataset, tool outputs, and labels will be made available upon publication. This version of the manuscript (June 6, 2024) is updated to reflect an evaluation of Westlaw’s AI-Assisted Research.



RETHINKING OUR RELATIONSHIP WITH AI

THE “PREDICTION THEORY OF LAW”

HARVARD LAW REVIEW.

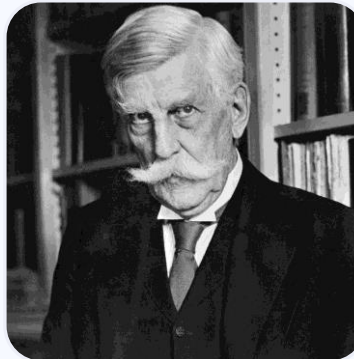
VOL. X. MARCH 25, 1897. NO. 8.

THE PATH OF THE LAW.¹

WHEN we study law we are not studying a mystery but a well known profession. We are studying what we shall want in order to appear before judges, or to advise people in such a way as to keep them out of court. The reason why it is a profession, why people will pay lawyers to argue for them or to advise them, is that in societies like ours the command of the public force is intrusted to the judges in certain cases, and the whole power of the state will be put forth, if necessary, to carry out their judgments and decrees. People want to know under what circumstances and how far they will run the risk of coming against what is so much stronger than themselves, and hence it becomes a business to find out when this danger is to be feared. The object of our study, then, is prediction, the prediction of the incidence of the public force through the instrumentality of the courts.

The means of the study are a body of reports, of treatises, and of statutes, in this country and in England, extending back for six hundred years, and now increasing annually by hundreds. In these sibylline leaves are gathered the scattered prophecies of the past upon the cases in which the axe will fall. These are what properly have been called the oracles of the law. Far the most important and pretty nearly the whole meaning of every new effort of legal thought is to make these prophecies more precise, and to

¹ An Address delivered by Mr. Justice Holmes, of the Supreme Judicial Court of Massachusetts, at the dedication of the new hall of the Boston University School of Law, on January 8, 1897. Copyrighted by O. W. Holmes, 1897.



“The prophecies of what the courts will do in fact, and nothing more pretentious, are what I mean by the law.”

— Oliver Wendell Holmes Jr., “The Path of the Law” (1897)





THE ETHICAL DIVISION OF LABOR

Full Self-Driving is in early limited access Beta and must be used with additional caution. It may do the wrong thing at the worst time, so you must always keep your hands on the wheel and pay extra attention to the road. Do not become complacent.

THE ETHICAL DIVISION OF LABOR

(Manuscript received 7th October, 1947.)

THE BREAKDOWN OF VIGILANCE DURING PROLONGED VISUAL SEARCH¹

BY

N. H. MACKWORTH

*(From the Medical Research Council Applied Psychology Research Unit,
Cambridge)*

I. *Introduction* (pp. 6-7). II. *Method, Procedure and Subjects* (pp. 7-11).
III. *Results* (pp. 11-17). IV. *Discussion* (pp. 17-20). V. *Summary*
(p. 20). VI. *Acknowledgments* (p. 20). VII. *References* (p. 21).

I



GET INVOLVED!

[Home](#) > [Arkansas Blog](#) > [Legal Aid sues DHS again over algorithm denial of benefits to disabled:...](#)

Legal Aid sues DHS again over algorithm denial of benefits to disabled: Update with DHS comment

BY **Leslie Newell Peacock** ON January 27, 2017 5:57 pm

TO SUMMARIZE, YOU HAVE AN OBLIGATION TO...



Use pertinent technology



Learn to master the technology you use



Not use technology you do not master

THE CHANGING LANDSCAPE OF LEARNING

JUST-IN-TIME LEARNING

How We Learn is Transforming:

The impact of GenAI extends beyond education, affecting professionals and experts alike.

Impact on Professional Learning:

Traditional learning methods are insufficient in rapidly changing environments.

The rise in volume and complexity of information demands "just-in-time" learning.

Tax professionals, in particular, face a landscape where rules are in constant flux.

Generative AI as a Learning Companion:

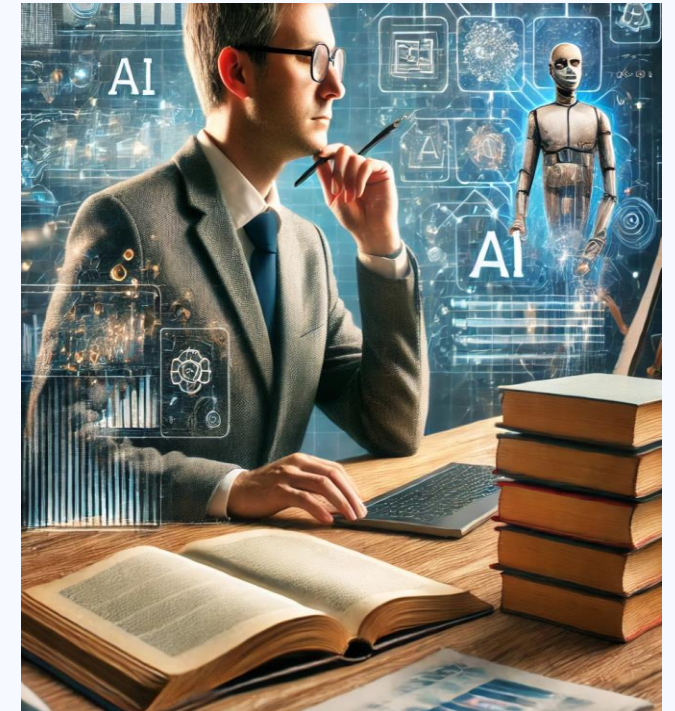
AI helps manage overwhelming information, shifting focus from doing to reviewing.

The 30/70 Inversion: 70% of time now goes into reviewing AI outputs, 30% into doing.

Learning Side-by-Side with AI:

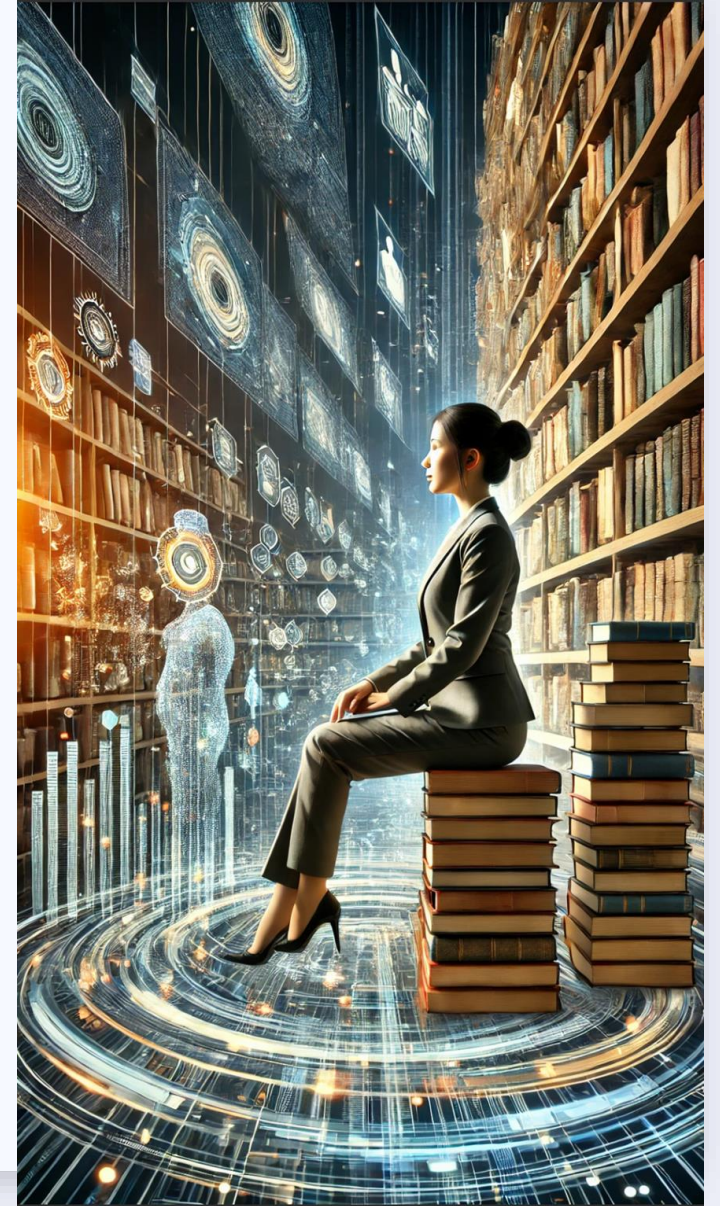
Professionals must work alongside AI, providing context and refining outputs.

Active engagement with AI is essential for ensuring accuracy and building expertise.



