AI Enhanced Asset Management

Harnessing Large Language Models for Strategic Advantage

June 2024

About Entis

Entis is a data-driven investment intelligence company that utilizes advanced AI to analyze global corporate dynamics. By converting complex qualitative data into clear, actionable insights, Entis supports more informed investment decision-making.

Established in 2018 from Deloitte Netherlands and acquired by APG Groep N.V., Entis now operates as an independent entity within the APG family. The company's focus is on delivering reliable data insights that enhance understanding of the investment landscape, grounded in transparency and precision.

Authors



Machiel Westerdijk, PhD

Founder and Managing Director of Entis



David Barber, PhD

Science Advisor of Entis, Director of the UCL Centre for Artificial Intelligence and Professor of Machine Learning at UCL.



Hans Op 't Veld, PhD

Principal Director Responsible Investment at PGGM

Summary

This whitepaper explores the transformative impact of integrating Advanced AI techniques, particularly Large Language Models (LLMs), on investment strategies. It discusses the role of data-driven decision-making in asset management and how AI's ability to process unstructured data provides the opportunity to obtain a competitive edge. It addresses developments in AI emphasizing the need for asset managers to understand these technologies.
Practical applications are discussed, including managing challenges such as biases, misinformation, infrastructural adjustments and making informed buy vs. build decisions. The whitepaper concludes with a discussion on the opportunities and risks associated with AI.

01 Introduction

Investment management fundamentally hinges on understanding asset characteristics, possessing more comprehensive information than competitors, and executing transactions at optimal prices. The advent of computers in the 1970s marked the beginning of a transformation in investment strategies, giving rise to systematic approaches such as quant investing, high-frequency trading, and smart beta strategies.

The integration of Artificial Intelligence (AI) into asset management is poised to both deepen and accelerate this transformation beyond what has been possible historically. For example asset managers are already shifting from manual document analysis to automated analysis using AI language models such as GPT allowing for a far more comprehensive and consistent evaluation of multiple data points. This frees up managers to focus on strategic level insights, gauging the impact of technologies on industries, spotting nascent sectors, devising novel quant alpha strategies, dissecting supply chain disruptions, broadening diversification with global trends, and leveraging tech-driven opportunities.
Those parties that know how to use these techniques are well positioned to seek out those investments that have the best prospects.

As the integration of AI accelerates, it also introduces challenges and ethical considerations. These include navigating an evolving regulatory landscape, ensuring the reliability of AI-driven decisions and the avoidance of biases, and addressing the varied pace of AI adoption across the industry.

By providing a nuanced understanding of AI's role, this paper offers insights into how the industry can navigate the possibilities new technology has to offer while managing potential risks.

The paper is structured as follows:

Chapter 2: Foundations of Systematic Approaches in Asset Management We discuss the long-standing application and evolution of systematic approaches in the field, highlighting how AI complements and enhances these methods. The chapter goes into how tasks traditionally performed by human analysts are increasingly being automated, exploring the benefits and challenges of deploying AI in these contexts.

Chapter 3: Developments in Artificial Intelligence We outline the fundamental understanding necessary to grasp AI's potential and its ongoing evolution. We emphasize how AI's progression is driven by faster computing capabilities and its changing nature . This awareness is important for developing future strategies and preparing for next-generation AI technologies.

Chapter 4: Strategic Deployment of LLMs in Asset Management Starting with a sketch of expected trends influenced by AI in asset management, we discuss practical strategies for implementing Large Language Models (LLMs). It covers strategic considerations for integrating AI technologies—whether to develop in-house or outsource. Discussions also focus on organizing infrastructure around AI to ensure scalable, secure, and effective integration within investment processes.

Conclusion: Navigating the AI revolution in asset management. We reflect on AI's transformative potential and the opportunities and risks, including ethical considerations and regulatory challenges. We emphasize that overcoming these challenges requires practical solutions in real-world implementations. Proactive engagement with AI, rather than relying on external narratives, is essential as this technology is poised to significantly disrupt the industry.

02 Foundations of systematic approaches in Asset Management

The internet era and subsequent technological advancements catalyzed the evolution of systematic strategies, with quantitative investing emerging as a dominant force. This period was characterized by mathematical models driving investment decisions, a trend that has only grown with advancements in computing power and data storage capabilities.

Key properties of systematic approaches:

- Rule-based decision making: application of predefined rules and algorithms, aiming to minimize human bias and leading to improved analytic consistency.
- Data-driven analysis: using quantitative data analysis based on statistical and mathematical models to inform decisions.
- Automation and scalability: enabling the execution of strategies across large asset pools.

A significant evolution in systematic approaches has been the integration of alternative data sources to complement (historic) financial data. Text and graphical data from news, sustainability reports, annual reports and supply chain data, have all become invaluable. These sources provide insights into market sentiment, corporate behavior and global trends, enabling more nuanced investment strategies that encompass the inclusion of more than traditional financial metrics.

A pivotal aspect of the evolution in systematic approaches is the convergence of technologies—AI, big data analytics, and unprecedented computational power. This synergy has enabled the processing of vast datasets and the development of more complex, adaptive investment strategies. This allows for a holistic view of the market, integrating diverse data types (structured and unstructured) and applying sophisticated algorithms that can learn and adapt over time.

These advancements have significantly enhanced quantitative strategies by enabling the quantification of properties such as innovation strength, employee skills, leadership strength, market position, and other intangibles that were previously beyond the scope of these strategies and were typically assessed qualitatively by investment teams. Simultaneously, especially in Europe, investment objectives are evolving as sophisticated analytics begin to integrate social and environmental contributions into decision-making processes. This shift, illustrated in Figure 1, is part of an ongoing evolution, with systematic strategies increasingly reflecting the incorporation of intangible factors in risk and return calculations.
This information, which is seen to be of a qualitative nature, is now being quantified through technology, which enables the consideration of these factors in a fashion equal to the financial metrics that are being used. Whereas investors used to solve for (financial) return and risk (only), the incorporation of this type of consideration in investment processes is complicating the optimisation of the investment solution.

These considerations provide an opportunity for investors to differentiate themselves and to capture the opportunities this additional stream of information provides. This is an attractive prospect for the asset management industry, which has been plagued by fee pressure over the past decade.

