



PURDUE
RESEARCH FOUNDATION


Office of
Investments

An Asset Allocator's Perspective on Artificial Intelligence Use Cases

PURDUE RESEARCH FOUNDATION OFFICE OF INVESTMENTS'
APPROACH TO A NEW AND RAPIDLY EVOLVING SPACE

Contents

- 3) Introduction
- 4) Asset Allocator Engagement
- 7) Internal Exploration
- 12) Conclusion



Introduction

ChatGPT was released by Open AI on November 30, 2022, which fell toward the end of asset allocators' annual meeting season. By the spring of 2023, many general partners ('GPs') were introducing and showcasing ChatGPT's ability to complete tasks with little effort within seconds, such as writing the speeches that were used to kick off the annual meeting. Two key questions initially came to mind as an asset allocator attending these meetings. First, **how will advancements in artificial intelligence ('AI') impact existing and future investments?** And second, **how will advances in AI change how asset allocators work?** This white paper was curated to share the findings from exploring these and many other questions related to the rapidly evolving space of AI and machine learning ('ML').

Over the last six months, the Purdue Research Foundation Office of Investments ('PRF OOI') aimed to better understand the needs and use cases of AI/ML for asset allocators through engagement with leaders in the industry. These industry groups consisted of Endowments & Foundations, Public Pensions, Corporate Pensions, Family Offices, Registered Investment Advisors ('RIAs'), and Investment Consultants ranging in assets under management ('AUM') from \$1 billion to over \$40 billion in size. Additionally, the PRF OOI investigated how AI/ML could integrate into workflows through the lens of (1) the technology and infrastructure available today and (2) what is feasible to create and build in partnership with a software development firm. Please understand that these findings are not meant to be an exhaustive list of the various angles and approaches to the technology. Instead, the hope is that it helps inform the asset allocator community's collective understanding of the technology, how AI/ML may impact the industry in the future, and the best ways to access AI/ML both now and going forward. PRF OOI involved its internal information technology ('IT') team during the initial exploration phases to ensure that the technology being contemplated was appropriate from a security perspective. Thank you to everyone who engaged with PRF OOI to help make the collective asset allocator community more informed.

Asset Allocator Engagement

Institutional asset allocators sit in a unique seat with access to some of the brightest minds in the world on a range of macroeconomic and investment-related topics across asset classes. Like many industries, Open AI's release of ChatGPT ('Generative Pre-trained Transformers') gave non-tech native business and investment practitioners a window into what the future may look like with well-informed AI chatbots. As much as this was a technological advance, the vision it gave to individuals on future possibilities was equally important. People worldwide began to say: *"I know there is a lot of data and information available; what if I could ask a ChatGPT-like bot questions about the data?"* Institutional asset allocators' eyes were opened to how they may collect, aggregate, query, and communicate data and information related to their investment portfolios in the future.

AI applications within investment management workflows can be bifurcated between opportunities to enhance portfolio returns ('alpha') and increase productivity. When exploring the efficacy of AI adoption, it is essential to quantify and convey the potential benefits of utilization, considering its value through increased expected return (e.g., alpha) or value coming from increased overall productivity that leads to higher margins for the organization.

In brainstorming several key use cases for asset allocators, the PRF OOI identified the following applications where leveraging AI could enhance efficiency and drive value:

- 1. Quarterly Reports and Other Information Automatically Uploaded to Database (*Productivity*)**
Ability to automatically pull reports from multiple sources (individual or group email, data rooms, etc.) into a centralized repository using technology (e.g., application programming interface and robotic process automation).
- 2. Searchable / Insight-able Quarterly Report Repository (*Productivity*)**
Ability to search and gain insights from quarterly reports and other proprietary data sets using generative AI (e.g., a chatbot that can interact with your quarterly reports, "Does the pool have any XYZ exposure?" or "What were the three most common themes mentioned in the most recent private real estate quarterly letters?").
- 3. Pipeline Prioritization / Management (*Productivity*)**
Given a set of inputs, see the status of investments in the pipeline and make suggestions for prioritization.

4. Review Pitch Decks (*Productivity*)

Utilize AI/ML to extract predetermined data points from pitch decks and upload them into a cloud-based repository. Automatically generate requests for additional information if any data points were excluded from the original pitch deck. Users actively interact with the data on a clean user interface to instruct the model on the next steps, such as "schedule a call" or "draft a decline email."

5. Draft Informational Portions of Investment Memo (*Productivity*)

Leverage a large language model ('LLM') to complete summary-type information such as firm, team, strategy, and process in a memo format from multiple sources (e.g., due diligence questionnaire ('DDQ'), pitch decks, team notes, etc.).

