# Collaborative Filtering and Contextual Bandits in Retail Finance: A Hybrid Recommender System for Personalised Financial Product Matching

*Dr. Paulo Leitão, Professor of Informatics, University of Minho, Portugal*

## 1. Introduction

Financial services are striving to adapt to rapidly changing consumer behavior and financial landscapes. For decades, financial services have been exclusively used by intermediaries, and customers have had to rely on advisors' judgment. With digitalization and advances in finance, services are becoming increasingly available to self-directed individual investors. As a result, the number of potential financial products, including combinations of investments and insurance with tax and estate planning, has increased dramatically. However, smart retail financial consumers demand personalized financial planning. It is critically important that we empower clients to consume both time-based and goal-based financial planning, especially goal-based planning over an extended horizon utilizing capital market simulators at a reasonable cost.

A new generation of powerful machine learning algorithms is said to be able to offer recommendations that are better approaches to solving an optimization problem than traditional actuarial pension products or an asset-only approach. Furthermore, the financial services industry is aggressively pursuing technology to enhance digital customer experiences, improve operational efficiency, and increase speed to market for new products and service offerings, displacing workers in those areas as a result. It is, therefore, crucial to understand the circumstances under which AI technologies and advanced machine learning systems can improve on broad modern portfolio thinking in personalized multi-period financial product recommendations. Social sciences and, increasingly, data scientists will need to understand what combination of portfolio risk might vary according to this person's risk tolerance, capacity, and willingness, and that person's human capital in the future.

### 1.1. Background and Significance

The global financial services industry has undergone significant transformation since the wide-scale adoption of technology. New financial products with varying features have been introduced to offer personalized financial services that cater to different individual needs. Today more than ever, we witness heightened interest in the development and application of advanced AI systems in the financial sector. These digital tools hold transformative potential for the industry. AI and machine learning can synthesize and analyze massive amounts of data, solve sophisticated quantitative problems, automate processes, and, in particular, improve decision-making for financial services companies and their customers.

For instance, AI-driven tools may be used for developing customized financial service products related to investment or credit risk ratings, improving the detection and prevention of financial crime, offering smart investment advice, consumer lending, or insurance-related products, and reducing costs related to compliance, risk management, and customer operations. Some 64 percent of U.S. consumers are willing to manage and control their financial life digitally, provided they have a platform that allows them to do it all. Another 44 percent of younger investors are drawn to robo-advisors, and 37 percent expect their financial future to be managed by AI.

Besides the extensive business practices, if technology in combination with AI is applied, scholars are acknowledging that the use of AI-driven machine learning systems today offers advanced modeling capabilities, data prediction, unsupervised learning, and optimization problems. It can easily offer a systematic solution process that can optimize the mathematical model. Most systems are user-friendly and can store personalized behavioral patterns of customers based on big data. The user interface of these digital investment and loan advisory tools is also consumer-centric. Nevertheless, an AI-driven decision-making tool or product is beneficial only if customers accept and use it. Hence, it should be customer-focused. The application of machine learning tools is drawing the attention of financial institutions in daily financial affairs, particularly in industrial and developing countries. Given the large potential for churn, the predictive model can be used to reduce the churn rate and to obtain new and improved strategies based on the predicted classes of customers.

**1.2. Objectives of the Study**

The increasing availability of financial transaction data opens up new opportunities for enhancing the personalized recommendation of financial products. The need for personalization lies in long-term savings decisions and short-term spending decisions, not in the decision on products to fulfill other financial needs. A key challenge for enhancing personalization, however, is the complexity of long-term savings decisions with the need to consider not only transaction data but also individual risk preferences and diverse future development scenarios. The aim of our study is to provide insight into opportunities created by AI techniques to overcome these complexities and to open up new horizons in the personalization of financial product recommendations.