Despite these developments, significant challenges persist for investors. Accessing reliable sustainability data remains difficult, and navigating the complex landscape of standards that define key non-financial metrics, such as environmental contributions and those deemed to be indicators of good corporate governance, continues to be problematic. Notable standards in this space include the EU Sustainable Finance Taxonomy, the United Nations Sustainable Development Goals (UN SDGs), and the OECD guidelines for multinational enterprises. These frameworks are partly overlapping and partly complementary – and sometimes even present conflicting criteria that complicate their incorporation into investment decisions.
In addition, systematic approaches face technical challenges such as a lack of transparancy, biases, the risk of overfitting, and need to rapidly adapt to market changes.

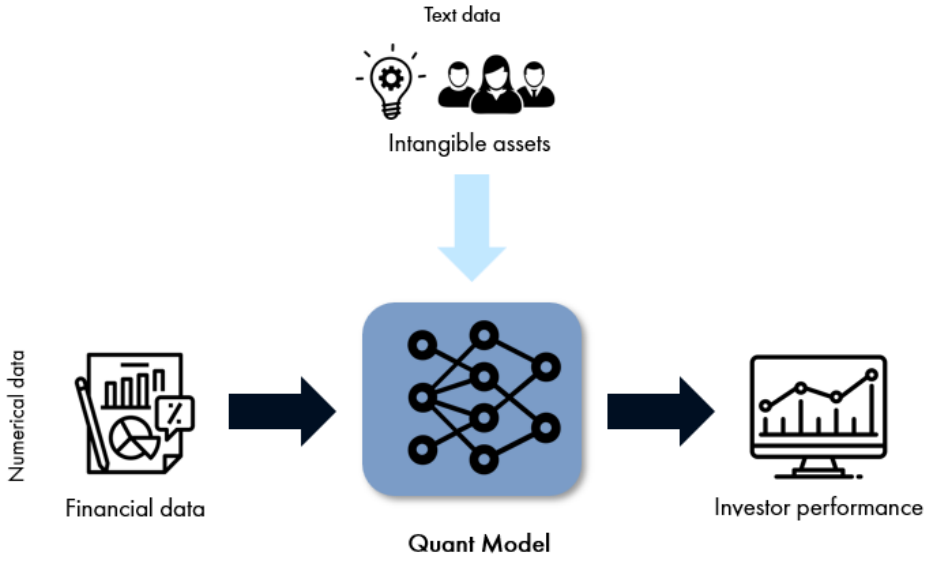


Figure 1: Integrating multi-dimensional data for enhancing quant alpha models: A synergy of traditional numerical data and unstructured information to gain improved stakeholder performance (which might include more than financial returns only).

Modern AI methods can enhance investment strategies by enabling the processing of diverse datasets and adapting quickly to new information. These methods promise significant advancements in quantifying intangible metrics and incorporating societal contributions from multiple standards. In the forthcoming chapters, we will go deeper into the latest developments in AI and examine practical strategies for integrating current large language models within asset management.

03 Developments in Artificial Intelligence

The current forefront of AI

In recent years, excitement has surged around AI advancements, spearheaded by systems like ChatGPT. These systems, falling under the umbrella of generative AI, are characterized by their ability to generate outputs such as visualizations, music, or texts based on textual prompts or descriptions. Note that generative AI differs from other AI systems, which generate output in highly structured formats, such as predicting chess moves, stock values, facial recognition or autonomous driving navigation.

The rise of generative AI, especially through the development of Large Language Models (LLMs), marks a transformative leap in human-computer interaction (see Appendix I for a brief explanation of systems like ChatGPT). These advancements significantly enhance our ability to process natural language and generate coherent content. Their applications span from serving as a combined chat-bot/super-encyclopedia/search engine to functioning as code generators and virtual software developers.

LLMs are particularly relevant for asset managers due to their capacity to access and extract valuable insights from vast quantities of textual data faster, more accurately, and in greater depth. The broader array of information found in texts like annual reports, scientific papers, and industry reports has traditionally been accessible only to human analysis, with all its inherent limitations. LLMs change this by making such textual data comprehensible to machines, enabling asset managers to analyze vast amounts of qualitative information efficiently, consistently and effectively.

One specific valuable application of LLMs is to use them for translating text into numerical data, allowing the semantic similarity between two texts to be expressed as a distance measure, see Figure 2. This functionality enhances the efficiency of processing large volumes of text and integrates seamlessly into automated quantitative analysis. In this way LLMs offer a concrete solution to incorporate intagible asset information into quant processes like we visualised earlier in figure 1.

