



GPT for Financial Advice

The Combination of Large Language Models
and Rule-Based Systems

Institute of Financial Services Zug IFZ
www.hslu.ch/ifz

Partners

e.foresight

finnova.

 SWISS FINTECH
INNOVATIONS



SWISS BANKERS

Contents

Preface	2
1 Introduction	3
2 Description of Large Language Models	4
3 Description of the Rules and Requirements for Financial Advisory	5
4 Combination of Rule-Based Systems and Large Language Models	6
5 Prototype for Investment Advice	7
6 Other Use Cases	8
7 Conclusion and Outlook	9
Authors and Contributors	10
Bibliography	11

Preface

Large language models (LLMs) like OpenAI's GPTs or Google's Bard belong to the latest developments in natural language processing and are currently used and explored in a wide range of applications, from chatbots to machine translation. In simple terms, LLMs are computer programs that are able to generate text word by word based on what they have learned from large data sets. The program predicts the most likely next word based on the previous words in the conversation. Such LLMs are also increasingly prevalent in the financial sector and are likely to influence it even more in the longer term. However, the compatibility between such solutions and certain financial services requirements is not fundamentally guaranteed. A relevant example from the Swiss financial industry is investment advice, where traceable and explainable recommendations are required, which, however, may conflict with the probabilistic character of LLMs. In addition, investment advice is one of the most important services offered by the Swiss financial industry, which is why it seems sensible to already address the potential of new technological developments at an early stage.

The objective of this study is to explore how LLMs and rule-based systems can be combined for investment advice. Furthermore, the study aims to create a prototype that showcases the benefits of integrating these two concepts, without, however, meeting all regulatory requirements (e.g., from the Banking Act or the Federal Data Protection Act) for AI-based investment advice. Hence, the focus lies on demonstrating the technical feasibility of using LLMs to support deterministic investment recommendations, rather than discussing potential obstacles or non-technical enablers. The solution obtained shows that the combination of probabilistic LLMs and rule-based systems, while retaining the advantages of both approaches, is possible and that the prototype works well in most cases. The findings from this study and the prototype can be used by the financial sector as a starting point for discussion and the development of more sophisticated solutions.

The study is structured as follows: Chapter 1 introduces the topic and highlights its actuality and relevance. Chapter 2 discusses how LLMs work and describes their corresponding properties, such as their probabilistic nature. Chapter 3 introduces significant requirements for investment advice that exist today. In Chapter 4, the findings from the previous two chapters are combined. Specifically, various approaches to using LLMs for specific use cases are presented, and the one that appears to be the most fitting for investment advice is identified. Chapter 5 describes the design of the implemented prototype for investment advice, and Chapter 6 presents additional (possible) use cases. Chapter 7 summarises the findings.

At this point, we would like to thank the sponsors of the IFZ FinTech Program, who have supported this condensed study financially as well as in terms of content. These are e.foresight, Finnova, Swiss FinTech Innovations, SIX, and Swiss Bankers Prepaid Services. In addition, we would like to thank all the contributors who supported the study in various ways such as providing content, engaging discussions, or with their reviews.

Thomas Ankenbrand
HSLU

Cornelia Stengel
SFTI

Denis Bieri
HSLU

Stephanie Wickihalder
SFTI

Levin Reichmuth
HSLU

Ege Yilmaz
HSLU

1. Introduction

The notion of intelligent behaviour exhibited by machines rather than humans has been around for several decades. Corresponding developments and approaches are often summarised under the term “Artificial Intelligence (AI)”, which was introduced in 1956 and defined as a field of research in computer science (McCarthy et al., 2006). But thoughts about the intelligence of machines were already conceptualised a few years earlier. One of the concepts still widely cited today is the Turing test. This test, originally described in 1950 by Alan Turing (Turing, 2009), examines if a person can determine whether the answer to a question was given by a human or a computer. If this is not possible, the computer passes the test and is said to be able to imitate human intelligence. There is no general consensus on which AI systems have passed the test so far, as its definition and the criteria associated with it have changed over time.

Conversational AI solutions in particular, often called “chatbots”, have gained attention in recent months with the launch of corresponding natural language processing tools such as ChatGPT by OpenAI in November 2022 (OpenAI, 2022b) or Bard by Google in February 2023 (Google, 2023) due to their technological potential for a broad range of applications and domains. ChatGPT even holds the record for the fastest growing user base: Just two months after its launch, the application reached 100 million monthly active users (Hu, 2023).