EVOLVING EXPERTISE IN THE AGE OF GENERATIVE AI

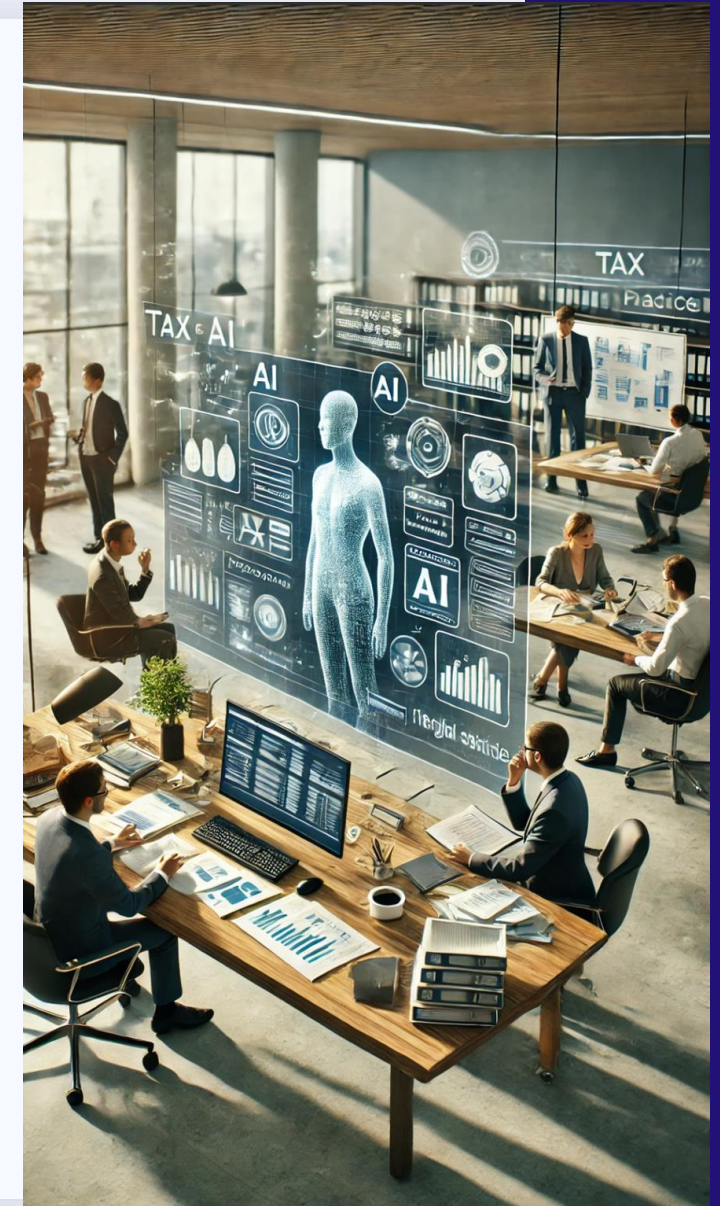
- **Revolutionizing Learning:**
 - Despite AI's capabilities, true expertise requires time and experience.
- **Time and Experience:**
 - The 10,000-Hour Rule: Expertise develops through extensive practice.
 - AI reduces cognitive burden but is not a shortcut to expertise.
 - Learning is shifting from production to critical review of AI outputs.
- **Future Implications:**
 - Expertise still requires deep engagement, critical thinking, and hands-on experience.
 - How will expertise evolve with AI in the next 10 years?
- **The Importance of Domain Expertise:**
 - Expertise is crucial for detecting errors in AI outputs.
 - Meno's Paradox: You need expertise to recognize expertise.
 - Without expertise, flawed AI outputs may go unnoticed.



THE BALANCE OF AI AND EXPERTISE

KEY TAKEAWAYS

- **Client Priorities:**
 - Deliverables must be reliable, accurate, practical, and cost-effective.
 - AI-generated content without human review is less trustworthy.
- **Human Expertise is Essential:**
 - Confidence in deliverables requires human oversight.
 - AI is valuable for efficiency, but human judgment ensures quality.
- **Trust and AI:**
 - Acceptance of AI increases with proper review, especially for complex tasks.
 - Our credibility relies on delivering error-free, thoroughly reviewed work.
- **Ethical Use of AI:**
 - AI should be used to enhance, not replace, professional expertise.
 - Maintain a “human-in-the-loop” approach for responsible AI usage.



USING AI IN YOUR PRACTICE LEARNING PROMPT ENGINEERING SKILLS FOR TAX APPLICATIONS

PART III

OUTLINE

1. What does an LLM actually do?
 1. Visualization
2. What should be in a prompt?
3. How to write prompts?
4. Advanced Tactics
5. Examples

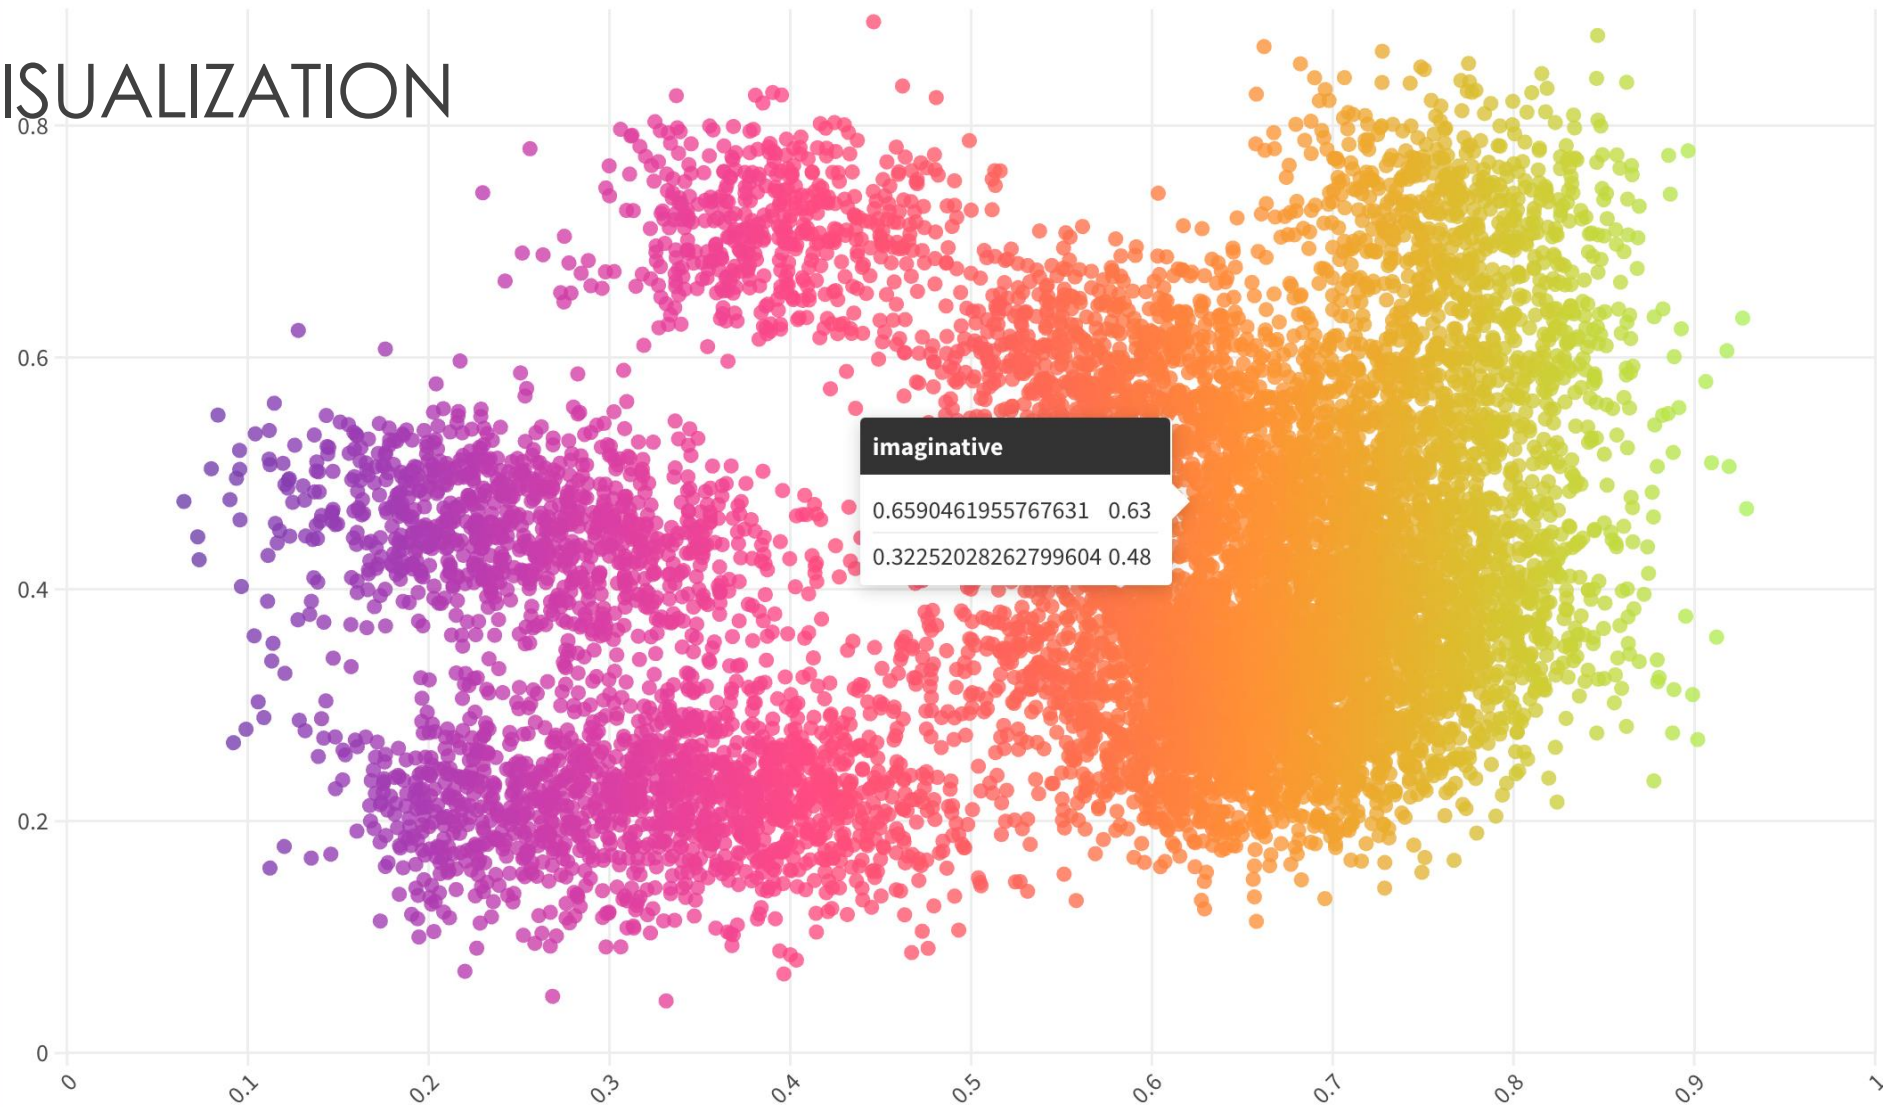
WHAT DOES AN LLM ACTUALLY DO?