1.2. Objectives of the Study Our study aims to clarify whether and under which circumstances transaction data can be used to make personalized financial product recommendations. In order to answer this general research question, we postulate and analyze the following specific central objectives: - We aim to compare various machine learning models regarding the precision of financial product recommendations. - We will investigate the preferences for interpreting and accepting the results and the data used for deriving recommendations based on machine learning models. In doing so, we want to determine whether information about the recommendation techniques impacts responses to the tailored recommendations provided. - We aim to assess the overall consequences of providing AI-driven financial product recommendations: the customer acceptance and satisfaction of personalized recommendations as well as the business performance facilitated by enhanced efficiency of provisioning customer service organizations. - As background research, we aim to investigate the decision-making behavior in long-term savings and how, and which decisions are taken by which customers. The results gained from answering our central objectives lead to conclusive answers to our main research question.

**2. AI in Financial Services**

AI is becoming increasingly popular in the financial services sector and within it has been recognized as a significant 'enabler.' Applications for AI in finance range from fraud detection and investigation to risk assessment and management, client services, and process automation. At present, we calculate that AI-based technologies have the ability to generate an increase in market performance by 30% by making financial advisories deliver proactive consultancy with advice acceptable to the customer precisely when they are most needed, which can be adequately adapted to their actual and timely necessities. This has been corroborated by the advent of AI Stock Advisor Tools that provide evaluation of key market performers' profitability and perform real-time monitoring of fluctuations and market elements.

Last but not least, another application of AI in finance is in terms of product recommendation and the accuracy thereof as per the unique individual customer. In fact, 80% of customers switched brands when recommended a product for their needs through AI-based recommendations. AI can distinguish in a way humans cannot the differences in data with people who will and will not default on the basis of content, timing, and tone. Adoption of AI-based advisory services provides the opportunity to revise the bank offer to current customers. Indeed, especially for current customers, having a proactive financial advisory service offered by the bank that is fully tailored to the client's individual situation (both past and present) makes the offer of services rather different from standard robo-advisory solutions. Given the fact that such service might even be partially based on advice on investing, customers could be prone to place their trust in AI-based recommendations made by the bank more than on purchase recommendations provided by wealth-based websites that may not be tailored to their specific situation, especially when it comes to secondary financial products. Thus, offering an efficient financial advisory service that leads to the customer purchasing products, owing to the accuracy of the proposed advice, is not the only way to have a positive business impact.

### 2.1. Overview of AI Applications in Finance

AI technologies are now being used extensively in the finance sector to support operations, analyze and manage various assets, and predict customer behavior and market fluctuations. These AI applications make use of various algorithms and models. Predictive modeling is used to predict outcomes, including fraud detection, while constraint satisfaction is used in search operations and portfolio selection. A wide variety of AI applications are included in investment banking applications and financial advisory. AI has seen significant progress in applications including portfolio optimization, personal advisories, and robo-advisors, which are used to manage and invest funds based on individual financial status and desired profit. Personal advisors and robo-advisors aim to provide personalized advice and help the user in managing financial transactions, which aids in customer engagement. There has also been significant progress in algorithmic trading, auto trading, and advanced credit scoring, which use different trading strategies and techniques.

When data increases in size and complexity, and investment opportunities expand, they also increasingly become difficult for fund advisors to manage efficiently. Many AI-based products and services have been introduced, including personal advisory systems and robo-advisors, to help with these tasks. Most pension funds and investors have adopted MV portfolios as a tool to help allocate their funds in different financial assets. Recent studies have shown numerous examples that AI systems can provide better investing solutions tailored to the customer's financial needs and investing goals and can achieve better risk-return portfolios. AI-based investment strategies can have a significant impact on the development of computational investing and have demonstrated competitiveness with other investing methods.

### 2.2. Benefits and Challenges of AI Adoption in Financial Services

(2. Benefits and Challenges of AI Adoption in Financial Services)

2.2. Benefits and Challenges of AI Adoption in Financial Services Financial institutions stand to gain numerous advantages from AI adoption. AI can enhance financial services by observing patterns in data that humans struggle to identify, enabling the automation of repetitive tasks, and thereby reducing the need for data analysis and assistance with decision-making. Fast processing and extensive analysis of data also enable more accurate risk assessments of loan and mortgage applications submitted by prospective customers. Moreover, AI algorithms can be used by financial institutions to build a personalized relationship with customers by recommending the most suitable financial products based on an individual's income, savings, and prior personal financial history, as well as other tailored services and insights.

