

Artificial intelligence tools like **ChatGPT** and **Google Gemini** generate a lot of excitement in executive suites and family offices alike. The promise of rapid insights, automated summarization, and advanced data parsing is alluring. Yet, a clear-eyed understanding of what these Large Language Models (LLMs) truly are — and are not — is critical to avoid costly missteps.

In this post, I'll unpack the key reasons why AI tools are *not* "true intelligence," explore specific risks like hallucinations, and discuss family office use cases where cautious optimism, combined with rigorous governance, can unlock real value.

Demystifying AI: Prediction, Not Reasoning

Here's a story that illustrates this perfectly: was shocked by the final bill.. One of the biggest misconceptions about LLMs such as ChatGPT and Google Gemini is the belief that they "understand" the text they generate. In reality, these models operate by identifying patterns in vast datasets and producing probabilistic predictions about what word or phrase should come next—not by reasoning or comprehending.

- **LLM Limitations:** At their core, LLMs generate content based on statistical relationships in training data, not on any form of conscious understanding.
- **Does Not 'Understand':** Unlike humans, these models lack self-awareness, context comprehension, and critical thinking.
- **Prediction, Not Reasoning:** Their outputs are best viewed as highly sophisticated guesses—often impressively coherent, but sometimes misleading.

This fundamental limitation means that outputs may appear authoritative even when inaccurate. For executives, recognising this distinction is crucial before placing undue faith in AI-generated content.

The Persuasive Danger of AI Hallucinations

One of the most challenging risks with AI-generated content is *hallucinations*—where the AI generates factually incorrect or fabricated information with high confidence. These hallucinations are particularly **persuasive** because the language is fluent and coherent.



Type of Hallucination Description Examples in Family Office Context Risks **Factual Hallucinations** Generated facts or data points that are completely false. Misstated private company valuations; incorrect regulatory dates. Misled investment decisions; compliance breaches. **Faithfulness Errors** Summaries or paraphrases that distort the meaning of original content. Misinterpreted private placement memoranda; inaccurate fund manager reports. Inaccurate due diligence outcomes; broken trust with stakeholders.

Because LLMs are trained on vast and diverse data sources, including internet text, they sometimes “make up” information when lacking sufficient data or when queried in ambiguous ways.

Executives must understand that an AI’s confident tone is not a guarantee of correctness. Instead, outputs should be treated as starting points requiring verification, especially when used for investment decisions or reporting to clients and auditors.

Family Office Use Cases: Where AI Shines — and Where Caution Is Needed

Family offices are increasingly experimenting with AI tools like ChatGPT and Google Gemini for a variety of tasks. Here are some areas where AI adds preliminary value, as well as inherent risks:

Preliminary Research and Idea Generation

- **Strengths:** Quickly scanning public information to generate lists of potential managers, strategies, or sectors.
- **Risks:** Information may be outdated or inaccurate, especially for niche private markets sectors.
- **Guardrails:** Utilize AI-generated outputs as “raw material” to guide human-led deep dives, not as sole decision input.

Parsing and Summarizing Complex Reports

- **Strengths:** Extracting key points from lengthy manager reports or regulatory filings to save analyst time.
- **Risks:** Faithfulness errors can omit or misrepresent critical details.
- **Guardrails:** Always cross-check AI summaries against original documents before sharing with stakeholders.

Handling Private Markets Data Scarcity

Private markets pose special challenges because data scarcity makes AI hallucinations more likely. Unlike public markets, where data is abundant and structured, private equity or venture fund information is often fragmentary and proprietary.

- **Risk Amplification:** Without sufficient training data, AI tools may fill information gaps with fabricated or inaccurate content.
- **Example:** An LLM may confidently fabricate past fund performance or investor commitments.
- **Mitigation:** Maintain human oversight and leverage internal proprietary data where possible.

Faithfulness Versus Factuality: Understanding Errors in AI Outputs

It’s important to distinguish two error types that affect AI output trustworthiness:

1. **Faithfulness Errors:** These happen when the AI output deviates from or distorts the original source material. For example, a summary that changes the meaning or omits important caveats.

2. **Factual Hallucinations:** When the AI fabricates completely false facts, numbers, or entities that don't exist in any reliable source.

Faithfulness errors can often be caught by cross-referencing source documents. Factual hallucinations are more subtle but usually become apparent upon trying to verify the output externally.



Practical Steps for Executives: What Would We Show an Auditor?

From my experience in family office technology risk governance, one simple framing helps executives vet AI adoption:

"What would we show an auditor if they asked us about this tool and its output?"

Any AI-driven process must be able to demonstrate:

- **Data Lineage:** Clear tracing of input documents or data sources feeding the AI output.
- **Human Oversight:** Documented steps showing how AI-generated results were reviewed and validated.
- **Risk Controls:** Defined processes to prevent or catch hallucinations, including spot checks or sampling.
- **User Training:** Ensuring team members understand LLM limitations and do not treat outputs as gospel.

Without these controls and documentation, relying on AI outputs for crucial investment or reporting decisions can expose the family office to reputational and regulatory risks.

Conclusion: AI Is a Powerful Tool, Not a Thinking Partner

Recent advancements with ChatGPT and Google Gemini demonstrate remarkable natural language generation capabilities. Still, these models do not provide "true intelligence" — instead, they perform pattern-based prediction without understanding, reasoning, or consciousness.

The persuasive fluency of AI-generated text masks real but subtle risks, including hallucinations and faithfulness errors, especially amplified in data-sparse private markets typical of multi-family office portfolios.

For canadianfamilyoffices.com executives evaluating AI tools, the key takeaways are:

- AI outputs should be viewed as proposals or drafts, not authoritative facts.

- Robust human review, validation, and governance are mandatory before relying on AI for critical family office activities.
- Framing AI use through compliance and audit standards helps build safe adoption cultures.
- Understanding that LLMs predict text without true “understanding” prevents misplaced trust and errors.

With these perspectives, family offices can responsibly harness AI capabilities to augment human expertise, rather than blindly trusting it to replace human judgment.