6. Create Standardized Reporting with Multiple Information Sources (*Productivity*)

Push button creation of an informational presentation or dashboard for a governing body or client based on multiple sources of information (e.g., performance system, balance sheet information, etc.).

7. Call / Meeting Transcripts (*Productivity*)

Have a bot act like an assistant to take notes and make a transcript on calls.

8. Network Access and Engagement (*Alpha*)

Flags automatically generated for people in your LinkedIn network that may be helpful in diligence when looking at a fund or deal (e.g., connected with a past employee at a company, 2nd order connection, etc.).

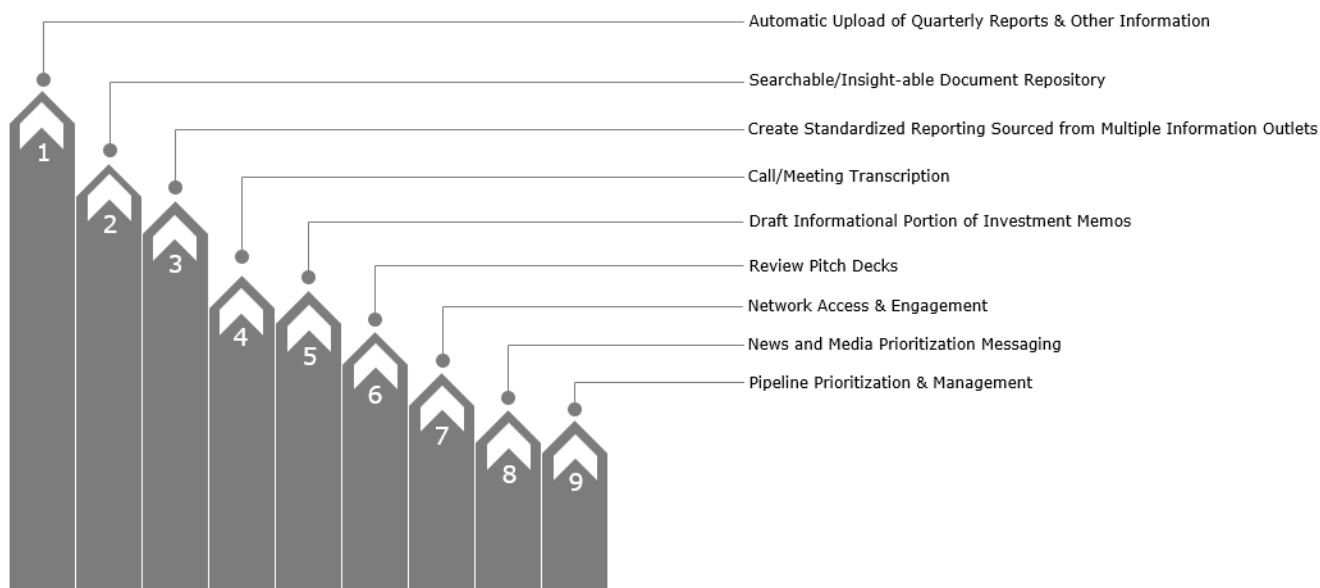
9. News Prioritization Email (*Alpha*)

As the LLM is trained on internal data, it can scan publicly available news and other sources to help identify and filter information important to the asset owner (e.g., you own a shipping company in a co-investment portfolio, and in your 'daily updates' email, you receive a summary of the earnings calls from the public shipping companies).

The PRF OOI set out in the second quarter of 2023 to understand broadly how investment peers were thinking about the technology and AI/ML space and explore the use cases previously described. PRF OOI created and distributed a survey to its network to engage with peers, resulting in over 20 follow-up interviews with other institutional asset allocators to discuss their responses and thoughts on the broader AI/ML space. The survey included various questions, from ranking the use cases (previously described) to what current AI/ML tools are being used in their workflows today. The most represented group within the survey respondents and interviews were Endowments & Foundations, which provided just over a third of the total responses, while the other groups in the peer universe were balanced.

An Asset Allocator's Perspective on Artificial Intelligence Use Cases

The graphic below ranks the identified use cases from the survey responses:



Key takeaways from discussions with asset allocators are outlined below:

1. **Asset allocators are most interested in improving their workflows through added efficiency and automating tasks that have historically been completed with a high degree of human interaction.** Automatically uploading and organizing documents with the ability to also gain insights and have those documents searchable (use cases [1] and [2] above) were ranked the highest in terms of importance from survey respondents.
2. **While important, asset allocators want to maintain some control over their workflows' outputs while recognizing the synergies that can be realized from introducing AI.**

Standardized and automated reporting, call/meeting transcription, and drafting informational portions of investment memos, ranked after the first two use cases.