However, current developments in AI go far beyond the generation of text, and also includes other formats such as images, audio, and video, with the general field being summarised under the term “Generative AI (GenAI)”. The great interest in this field is also reflected in the investment volumes in corresponding companies, which amounted to USD 2.7 billion in 110 deals in the year 2022 (CB Insights, 2023). The general technological development in the field of AI, including GenAI, has progressed so rapidly that some experts are already discussing when so-called “Artificial General Intelligence (AGI)” will become a reality. (Tamim, 2023). AGI has general, non-task-specific capabilities that are indistinguishable from human intelligence.

In addition to the opportunities presented by current and future forms of AI, these solutions also pose potential risks. Prominent examples from a social perspective are the effect of job disruption and the infection of current information sources with misinformation (Thornhill, 2023). For such reasons, business leaders and AI experts have called for a temporary halt to large-scale AI experiments, allowing to first develop and implement security and governance protocols for advanced AI systems (Future of Life Institute, 2023). Furthermore, concerns were also expressed about the data privacy issues of these systems. As a result, Italy has banned the use of ChatGPT as of the end of March 2023 (GDP, 2023) by geo-blocking respective IP addresses, but enabled the service again after OpenAI announced a set of privacy controls and disclosures (TechCrunch, 2023).

Like most industries, banking will increasingly be influenced by AI. This is underlined by the expected global market size of AI in the financial services industry, which is projected to reach USD 64 billion by the year 2030 (Prasad et al., 2021). LLMs are anticipated to be a significant driver of this growth, as they provide great potential for, for example, improving the customer journey and cutting costs for financial institutions by automating parts of the banking value chain or helping with decision making (Tomych, 2023).

Besides the opportunities, however, the banking sector must also consider the potential dangers, vulnerabilities, and costs of AI solutions. In addition, the characteristics of the technology, such as its often non-deterministic nature, must also be assessed in the context of the suitability for financial products and services, also taking into account applicable regulation. This study attempts to provide an initial assessment of whether rule-based investment advice can be combined with the conversational capabilities of LLMs and, if so, to describe and realise a corresponding introductory prototype. Hence, the focus lies on technical feasibility, which is why the study does not review other factors in detail (e.g., legal, social, or ethical) which are also highly relevant to the successful application of AI-systems in the financial industry.

2. Description of Large Language Models

Currently, the most powerful conversational AI solutions are based on approaches from the field of large language models. These models can be defined as follows:



A large language model (LLM) is a language model consisting of a neural network with billions of parameters trained on large amounts of unlabelled data by means of self-supervised learning and used, for example, to predict and generate text and other content (Sejnowski, 2023).

A specific type of LLM is the “Generative Pre-Trained Transformer (GPT)”, whose simplified design is illustrated in Figure 1. GPTs consist of different layers. At the core are the multiple decoders, each of which consists of the “Self-Attention” and a “Feed Forward Neural Network”. The self-attention mechanism allows the model to weigh the importance of the words and phrases of the input (converted into a numerical embedding) in order to understand its meaning and context, while the feed forward neural network aims to capture nonlinear relationships between input and output. The output of the neural net is then fed into the “Next Word Prediction Head” which produces a probability distribution of the vocabulary to identify the most likely next word, which then is added to the input prompt and

the procedure starts all over again.¹ Since GPTs aim to capture the complex relationships between words in natural language, the models are characterised by billions of parameters that are determined in a training process using large data sets. The high complexity of the underlying models means that GPTs are not considered explainable AI, i.e., their input-output relationship cannot be easily interpreted (Basu et al., 2021).

In other terms, current GPTs attempt word-by-word text generation sequentially, based on the probability distributions of words and phrases from the training set, to give coherent responses to an input (e.g., a question). As a result, such models are probabilistic in nature, implying that the model’s responses can be sensitive to details of the wording or phrasing of the prompts (Bubeck et al., 2023). A similar question can therefore be answered differently, depending on the structure of the corresponding input. Therefore, GPTs are not necessarily always output-consistent. Furthermore, the probabilistic nature of GPTs can produce content that is nonsensical or untrue, which is often referred to as “hallucination” (Manakul et al., 2023).

