FAR-AI: A Modular Platform for Investment Recommendation in the Financial Domain

Javier Sanz-Cruzado[0000−0002−7829−5174],
Edward Richards[0009−0008−1130−3767], and
Richard McCreadie[0000−0002−2751−2087]

University of Glasgow, Scotland, UK

Abstract. Financial asset recommendation (FAR) is an emerging sub-domain of the wider recommendation field that is concerned with recommending suitable financial assets to customers, with the expectation that those customers will invest capital into a subset of those assets. FAR is a particularly interesting sub-domain to explore, as unlike traditional movie or product recommendation, FAR solutions need to analyse and learn from a combination of time-series pricing data, company fundamentals, social signals and world events, relating the patterns observed to multi-faceted customer representations comprising profiling information, expectations and past investments. In this demo we will present a modular FAR platform; referred to as FAR-AI, with the goal of raising awareness and building a community around this emerging domain, as well as illustrate the challenges, design considerations and new research directions that FAR offers. The demo will comprise two components: 1) we will present the architecture of FAR-AI to attendees, to enable them to understand the how's and the why's of developing a FAR system; and 2) a live demonstration of FAR-AI as a customer-facing product, highlighting the differences in functionality between FAR solutions and traditional recommendation scenarios. The demo is supplemented by online-tutorial materials, to enable attendees new to this space to get practical experience with training FAR models. VIDEO URL

Keywords: Recommendation · Finance · Information Retrieval · Machine Learning.

1 Introducing Financial Asset Recommendation (FAR)

Financial organisations such as Banks, Fund Operators, and Fintech companies are undergoing a digital transformation, driven by the need for 24-7 online services, as well as demand for more effective automated analytic and artificial intelligence tooling to remain competitive in the market and serve the growing number of citizens requiring financial products [6]. The investment market in particular is experiencing significant disruption, as organisations transition from a model where only certified financial advisors serve customers, to mixed models where decision making by advisors are supported by automated services, or in some cases such services are used to directly serve customers with investment advice[1]

At its core, providing investment options to a customer is a ranking task, where the goal is to rank a set of assets (from one or more markets) based on their suitability for investment, conditioned on the current customer [2]. We refer to this ranking task as Financial Asset Recommendation (FAR). However, what makes FAR challenging is that the suitability of an investment is primarily derived from external factors, such as the short or long term market returns, the value of the currency used in the trading process, and the impact of governmental regulations or global events [7]. In addition to these external factors, FAR systems need to consider the customer circumstances, like the alignment of the recommendations with their financial risk tolerance and preferences [4]. As such, the types of solution built for FAR differ markedly from those in more traditional recommendation domains (such as movie [1] or product recommendation [5]).

More specifically, there is a stronger focus on item (asset) modelling using a diverse set of information sources. Effective solutions need to have the ability to model financial pricing data to distinguish high performing or under-valued assets. Numerical and factual data should be extracted from (textual) financial reports. Meanwhile fundamental information about asset categories is needed to enable appropriate filtering of assets based on the needs of the user [3]. Indeed, on the customer side, solutions need a means to dynamically adapt to the risk appetite of each individual customer, either by categorizing assets by risk, or developing risk-aware recommendation models. Furthermore, investments are not for forever, at some point the customer needs to realize the profits from their investment. Solutions need to consider how long the customer is expecting to invest for (the investment horizon), and rank assets according to their expected profitability for those different horizons, where the outcome of a recommendation may not be known for months or years [4].

FAR is a recommendation domain that is still its infancy, but is of growing interest both to researchers and industry. Indeed, we have been seeing a growing number of shared tasks/data challenges (e.g. FinNLP 2023 and FinArg 2024) and workshops (e.g. FinRec and FinIR) in this space over the last couple of years. The primary goal of this demo is to raise awareness of FAR as an interesting and fertile area for future researchers to work on, support community building around this emerging domain, as well as illustrate the challenges requiring further research.

2 Demonstration

In this demonstration we will present a modular platform for financial asset recommendation, referred to as FAR-AI, which illustrates how to tackle these challenges. FAR-AI is a suite of modular micro-services with an associated front-end that has been developed over the prior two years. It started research and

[1] Recommendation services are often referred to as robo-advisors in this context.

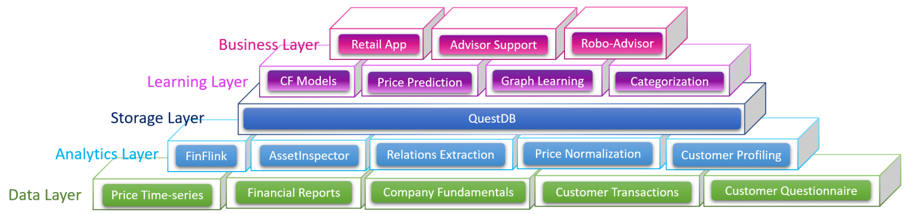

Fig. 1. FAR-AI Architecture

development within the European Commission-funded Flagship Horizon 2020 Infinitech project, and was later extended to additional markets and with new functionality through Impact Accelerator funding via UKRI. The overall architecture of FAR-AI is divided into five architectural layers, each comprised of different microservices, as shown in Figure 1. We summarize each layer below:

- **Data Layer**: Provides a series of plug-in data ingestion services to collect information about assets, the market as a whole, customers and companies.
- **Analytics Layer**: Contains services that convert the raw input data into a form that is suitable for learning. Notable micro-services within this layer include: FinFlink, which generates customisable technical indicators from pricing time-series in real-time; AssetInspector, which collects asset descriptions and metadata for financial assets from various web sources; and relation extraction, which analyses textual financial documents and extracts factoids to be stored within a knowledge graph. This layer also is responsible for cleaning and normalising both pricing and customer data.
- **Storage Layer**: Provides persistent storage of both the raw data and outputs from the analytics layer. All data here is stored as a time-series, such that system state for past time-points can be simulated on-demand, to better facilitate back-testing.
- **Learning Layer**: This layer is responsible for batch-training the various supervised models used by the upper business layer. This includes the training of traditional collaborative filtering models (based on transaction data), as well as price-change predictors for various investment horizons and encoders for relational knowledge graphs.
- **Business Layer**: The final business layer hosts the user-facing services for different use-cases.