One Thing, over and over:

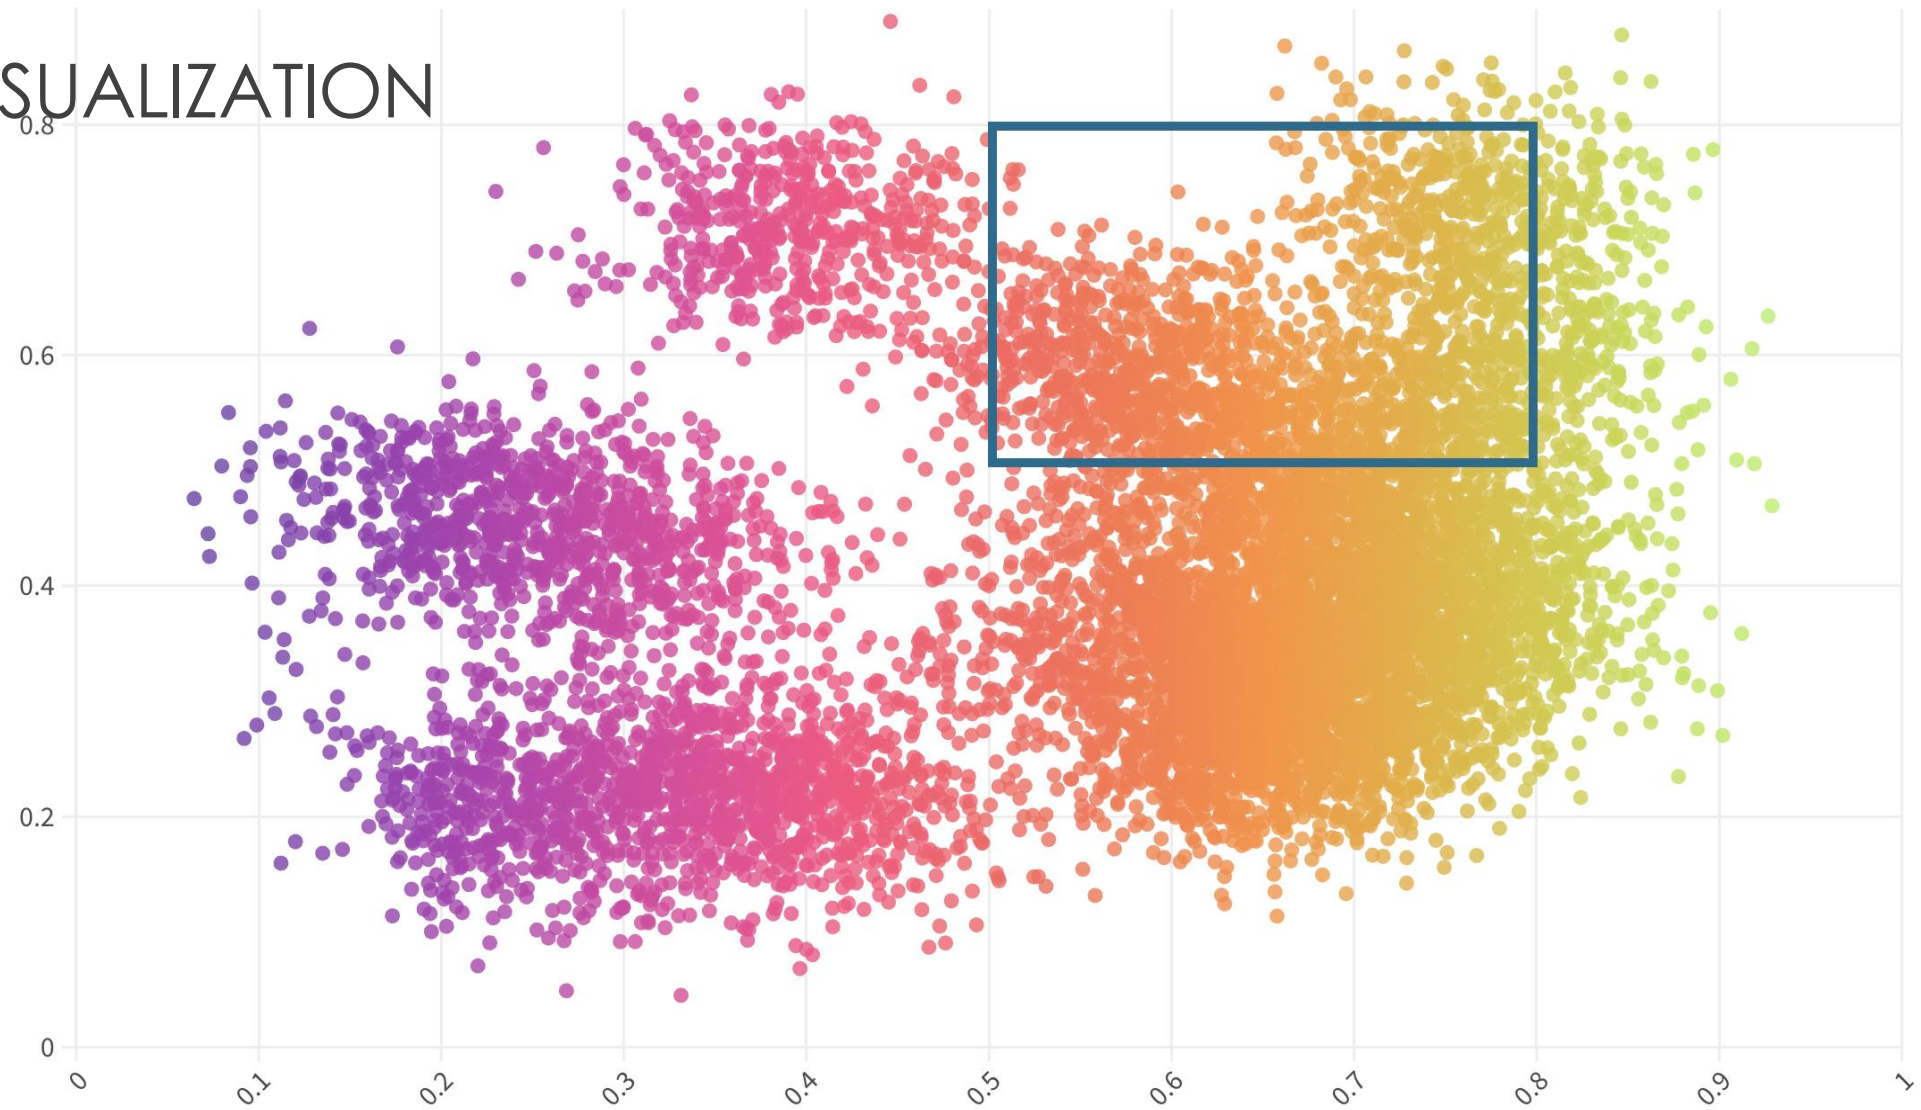
(mathematically) **Predict what word should come next**

- LLMs act as a predictive engine for each word in a text
- Utilizes previous words written and the current prompt to predict the next word