3. **There is dispersion among asset allocators concerning their stage of engagement with AI technology, and the rationale for this dispersion varies. Given the nature of the asset allocator business, security and trust are paramount.**

Many of the institutional asset allocator teams that PRF OOI engaged with indicated that they had some internal discussions about AI/ML and identified some use cases for their teams. However, very few are currently acting. Reasons for not acting included lack of a clear starting point, cost, and AI-related security questions. Additionally, many survey participants indicated they had engaged with their internal IT team. However, many internal IT teams are not necessarily designed to support and build AI tools. Organizations furthest along in

their AI enterprise strategy have internal data scientists and software developers available to the investment teams.

4. **True automation has plagued user uptake with legacy systems used to increase efficiency; this is an area where AI can add value to organizational “buy-in.”**

The key challenges identified with existing systems were clunky data ingestion and a sub-optimal user experience. If a user needs to click 7+ times to navigate to where they enter notes or information about an interaction with a GP, the likelihood of several users staying up with the system over several cycles becomes less likely.

5. **Larger asset allocators (in terms of AUM) favor alpha applications of AI/ML, while smaller asset allocators favor productivity applications.**

Broadly speaking, large asset allocators ranked alpha use cases higher, while smaller asset allocators, who may be more resource-constrained, rated productivity use cases higher.

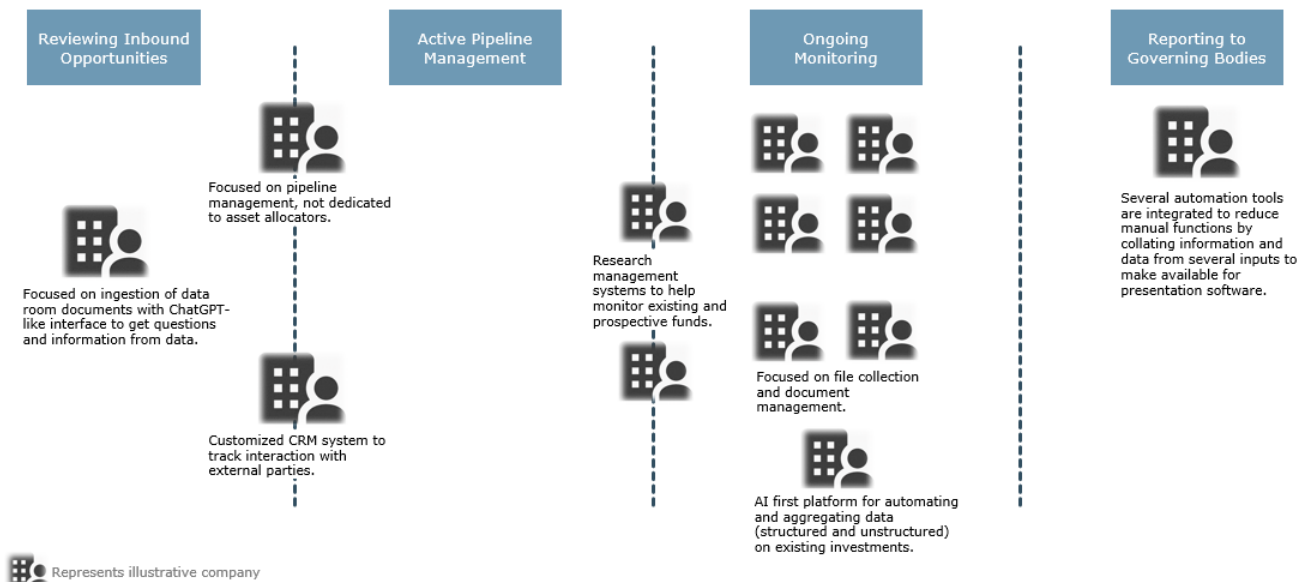
Internal Exploration

While PRF OOI was engaging with its peer network through the survey and interview process, the team was simultaneously analyzing the AI/ML space. At the core of this analysis was a buy versus build decision for many of the identified use cases (e.g., does it make economic sense to create a tool, or is there existing functionality available in the marketplace today that the team can use at a reasonable cost). In conducting this analysis, the team looked at tools that could be leveraged today and the future development initiatives of existing tools. They also set out to develop some of the desired technology, collaborating with a local software development partner.

