

White Paper Report

Hey AI, Give Me Investment Advice



How can I help you today?



Hallucination Mode High ▾



The use of artificial intelligence through large language models (LLMs) such as ChatGPT, Gemini, and Claude are quickly being integrated into everyday life. It is estimated that there are almost 2.5 billion active users between ChatGPT, Gemini and Claude, up from ~1.5 billion in 2025.¹ Users are finding recipes, planning travel, summarizing data, formulating emails, and doing research. However, we are also witnessing a growing trend of investors using LLMs to assist with basic financial tasks (e.g., budgeting, taxes) as well as picking stocks, creating asset allocation models and analysis of current investment allocations.

In this quarterly white paper, we want to help investors understand how LLMs actually work, where they can be useful, and where they fall short, particularly when it comes to investment decisions. Our goal is not to discourage the use of new technology. It is to give investors a clearer picture of what these tools can and cannot do and explain why we believe constructing and managing a long-term asset allocation continues to require human expertise, client knowledge, and ongoing accountability.

Language Models, Not Math Models

At their core, LLMs are exactly what their name suggests, a model that is built to work with language, not numbers. To understand where they belong in investment decisions, it helps to start with what these tools do. A large language model is a text predictor. It has been trained on vast amounts of writing from books, articles, websites, blogs and online forums. From that material, it learns the patterns of how words tend to follow other words across different contexts.

When you ask an LLM a question, it is not thinking, calculating, or looking up facts. It is choosing the most likely next words based on what it has seen before. The response sounds like a thoughtful answer because the model has learned what thoughtful investment advice sounds like, not because it has any of its own. Building a long-term asset allocation model is a math problem. An LLM is not built to solve math problems, only to produce language that sounds like the answer to one.

It Looks Like Optimization. It is Not!

The difference between language and math becomes clear when you look at how a real investment portfolio is built. When investors and advisors talk about portfolio optimization they are usually referring to mean-variance optimization. This is a mathematical method developed by Harry Markowitz in the 1950s that earned him a Nobel Prize, and it is still the foundation of how most professional portfolios are constructed today.

The method takes three inputs: expected returns, expected volatility, and how correlated the investments

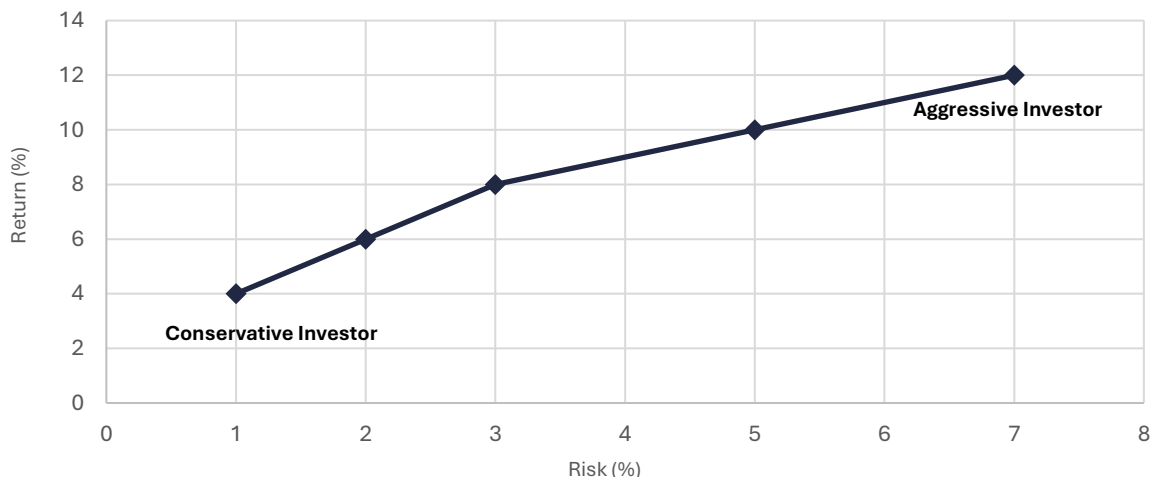
are expected to be in the future. From those inputs, it produces a range of portfolios that aim to deliver the best possible return for a given level of risk. This is known as the efficient frontier (**Chart 1**). Then a Financial Advisor selects the portfolio on that range that best fits the investor's long-term goals, time horizon, and ability to tolerate losses. The result is an allocation built on math and shaped by judgment about the investor. An LLM cannot generate the inputs the math requires, it does not run the calculation and cannot make the judgment that has to follow.