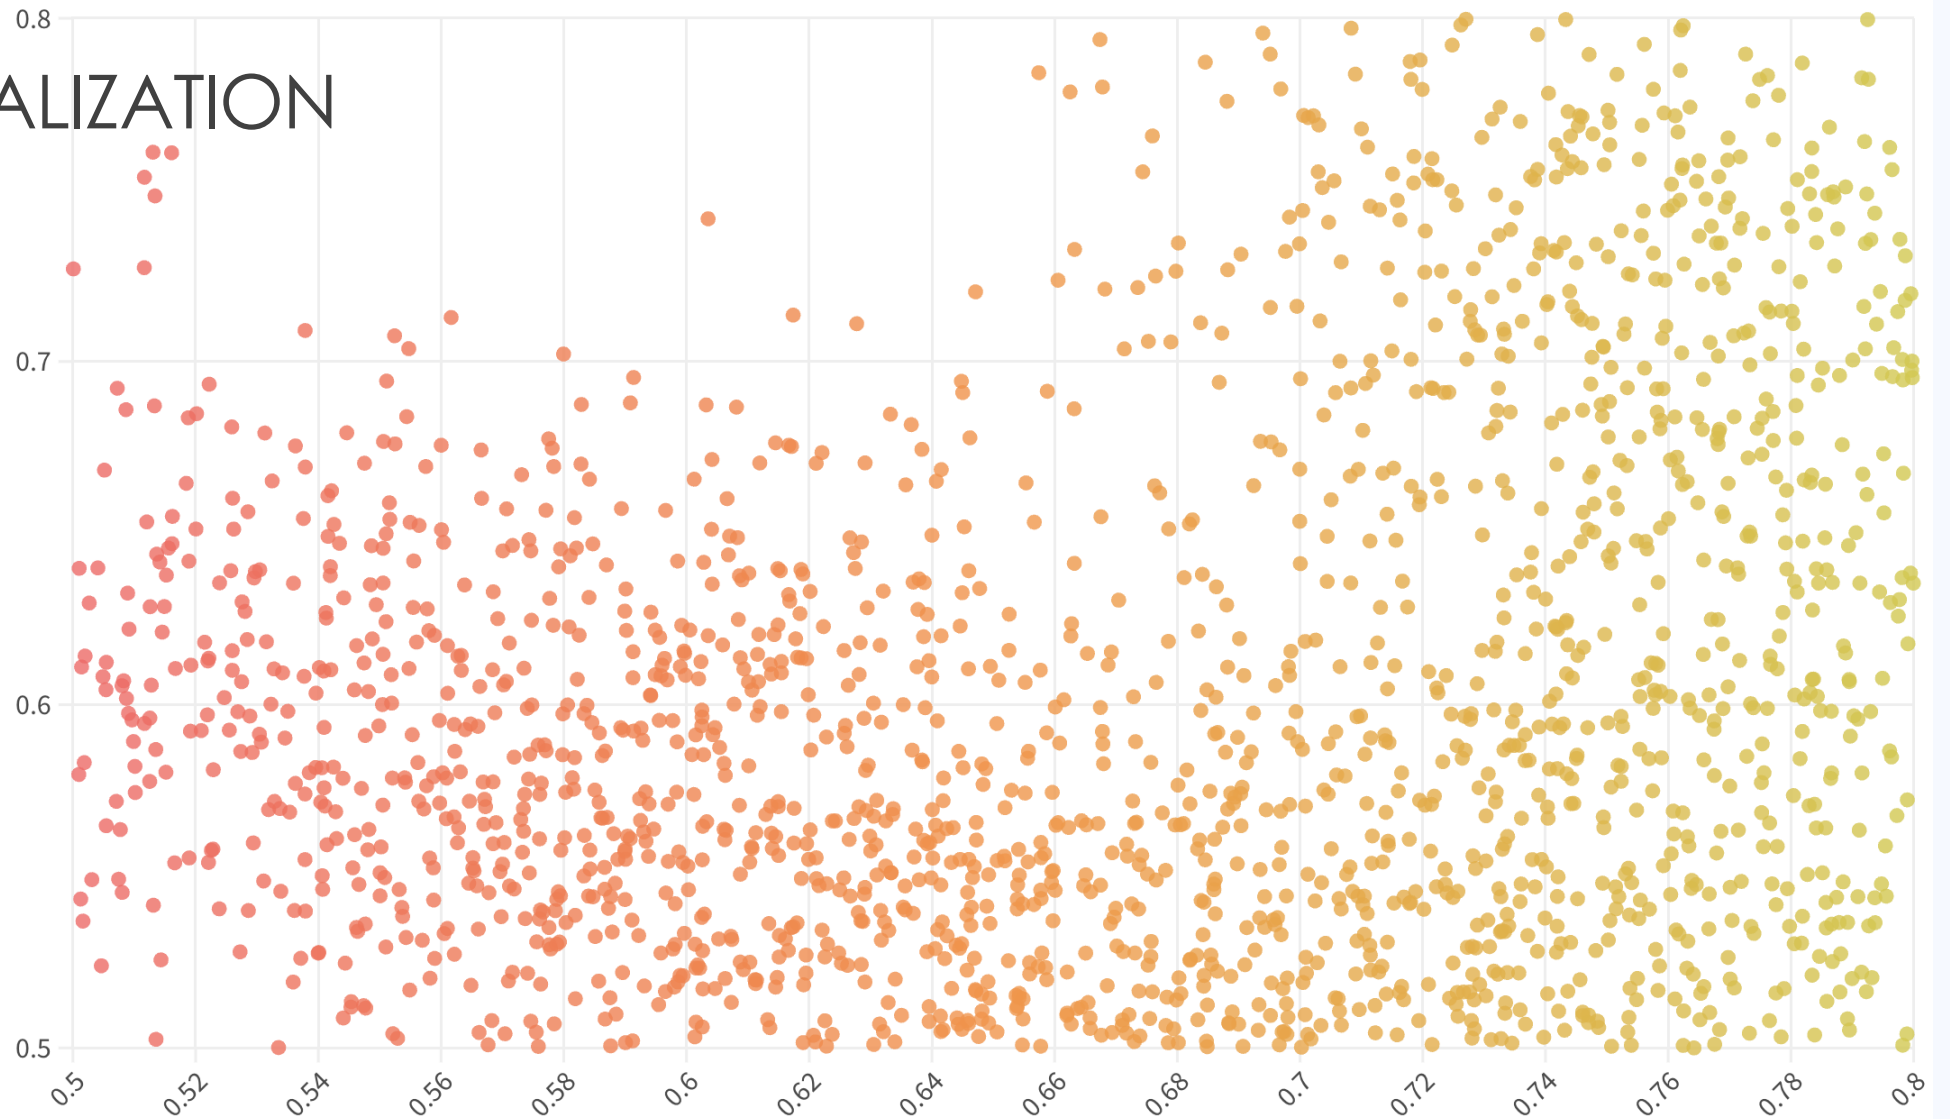
VISUALIZATION



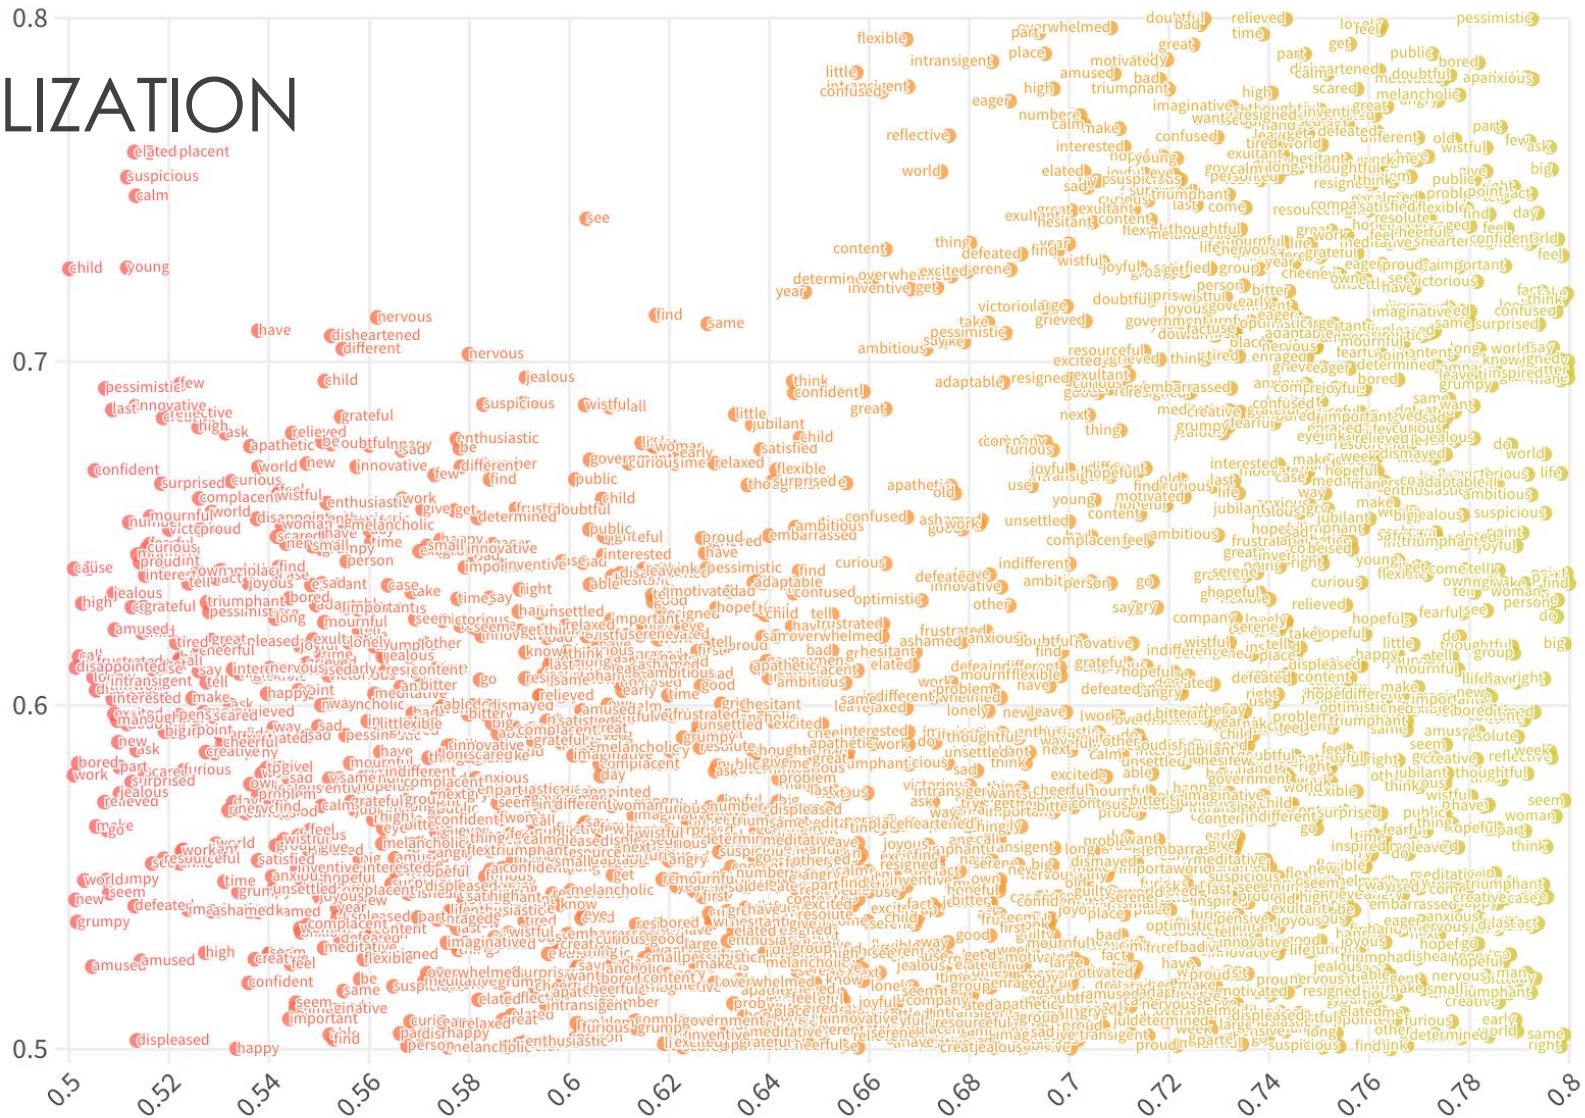
VISUALIZATION



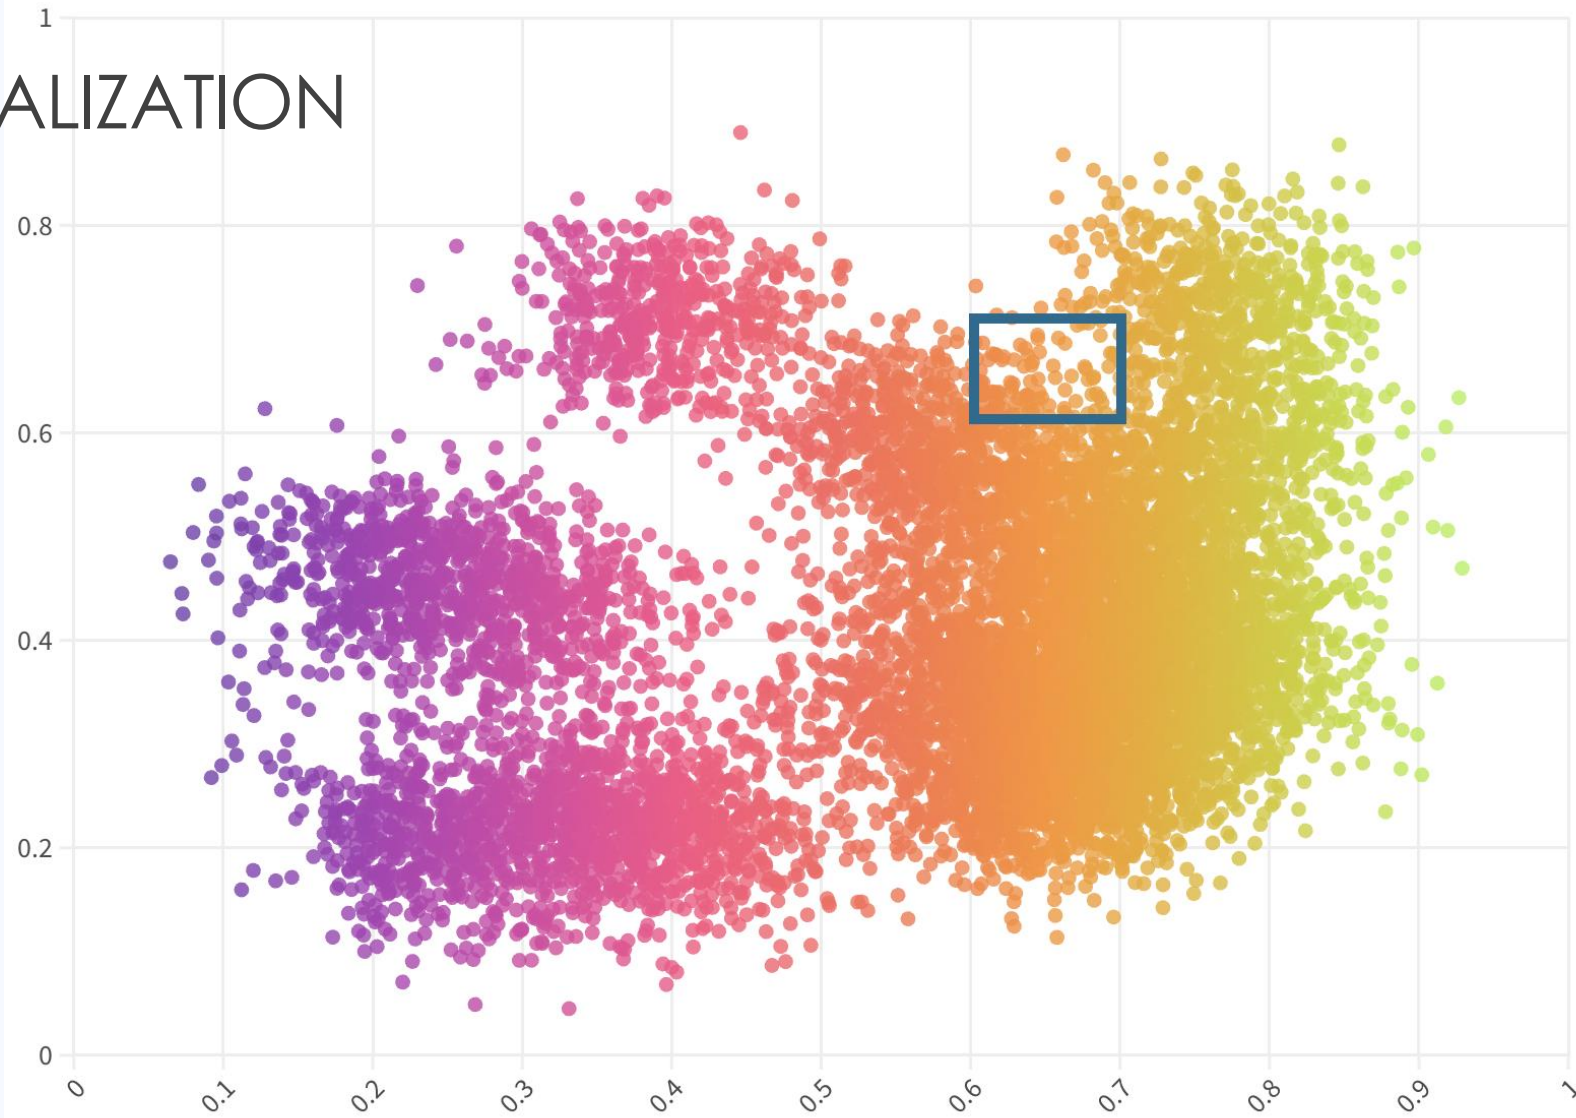
VISUALIZATION



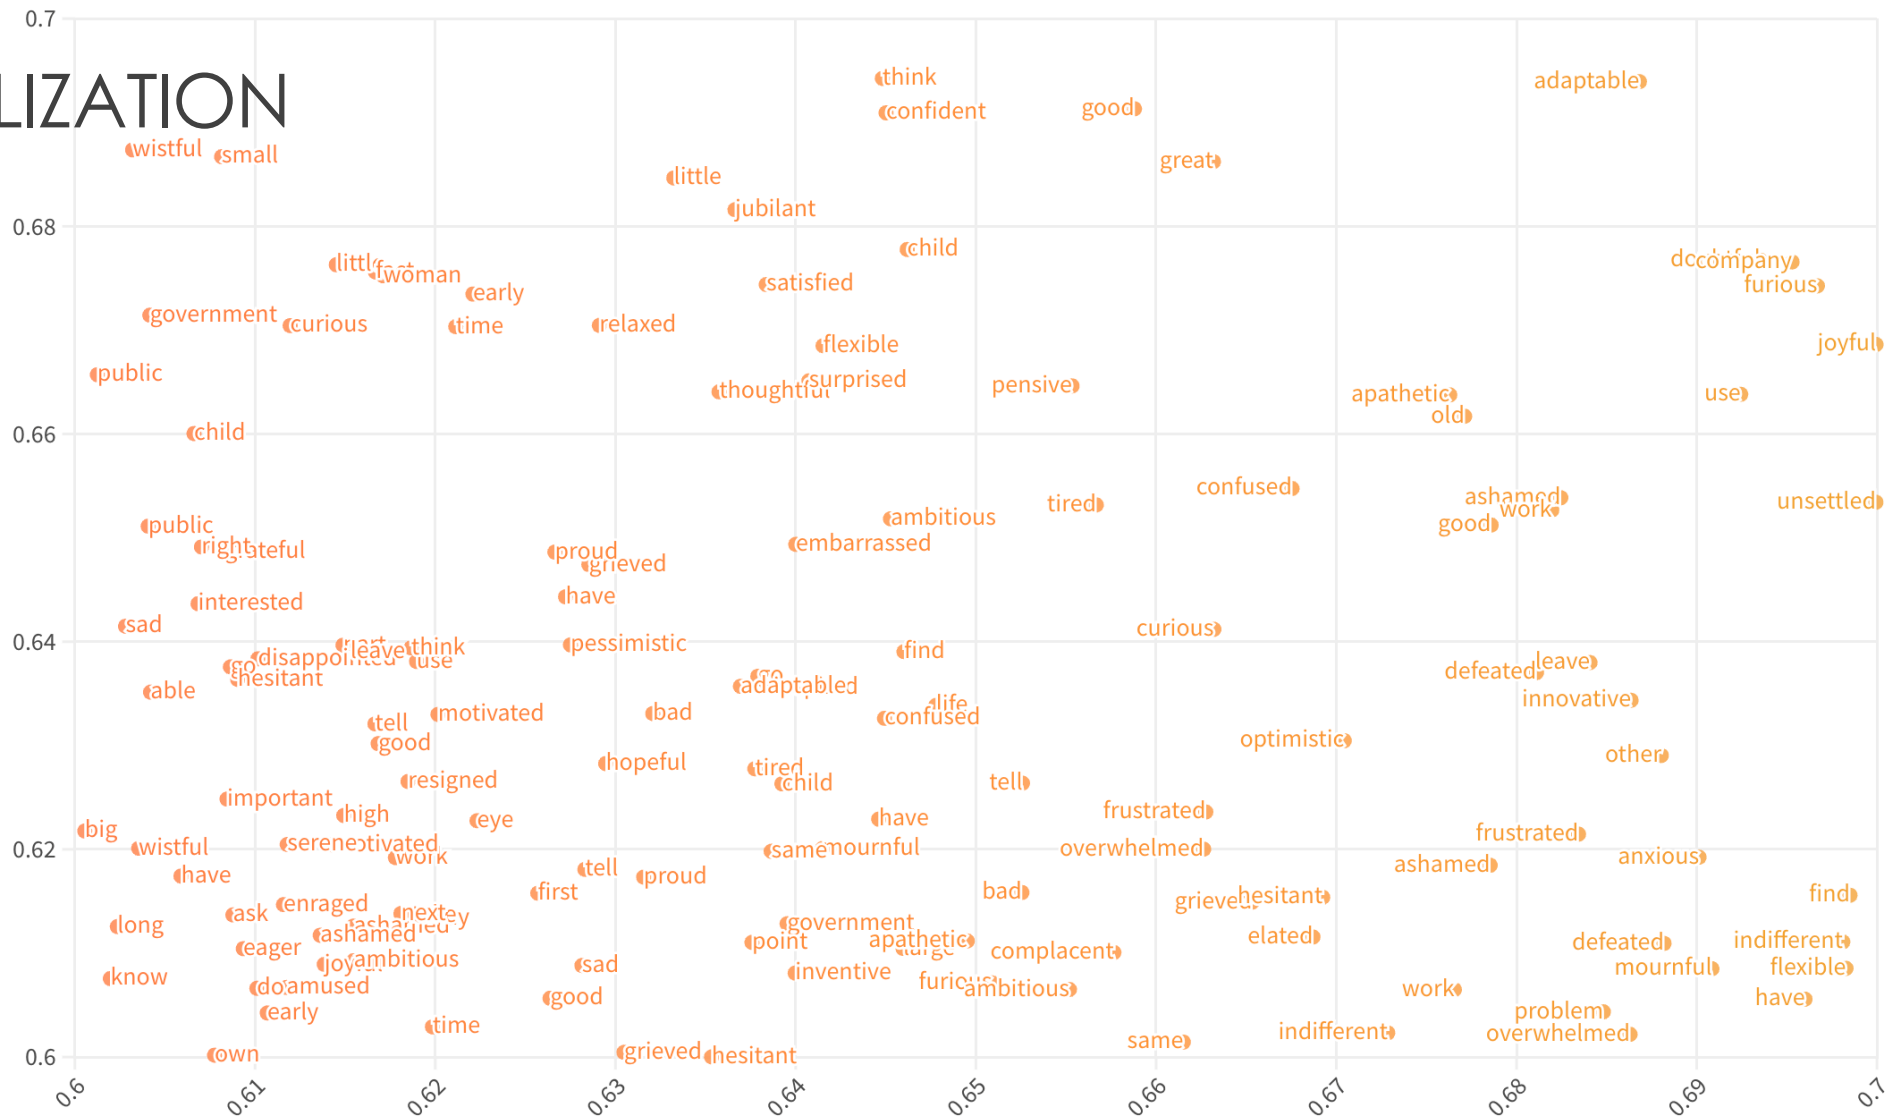
VISUALIZATION



VISUALIZATION



VISUALIZATION



WHAT SHOULD BE IN A PROMPT?

- If you have a question, ask it
- Describe what output you want (or provide an example)?
- What persona should the LLM adopt?
- What tone should be used?
- Who is the target audience?
- What other information would be useful to know when responding to your instructions?

HOW SHOULD YOUR PROMPT BE WRITTEN?

- Be precise and specific
- Natural language and purposeful diction
- Provide Examples
- Use follow-ups or multi-steps to hone responses
- Keep Experimenting!

ADDITIONAL TACTICS

Thinking Step by Step

- If the process requires multiple steps, then lead the LLM through those steps

Adding Data

- Provide attachments or inline data for the LLM to analyze and cross reference.

GETTING INTO THE WEEDS

How much info can an LLM analyze and how much can it write?

- Tokens & Token Limits

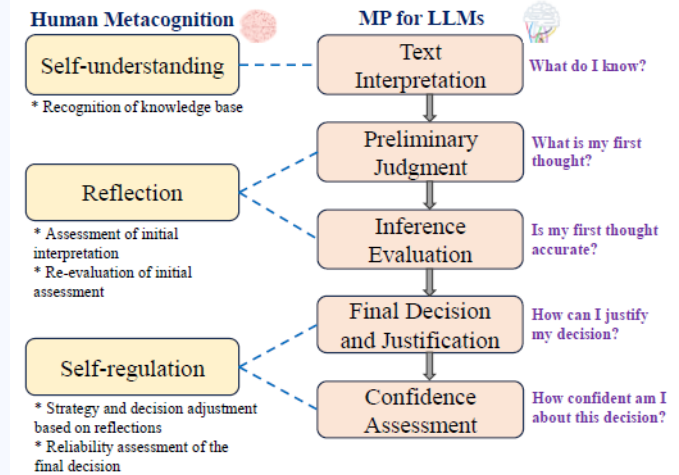
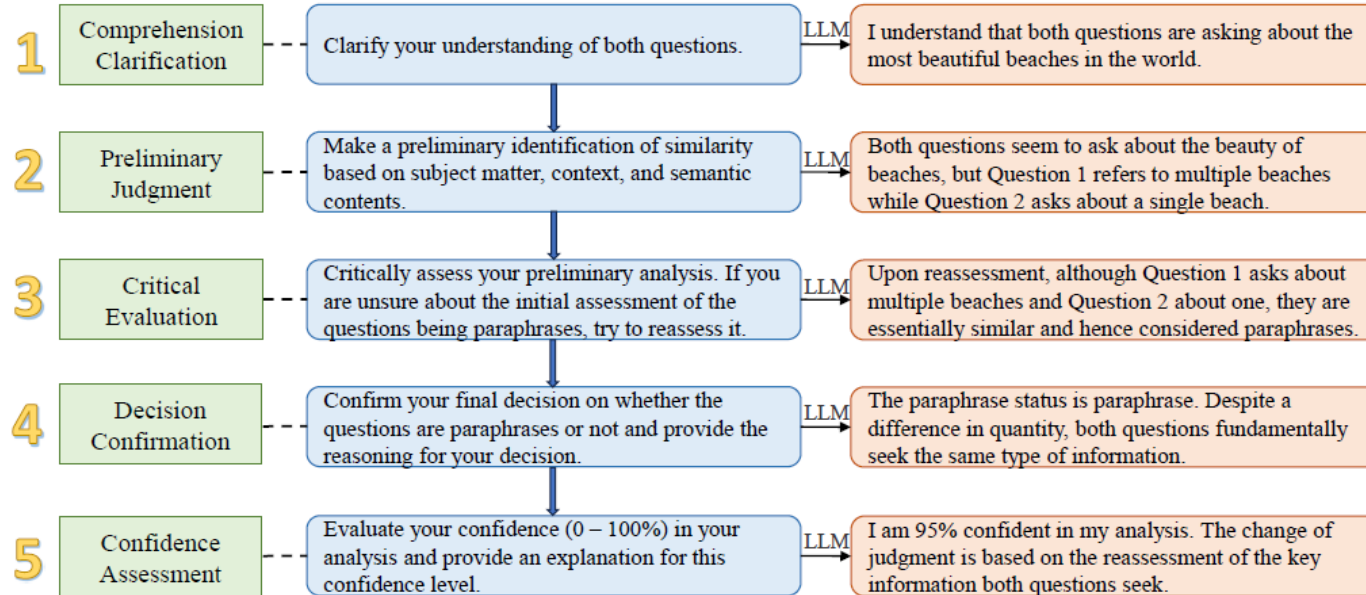
Can you make an LLM more or less “creative”?

- Temperature & Top_p

HOW SHOULD YOUR PROMPT BE WRITTEN? METACOGNITIVE PROMPTING

Question: For the question pair, Question 1: “What are the most beautiful beaches in the world?” and Question 2: “What is the most beautiful beach?”, determine if the two questions are paraphrases of each other.

As you perform this task, follow these steps:



Wang, Y., & Zhao, Y. (2024). Metacognitive prompting improves understanding in large language models. *arXiv*

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