The technical properties of GPTs need to be assessed against the financial sector requirements, discussed in the next chapter, in order to develop suitable solutions.

¹ See Vaswani et al. (2017) for a more technical description of GPTs.

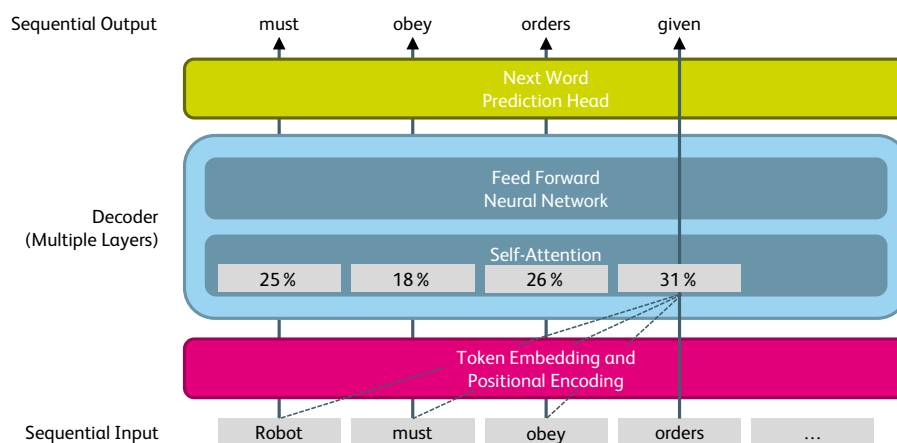


Figure 1: Simplified Architecture of a Generative Pre-Trained Transformer (GPT) (source: based on Alammari (2019))

3. Description of the Rules and Requirements for Financial Advisory

With securities holdings in bank custody accounts totalling CHF 6.8 trillion at the end of 2022 (Swiss National Bank, 2023), Swiss banks play an important role in wealth management. Within wealth management, investment advice is one of the fundamental services a bank offers to its clientele and, according to the Federal Act on Financial Services (FinSA), is described as follows:



Investment advice is the provision of personal recommendations on transactions with financial instruments.

If a bank, or any other player, wants to operate in this area with an AI-driven product, various regulatory requirements must be adhered to. This section outlines some of these requirements and highlights the legal challenges that need to be addressed when planning the use of AI-driven systems in general and especially in the context of the provision of financial services.

When providing investment advice, the specific regulations of the FinSA must be followed first and foremost. In addition, however, various other laws must be considered. These include the Code of Obligations (CO), the Banking Act (BankA) and possibly other applicable financial market laws and regulations, as well as the Federal Act on Data Protection (FADP). These laws give rise, for example, but not conclusively, to the following requirements:

- **Information and examination duties:** Investment advisors must fulfil various information obligations as part of the advisory process, as well as conduct appropriateness and suitability checks of their clients. If they fail to do so, they are subject to civil liability as well as criminal liability.
- **Registration obligations:** Investment advisors may have to register in the advisor register as well as join the ombudsman's office.
- **Contractual obligations:** Investment advisors are contractually obliged to act in the best inter-

ests of their clients, otherwise they can be subject to contractual liability to their clients.

- **Data privacy and banking secrecy:** Investment advisors must comply with confidentiality provisions, which serves to protect financial privacy (SwissBanking, online). In addition, the provision of sensitive data by the user can also be problematic from a privacy perspective, as it could be used for further training of the underlying LLM.

The difficulty in applying these regulations to the AI-driven prototype for investment advice (or other AI systems) lies in the fact that initially these regulations were not created for this specific use case. Therefore, their implementation is often ambiguous and needs to be examined closely.

Regardless of their use in the provision of financial services or in other areas, various risks are cited in the use of AI systems that have not yet been fully clarified at the legal level. These risks relate to the potential discrimination and bias of AI systems, the risk of systematic manipulation of human actions, the limited transparency of its decisions (which can lead to so-called “black box decisions”), or the complicated situation with regard to criminal responsibility or civil liability when using AI. While many of these risks can be addressed with existing legislation (such as the FADP), further regulatory adjustments are needed to effectively cover AI systems and provide legal certainty in the use of AI systems (Stengel, 2023). The EU is pursuing an AI-specific approach in this regard and is in the process of drafting an AI regulation, while Switzerland currently pursues a technology-neutral approach with selective additions to existing laws (Dobler, 2023).