Chart 1

Sample Efficient Frontier Portfolio

This data is for representational purposes only to show that as an investor takes on more risk, they should receive more return.

Data Source: Verdense Capital Advisors.



Built on the Past, Not the Present!

Even setting aside the math problem, there is a second limitation when using an LLM for portfolio advice. The model is always working from yesterday's information. Each LLM has a training cutoff date (**Table 1**), which is the date after which it has seen no new information. Interest rate changes, earnings results, geopolitical events, and shifts in valuations that occurred after that date are all unknown to the model.

When an investor asks an LLM for an investment recommendation today, the answer may be based on a view of the world that is six months, a year, or several years out of date. The model does not bridge that gap,

and it answers with the same confidence whether the data behind it is current or stale.

That matters for long-term asset allocation in two ways. First, any investment decision should start with where markets are right now, and the LLM may not know that. Second, the assumptions baked into the model's training may no longer hold. For example, what bonds were expected to do, how different investments tended to move together, what valuations once looked reasonable can shift quickly in ways the model has not seen. The output is current, but the thinking is not. The response is written with the same confidence either way, and the investor has no way to tell the difference.

Table 1

How current is the answer when you ask AI?

Flagship and default free models as of June 2026. Each company runs several models at once, and free tiers often run a different version than paid, so cutoff dates shift with every release. Even the most current model was trained on a view of the world that is months, or longer, out of date.

AI Model	Training Cutoff	How Far Behind Today
ChatGPT GPT-5.5	August 2025 ²	About 10 months
Claude Opus 4.8	January 2026 ³	About 5 months
Gemini 3.1 Pro	January 2025 ⁴	About 17 months

A Second Opinion That is Not One!

One trend we have seen is investors using LLMs as a second opinion on their current portfolio, often by pasting their current holdings into the LLM and asking the model to evaluate it. We realize that it could make sense if the model has access to the actual holdings, the feedback should be more useful.

Unfortunately, giving the model more data does not change what the model is. It is still not performing real optimization, its information may still be out of date, and the response still reads like a real analysis while no real analysis has been done.

What the model still does not see is everything that truly shapes a good investment allocation. It does not know the investor's tax situation, liquidity needs, whether they have concentrated stock positions, or how that investor tends to behave when markets fall. The output may look like a thoughtful second opinion, but it is built on the same limitations explained earlier.

There is also a practical concern worth monitoring. When portfolio details and personal information are pasted into an LLM, that information may be stored on a third-party server and may be retained, used for training, or exposed in a security breach.

Where the Information Comes from Matters

LLMs are trained on the internet, and most of the investment content online is written for retail audiences. This includes trading forums, personal finance blogs, marketing material, and outdated articles that were never updated. Peer-reviewed academic finance, institutional research, and the work that professional investors rely on makes up a much smaller share of what the model has seen. A lot of that work sits behind subscriptions, including Bloomberg, FactSet, and Morningstar Direct.

To the model, a Reddit thread and a peer-reviewed paper are both just text, and they sit in the training data at roughly the same weight. A useful reminder of how retail investors can move markets came in January 2021 when a coordinated effort on Reddit's WallStreetBets

forum sent GameStop's share price up more than 1,600% in a matter of weeks. The stock then collapsed and many investors who bought late were left with significant losses. The posts, the hype and the rationalizations from that period now sit in the same training data the model pulls from when an investor asks it about portfolio construction.