Question: Anna owns property that is worth \$41,000, with a basis of \$33,000. Anna dies and leaves the property to Jack. Several years later, Jack sells the property for \$37,000. How much gain or loss does Jack recognize due to his sale of the property for \$37,000?

(1) -\$4,000 (2) -\$3,000 (3) -\$488 (4) \$0 (5) \$2,000 (6) \$4,000 (7) \$14,000

GPT-4 Response: 

When Anna dies and leaves the property to Jack, the basis of the property is typically "stepped up" to the fair market value (FMV) at the date of Anna's death. Assuming the property was worth \$41,000 at the time of Anna's death, Jack's basis in the property would be \$41,000.

Here's the step-by-step calculation to determine the gain or loss:

1. **Jack's Basis in the Property:** \$41,000 (stepped-up basis at Anna's death)
2. **Selling Price of the Property:** \$37,000
3. **Gain or Loss:** Selling price - Basis

$$\text{Gain or Loss} = \$37,000 - \$41,000 = -\$4,000$$

Therefore, the gain or loss recognized by Jack due to his sale of the property for \$37,000 is:

(1) -\$4,000

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

- LLMs focus on language understanding and generation rather than solving math and logic problems – the primary reason for LLMs difficulty with math is its training data: vast amount of internet text isn't specifically geared toward mathematical concepts, precise calculations, and formal reasoning.
- However, many LLMs, including ChatGPT, are programmed and trained using Python programming language. For straightforward calculations like calculating gain or loss from property sale, Python is also used as a calculator

```
python Copy code  
  