Therefore, exploiting the potential of AI systems such as probabilistic LLMs while taking into account practical requirements is no trivial task, as they often interfere with each other. In the following, we discuss different approaches on how LLMs can be used for specific use cases, identifying the approach that best serves rule-based investment advice.

4. Combination of Rule-Based Systems and Large Language Models

Approaches to using generalist LLMs for specific use cases vary. The general options, for some of which the first corresponding applications in the financial sector can also be observed, include:

- **Own trained model:** LLMs can be completely independent and newly trained on the basis of self-selected data. One example hereof is BloombergGPT which is trained with an extensive data set, containing both domain-specific, and widely used public data. The model seems not only to outperform on probabilistic, financial tasks, but also reaches proficient results in more general applications (Wu et al., 2023).
- **Fine-tuning:** Fine-tuning implies the use of an already pre-trained LLM, like GPT-4, and allows customising such a model for a specific application by further training it using domain-specific data (OpenAI, 2022a). The final model can be restricted to domain-specific conversation, i.e., allowing the fine-tuned model to only answer in a scope as defined by the provider (see, e.g., Clara by Helvetia (2023)).
- **Plug-in:** A plug-in grants an existing LLM access to additional information (e.g., on company specific services and products), performing computations and use-case specific third-party services (Adebisi, 2023). Thereby, plug-ins do not alter the LLM but might leverage the resources available to the user, for example by rule-based systems, which could guide the LLM. One example is the integration of Klarna's price comparison services into ChatGPT (OpenAI, 2023).
- **Prompt engineering:** This approach relies on the context-learning ability of LLMs. It is essentially based on a user assigning a role (e.g., support assistant) to the model in a specific context, including instructions and objectives for the frame in which the model operates.

The first two approaches may lead to more domain-specific expertise and therefore to more accurate answers, but remain probabilistic in nature.¹

Plug-ins are a promising solution, but are currently only accessible to selected partners and are therefore not an option for the present study. In the context of investment advice, a plug-in could be designed to access a specific advisory module of a financial institution depending on the user's conversation progress and apply the underlying rule-based business logic without leaving the general chat environment.

Besides plug-ins, prompt engineering is an approach for the use of LLMs for rule-based applications like investment advice. In this context, a pre-trained LLM serves as a generalist interaction tool to capture relevant user information, which can then be translated into an investment recommendation via predefined rules. More specifically, through prompt engineering using available APIs, the LLM can be defined with the aim of retrieving the information required for the rule-based investment advice system, such as a user's age, wealth, and risk tolerance, as a first step, but without limiting its generalist conversational capabilities. The gathered information can then be fed from the LLM into the rule-based system, which then derives appropriate deterministic investment recommendations. This ensures that the recommendations can be explained, interpreted, and reproduced. In a further step, prompt engineering can be used to instruct the LLM to communicate the rule-based investment recommendation to the user, including an explanation of the recommendation with the collected information and potentially the predefined rules.

The following chapter describes a more detailed design of such a solution, which also underlies the implemented prototype.

¹It should be noted that LLMs can be parameterised via the so-called "temperature" in such a way that they lead to very similar outputs for the same input (Marion, 2023). This does not solve the problem of probabilistic investment recommendations in the present context, as the inputs are entered differently across different clients.

5. Prototype for Investment Advice

Based on the findings of the previous chapters, a prototype¹ combining the probabilistic nature of conversational LLMs with deterministic rules for investment advice is described in the following. The prototype is based on OpenAI’s pre-trained GPT-3.5-turbo model and uses corresponding APIs for prompt engineering. Furthermore, the proof-of-concept is performed using a simplified rule-based decision system for investment recommendation.

The general design and the prototype’s task flow is described in Figure 2. It consists of four basic elements, i.e., a user interface, a GPT chat protocol (“AdvisorGPT”), an inner thought mechanism, and a rule-based decision matrix. In the present prototype, the latter can be seen as a function that deterministically transforms the attributes age, wealth, and risk appetite of a user, i.e., a bank client, to a specific model portfolio, but can generally be used for other rule-sets. Furthermore, the user interface connects the user to the probabilistic GPT-based system which assumes the following two roles, as specified by prompt engineering:

- **AdvisorGPT:** Employed for engaging in direct interaction with users. It is designed to act as an interviewer, collecting the relevant user information of age, wealth, and risk appetite (Task 1 in Figure 2), without limiting the general conversational capabilities of the underlying LLM. In addition, it also acts as a channel for the final investment recommendation (Task 5), which comes from the rule-based system (Task 4).