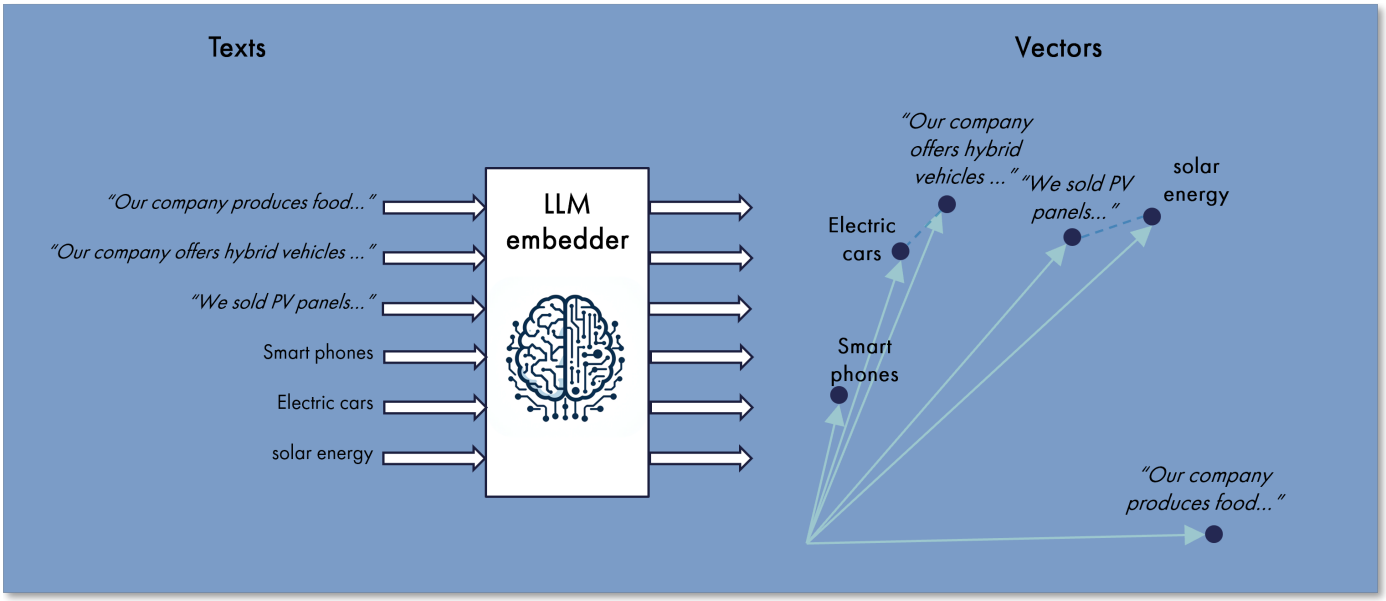


Figure 2: Transforming text to vectors with LLM embedders - The embedder component of a large language model can be used separately as a text analysis tool, capable of gauging the relevance of a business document to a given topic. This is achieved by converting text into a numerical format, thus quantifying its content into vectors. Each vector is an array of numbers reflecting various textual features—themes, sentiment, and keywords. Transforming text into numerical vectors facilitates precise, efficient comparisons. For example, the similarity between the vector profile of a text in a business report and that of an investment thesis on renewable energy technologies can be readily evaluated. The numerical proximity of these vectors indicates topical relevance. Such a mechanism allows for rapid review of extensive documents, trend identification, and relevance assessment, enhancing decision-making through a blend of quantitative and qualitative content analysis methods.

Fundamental tasks and functions that leverage the unique strengths of Large Language Models:

- Natural Language Processing (NLP): This enables automated analysis of vast textual data, such as annual reports and industry reports. It offers a deeper understanding and comprehensive integration of investment opportunities and risks, enhancing decision-making.
- Strategic AI Conversations: LLMs enable strategic dialogue acting as expert assistants to help investors refine their strategies, generate new ideas, and sharpen their competitive edge in financial markets.
- Comprehensive Knowledge Base: Large Language Models function as dynamic knowledge bases, providing immediate access to extensive information on companies, industries, science and technology, and more. This capability is invaluable for improving investment decisions, even in unfamiliar domains. These models encapsulate a wealth of knowledge, having been trained on a broad array of sources including Wikipedia, scientific articles, books, and more, allowing users to query this information in natural language rather than relying on traditional SQL database queries.
- Code Generation and Software Development: Automating coding tasks from creating quant models to portfolio analysis tools, boosting efficiency and reducing costs.

Synergistic advances in AI

Recent AI developments are not sudden but the result of long-term trends. Understanding these trends is essential—not only to predict where AI is heading, but to recognize that today's innovations are just the beginning. For business leaders and strategists, grasping these developments is key to preparing for an uncertain future. The figure below describes the three main forces driving AI advancements: The accelerated increase in compute power, breakthroughs in AI architectures and algorithms and increased volumes of data on which AI models are trained.

The confluence of these factors results in a massive acceleration in the ability to intelligently synthesize and leverage data.



Compute Power

The advent of GPUs (Graphics Processing Units) and TPUs[1] (Tensor Processing Units) has revolutionized the power available to train and use AI systems (NVID-IA's fastest computers in a single rack now exceeding the petaflop[2] limits (https://www.nvidia.com/en-us/data-center/gb200-nvl72/) for deep learning performance). These AI-optimized architectures are designed to efficiently ex-ploit the parallel nature computations at the heart of artificial neural networks (which underpin deep learning – the main technique behind modern AI), significantly enhancing AI's processing capabilities.



Algorithmic breakthroughs

Transformer-based models like ChatGPT, have significantly improved AI's natural language processing and generation capabilities. These models analyze text as sequences of 'tokens' — essentially, characters — which are converted into numerical data to capture linguistic features and meanings. This process allows for accurate predictions of subsequent tokens, enhancing the model's ability to generate coherent and contextually relevant responses. One challenge with models like GPT, however, is their scalability: Processing longer texts, such as an entire textbook, is currently impractical. Ongoing research to create more efficient and scalable AI algorithms promises to address the constraints of current models.



Training data volume

The race to develop superior AI models is marked by a significant increase in the volume of data (>10 trillion words for training current LLMs) used for training, alongside the computational power and the number (>1 trillion in current LLMs) of parameters within these models. This trend is evident in the emergence of models like GPT-4, Gemini, and Claude 3, each claiming to surpass its predecessors in various cognitive tasks.

Figure 4: The driving forces behind the AI revolution. Notes to the figure: [1]TPUs are computer chips designed for high volume low precision computation specifically designed for machine learning applications. see https://en.wikipedia.org/wiki/Tensor_Processing_Unit; [2]A petaflop is a measure for a computer's performance. One petaflop is the capability to do a 1,000 trillion (or one quadrillion) calculations per second on floating point numbers i.e. numbers with decimal points e.g. '5.4518'.

While the driving forces boosted AI development, they do not directly address the expected utility for specific sectors like asset management. Historically, powerful AI systems such as dedicated chess computers represented the forefront of AI, yet offered limited value in other applications such as asset management. Therefore, understanding the evolution of AI intelligence is crucial, reminding us that systems like ChatGPT are waypoints, not endpoints, in AI's ongoing development.

The rapid advancement in AI, particularly evident in LLMs, reflects enhancements not just in the amount of compute power but also in how AI models use their specific architectural design to achieve smarter outputs. For instance, the progression from GPT-3 to GPT-4 isn't merely about increased computation; it involves increasing the number of a model's parameters and feeding it more training data, significantly boosting the system's intelligence.

High-frequency trading systems, despite their relative simplicity, also benefit from both faster processing and larger datasets, which directly correlates with their ability to capitalize on rapid market changes. Whilst quantitative models in finance can handle larger data sets, their cognitive improvements are capped by the fundamental limitations of their algorithmic designs.

