

In the evolving world of artificial intelligence, tools like **ChatGPT** and **Google Gemini** have introduced unprecedented capabilities in natural language processing and information synthesis. They can help family offices with preliminary research, parsing complex reports, and generating summaries. However, these generative AI systems also produce *hallucinations* — confident but incorrect or fabricated statements — creating significant risks, especially in data-scarce domains such as private markets.

Understanding AI Hallucinations: What Does "Grok" Mean Here?

Before diving into failure modes, it's important to clarify what "grok hallucinations" implies. *Grok* is often used informally to mean deeply understanding or intuitively grasping something. When applied to AI, "grok hallucinations" emphasizes how advanced AI models can produce outputs that appear deeply knowledgeable and coherent — in other words, they seem to *grok* the subject — but are in fact fabricating or misleading.



This presents a unique challenge: the AI output feels like a confident expert's insight, yet it contains errors that may not be obvious at first glance.

Why Are AI Hallucinations Dangerous for Family Offices?

Family offices rely heavily on accuracy and trustworthiness when dealing with ultra-high-net-worth (UHNW) clients. Typical AI use cases in our domain include:

- Preliminary research summaries on sectors, funds, or investment themes
- Parsing and synthesizing richly detailed private equity or venture capital reports
- Generating concise executive summaries for committees and advisors
- Data extraction and normalization from diverse sources for reporting

However, in these scenarios, AI hallucinations can lead to:

- **Confident fabrication:** The AI confidently asserts information that isn't grounded in data, leading to false assumptions.
- **Misleading summaries:** Summary texts that omit critical caveats or misinterpret reported facts.

- **Overreliance on incomplete or scarce data:** Especially in private markets, where data scarcity increases error risk.

These issues can impact manager due diligence, reporting accuracy, investment decisions, and ultimately family trust.

The Role of Data Scarcity in Private Markets

One of the key [get more info](#) drivers of hallucination risk in family office technology is the nature of the private markets. Unlike public markets with abundant, standardized data and regulatory filings, private market data is:

- Scattered across non-standardized reports and private communications
- Often confidential or incomplete, with limited historical transparency
- Highly heterogeneous in format, requiring sophisticated interpretation

AI tools like ChatGPT or Gemini, which are primarily trained on broad internet corpora and public datasets, have less footing in these specialized domains. As a result, they may “fill in the gaps” by fabricating plausible-sounding but incorrect information—a classic hallucination mode.

Common Failure Modes of Grok AI Hallucinations

Understanding the common failure modes helps us set appropriate guardrails. Here are the key types:

1. Confident Fabrication

This is the hallmark failure mode: the AI outputs specifics such as fund names, dates, performance metrics, or strategic assertions with an air of certainty—but these facts are invented. For example, an AI might state that a private equity fund achieved a 25% IRR over 5 years, when in reality no such figure exists or has been reported.

2. Misleading Summaries by Omission

Summaries that exclude important nuance or caveats can mislead decision-makers. For example, parsing a due diligence report, the AI might leave out references to key risks identified by a manager or regulatory issues that were explicitly flagged.

3. Overgeneralization and Overconfidence

The AI model may generalize from partial data and present it as comprehensive truth. Statements such as “All funds in this sector have outperformed benchmarks” reflect dangerous overconfidence that can misdirect investments.

4. Mixing Fact and Fiction

Sometimes hallucinations appear embedded within largely factual text, making them especially hard to spot. For example, explaining a manager’s strategy correctly but inventing unknown portfolio companies or deal structures.

5. Misinterpretation of Technical or Domain-Specific Terms

Private markets involve jargon, complex structures, and legal terms. AI may misparse these, confusing waterfall structures, carried interest terms, or fund vintage definitions, leading to incorrect interpretations.



Faithfulness vs Factual Hallucinations

It is useful to distinguish two related but different issues:

- **Faithfulness:** Does the AI accurately represent the input data or source text? For example, faithfully summarizing a due diligence report without inventing details or adding bias.
- **Factual accuracy:** Are the facts themselves true and verifiable in the external world? Even a faithful summary can still fail if the underlying data has inaccuracies.

Many AI hallucinations fall into the category of *unfaithfulness* where the AI creates statements that are inconsistent with the input provided. This is particularly damaging in family office contexts where the goal is to augment human intelligence, not replace it with unvetted output.

What About ChatGPT and Google Gemini? How Do They Compare?

Both ChatGPT and Google Gemini exemplify cutting-edge large language model (LLM) technology; however, their hallucination profiles reflect underlying training data, architecture, and governance approaches:

Aspect	ChatGPT	Google Gemini
Training Data	Mostly internet text, plus supervised fine-tuning.	Google claims integration with their search index and multimodal data, which may improve recency.
Hallucination Tendencies	Known to confidently fabricate details, especially when asked about niche or private data. Early reports are mixed; some improvements claimed but hallucinations persist in complex domains.	OpenAI has introduced fact-checking plugins and "GPT-4 Turbo" which attempts to improve factuality. Google integrates with real-time search and sources, but transparency on hallucination mitigation is limited.
Governance Features	OpenAI has introduced fact-checking plugins and "GPT-4 Turbo" which attempts to improve factuality.	Google integrates with real-time search and sources, but transparency on hallucination mitigation is limited.
Use in Family Offices	Widely used for draft summaries, but requires human vetting and sources cited. Still emerging, promising but unproven at scale in private market contexts.	

Mitigating the Risks of Grok AI Mistakes in Family Office Workflows

Based on years of experience managing ultra-high-net-worth family reporting and due diligence, here are best practices to handle hallucinations:

1. **Always ask: “What would we show an auditor?”** If the AI output does not have a credible, verifiable trace back to source documentation, do not rely on it.
2. **Use AI as a starting point, not an oracle.** Treat summaries or preliminary research as drafts that must be checked by domain experts.
3. **Maintain a “things AI got wrong” log.** Collect examples of hallucinations encountered in your workflows to educate users and improve prompts.
4. **Combine AI tools with structured data governance.** For example, parse fund reports into databases with human validation before feeding to AI summaries.
5. **Request citations and source attributions where possible.** Models or tools that link to underlying documents improve trust and traceability.
6. **Train teams on common failure modes.** Educate investment analysts and clients about hallucination types and indicators.

Conclusion: Balancing Innovation with Caution

AI generative tools like ChatGPT and Google Gemini are transformative, enabling family offices to process information faster and generate insights that were previously laborious. Yet their **confident fabrications** and **misleading summaries** require sober recognition and controls.

By understanding the common hallucination failure modes — especially in data-scarce private markets — and embedding robust verification processes, family offices can grok AI’s potential without falling victim to its pitfalls. This reminds me of something that happened wished they had known this beforehand.. Until AI can fully explain its reasoning and reliably cite all facts, human judgment and skepticism remain indispensable.