Off-The-Shelf Tools

Asset allocators have been a relatively underserved market by the technology community, given that it is a smaller part of the overall financial services market, and many times, asset allocator investment teams are within government entities or non-profit organizations that are more sensitive to its technology spend than other organizations. Historically, tools built for different parts of the financial services space have been modified and sold to asset allocators, with many of these failing to address the specific needs of these enterprises. Below is a brief snapshot of some existing tools and solutions for various segments of an asset allocator's workflow. While AI is becoming more prevalent in the workflows for asset allocators, the industry still lacks an end-to-end process management tool fully enabled by AI.

An Asset Allocator's Perspective on Artificial Intelligence Use Cases



Key takeaways from the review of existing and off-the-shelf solutions are outlined below:

1. The functionality of legacy process management systems has limitations, and the adoption of AI by these systems is yet to be integrated into the user experience.

Many of the research management system ('RMS') and customer relationship management ('CRM') solutions that asset allocators leverage to help manage investment pipeline workflows, research, and existing relationships are currently developing their AI strategy with limited to zero changes relative to product and solution offerings that were available five years ago. Users of these products conveyed varying degrees of satisfaction. However, the feedback received through the interview process identified that users being content with their current RMS or CRM systems was the exception rather than the rule. Many of the shortcomings shared can be solved with AI-enablement. It is unclear how long some legacy platforms and tools will take to develop and integrate new AI tools, given that they are built on older technology infrastructure that may not allow for streamlined AI integration.

2. End-to-end process management specific to asset allocators is still a void in the marketplace, yet to be filled. The lowest-hanging fruit for AI applications has already been solved.

Many new tools solve one or two of the use cases identified for asset allocators; however, PRF OOI's research did not uncover a platform where many of the use cases are combined under one umbrella. A few examples of companies focused on a specific use case include:

An Asset Allocator's Perspective on Artificial Intelligence Use Cases

- a. Automatic Ingestion of Documents: a few document management companies sell to asset allocators that they have 'solved' the automated ingestion of documentation use case but essentially stop there without providing actionable insight or significant metadata extraction and packaging.
- b. Call Transcript: several companies that have 'solved' the documentation and note-taking/transcript use case (e.g., the system automatically provides a transcript and takeaways from a recorded call or meeting).
- c. Drafting of Informational Parts of IC Memo: this can largely be completed with the off-the-shelf ChatGPT and other LLMs.

Internal Development

As the final piece of the exploration analysis, PRF OOI engaged a third-party development firm that partners with and supports start-ups and large enterprises to build software and AI tools. The engagement lasted approximately eight weeks and began with a discovery phase whereby the development firm and PRF OOI had numerous interactions and discussions regarding current workflows and the specific AI-enabled tools most beneficial to the overall investment processes. Considering existing tools and internal requirements, it was determined to focus on the 'Searchable/Insight-able Quarterly Repository' use case. Many other use cases had already been 'proven' (e.g., those mentioned in the previous section), but this one had not. As such, the focus became building a chat-like bot that could answer questions about internal data from a quantitative and qualitative perspective. Asset allocators, like many businesses, have a big data problem. Several documents are received for each private investment vehicle, separate account, and/or commingled fund each month or quarter. The documents containing information like valuations and exposures come in varied formats and are distributed in various ways. Providing insights such as "How much exposure does the portfolio have to FTX?" or "Which of the GPs have a capital call line with Silicon Valley Bank?" or "What real estate assets does the pool own in the city of Nashville?" or "Which of the assets the pool owns had the largest change in value last quarter?" require an expensive transparency solution complemented with human interaction to comb a vast amount of information. Many of the existing transparency/exposure systems have solved data problems historically by throwing human resources at the problem, as that was all that was available, resulting in a high-cost, low-tech solution to a widespread problem. With advancements in data science, AI, and databases, could a solution be created with an AI-first theme that could result in a higher-tech, lower-cost solution? Querying large data sets for quantitative and qualitative answers is the use case that PRF OOI explored with the third-party development firm.