While current LLMs excel in certain tasks, they often falter in areas requiring complex logical, symbolic, or commonsense reasoning. Critics note that LLMs, at their core, function as advanced token predictors, lacking a genuine understanding of the concepts or the physical world they reference. Such limitations can lead to decision errors when naively applying these systems in sectors such as asset management.

Interestingly, LLMs display forms of emergent behaviour; for example, they understand the structure of human and computer languages without explicit training. Furthermore, AI-generated video from textual descriptions (see e.g. https://sora.aitubo.ai/), can accurately depict physical interactions hinting at an underlying, albeit not explicitly programmed, grasp of physical principles. Thus, it should not be surprising if future versions of GPT show improved performance in abilities such as logical reasoning.

Another promising direction in AI is the development of hybrid AI systems that integrate specialized modules for specific tasks. For instance, combining modules for logical reasoning with LLMs could yield systems that have capabilities exceeding human reasoning and at a much larger scale, with vastly superior knowledge and reasoning capabilities.

As we look to the future, the implications for asset managers of more advanced AI systems are profound. Currently, LLMs are primarily used as preprocessing tools that convert textual information into numerical data, which can then be incorporated into traditional quantitative models. However, we are beginning to see the potential for much more integrated systems. These future systems could seamlessly reason over both textual and numerical data, combining reasoning and language interpretation skills with statistical forecasting and simulation capabilities to make forward-looking predictions.

04 Strategic Deployment of Large Language Models in Asset Management

The rapid evolution of AI presents both challenges and opportunities in asset management. Several key developments, while interconnected, each offer unique contributions that together are poised to reshape the industry:

1. Incorporating Unstructured Data: As illustrated earlier in figure 1, AI is revolutionizing investment analysis by integrating qualitative data into quantitative frameworks. This approach enables a nuanced analysis that blends traditional investment expertise and modeling with the speed of quantitative strategies, revealing insights at a previously unattainable scale. The ability to ingest and process a variety of non-traditional data sources greatly enhances the scope for strategies to benefit from this data.

2. Managing Complex Investment Problems: AI enhances traditional investment optimization by incorporating both sophisticated constraints and objectives into the decision-making process. This includes the dynamic integration of real-time, unstructured data alongside traditional financial metrics. While this builds on the capability to handle unstructured data, it specifically addresses the complexity and holistic nature of optimizing portfolios in real-time.

3. Democratizing Portfolio Management: AI is simplifying sophisticated portfolio management by enabling individual investors to define their risk preferences and receive tailored recommendations, akin to having a personal financial advisor.

4. Enhancing Market Efficiency: By processing and acting on information swiftly, AI is expected to increase market efficiency. Micro-inefficiencies become easier to arbitrage on, which means that information is included in the pricing of instruments more quickly. This point highlights the impact on market dynamics and the need for agile strategies to keep up with faster information processing.

5. Accessing Customized Investment Products: AI enables the creation of tailored investment products, such as customized thematic funds. These products align closely with personal or ethical preferences, offering more personalized options beyond traditional index funds and ETFs. This emphasizes the creation of new, tailored financial products catering to individual needs.

6. Providing Consistent Investment Decisions: AI can replace discretionary, more subjective portfolio management decisions with a more consistent systematic application of qualitative information and increased context awareness. This both reduc-

1. Examples of hybrid AI architectures are hierarchical mixtures of experts, multi-modal learning, neuro-symbolic AI, en-semble methods, and LLMs combined with reinforcement learning architectures like used in alphaGo see e.g. https://github.com/microsoft/Everything-of-Thoughts-XoT.

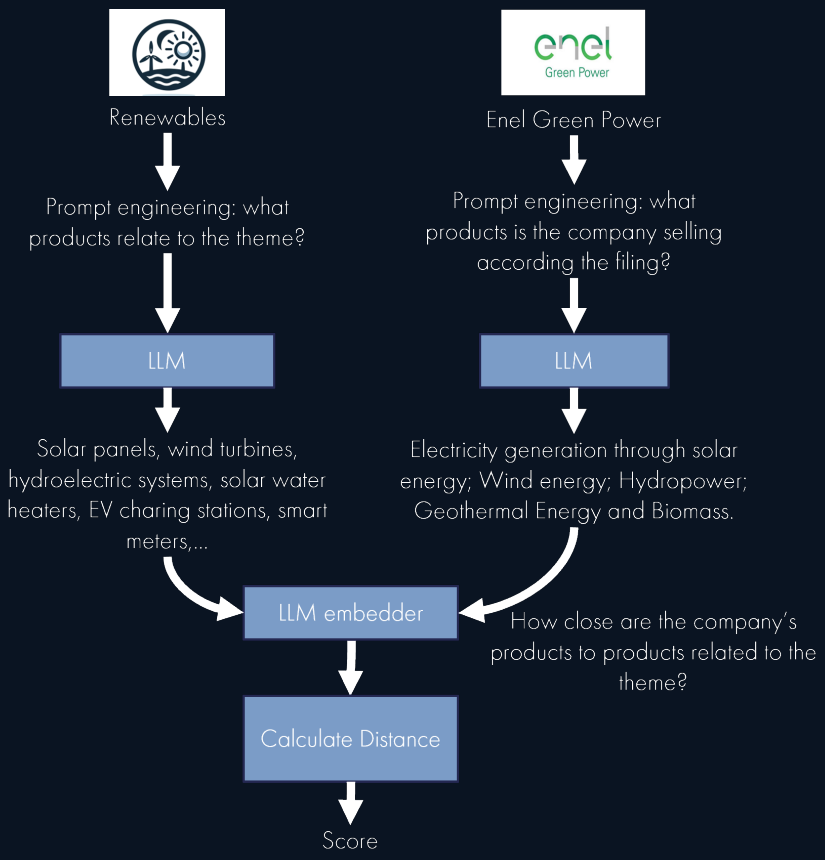
The extent to which asset managers keep up with technological advancement will increasingly prove to be a distinguishing factor for success. Even though the access to basic tooling is now becoming more widespread (e.g. ChatGPT and low-coding/no-coding solutions), more advanced uses of the opportunities AI offers are by no means commonplace. It is therefore important for asset managers to determine their strategy with respect to onboarding, maintaining and improving their AI capabilities.