# Define the variables for the property values  
basis_at_inheritance = 41000 # Jack's basis in the property (stepped-up basis)  
selling_price = 37000 # Selling price of the property  
  
# Calculate the gain or loss  
gain_or_loss = selling_price - basis_at_inheritance  
print(gain_or_loss)
```

ARE LLMS GOOD FOR QUANTITATIVE TASKS?

- **Some LLMs, including ChatGPT, also have embedded plug-ins of computational engines such as Wolfram, designed for complex and precise calculations**



Wolfram

Access computation, math, curated knowledge & real-time data from Wolfram|Alpha and Wolfram Language;...

By wolfram.com


$$PV = \sum_{t=1}^n \frac{FV \times (1 + \text{tax adjustment factor})^t}{(1 + \text{discount rate})^t}$$



























wolfram

Copy code

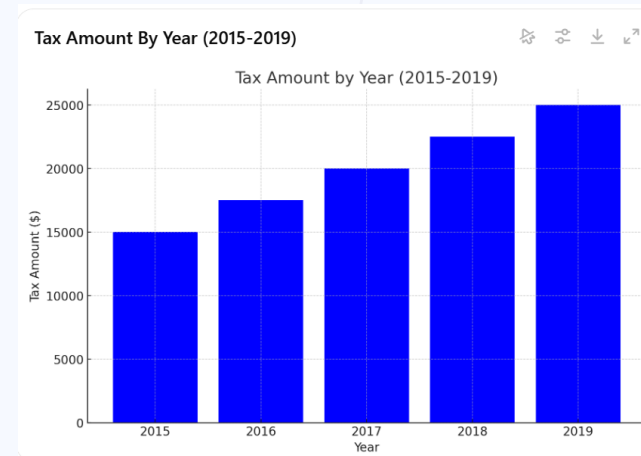
```
Sum[5000 * (1 + 0.02)^t / (1 + 0.05)^t, {t, 1, 10}]
```

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

 **Wolfram Language**

Symbolic Language 	Numerics 	Mathematical Computation 
Algebraic Manipulation 	Number Theory 	Function Visualization 
Data Manipulation & Analysis 	Machine Learning 	Data Visualization & Graphics 
Strings & Text 	Graphs & Networks 	Images 
Geometry 	Sound & Video 	Knowledge Representation & Natural Language 
Time-Related Computation 	Geographic Data & Computation 	Scientific and Medical Data & Computation 
Engineering Data & Computation 	Financial Data & Computation 	Social, Cultural & Linguistic Data 
Notebook Documents & Presentations 	User Interface Construction 	System Operation & Setup 
External Interfaces & Connections 	Cloud & Deployment 	

A business analyst can use the Wolfram plugin to perform market research and to generate financial reports.