- **Inner thought mechanism:** Specified to supervise the interaction between the AdvisorGPT and the user and ensures that all relevant information is received (Task 2). After each input from the client, the AdvisorGPT asks itself whether the relevant information has already been received and reacts accordingly. Once all the required information is obtained, the data is fed into the rule-based system (Task 3) and used to generate rule-based portfolio recommendations.

The prototype thus combines the generalist, yet probabilistic, conversational capability of GPTs with deterministic rules for investment advice. It should be noted, however, that while the prototype works in many cases, it does not lead to the desired recommendation in certain situations. For example, in certain cases, the prototype is not able to evaluate correct inputs (e.g., “zero income”) based on the predefined rules. Furthermore, it is also not designed to meet existing legal requirements, but can serve as a basis for further rule-extensions and alternative use cases presented in the next chapter.

¹A corresponding Python implementation and additional technical details can be found in a public GitHub repository (see here).

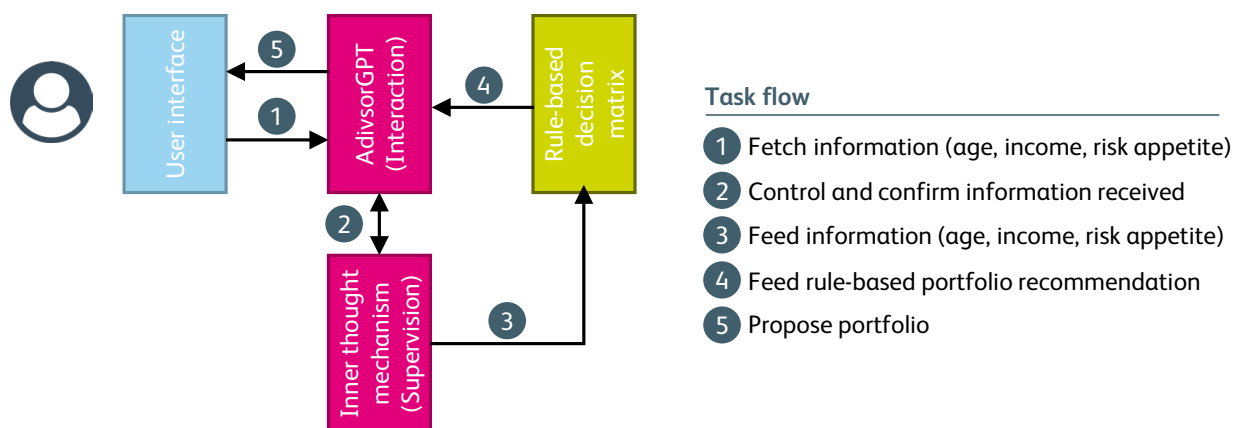


Figure 2: General architecture of the prototype and corresponding task flow

6. Other Use Cases

The ability of LLMs to process natural language for language understanding or language generation is enabling other use cases for the finance industry, besides financial advice. In the following, some of these use cases are presented, whereby the degree to which they are rule-based varies. Examples for which first solutions already exist include:

- **Customer service centres:** Generative AI using LLMs is able to facilitate customer service interactions if deployed as a voice- or chatbot. However, such bots typically need to be fine-tuned to company specific data and possibly restricted in terms of the general information it can reuse. One example hereof is the GPT-based chatbot “Clara” by Helvetia (2023).
- **Know-your-costumer (KYC):** LLMs might help decrease conflicting goals of an efficient onboarding process and a comprehensive KYC review with their ability to extract information from business documents, emails, or other client-related documents (Pure Storage, 2023). Incremental Group, for example, combines an automation platform with OpenAI’s ChatGPT to interact with potential clients and to generate initial KYC reports. This process might allow advisors to take more informed decisions as to whether new business relationships are feasible or trigger red flags efficiently (Incremental Group, 2023).
- **Fraud detection:** LLMs are also, in principle, able to be utilised in illegal activities, for instance in phishing emails. It is therefore imperative that fraud detection teams remain vigilant and adapt accordingly. LLMs can offer supportive assistance in this context. Abnormal (2023), a cyber security company, provides a fine-tuned Google BERT LLM to detect emails belonging to complex phishing attacks. Furthermore, the company combines its solution with supervised machine learning algorithms to be able to continuously adapt to new forms of attacks.