Example – Using Large Language Models to Quantifying companies' similarity to investment themes

An example of AI's capability to create sophisticated investment products can be seen at Entis, where AI-enabled solutions swiftly derive quantifiable metrics to score publicly listed companies on any investment theme. Whether it's the future of mobility, healthcare innovations, or other emerging sectors, the system utilizes LLMs and other NLP methods to precisely evaluate companies' exposure to these themes based on revenue and patent activities. Once established, our system continuously monitors these companies, updating their thematic scores daily.

These precise thematic exposure metrics facilitate the creation of thematic indices, which can serve as the foundation for specialized investment funds. This approach not only offers investors targeted investment opportunities but also enhances the adaptability and responsiveness of financial products to market trends and technological advancements.

A simplified* picture of Entis' process to apply prompt engineering and LLMs to quantify the similarity between a theme and company texts.



*Note that while the application of LLMs to quantify thematic exposure appears straightforward, the practical implementation involves significant challenges. Ensuring the accuracy and reliability of the data is crucial, as the quality of input directly impacts the system's output. Rigorous validation processes are also necessary to maintain the integrity of the thematic scores. These processes include continuous verification of data sources and frequent updates to the models to align with the latest market developments and technological innovations. We will explore these challenges and the strategies to address them in more depth in the following chapter, providing a more detailed insight into the complexities of using AI to drive investment decisions

Managing Risks

As AI technologies reshape asset management, they introduce risks and practical implications that necessitate careful navigation. Regulatory requirements and new legislation will be instrumental in mitigating these risks, ensuring that AI's transformative power is gainfully harnessed. The challenges of complying with comprehensive regulatory standards suggest that large-scale systematic investment professionals, such as asset managers and index providers, will likely remain influential. These established players have the expertise and infra-structure to navigate complex regulatory environments, which can be a significant barrier for new-comers. The financial sector places a high premium on safety and reputation, areas where established firms have a considerable advantage.
However, innovation and competition, particularly from tech giants with advanced AI models, will drive rapid adaptation within the financial sector. The speed of these changes is not to be underestimated.

Identified Risks and Mitigation Strategies

Mastering the skill of 'prompt engineering'—the art of crafting precise queries and providing the right contextual information (see Appendix II) is critical for leveraging LLMs effectively. This approach helps constrain AI to use information pertinent to specific queries, resulting in higher quality responses and reduced 'noise'. Additionally, requiring evidence from an LLM for its conclu-sions can mitigate risk and improve interpretability and confidence in decision making. Appendix II provides examples of prompt engineering and guidelines on how to enforce evidence require-ments in model responses. Validating outputs through indepen-dent models and cross-validation techniques further supports the robustness of findings.

The historical preference for rule-based systems or human judgment in investment decisions often stems from their transparency and perceived reliability.

Traditional rule-based methods might define investment portfolios based on inclusion rules such as market capitalization thresh-olds, industry, and country exposures.
However, AI models, particularly those based on neural networks like LLMs, do not readily offer clear insights into their deci-sion-making processes. To enhance transparency, it is crucial to demand evidence and explanations from AI systems, like justifi-ca-tions provided by human experts. By integrating AI's cognitive capabilities into these rule-based frameworks, decisions can adhere more closely to predefined criteria, thus enhancing both compliance and transparency.

Table 1 presents an overview of some of the key risks with utilizing LLMs in asset management, alongside practical tips and mitigation strategies.

Enhancing Efficiency and Effectiveness in AI Utilization

The successful integration of LLMs into asset management requires strategic decision-making around cost management and the procurement of necessary expertise. Deciding whether to buy the required expertise, collaborate with others, or initiate a propri-etary solution is crucial. This requires careful cost-benefit analysis.

Strategic AI integration decisions for Asset Managers:

- Cost-effective AI integration: Directly training LLMs from scratch is expensive and unlikely to yield better results than utilizing state-of-the-art models developed by leading AI organizations. A proprietary approach requires significant investment, both in terms of financial and human resources as well as in governance.
- Fine-Tuning and Model Utilization: A more practical approach involves fine-tuning an existing state-of-the-art model to tailor it to specific needs, such as analyzing specialized data relevant to sustainable investing. Fine-tuning involves making small adjustments to these pre-trained models to improve their performance on specific datasets, which is generally more cost-ef-fective and less complex than building a model from the ground up. While this process can enhance model performance for particular tasks, it requires a nuanced understanding of AI technologies. Asset managers should assess their team's capabili-ties and consider partnerships or specialized training to acquire the necessary skills for effective fine-tuning.

- Prompt engineering (see appendix II): The effectiveness of LLMs can be significantly influenced by the skill of prompt engineering. Effective prompt engineering can yield superior results for reasonable costs, making it a critical skill for asset managers.
- Selective model component use: For certain tasks, such as text classification, using specific components of LLMs (e.g., the embedder component, see figure 2) can offer more precise and cost-effective solutions.