The risk grows when you consider how investors use these tools. Most ask follow-up questions until the model agrees with them. Remember, the model is built to be helpful and agreeable. What you end up with is confirmation bias dressed up as investment advice, an answer that feels validated when the user has talked the model into it.

What you end up with is confirmation bias dressed up as investment advice.

Hallucinations and the Herd

Another limitation worth noting is that LLMs can produce what the industry calls hallucinations. These are confident, specific-sounding answers that are simply wrong. This may include made-up fund symbols, expense ratios that are inaccurate and historical returns that were not real. An LLM model delivers all of it with the same confidence as real information making the errors difficult for an everyday user to spot.

LLMs Do Not Know You

Know Your Client is a foundational principle of investment advising. It is a regulatory requirement and the starting point for any portfolio decision that has a chance of being right for a client. An LLM model does not know an investor, it only knows what was typed in the chat box that day, which is rarely the full picture.

Real portfolio management is also an ongoing job, not a one-time recommendation. A real advisor monitors the portfolio over time, decides when to rebalance, looks for opportunities to harvest tax losses, and coordinates with the investor's CPA and estate attorney. An LLM

Now consider that millions of investors are asking the same handful of models the same kinds of questions and acting on similar answers. We have seen herd behavior move markets before without any help from AI. Add a tool that pushes flawed analysis out to everyone at once and you have a recipe for amplified volatility in markets that are already prone to it.

produces an answer in a single moment and then forgets that the conversation happened.

Perhaps most importantly, an LLM is not a registered investment advisor. It is not regulated and it has no fiduciary duty to act in your best interest. A professional advisor does, and that obligation shapes every recommendation.

What is an AI hallucination? When an AI system produces information that sounds confident and plausible but is wrong or made up.

How Verdence Uses AI

It may be easy to read this white paper as a broad criticism of artificial intelligence. That is not our view. AI plays a meaningful role in the work we do, and we believe it will continue to drive efficiency in ways that translate into a better experience for our clients. The question is not whether to use it, but where does it add value and where does it not belong?

From an investment strategy perspective, LLMs help our team move faster through research and pull insights from large data sets in less time than before. It can summarize dense regulatory filings, earnings reports, and other long documents, which lets us spend more time thinking and less time reading. LLMs do not form our views or draw conclusions, but they help organize thoughts and surface questions worth asking. The arguments, the house view, and the final positions are our own.

In our work with clients, AI helps us prepare for meetings by organizing notes and gathering relevant context we might otherwise miss. It is useful for stress-testing scenarios and modeling outcomes we want to walk through together. It can also support routine operational tasks that allow our team to focus more of their time on the work that matters.

What we do not use AI for is just as important. We do not use it to build long-term asset allocations, generate investment recommendations, or replace the judgment that sits behind every decision we make for our clients. That work continues to require human expertise, real knowledge of the client, and the ongoing accountability that defines a fiduciary relationship.

This is what we believe is the right use of this growing technology. AI belongs as a tool that supports work, not one that replaces the judgment and the relationships behind it.

AI belongs where it organizes thinking, summarizes dense material, and surfaces questions worth asking, but the judgment, conclusions, and final positions stay human.

Our View

Artificial intelligence is here to stay, and its capabilities will continue to grow. For investors, the right response is not to avoid it, but to understand it. LLMs are remarkable tools for many tasks, but they are not built for the work of constructing a long-term investment plan. They cannot perform the math, do not have current market information, do not know the investor, and carry no responsibility for the outcome of any recommendation they produce.

A better question to ask is not whether AI should be involved in investment decisions, but how? AI can help advisors work faster, dig deeper, and spend more time on the parts of the job that genuinely require human judgment. That is where we believe it belongs, and that is how we use it at Verdense.

If you have questions about how we use technology in our work, or how we think about your portfolio, please reach out to your advisor.