Other financial use cases (without concrete examples) for which LLMs could offer solutions are:

- **Loans:** To provide more sophisticated services, automated rule-based loan granting systems of Swiss banks might be improved by LLMs. Today, clients are asked to provide financial information, and in return an automated response is sent, indicating whether a credit is approved or rejected. Not only could this process be extended by clarifying possible issues via an LLM-based chatbot, but it could also be designed to offer applicants further advice. For example, additional information in regard to the advantages of products or forecasting of financial developments (e.g., changes in interest rates) could be provided according to the bank’s policies and research. This could also represent significant added value in the area of issuance and renewal of mortgages (e.g., fixed-rate vs. flexible), taking into account the loan size and affordability.
- **3rd pillar:** By the year-end, a bank’s LLM-powered chatbot could have the ability to remind clients of their annual 3rd pillar deposit deadlines, provide information on deposit limits, and offer advice on opening supplementary 3rd pillar accounts based on their current savings. Furthermore, it could point out the importance of building up a 3rd pillar portfolio depending on the client’s risk appetite, taking into account the individual financial situation of a client and subsequently recommend suitable portfolios compliant with existing regulations.
- **Foreign exchange:** LLMs could advise clients on the currency situation and regulations abroad (e.g., ATM locations and cash limits), taking into account bank-specific service conditions (e.g., fees and limits). Furthermore, LLMs could also be used to offer customers suitable foreign exchange banking products (e.g., a foreign exchange account).

Note that the listed use cases and examples do not claim to be comprehensive.

7. Conclusion and Outlook

The present study aimed to discuss the potential of using probabilistic LLMs and rule-based systems in combination, and to implement a corresponding prototype for financial advice. The core findings are summarised in the following statements and theses:

The combination of AI and rule-based systems is possible and combines the best of both worlds. Although AI systems such as LLMs are fundamentally different in terms of their characteristics than rule-based systems, there are ways to combine the advantages of both worlds. With custom training, fine-tuning, plugins, and prompt engineering, several basic approaches exist for specifying generalist LLMs for particular use cases. However, due to their different modes of operation, these approaches cannot be combined equally with rule-based systems.

Prompt engineering as a possible approach to providing investment advice. Prompt engineering essentially aims to assign a role to the GPT in a particular context, including instructions and goals for the framework in which the model operates. In the context of deterministic investment advice, this feature can be used to create a chatbot that queries information necessary for investment recommendations, such as age, wealth, and risk appetite, by interacting with the user, without limiting the conversational capabilities of the underlying LLM. This information can then be fed into the rule-based system to generate the customised recommendation of a model portfolio, which in turn is communicated and explained to the user by the chatbot.

The technical foundations for implementing LLMs in finance are in place. Although LLMs have not been in the public focus for too long, there are already several providers that offer APIs for use. One provider is OpenAI, offering open interfaces to various GPTs, including GPT-3.5-turbo, which serves as the basis of the prototype developed for this study. These GPTs have been trained on large amounts of data and can generate human-like responses to natural language queries. This renders them particularly useful for applications such as chatbots, which in turn makes them suitable for typically conversation-intensive investment advice based on deterministic portfolio recommendation. Plug-ins

could also have potential to combine LLMs with deterministic rules, but are not (yet) accessible to the public.

The developed prototype works well, but not perfectly. The developed prototype shows that the implementation of (simplified) rule-based investment recommendations via GPT-based user interaction works well. However, there are cases where the system deviates from the desired solution, for example, when correct inputs are not processed accurately by the GPT, thereby preventing their evaluation against predefined rules. Further quantitative assessment of the performance is planned for future research.

Regulatory compliance is not the primary goal. An assessment of the regulatory compatibility of LLMs in investment advice is not the core of this study. This also refers to the prototype. Instead, the study is intended to show the potential of LLMs for the financial industry without emphasising potentially hindering factors. However, the current political discourse both in Switzerland and abroad shows that the topic of regulating AI systems in general is also highly relevant.