	Mitigation Strategy / Tips
Reliability and Quality concerns	**Provide the LLM with the latest information.** When seeking insights on specific, time-sensitive matters, incorporate the most recent data into your prompt. This ensures that the LLM's analysis is relevant and based on the latest information.
	Critically Evaluate LLM Responses: Validate responses against current data to avoid reliance on outdated or inaccurate information.
	Apply Human Oversight: Ensure decisions are understandable and grounded in expert judgment to improve reliability.
	Distinguish Reasoning from Knowledge: Recognize that LLM knowledge is extensive yet static, confined to its last training. Use LLMs for processing and analyzing information, not as a source for current facts.
	Address Bias and Misinformation: Actively use LLMs for text comprehension and interpretation rather than as sources of knowledge. Be aware that LLMs require contextual knowledge about language and entities, which can be imperfect. For example, differentiate between asking LLMs for growth expectations based on general knowledge versus specific text-based queries where external information is not used.
Privacy and security concerns	**Secure IT Infrastructure:** Ensure robust security measures are in place to protect the IT infrastructure handling AI models.
	Manage Sensitive Data: Avoid working with privacy-sensitive data unless absolutely necessary and ensure compliance with privacy regulations.
	Verify AI Vendor Privacy Safeguards: Regularly check the privacy measures provided by AI vendors to ensure that prompts and responses remain confidential and are not used or accessed by third parties.
Black box decision making	**Demand Explanations:** Require clear justifications for the LLM's reasoning to increase transparency.
	Combine Insights with Human Oversight: Validate LLM insights through expert review to ensure accuracy and relevance.
	Generate Hypotheses for Further Investigation: Use LLMs as a starting point for deeper expert analysis.
	Integrate Within Rule-Based Frameworks: Align LLM outputs with established criteria to enhance compliance and transparency.
Lack of source information	**Provide Original Source Data:** Supply LLMs with direct access to original data sources to enhance insight accuracy.
	Establish Controlled Data Systems: Develop systems to manage and control data quality systematically.
	Acknowledge and Supplement Limitations: Recognize the limits of LLMs and augment with data from public and commercial providers as needed.
	Identify Data Gaps: Utilize LLMs to pinpoint missing information and guide data acquisition strategies.
	Addressing Data Unavailability: Plan for scenarios where critical data, like a company's latest annual report, is unavailable. Develop strategies to obtain necessary documents or identify alternative sources.

Table 1: Mitigating risks in leveraging Large Language Models for asset management - This table outlines key risks associated with the use of Large Language Models (LLMs) in asset management and provides tips and mitigation strategies to address these challenges.

Organizing the Infrastructure Around AI Models

Working with AI models, especially with the LLMs accessible through chat interfaces, is relatively straightforward and allows for easy updates when new versions are released. However, an often underestimated challenge is the need to invest in organizing the surrounding infrastructure, as illustrated in Figure 5, effectively.

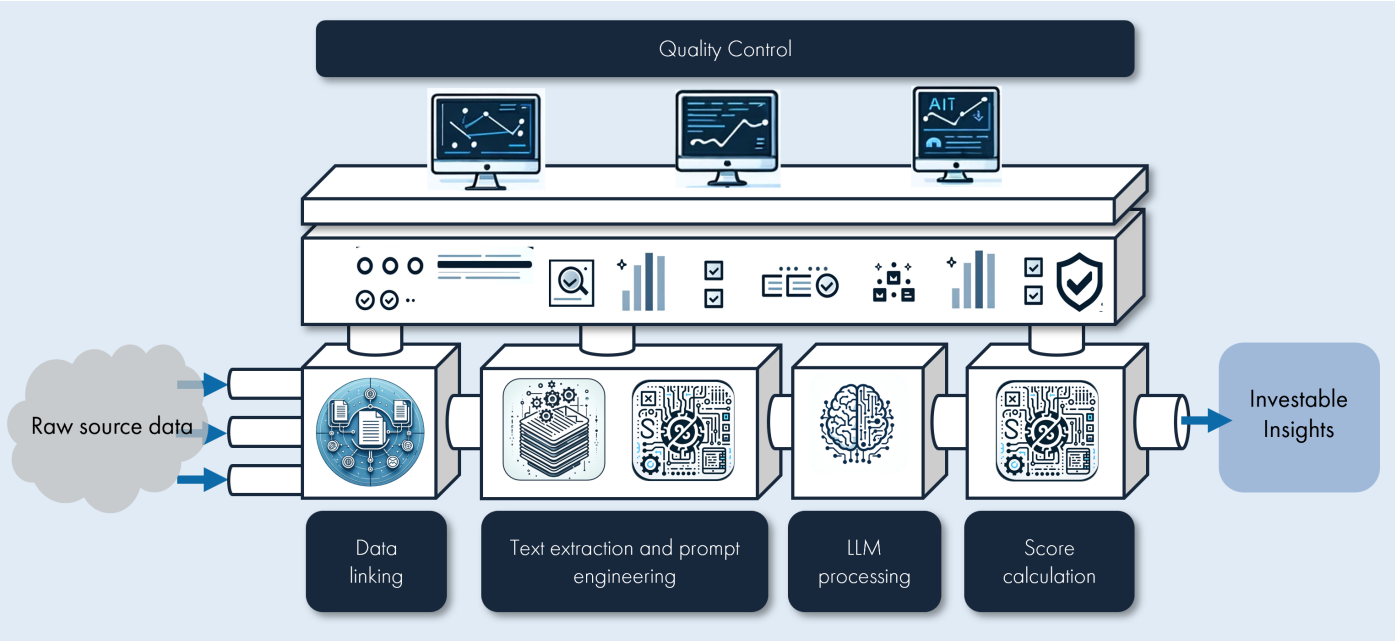


Figure 5: From raw source data to investable insights - Data processing infrastructure for utilizing AI models such as LLMs. Schematic overview.

Reliable AI deployment depends on providing accurate data and formulating precise queries through prompt engineering. Steps essential for validation include verifying responses against reference data and validation by human experts.
Setting up a robust data intake process for LLMs requires significant investment, especially for scalable, highly automated solutions aimed at generating insights across sectors, countries, and covering thousands of companies.
For instance, extracting insights from bulk annual reports using LLMs presents several challenges before even engaging with the models. Reports must be automatically sourced and correctly matched to the corresponding stocks or bonds, ensuring they are not mistakenly linked to subsidiary companies or other related but incorrect entities. While asset managers already spend significant effort solving data linking challenges, utilizing LLMs requires linking these new, often untagged text sources to identifier systems like SEDOL, CUSIP, and ISIN, which is more complex due to these documents not being directly associated with such identifiers.

Once data is linked, the next hurdle is determining which information to present to the LLM. Asking questions based on an entire annual report is impractical; these documents are too lengthy and mixed with text, tables, and figures, which current LLMs struggle to process. Identifying key sections relevant to specific interests, such as business segment descriptions, may require specialized text extraction software, posing a difficult buy-or-build decision. Commercial software may not precisely fit your needs, suggesting a custom solution might be more effective to parse and ingest the data.
After extracting and verifying the relevant text and data, the next step involves crafting well-developed questions to ensure the LLM delivers accurate and usable results without drawing from unwanted sources or providing hallucinated content. As previously explained in the section on prompt engineering, these queries must be meticulously designed to ensure that responses are formatted for further automated processing, such as quantitative analysis.

