

ThinkTwenty20's "Twenty Rules for AI for Financial Professionals": Alpha Version – Part 2

By Eric E. Cohen, CPA

AI > GenAI; some GenAI > other GenAI

This column continues an ongoing series of postings to develop helpful guidance for financial professionals related to artificial intelligence. It is developing a list of guidelines and advice, with the hope that we can collaboratively make some of them more organized and permanent. If you missed the introduction to the series, with the first five suggestions, focused on things to know before you start providing input to, or read and share the output from, a generative artificial intelligence (GenAI) tool, you will want to read that post, which began with five areas:

- Confidentiality: Don't type anything into an AI that you would not want made public.
- Skepticism: Don't automatically trust anything coming from an AI without review.
- Diversification: Don't put all your eggs (Algs?) in one basket.
- Compliance: Consider how any output might comply with industry and ethical regulations and standards.
- Transparency: Be careful to consider when you need to disclose your use of these tools.

This time, I'd like to focus on one piece of guidance:

- **Tool selection: Generative AI may not be the right AI for the job; your chosen GenAI may not even be the best GenAI for the job.**

To some, this may seem like a specialization of the Diversification principle; while choosing and sticking with one GenAI tool may be simple (related to licensing, training and other very practical issues), sticking with one GenAI for all tasks is in many other ways suboptimal. While GenAI has the low barrier to entry and perceived minimal costs, it's not always the right tool for the jobs a financial professional is engaged in.

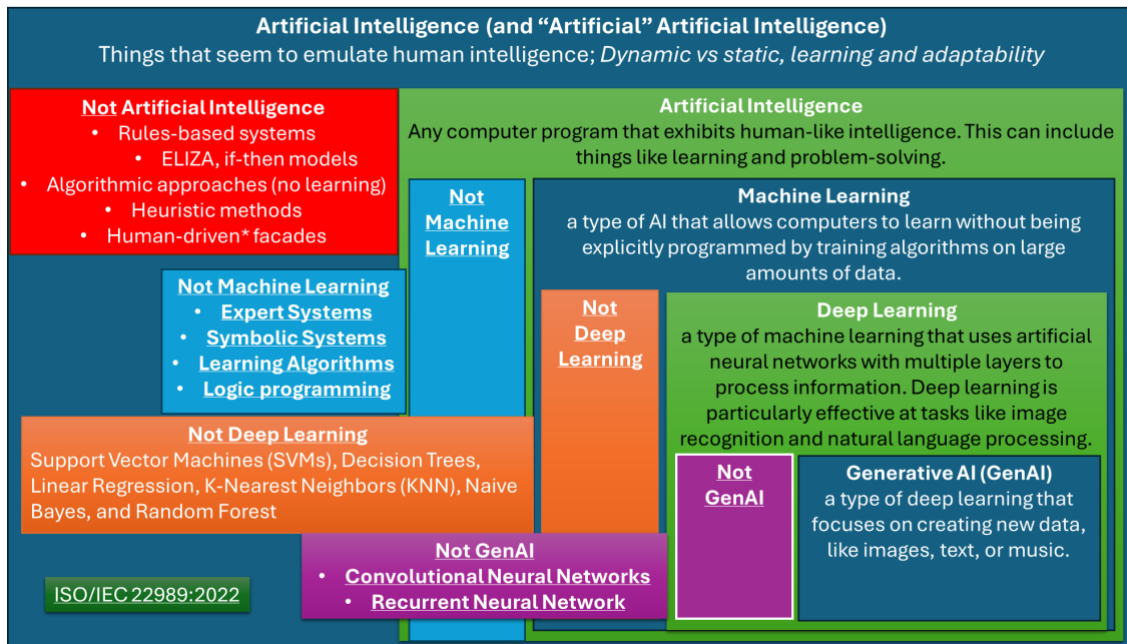


Figure 1: GenAI in the world of AI

Figure 1 is an illustration of the world of things called "artificial intelligence." With people throwing billions of dollars at anything with "AI" in the description, and with the craze around GenAI, the business community has looked to GenAI for solutions to their problems, even though another tool in the AI toolkit may be more appropriate.

In the illustration, there are areas such as “Not Machine Learning,” “Not Deep Learning” and “Not GenAI.” That is not to say those methods are not valuable or important; in fact, those may be the better tool for the tasks needed in the enterprise. In accounting and audit, where analysis and prediction may be appropriate to task, I asked Gemini about some machine learning tasks that might be useful:

“Supervised learning for anomaly detection: This technique trains AI models to recognize unusual patterns in financial data. These anomalies could indicate potential fraud or errors. By automatically flagging these transactions, accountants and auditors can focus their time on investigating the most suspicious activity.

“Unsupervised learning for clustering: This lets the AI group similar transactions together. This can help identify trends or areas of high risk. For example, unsupervised learning might cluster purchases from vendors not previously used by the company, which could warrant further investigation.

“Reinforcement learning for optimizing audit procedures: This allows AI to learn and improve its auditing approach over time. Imagine an AI system that gets rewarded for identifying high-risk areas during audits. Over time, the AI would become better at prioritizing tasks and performing efficient audits.”

My illustration in Figure 1 doesn’t reflect that not all GenAI offerings are the same. In fact, different interfaces within the same GenAI offering may offer challenges or value to your tasks.

Text input limitations: Large Language Models, for text-to-text, may have differences in the amount of textual input you can provide. Microsoft’s Copilot has its standard input, which until recently was limited to 2000 characters, and Notebook interface, limited to 18000 characters. This blog itself is already more than 4500 characters. Compare that with others that may accept 2000 tokens (tokens are words or parts of words, so 2000 tokens could be five times or more than the input of 2000 characters). I can’t paste this content into Copilot; I can paste it into Claude.

Other input limitations: Moving beyond the simplest typing of text into the input area, the various GenAI solutions are inconsistent with how to feed information in. Some let you drag-and-drop, some let you reference a URL (web link), others require a file upload and may limit the types of files you can upload.

Financial professionals use PDF (Adobe Acrobat) files extensively. Those files may be on a website, on your local drive, on a Google drive, on a Microsoft OneDrive or on other sources. Some PDFs make their text available easily; others will require optical character recognition to read.

I received a PDF file with locked content. I could not find a tool that would upload it. I could upload it to Google Drive, however, open it as a Google Doc, and it performs OCR in the process; I could then copy the text into the tools for evaluation.

Similarly, we use video more and more, and Youtube is right up there as an information source. Although Gemini (web interface) works with Youtube files (same company), I asked Gemini to summarize a one-hour video by time stamp. The response did not seem correct, so I copied and pasted the transcript. It was too large for (free) Claude (by only 4%). I then pasted it to Gemini 1.5 Pro in the Google AI Studio interface. It was around 27,000 tokens, and it processed the content without issue. Claude Pro can take 200k tokens (around 350 pages of text); the limits of the free version “depend on current demand.” I guess $27,000 * .96$ is around 26,000 tokens at the time for Claude.

The call to action here is to keep an open mind for tooling. Don’t let GenAI, and specifically, don’t let a specific GenAI, be your go-to for every task.