AI can benefit financial institutions by being applied to enhance their operational efficiency, for example, in reducing the time taken to sanction a loan or process a mortgage application. This will help them lower their operational costs and deliver profit. The primary motivation for greater use of AI among financial institutions was to identify efficiencies and cost reductions for their operational systems. While operational systems are being automated, the personal touch aspect of a client relationship can be enhanced and personalized, as human capital currently supporting back-end operations could become increasingly directed towards sales and customer care services; in turn, automating the system to be more personalized in all respective back-end services and sales. In contrast to the numerous benefits of AI systems, the adoption of AI among financial institutions also raises several challenges and potential risks. These include the additional financial investments required to establish and maintain AI systems, and the costs to integrate these with existing systems. There also remain additional costs and time to be invested in training employees to use the AI system and to keep it operational and relevant. Further regulatory compliance issues that need to be overcome include ensuring that the AI system and the algorithms it uses comply with customer appreciation laws and regulations, as well as ethical standards. Additionally, there is a societal unease potentially causing a loss of trust in financial institutions, due to such factors as data being misconstrued. The modern customer can become concerned that algorithms were not programmed correctly, are biased, or were never intended to help with financial planning. This could lead to an increased chance of a data breach, thereby increasing risk to financial institutions. There is also a risk that they may not update when necessary, and may produce forecasts based on older and not up-to-date training data. There are further challenges in dealing with wealthier clients.

**3. Machine Learning Models for Personalizing Financial Products**

Personalizing financial products is an emerging frontier in FinTech that aims to offer desirable products to consumers. Machine learning models are being developed to operate individually with consumers and to predict products to offer to each user. Two main types of machine learning models are applied for several FinTech services: supervised learning models and unsupervised learning models. Supervised learning models aim to predict an interest or a variable from a consumer-based massive dataset containing many distinct user features, habits, and behaviors. Regression, support vector machines, decision trees, k-nearest neighbors, and so forth can be utilized to train the models to offer the predictions. These techniques can also be ensemble, performed with operational failure indexes or introducing superior channels to test, and so forth.

Likewise, unsupervised learning models aim to evaluate such datasets possessing no response or no target. These techniques are especially beneficial when target variables are not obtainable, which may occur sometimes in the simulation of demand for diverse, brand-new financial or non-financial products, for instance. These machine learning models inform business negotiators and decision-makers about financial behaviors, like which economic segments perform optimally, how people behave and take action, stock exchanges, interest rates, and the fluctuations in other socio-economic variables. Additionally, they can be trained to discover whether a user might abandon a product or not. For instance, famous and innovative algorithms, including k-means, GMM, PCA, t-SNE, and UMAP, can be utilized for distinct financial services. Both types of these algorithms are suggested to be fed with high-quality and accurate user data. Lastly, it is critical to properly preprocess the datasets as there may be redundant noise or inconclusive features in large financial systems. Furthermore, the results of these machine learning models are not one-value predictions. User interaction with such a personalized financial system significantly influences the success rate of the user-product communications.

### 3.1. Types of Machine Learning Models

Machine learning models can be divided into groups based on the mechanism by which they capture and serialize information for prediction. Regression models numerically order different events based on the magnitude of their impact. In a piecewise linear approach such as decision trees, each split provides a different numerical order of impact for specific events. A forest of such trees becomes a non-linear sum that captures a wide variety of different combinations of inputs. While neural networks can capture different causes for similar events, they are difficult to visually interpret because each input is combined in every part of their non-linear structure. This categorization is crucial for understanding how to use these models for personalizing financial product recommendations. For example, whereas regression and forests require one input for every different pattern of making a decision, neural networks can focus on each input in different ways for each main decision. For simple financial challenges, a regression model with carefully chosen inputs may suffice, whereas a neural network that captures diverse complex patterns of consumer behavior is necessary for large and complex applications. The main strengths and weaknesses of each type of model will directly affect the quality of personalized product recommendations and the ease with which they are embedded into the wider economic environment. Innovations in machine learning research are producing increasingly complex models for financial decision-making that can capture a wide variety of different patterns of behavior. Cash-on-card and PFM aggregation services use linear regression to examine the weighted significance of each input in, for example, the level of spending or balance. They use various combinations of inputs to generate many different versions of the same linear model. Cash-on-card credit scorers issue credit cards as a loan product to subprime consumers by replacing a key input with a substitute that reduces the perceived likelihood that the borrower will default. Use such a regressional approach to make an inventory-backed personal loan product available via selected retailers for consumers who have underlying savings.