This overview highlights that leveraging AI models like LLMs for large-scale, automated applications—essential for managing extensive investment portfolios or maintaining thematic indices—requires more than just having powerful models. It necessitates a well-organized process surrounding the AI systems. Asset managers also face the buy-or-build decision: purchasing investable insights from parties that have already invested in such platforms versus creating proprietary systems. The critical question remains: What competitive advantage does an asset manager retain if they base decisions on insights available to others?

The above underlines the importance of a clear strategy. Decisions have potentially far-reaching consequences with respect to the product offering of asset managers as well as to the skillsets required from personnel. Investment professionals may increasingly have to be well-versed in AI in order to keep up with the knowledge needed, and their processes will have to adapt accordingly. This will add to the changes which are already visible, going from idiosyncratic to rule-based investment. Likely, this will not only affect the public markets side of the asset management market -- private market investments will also be impacted by this change. As professionals in this market are less exposed to large datasets, the consequences may be even more pronounced as the gap between traditional and innovative approaches is likely to be larger.

05 Conclusion: Navigating the AI revolution in asset management

This whitepaper has explored the transformative impact of AI on asset management, driven by rapid advancements in computational power, algorithmic innovations, and increasing data availability. These developments have significantly enhanced AI's ability to process and analyze complex information, thus reshaping investment strategies and operational frameworks within the industry.

Despite its potential, the rapid evolution of AI also introduces substantial challenges and risks, such as biases, reasoning flaws and transparency issues, which demand careful management. Effective integration of AI goes beyond technology adoption, requiring a deep understanding of their potential and limitations.

We have emphasized the need to build a robust infrastructure around AI technologies, which necessitates a strategic overhaul of organizational processes to harness AI's full capabilities effectively. While this presents unparalleled opportunities for asset managers, it also challenges them to keep up with the rapid pace of technological advancement.

AI's strength in asset management lies in its enhanced ability to identify true value, from recognizing promising technologies to spotting well-managed companies poised for impact. It improves asset managers' capacity to understand the interconnectedness of business activities and to manage uncertainties more effectively. AI facilitates operations on a larger scale and adapts to personal preferences, making sophisticated investment strategies more accessible, even to those with minimal programming and mathematical skills. Thoughtful AI integration seeks not only to improve financial outcomes but also to drive substantial societal benefits by channelling capital towards ventures that genuinely advance global progress.

Still, AI is not a panacea. As initial excitement settles, its practical limitations become clearer. Asset management professionals must adopt a balanced view, acknowledging that while AI can significantly enhance decision-making and efficiency, it must be handled with care to prevent overreliance on automated systems that are left unchecked.

Looking forward, AI will continue to drive innovation in financial products, making them more customized and responsive to global shifts and investor preferences.

As we approach this AI-driven transformation, asset management firms that engage proactively with these technologies, strategically plan, and dynamically adapt, will not just survive but thrive. Industry leaders should aim to lead change, ensuring that AI's integration into asset management maximizes benefits for all stakeholders in the financial ecosystem.

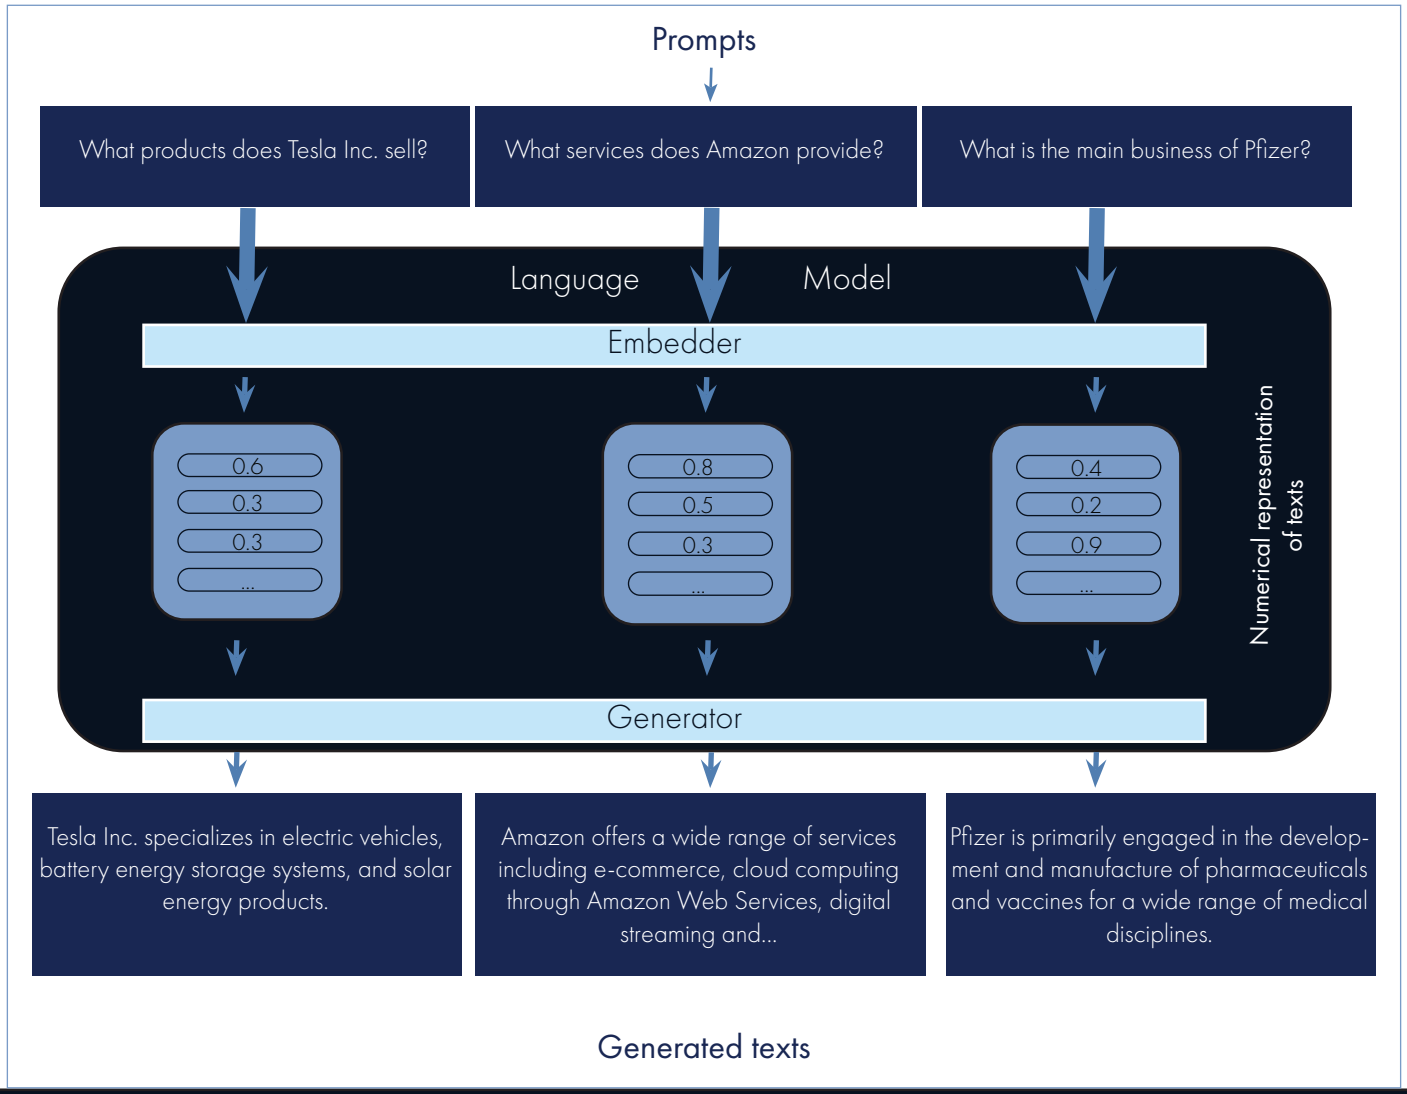
Appendix I - Understanding Large Language Models, Transformers, and Generative AI