The demo will comprise two components. First, via an associated poster we will present the architecture of FAR-AI to attendees, to inform them about the underlying ingestion, analytics and learning layers that are needed to make effective FAR solutions. Second, we will present a live demonstration of FAR-AI as a customer-facing product, illustrating a real-world use-case for the technology, namely automatic personalised portfolio generation for financial advisors.

Demo Use-Case - Financial Advisor Support: Within Banks and other financial institutions, financial advisors meet with clients to help them develop investment portfolios. The goal here is to produce a set of assets weighted by

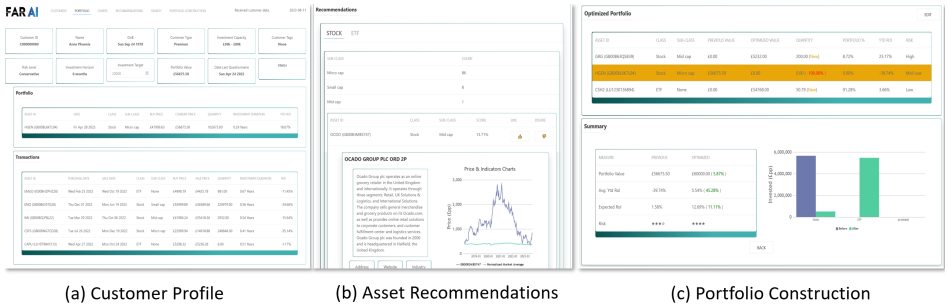

Fig. 2. FAR-AI Financial Advisor Support UI

investment value for a client, where the assets should be both profitable (given a target investment horizon), and reflect the client's risk appetite while minimising risk overall. To enable this use-case, FAR-AI provides the following functionality: 1) customer profile and portfolio visualisation; 2) supervised learning of both transaction based and price change prediction-based asset recommenders; 3) financial asset search; 4) automatic portfolio construction and optimisation using the asset recommender models. The demo interface is illustrated in Figure 2.

3 Target Audience

The primary audience for this demo is PhD students and researchers working in the areas of information retrieval and recommendation, who might be interested in the finance space as a new domain to explore. The finance domain is a growing research area and we believe that not only may there be existing ECIR attendees who might wish to adapt their research and technologies to this domain, but that there are exciting opportunities for new foundational research in this space for the next generation of students to explore. Indeed, recent relevant workshops include: FinRec at RecSys 2022, as well as FinIR and KDF at SIGIR 2020 and 2023.

4 Open Source Materials

While the full FAR-AI platform is not currently open-source we provide supplemental tutorial materials to enable attendees new to this space to get practical experience with training FAR models:

- [Tutorial iPython Notebook] Price Change Prediction
- [Tutorial iPython Notebook] Learning from Financial Transactions
- [Tutorial iPython Notebook] Financial Knowledge Graph Embedding

Acknowledgements: This work was in part carried out within the Infinitech project which is supported by the European Union's Horizon 2020 Research and Innovation programme under grant agreement no. 856632. Subsequent platform development was also financially supported via Engineering and Physical Sciences Research Council (EPSRC) Impact Accelerator, part of UK Research and Innovation (UKRI) with grant ref. number EP/X5257161/1.

References

1. James Bennett and Stan Lanning. The Netflix Prize. In *Proceedings of KDD Cup and Workshop 2007*, pages 3–6, San Jose, California, USA, 2007.
2. Dominik Jung, Verena Dorner, Florian Glaser, and Stefan Morana. Robo-advisory. *Business and Information Systems Engineering*, 60:81–86, 2018.
3. Richard McCreadie, Konstantinos Perakis, Maanasa Srikrishna, Nikolaos Droukas, Stamatis Pitsios, Georgia Prokopaki, Eleni Perdikouri, Craig Macdonald, and Iadh Ounis. Next-Generation Personalized Investment Recommendations. In *Big Data and Artificial Intelligence in Digital Finance: Increasing Personalization and Trust in Digital Finance using Big Data and AI*, pages 171–198. Springer, 2022.
4. Javier Sanz-Cruzado, Richard Mccreadie, Nikolaos Droukas, Craig Macdonald, and Iadh Ounis. On transaction-based metrics as a proxy for profitability of financial asset recommendations. In *Proceedings of the 3rd International Workshop on Personalization & Recommender Systems in Financial Services (FinRec 2022)*. Seattle, Washington, USA, 2022.
5. Brent Smith and Greg Linden. Two decades of recommender systems at Amazon.com. *IEEE Internet Computing*, 21(3):12–18, 2017.
6. John Soldatos and Dimosthenis Kyriazis, editors. *Big Data and Artificial Intelligence in Digital Finance*. Springer, 2022.
7. Dávid Zibriczky. Recommender Systems meet Finance: a Literature Review. In *Proceedings of the 2nd International Workshop on Personalization & Recommender Systems in Financial Services (FinRec 2016)*, pages 3–10, Bari, Italy, 2016.