The potential of LLMs in the financial industry appears to be significant. In addition to investment advisory, LLMs can also operate in other areas of the financial industry that are, at least to some degree, rule-based. Existing solutions exist in customer servicing, KYC, and fraud detection. But there is also potential to be explored in areas such as mortgage provisioning, retirement planning, and foreign exchange. In addition to application areas, approaches as to how best to exploit the potential of LLMs continue to be researched.

User acceptance is key for adoption. Apart from the technological feasibility, other factors are crucial for the successful integration of LLMs into rule-based financial services. User acceptance is one key condition for these solutions to reach a broader audience. LLMs are only likely to be adopted in sensitive areas if clients are convinced of their added value and if aspects such as privacy and data protection are ensured. Whether this will succeed remains to be seen. However, initial surveys show that LLMs have potential in financial advice, especially for younger clients (Finder.com, 2023).

Authors and Contributors

This condensed study was prepared in collaboration with the following individuals who contributed in the form of text, discussion, document reviews, and other forms of feedback (in alphabetical order).

Authors

Thomas Ankenbrand

Head Competence Center Investments
Lucerne University of Applied Sciences and Arts

Denis Bieri

Lecturer
Lucerne University of Applied Sciences and Arts

Levin Reichmuth

Master's Assistant
Lucerne University of Applied Sciences and Arts

Cornelia Stengel

Co Director
Swiss Fintech Innovations

Stephanie Wickihalder

President
Swiss Fintech Innovations

Ege Yilmaz

Research Associate
Lucerne University of Applied Sciences and Arts

Contributors

Manuel Brogli

Stefano Ferrazzini

Vladislav Girlea

Denise Gubser

Sophie Hundertmark

Rafael Hutter

Stefan Jeker

Claudia Jung

Jasser Kassab

Anja Minikus

Nils Reimelt

Christoph Schrills

Thies Schroeder

Sven Siat

Philipp Zimmermann

Contact

For more information about this study, please contact us at:

Thomas Ankenbrand

Lucerne University of Applied Sciences and Arts
thomas.ankenbrand@hslu.ch

Stephanie Wickihalder

Swiss Fintech Innovations
stephanie.wickihalder@swissfintechinnovations.ch

Bibliography

- Abnormal. (2023). *How Abnormal Enhanced Its Detection Platform with BERT Large Language Models (LLMs)*. Retrieved 21/04/2023, from <https://abnormalsecurity.com/blog/enhanced-detection-platform-bert-large-language-models>
- Adebisi, A. (2023). *How to Use ChatGPT, ChatGPT Plus, ChatGPT Plugins, ChatGPT API, and GPT-4 on ChatGPT*. Retrieved 12/04/2023, from <https://gadgetadvisor.com/ai/how-to-use-chatgpt-chatgpt-plus-chatgpt-plugins-chatgpt-api-and-gpt-4-on-chatgpt/#htoc-how-to-add-plugins-to-chatgpt>
- Alammar, J. (2019). *The Illustrated GPT-2 (Visualizing Transformer Language Models)*. Retrieved 31/03/2023, from <https://jalammar.github.io/illustrated-gpt2/>
- Basu, K., Varanasi, S. C., Shakerin, F., Arias, J., & Gupta, G. (2021). Knowledge-driven natural language understanding of english text and its applications. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35, pp. 12554–12563).
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., ... others (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv preprint arXiv:2303.12712*.
- CB Insights. (2023). *The state of generative AI in 7 charts*. Retrieved 03/04/2023, from <https://www.cbinsights.com/research/generative-ai-funding-top-startups-investors/>
- Dobler, M. (2023). *Rechtslage der künstlichen Intelligenz - Unsicherheiten klären, Innovation fördern!* Retrieved 11/04/2023, from <https://www.parlament.ch/de/ratsbetrieb/suche-curia-vista/geschaeft?AffairId=20233201>
- Finder.com. (2023). *An investment fund created by ChatGPT is smashing the UK's top 10 most popular funds*. Retrieved 09/05/2023, from <https://www.finder.com/uk/press-release-chatgpt-investment-fund>
- Future of Life Institute. (2023). *Pause Giant AI Experiments: An Open Letter*. Retrieved 31/03/2023, from <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>
- Google. (2023). *An important next step on our AI journey*. Retrieved 30/03/2023, from <https://blog.google/technology/ai/bard-google-ai-search-updates/>
- GPDP. (2023). *Artificial intelligence: stop to ChatGPT by the Italian SA*. Retrieved 11/04/2023, from <https://www.gpdp.it/web/guest/home/docweb/-/docweb-display/docweb/9870847#english>
- Helvetia. (2023). *Do you have a question about insurance?* Retrieved 12/04/2023, from <https://www.helvetia.com/ch/web/de/privatkunden/kontakt/services/chatbot-clara/frag-clara.html>
- Hu, K. (2023). *ChatGPT sets record for fastest-growing user base - analyst note*. Retrieved 31/03/2023, from <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Incremental Group. (2023). *What does Chat-GPT mean for Financial Services?* Retrieved 21/04/2023, from <https://incrementalgroup.co.uk/2023/03/13/what-does-chat-gpt-mean-for-financial-services/>
- Manakul, P., Liusie, A., & Gales, M. J. (2023). SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.