1. <https://seekingalpha.com/news/4529731-chatgpt-reaches-900m-weekly-active-users-but-gemini-adoption-continues-to-heat-up-report>, ChatGPT Stats May 2026 - Usage, Market Share, and More. - fatjoe., Claude AI Statistics (2026) - Active Users, Revenue & Growth
2. <https://help.openai.com/en/articles/6825453-chatgpt-release-notes> | ChatGPT — Release Notes June 1, 2026
3. <https://platform.claude.com/docs/en/about-claude/models/overview> | Release Notes April 2026
4. <https://docs.cloud.google.com/gemini-enterprise-agent-platform/models/gemini/2-5-pro> | Release Notes June 1, 2026

If you have questions about how we use technology in our work, or how we think about your portfolio, please reach out to your financial advisor.

Disclaimer:

© 2026 Authored by **Megan Horneman**, Chief Investment Officer, Verdense Capital Advisors, LLC. Reproduction without permission is not permitted.

This material was prepared by Verdense Capital Advisors, LLC ("VCA" or "we", "our", "us"). VCA believes the information and data in this document were obtained from sources considered reliable and correct and cannot guarantee either their accuracy or completeness. VCA has not independently verified third-party sourced information and data. Any projections, outlooks or assumptions should not be construed to be indicative of the actual events which will occur. These projections, market outlooks or estimates are subject to change without notice. This material is being provided for informational purposes only and is not intended to provide, and should not be relied upon for, investment, accounting, legal, or tax advice. Past performance is not a guarantee of future results. Different types of investments involve varying degrees of risk, and there can be no assurance that the future performance of any specific investment, investment strategy, or product or any non-investment related content, made reference to directly or indirectly in these materials will be profitable, equal any corresponding indicated historical performance level(s), be suitable for your portfolio or individual situation, or prove successful. You should not assume that any discussion or information contained in this report serves as the receipt of, or as a substitute for, personalized investment advice from VCA. Alternative investments are designed only for sophisticated investors who are able to bear the risk of the loss of their entire investment. Investing in alternative investments should be viewed as illiquid and generally not readily marketable or transferable. Investors should be prepared to bear the financial risks of investing in an alternative investment for an indefinite period of time. Due to various factors, including changing market conditions and/or applicable laws, the content may no longer be reflective of current opinions or positions. All indexes are unmanaged, and you cannot invest directly in an index. Index returns do not include fees or expenses. Sector Watch Use of this website is intended for U.S. residents only. Any recommendation, opinion or advice regarding securities or markets contained in such material does not reflect the views of Verdense Capital, and Verdense Capital does not verify any information included in such material. Verdense Capital assumes no responsibility for any fact, recommendation, opinion, or advice contained in any such research material and expressly disclaims any responsibility for any decisions or for the suitability of any security or transaction based on it. Any decisions you may make to buy, sell, or hold a security based on this research will be entirely your own and not in any way deemed to be endorsed or influenced by or attributed to Verdense Capital. It is understood that, without exception, any order based on such research that is placed for execution is and will be treated as an UNRECOMMENDED AND UNSOLICITED ORDER. Further, Verdense Capital assumes no responsibility for the accuracy, completeness, or timeliness of any such research or for updating such research, which is subject to change without notice at any time. Verdense Capital does not provide tax, or legal advice. Under no circumstance is the information contained within this research to be used or considered as an offer to sell or a solicitation of an offer to buy any particular investment/security. Fixed income securities are subject to increased loss of principal during periods of rising interest rates. Fixed income investments are subject to various other risks including changes in credit quality, market valuations, liquidity, prepayments, early redemption, corporate events, tax ramifications and other factors. Lower rated securities are subject to greater credit risk, default risk, and liquidity risk. Commodity-related products, including futures, carry a high level of risk and are not suitable for all investors. Commodity-related products may be extremely volatile, illiquid and can be significantly affected by underlying commodity prices, world events, import controls, worldwide competition, government regulations, and economic conditions, regardless of the length of time shares are held. Data is provided for information purposes only and is not intended for trading purposes. Verdense Capital shall not be liable for any errors or delay in the content, or for any action taken in reliance on any content.