Understanding Large Language Models, Transformers, and Generative AI

In the world of artificial intelligence, understanding natural language—how we communicate through speech and writing—is a complex challenge. To address this, researchers have developed advanced models capable of not only understanding but also generating human-like text. Here's a straightforward breakdown:

- Large Language Models (LLMs): These are powerful tools designed to process and generate human language. Think of them as highly advanced programs that can understand and produce text based on the training they've received from a vast database of written language.
- Transformer Models: LLMs often use a special architecture known as "transformers." These are designed to handle sequences of data (like sentences) and are exceptionally good at understanding the context and relationships between words in a sentence. Transformers are a type of deep neural network, which means they perform their tasks through layers of processing that help determine the meaning of text and predict what comes next.
- Generative AI Models: Within the world of LLMs, some are specifically geared toward generating text. When you ask a question, the model generates an answer based on patterns and examples it has learned during training. It does this by transforming the text of your question into numerical data (vectors), which capture the deeper characteristics of the text. The AI then uses these vectors to generate a response, predicting which words (or tokens) are likely to follow.

How It Works: Imagine feeding a few sentences into a "box" (the language model). Inside this box, the text is converted into numbers that uniquely represent the content and meaning of those sentences. This numerical transformation allows the model to compare different texts and understand how similar or different they are. Similarly, when generating text, the model uses these



Appendix II - Prompt Engineering in Asset Management

Simple example

Scenario: An asset manager wants to evaluate the investment potential of a technology company, specifically looking at its financial health and market growth opportunities.

Basic Prompt:
- **Prompt**: "Tell me about Apple's financial health."
- **Potential AI Response:** "Apple Inc. is financially stable."

This basic prompt is too vague and the response is overly broad, lacking actionable insights.

Engineered Prompt:

- Prompt: "Provide a detailed analysis of Apple's financial health based on the latest quarterly earnings, focusing on revenue growth, profit margins, debt levels, and cash flow. Compare these metrics against the industry average and highlight any financial risks that could affect its investment potential over the next year."
- Expected AI Response: "In the latest quarterly report, Apple Inc. reported a revenue growth of 6%, which is above the industry average of 4%. The profit margin stood at 21%, consistent with the sector. Debt levels have increased slightly from the previous quarter, now representing 30% of the asset value, but cash flow remains robust at $15 billion, indicating good liquidity. Potential financial risks include increased borrowing costs and potential supply chain disruptions in Asia, which could impact profitability margins over the next year.

Systematic prompt engineering - Example of Building an Optimized Prompt

Let's refine the example with Apple's financial analysis to demonstrate the iterative refinement towards an optimized prompt:

- Initial Prompt: "Provide a detailed analysis of Apple's financial health based on the latest quarterly earnings, focusing on revenue growth, profit margins, debt levels, and cash flow."
- AI Response: Gives basic figures without page references or sufficient breakdown between different types of cash flows.
- Refined Prompt: "Provide a detailed analysis of Apple's financial health based on the latest quarterly earnings. Focus on revenue growth, profit margins, debt levels, and differentiate between operating cash flow and free cash flow. Include page references from the earnings report for all data mentioned."
- AI Response: Responds with detailed metrics including specific types of cash flow and page references.
- Evaluation and Further Refinement (if necessary): Check if all financial metrics are covered as required, the accuracy of page references, and the clarity of financial terms used. If further specifics are needed, refine the prompt again.
- Optimized Prompt (Final): Once the prompt reliably yields comprehensive, accurate, and fully referenced data, finalize it for storage: "Analyze Apple's quarterly financial health detailing revenue growth, profit margins, debt levels, operating and free cash flow, specifying page references for each metric in the report."

Utilizing the Optimized Prompt

This final, optimized prompt can then be saved in a prompt library and used as a standard query format for similar financial analysis tasks across different companies or quarterly reports. This systematizes the process, allowing for scalable and efficient financial analyses using AI, with each prompt tailored to retrieve the maximum amount of relevant and verifiable information.

Contact Information

Entis b.v.

Papendorseweg 100
3528 BJ Utrecht
The Netherlands

info@entis.ai

+31(0)30 7991140

Visit the website for more information: www.entis.ai