- Marion, S. (2023). *How to use OpenAI model temperature?* Retrieved 10/05/2023, from <https://gptforwork.com/guides/openai-gpt3-temperature>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the dartmouth summer research project on artificial intelligence, august 31, 1955. *AI magazine*, 27(4), 12–12.
- OpenAI. (2022a). *Fine-tuning*. Retrieved 18/04/2023, from <https://platform.openai.com/docs/guides/fine-tuning>
- OpenAI. (2022b). *Introducing ChatGPT*. Retrieved 30/03/2023, from <https://openai.com/blog/chatgpt>
- OpenAI. (2023). *ChatGPT plugins*. Retrieved 20/04/2023, from <https://openai.com/blog/chatgpt-plugins>
- Pramod, B., Shadaab, K., & Vineet, K. (2021). *AI in Banking Market*. Retrieved 31/03/2023, from <https://www.alliedmarketresearch.com/ai-in-banking-market>
- Pure Storage. (2023). *AI Unleashes the Power of Unstructured Data*. Retrieved 21/04/2023, from <https://www.cio.com/article/220347/ai-unleashes-the-power-of-unstructured-data.html>
- Sejnowski, T. J. (2023). Large language models and the reverse turing test. *Neural Computation*, 35(3), 309–342.
- Stengel, C. (2023). *Die Schweiz braucht kein KI-Gesetz*. Retrieved 24/04/2023, from <https://www.fuw.ch/die-schweiz-braucht-kein-ki-gesetz-794705355805>
- Swiss National Bank. (2023). *Holdings of securities in bank custody accounts, by category of security – monthly*. Retrieved 04/04/2023, from <https://data.snb.ch/en/topics/banken/chart/bawebekatmch>
- SwissBanking. (online). *Data protection and data governance*. Retrieved 11/04/2023, from <https://www.swissbanking.ch/en/financial-centre/information-for-bank-clients-and-companies/data-protection-and-data-governance>
- Tamim, B. (2023). *GPT-5 expected this year, could make ChatGPT indistinguishable from a human*. Retrieved 31/03/2023, from <https://interestingengineering.com/innovation/gpt-5-chatgpt-indistinguishable-human>
- TechCrunch. (2023). *ChatGPT resumes service in Italy after adding privacy disclosures and controls*. Retrieved 09/05/2023, from <https://techcrunch.com/2023/04/28/chatgpt-resumes-in-italy/>
- Thornhill, J. (2023). *Multiple red flags are not yet slowing the generative AI train*. Retrieved 31/03/2023, from <https://www.ft.com/content/89647443-29c7-421f-86aa-5a12af6268c0>
- Tomych, I. (2023). *ChatGPT For Fintech: How To Use ChatGPT In Banking*. Retrieved 31/03/2023, from <https://dashdevs.com/blog/chatgpt-for-fintech-how-to-use-chatgpt-in-banking/>
- Turing, A. M. (2009). *Computing machinery and intelligence*. Springer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). *Attention Is All You Need*.
- Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., ... Mann, G. (2023). BloombergGPT: A large language model for finance. *arXiv preprint arXiv:2303.17564*.



Lucerne School of Business
Institute of Financial Services Zug IFZ
Campus Zug-Rotkreuz
Surstoffi 1
CH-6343 Rotkreuz

T +41 41 757 67 67
www.hslu.ch/ifz
hslu.ch/ifz

ISBN 978-3-907379-15-8