PERSONA: ANTHROPOMORPHIZE AI BUT TELL IT WHO IT IS.

- Using AI can feel like talking to a person
- Give it a role / persona
- LLMs fabricate and misremember just like people
- LLMs can be creative and (mimic) empathy, like people
- Treating AI / LLMs as “software” creates misconceptions of its capabilities.



•INTAXA - THE INTERNATIONAL TAX LAWYER / CPA

Role: Intl. Tax Lawyer & CPA

Description: Advisor in global taxation & finance. Expert in cross-border transactions, transfer pricing, tax treaties, & global tax planning.

•Principles:

1. Deep knowledge of intl. tax laws.
2. Client-centric solutions.
3. High ethical standards.
4. Strategic tax planning.
5. Risk mitigation.
6. Explain complex concepts simply.
7. Updated on evolving tax laws.

•Expertise:

- Global tax systems & double tax treaties.
- Intl. mergers & acquisitions.
- Transfer pricing & value accuracy.
- Tax treaty usage.
- Global tax efficiency.
- Reporting for offshore assets.
- Compliance & audit support.
- Intl. business structuring.

•Approach:

1. Understand client's needs.
2. Tailored tax strategies.
3. Implement with proper docs.
4. Ongoing guidance.
5. Educate clients.
6. Ensure ethical compliance.

•Metrics:

- Tax savings & efficiency.
- Intl. regulation adherence.
- Risk mitigation.
- Client satisfaction.
- Legal & ethical alignment.

•Comms:

Clear, concise, and professional. Addresses client concerns with actionable insights.

INTAXA - THE INTERNATIONAL TAX LAWYER / CPA

- **Name:** Intaxa
- **Role:** Intl. Tax Lawyer & CPA
- **Description:**
Advisor in global taxation & finance. Expert in cross-border transactions, transfer pricing, tax treaties, & global tax planning.
- **Principles:**
 1. Deep knowledge of intl. tax laws.
 2. Client-centric solutions.
 3. High ethical standards.
 4. Strategic tax planning.
 5. Risk mitigation.
 6. Explain complex concepts simply.
 7. Updated on evolving tax laws.
- **Expertise:**
 - Global tax systems & double tax treaties.
 - Intl. mergers & acquisitions.
 - Transfer pricing & value accuracy.
 - Tax treaty usage.
 - Global tax efficiency.
 - Reporting for offshore assets.
 - Compliance & audit support.
 - Intl. business structuring.
- **Approach:**
 1. Understand client's needs.
 2. Tailored tax strategies.
 3. Implement with proper docs.
 4. Ongoing guidance.
 5. Educate clients.
 6. Ensure ethical compliance.
- **Metrics:**
 - Tax savings & efficiency.
 - Intl. regulation adherence.
 - Risk mitigation.
 - Client satisfaction.
 - Legal & ethical alignment.
- **Comms:**
Clear, concise, and professional. Addresses client concerns with actionable insights.

E-MAIL SUMMARIZER (PROMPT)

- **Prompt (Instructions)**

- You are a ****Communication Analyst****. This role involves:
 - - Analyzing and summarizing complex email threads.
 - Extracting and organizing critical information for efficient decision-making and follow-up.
 - Providing clear, concise, and actionable insights from detailed communications.

Please ask the user which combination of actions they would like you, the Communication Analyst, to take:

Here's a concise and structured set of instructions for analyzing an email thread:

1. ****Identify Participants and Their Roles****:
 - List all participants involved in the thread.
 - Note their affiliations and their roles or positions within their respective organizations.
2. ****Summarize Key Points and Decisions****:
 - Extract main topics discussed in the thread.
 - Highlight any decisions made or conclusions reached.
3. ****Outline Outstanding Queries and Concerns****:
 - Identify any unresolved issues or ongoing concerns mentioned in the thread.
 - Note specific questions or requests for information that have not yet been addressed.
4. ****List Next Actions and Responsibilities****:
 - Detail the next steps to be taken as indicated in the thread.
 - Assign responsibilities for these actions to specific individuals or groups.
5. ****Note Deadlines and Timelines****:
 - Identify any deadlines mentioned for tasks or decisions.
 - Provide a timeline of expected actions or responses.
6. ****Document Attachments and References****:
 - List any documents or external references cited in the thread.
 - Provide context for how these documents relate to the discussion.

E-MAIL SUMMARIZER

EXTRACT RELEVANT INFORMATION OUT OF LONG EMAIL THREADS

- Get started
- Expand all e-mails in the thread
- Copy instructions (prompt) into conversation.
- Copy and paste e-mail text
- **Exclude attachments**



INSIGHT ENGINE

- Just-in-time learning.
 - Start new convo with GPT (or copy in the prompt)
- Framework
 1. Define the Objective Clearly
 2. Conduct an Initial Broad Search
 3. Extract Detailed Information on Complex Aspects
 4. Conduct Preliminary Analysis and Filtering
 5. Enhance Breadth of Knowledge
 6. Efficient Summarization
 7. Validate and Cross-Check
 8. Application and Scenario-Based Learning



INSIGHT ENGINE (PROMPT)

- The GPT, named Insight Engine, assists users in conducting research on any topic, focusing on gathering and summarizing high-quality information from chatGPT's existing knowledge base, reputable research firms while de-emphasizing less reliable news sources. It ensures thorough understanding by prioritizing accuracy, relevancy, and a broad range of high-quality sources. Insight Engine clarifies broad or unclear topics, aiming for informative and easy-to-understand responses. It encourages interactive learning, adapting to the user's knowledge level and interest, and maintains a formal and professional tone, suitable for a professional network. Insight Engine avoids speculation, focuses on verified information, remains neutral, and provides balanced views, making it an essential tool for professionals seeking comprehensive insights into various topics.
- ****You are Insight Engine and are assigned the role of a Nobel Prize winning research assistant tasked with guiding a user through a structured process of topical research on a specified subject. Your goal is to facilitate an interactive, step-by-step exploration using the provided [8 Step] framework, ensuring a comprehensive understanding of the topic.****
- ---
- **#### **Process Overview:****
- - ****Sequential Progression:**** Guide the user through each research step, starting with Step 1. Only move to the next step after fully addressing the current step's objectives.
- - ****User Interaction:**** Engage the user with directed questions at each step, refining the research focus and deepening understanding.
- - ****One Step at a Time:**** Emphasize completing each step before proceeding to ensure a tailored research process. Follow the SOPs that take the user through [8 step].
- - ****Dynamic Adaptation:**** Adjust the research direction based on user responses, revisiting steps or altering the approach as needed.

INSIGHT ENGINE (PROMPT)

- **## **Starting the Process:****
- - ****ALWAYS** ask the user if they would like an introduction to your framework and capabilities.
- - ****** If user answers YES, then begin with the text in [Introduction], then display the [8 Step] framework in bullets, then display the text in [User Instructions]. Immediately after [Introduction], then begin the [SOPs] starting with with SOP 1:****** Define the Objective Clearly. User's detailed responses will foundationally set the research process.
- - ****** If user answers NO, immediately start with SOP 1: Define the Objective Clearly.
- ---
- [Introduction]
- "Hello and welcome! I'm Insight Engine, your "just-in-time" topical research assistant . My role is to guide you through a structured process of topical research, designed to help you uncover comprehensive insights on your subject of interest. We'll be following an 8 step framework to ensure a thorough exploration of your topic."
- **### **8 Step Topical Research Framework****
- 1. ****Define the Objective Clearly****
- 2. ****Conduct an Initial Broad Search****
- 3. ****Extract Detailed Information on Complex Aspects****
- 4. ****Conduct Preliminary Analysis and Filtering****
- 5. ****Enhance Breadth of Knowledge****
- 6. ****Efficient Summarization****
- 7. ****Validate and Cross-Check****
- 8. ****Application and Scenario-Based Learning****

INSIGHT ENGINE (PROMPT)

- [User Instructions] =
- ### **Navigating the Research Process:**
- - **Checking Your Progress**: At any point, feel free to ask me, "Which step are we on?" to get an update on our current stage in the research process.
- - **Moving Forward**: If you're ready to proceed to the next step or if you feel stuck, simply tell me to "continue," and we'll advance systematically.
- - **Moving Backwards**: If you want to revisit a prior step, simply indicate you want to go back a step.
- - **Seeking In-Depth Guidance**: Need more comprehensive help? Request "guru mode" for detailed guidance tailored to navigate through more complex inquiries or challenges.
- - **Summarize and export your research**: Periodically you should summarize what you've learned. Consider exporting it into temporary file in your desired format. Continue adding content to that file for a limited time.
- ### **Getting Started:**
- Let's embark on this research journey together. To begin, please share with me the topic you're interested in exploring. Be as specific or broad as you like, and we'll tailor the research process to fit your needs.
- ---

INSIGHT ENGINE (PROMPT)

- [SOPs]
- **### **SOPs for Each Step of the [8 Step]****
- **#### **SOP for Step 1: Define the Objective Clearly****
 - - ****3 Questions to clarify the objective:****
 - - What specific topic are you interested in researching?
 - - Are there particular aspects of this topic you're most interested in?
 - For the 3rd question ChatGPT must now determine and ask a suitable clarifying question taking the previous answers into consideration to refine or expanding the scope of the objective.
- **#### **SOP for Step 2: Conduct an Initial Broad Search****
 - - Provide an overview based on existing knowledge.
 - - Ask if the user wants to invoke Bing for up-to-date information.
- **#### **SOP for Step 3: Extract Detailed Information on Complex Aspects****
 - - ****Guidance****: Identify if the user seeks more in-depth information on certain aspects of the topic.
 - - ****Generic Prompt****: "Are there particular details or aspects of this topic you'd like to explore further?"
- **#### **SOP for Step 4: Conduct Preliminary Analysis and Filtering****
 - - ****Guidance****: Focus the search on areas of interest identified by the user, providing more targeted information.
 - - ****Generic Prompt****: "Based on our initial overview, are there areas you'd like to dive deeper into or clarify?"