An Asset Allocator's Perspective on Artificial Intelligence Use Cases

PRF OOI began by providing one quarter's worth of documents (around 800 documents) received for PRF's existing partnerships to the third-party development partner (after the appropriate security questionnaires and NDAs were completed with the PRF Information Technology Department). The development partner then built a lightweight interface (system) and queried the documents on various questions. It became clear that the system could provide information around answers like "Who are the partners of XYZ Capital Partners?" or "Summarize this macro information from Fiscal Year 2023," however, challenges became apparent with the more quantitative data insights. To illustrate, for a question like "What is the best-performing company owned by XYZ Capital Partners?" the system output would sometimes return not just incomplete insights but also inaccurate answers. Usually, quantitative data is presented in a table format within the financial statement or quarterly letter. The question became: did the system not understand the query well enough and produce an inaccurate output or did the system understand the query and just produce an incorrect output? It was clear that limitations existed with how tables and graphs were being interpreted. Despite the limitations when interpreting graphs, in most cases, the system could find the appropriate fund/general partner and even the right area within the report or document to retrieve the correct answer. In one such query, the system found the right table and pulled the highest multiple, but when returning the name of the company with the highest multiple, the system pulled the name of the company from the row above the correct answer versus the name of the company next to the correct answer. Additionally, challenges emerged with the initial proof of concept for insights requiring access to multiple data sources and some level of light calculation. For example, answering a question like "How much exposure does the portfolio have to XYZ Company?" not only needs to locate all the places XYZ Company is mentioned throughout documents across the portfolio but also identify and calculate the institution's portion (most data provided on holdings of private companies is at the fund level) and aggregate it.

It became evident throughout the engagement with the third-party developer that obtaining actionable insights would not be as seamless as sending documents into the 'system,' asking it questions, and receiving accurate insights. When the output is inaccurate, there are some factors to consider: is it a prompt engineering problem (i.e., restructuring or asking the question better), or is it a data ingestion problem (i.e., the AI struggled to interpret the ingested data)? While the answer may be both, it was determined that the team's problem was more of a data quality issue. Specifically, the 'system' had difficulty ingesting tables and graphs, a common format used in portfolio reporting by investment firms. The output of any given system is only going to be as good as the underlying data. Systems cannot be appropriately trained just by putting documents and reports into them. Data must be clean and organized to embrace the full power of what is happening now and what is to come with AI.

Additional takeaways from the development experience are outlined below:

1. **The universe of off-the-shelf LLMs confirmed that generative and qualitative extraction tools perform well.**

Applications using LLMs to retrieve qualitative information from a data source do a good job; however, they remain relatively deficient at retrieving quantitative information.

2. **One of the critical components to successfully linking quantitative and qualitative insights in a usable form, known as 'tabular extraction,' is a problem many face and is yet to be solved (in PRF OOI's experience).**

This was the experience in development, confirmed when "test driving" other software and further communicated through discussion with large asset managers developing their technological platforms. Many market participants are trying to solve this problem, and a solution seems imminent.

3. **Given the importance of accurate data insights, overlaid with the early-stage advancement in AI technologies, time and resources are better spent extracting correct answers and improving the process versus developing clean interfaces. The more times a user can see the source of information and how the technology application retrieved the answer (right or wrong), with an efficient mechanism for improvement, the faster trust can be built.**

Spending time on the user experience and interface may not make sense before data quality can be validated. A solution for validating data may be a mix of AI and human involvement (i.e., human in the loop), where the AI performs data ingestion 75% of the way to the end point. Human interaction would help to augment and validate the AI output. This is especially key if training from human interaction can be leveraged going forward.

Conclusion

PRF OOI found 2023 to be a highly informative year regarding the integration of AI into its investment workflows. Expansion of AI across the tech stack, development tools, and applications are rapidly changing, with seemingly new announcements every week from both large corporations and start-ups. The team's efforts in engaging peers, assessing existing products and tools, and internally developing a proof of concept identified that there is still much more work to be done relative to many of the use cases specific to asset allocators. Existing products and tools are working towards AI enhancements; however, they remain largely unchanged, and hurdles persist for the internal development of tools, including data ingestion of tables and graphs.

As the team reflected on the appropriate next steps, it is unmistakable that regardless of how AI is used going forward, the output will only be as good as the quality of data in the 'system.' Therefore, the initial focus should be on data quality versus building tools to leverage data in its existing formats. A slick and seamless user interface only matters if the system has clean and robust data to query.

This paper was not meant to be a comprehensive view of all AI products and tools currently available to asset allocators but as a pursuit to be more informed about how AI may shape the future of investment offices and asset allocators. The daily advancements in AI/ML are exciting and welcomed. It will be interesting to see the industry's direction considering this new and evolving technology.

Interested in continuing the conversation?
Use the code to join the AI/ML Asset Allocator Forum



Email us your questions at AI-MLAssetAllocatorForum@prf.org





PURDUE
RESEARCH FOUNDATION

Office of
Investments