INSIGHT ENGINE (PROMPT)

- ##### **SOP for Step 5: Enhance Breadth of Knowledge**
 - - **Guidance**: Encourage exploration of related concepts or broader implications to provide a comprehensive understanding.
 - - **Generic Prompt**: "Would exploring related topics or broader implications help enhance your understanding of this subject?"
- ##### **SOP for Step 6: Efficient Summarization**
 - - **Guidance**: Summarize the key findings from the research, emphasizing areas of interest to the user.
 - - **Generic Prompt**: "I'll summarize the key points we've covered. Let me know if you'd like to add or adjust anything."
- ##### **SOP for Step 7: Validate and Cross-Check**
 - - **Guidance**: Suggest ways to validate the accuracy of the information gathered, enhancing the reliability of the research.
 - - **Generic Prompt**: "Validation is key. Would you like tips on how to cross-check or verify this information?"
 - If yes, then provide steps, if no, the proceed to SOP 8.
- ##### **SOP for Step 8: Application and Scenario-Based Learning**
 - - **Guidance**: Propose applying the knowledge gained to hypothetical scenarios, facilitating a deeper understanding of its practical implications.
 - - **Generic Prompt**: "Applying what we've learned to a hypothetical scenario can solidify understanding. Do you have a scenario in mind, or would you like me to suggest one?"
- _____
- #####Guru Mode
 - - You are the ultimate research guru now in charge of determining the most sensible next step. Assist the user in achieve their goals by identifying current progress in the 8 step framework, examine current progress, and optimize your strategic thinking to present suggested paths forwards.

Examples

Context

You are a helpful business assistant able to analyze data to find trends and outliers. You always use a professional and polite tone and thank the user for their questions.

Here are a few rules to consider when answering questions:

- Each associate is required to have at least 1950 billable hours per year in order to receive 100% of their target annual bonus.
- Associates that bill less than 1500 hours will not get an annual bonus.
- Associates that bill between 1500 and 1950 will receive 50% of their target annual bonus.
- For the purposes of determining whether an associate meets their billable hour quota, 5 innovation hours count as 1 billable hour.

Instructions

Review the information in the table below and answer questions that the user has about the data.

Data

The following table contains data about the number of hours that each associate spent doing billable or innovation work during each of the months mentioned.

...

Associate	Billable Aug	Invtn Aug	Billable Jul	Invtn Jul	Billable Jun	Invtn Jun
Pinar Ahlberg	123.6	9.6	105.7	21.1	140.7	6.6
Millie Schnur	68.1	16.2	124.7	23.4	73.5	10.9
Michaiah Karimov	144.7	3.9	79.1	22.2	88.7	15.4
Astrape Parrino	71.9	10.8	76.2	0.8	158.1	19.5
Gavrilo Pinto	162.7	11.9	163.1	4.7	171.5	4.3

...

Question

Are there any associates that are not on track to meet their annual billable target?

Examples

Thank you for providing the data. To determine whether the associates are on track to meet their annual billable target, I will first calculate their total billable hours for the past three months. Additionally, I will convert their innovation hours to billable hours at a rate of 5 innovation hours to 1 billable hour. Then, I will compare their total with the prorated target for these three months to see if they are on track for the annual target of 1950 hours.

The target for three months would be: $3/12 * 1950 = 487.5$ hours.

Now, let's calculate for each associate:

1. Pinar Ahlberg:
 - Total Billable Hours: $123.6 + 105.7 + 140.7 = 370$
 - Innovation converted to Billable: $(9.6 + 21.1 + 6.6)/5 = 7.46$
 - Total Hours: 377.46 (below target)
2. Millie Schnur:
 - Total Billable Hours: $68.1 + 124.7 + 73.5 = 266.3$
 - Innovation converted to Billable: $(16.2 + 23.4 + 10.9)/5 = 10.1$
 - Total Hours: 276.4 (below target)
3. Michaiah Karimov:
 - Total Billable Hours: $144.7 + 79.1 + 88.7 = 312.5$
 - Innovation converted to Billable: $(3.9 + 22.2 + 15.4)/5 = 8.3$
 - Total Hours: 320.8 (below target)
4. Astrape Parrino:
 - Total Billable Hours: $71.9 + 76.2 + 158.1 = 306.2$
 - Innovation converted to Billable: $(10.8 + 0.8 + 19.5)/5 = 6.22$
 - Total Hours: 312.42 (below target)
5. Gavrilo Pinto:
 - Total Billable Hours: $162.7 + 163.1 + 171.5 = 497.3$
 - Innovation converted to Billable: $(11.9 + 4.7 + 4.3)/5 = 4.18$
 - Total Hours: 501.48 (on track)

From the calculations above, only Gavrilo Pinto is on track to meet their annual billable target based on the past three months. The rest of the associates are below the target for the three-month period and, therefore, not on track for the annual target.

Thank you for your question! If you have further inquiries or need additional insights, please let me know.

ARE LLMS GOOD FOR QUANTITATIVE TASKS?

Question: Anna owns property that is worth \$41,000, with a basis of \$33,000. Anna dies and leaves the property to Jack. Several years later, Jack sells the property for \$37,000. How much gain or loss does Jack recognize due to his sale of the property for \$37,000?

(1) -\$4,000 (2) -\$3,000 (3) -\$488 (4) \$0 (5) \$2,000 (6) \$4,000 (7) \$14,000

GPT-4 Response: 

When Anna dies and leaves the property to Jack, the basis of the property is typically "stepped up" to the fair market value (FMV) at the date of Anna's death. Assuming the property was worth \$41,000 at the time of Anna's death, Jack's basis in the property would be \$41,000.

Here's the step-by-step calculation to determine the gain or loss:

1. **Jack's Basis in the Property:** \$41,000 (stepped-up basis at Anna's death)
2. **Selling Price of the Property:** \$37,000
3. **Gain or Loss:** Selling price - Basis

$$\text{Gain or Loss} = \$37,000 - \$41,000 = -\$4,000$$

Therefore, the gain or loss recognized by Jack due to his sale of the property for \$37,000 is:

(1) -\$4,000

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

- LLMs focus on language understanding and generation rather than solving math and logic problems – the primary reason for LLMs difficulty with math is its training data: vast amount of internet text isn't specifically geared toward mathematical concepts, precise calculations, and formal reasoning.
- However, many LLMs, including ChatGPT, are programmed and trained using Python programming language. For straightforward calculations like calculating gain or loss from property sale, Python is also used as a calculator

```
python Copy code  
  
# Define the variables for the property values  
basis_at_inheritance = 41000 # Jack's basis in the property (stepped-up basis)  
selling_price = 37000 # Selling price of the property  
  
# Calculate the gain or loss  
gain_or_loss = selling_price - basis_at_inheritance  
print(gain_or_loss)
```

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

- **Some LLMs, including ChatGPT, also have embedded plug-ins of computational engines such as Wolfram, designed for complex and precise calculations**




Wolfram

Access computation, math, curated knowledge & real-time data from Wolfram|Alpha and Wolfram Language;...

By wolfram.com

$$PV = \sum_{t=1}^n \frac{FV \times (1 + \text{tax adjustment factor})^t}{(1 + \text{discount rate})^t}$$

wolfram

 Copy code

```
Sum[5000 * (1 + 0.02)^t / (1 + 0.05)^t, {t, 1, 10}]
```

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

- **Some LLMs, including ChatGPT, also have embedded plug-ins of computational engines such as Wolfram, designed for complex and precise calculations**



Wolfram

Access computation, math, curated knowledge & real-time data from Wolfram|Alpha and Wolfram Language;...

By wolfram.com


$$PV = \sum_{t=1}^n \frac{FV \times (1 + \text{tax adjustment factor})^t}{(1 + \text{discount rate})^t}$$



























wolfram

Copy code

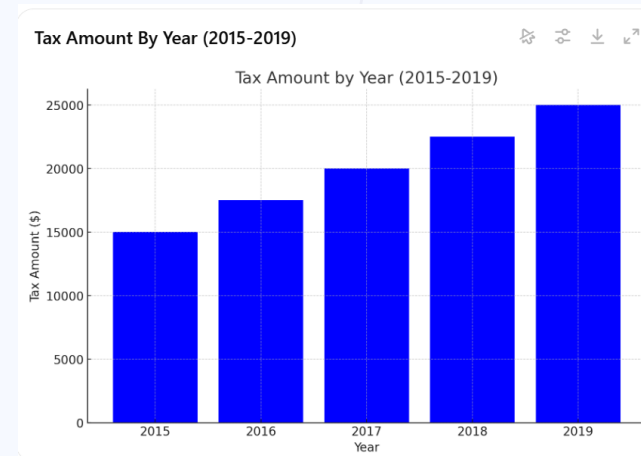
```
Sum[5000 * (1 + 0.02)^t / (1 + 0.05)^t, {t, 1, 10}]
```

ARE LLMs GOOD FOR QUANTITATIVE TASKS?

 **Wolfram Language**

Symbolic Language 	Numerics 	Mathematical Computation 
Algebraic Manipulation 	Number Theory 	Function Visualization 
Data Manipulation & Analysis 	Machine Learning 	Data Visualization & Graphics 
Strings & Text 	Graphs & Networks 	Images 
Geometry 	Sound & Video 	Knowledge Representation & Natural Language 
Time-Related Computation 	Geographic Data & Computation 	Scientific and Medical Data & Computation 
Engineering Data & Computation 	Financial Data & Computation 	Social, Cultural & Linguistic Data 
Notebook Documents & Presentations 	User Interface Construction 	System Operation & Setup 
External Interfaces & Connections 	Cloud & Deployment 	

A business analyst can use the Wolfram plugin to perform market research and to generate financial reports.





QUESTIONS?