### 3.2. Data Sources and Preprocessing

3.2 Data Sources and Preprocessing. One of the cornerstones of any machine learning-based financial recommendation or allocation system is the choice of data. The choice of technologies or models depends on the available data. To build a general system accommodating a wide range of personal finance-related domains, the following data will be required. Transaction Data: Transaction data showing account movements over various periods is of paramount importance. Demographic Information: Age, sex, education, etc., can profoundly impact financial decisions and hence should be taken into account. Behavioral Data: Subscriptions, routine expenses and incomes, usage patterns are necessary to build more personalized models and recommendations. Due to the sensitivity of financial data, careful consideration must be given to ways to obtain a diverse dataset reflecting as many scenarios as possible.

Models and algorithms are just tools for determining the model parameters that can apply to data. Ensuring that the data is of high quality, adequate, and relevant is an essential part of the process. Raw data collected from banks or open spending sources contain a lot of errors that can ultimately alter the predictions or recommendations provided at the model's end. The collected data initially goes through the data preprocessing phase, which includes tasks such as data cleaning, data normalization, and finally feature selection for modeling that turns raw data into a reliable, usable, and effective data source. The success of machine learning models and algorithms depends largely on the quality of the data used in model training and fitting. To summarize, without complete, accurate, and relevant data, any developed model will provide inaccurate insights and possibly harmful recommendations. Knowing where to limit or make assumptions to the model based on the experienced deficiencies or shortcomings of the data is also important, as it offers a way to establish a boundary reflecting a potentially fickle estimate.

**4. Case Studies and Success Stories**

There are many organizations that have been able to leverage artificial intelligence effectively to provide revolutionary new financial product recommendation solutions.

Startups

Guruvest and Elephant Data are two particularly interesting examples of organizations that have developed new data-driven recommendations for fund products. Guruvest seemed to make strong recommendations for funds that would exhibit strong performance, but according to the company, ceased operations in the summer of 2012. The website is no longer active and appears to have been removed.

Large established companies

Nordnet has implemented an artificial intelligence system that customizes the front page of its website for each unique user. This personalization is based on the system's calculation of a variety of interesting news, content, and recommendations. The system relies on machine learning to better serve the customers who use it. The insights Nordnet gathers from users who interact with the AI are in turn used to inform the product managers, designers, and front-end developers who design the trading and finance services that Nordnet delivers.

IBM has shared a case study of its collaboration with a bank in Colombia, known as Colpatria Bank. The bank used Watson to create a chatbot that identifies good personal loans. It functions by asking clients why they need money and their plans after taking out a personal loan. The questions are aimed at getting a deeper understanding of the client's needs and lifestyle, and each answer helps the bank to determine which product to recommend. This is a good example of the usage of collaborative filtering and the importance of social data in the recommendation process. It is dogma today that you should always consider other customers' ratings and usage patterns in a recommendation context. The approach used by Colpatria Bank to deliver good recommendations led to a 30% increase in the number of engagements with new clients and an increase in client satisfaction by 99%, mainly due to the rise in perceived ease of use. For instance, the cycle of awareness is reduced as more users who trust the system use Watson. In the end, the recommendation engine is used to send personal loan offers to website users.

**4.1. Real-world Examples of AI-driven Financial Product Recommendations**

This annotated review has detailed how recommendation systems, rankings, and matching algorithms can be used in the design of financial institutions' front-end architectures to personalize channels and products and A/B test them systematically. It also mentioned how organizations have used robust algorithms to filter customers internally and design appropriate products for each. This subsection reviews some recent case studies showing diverse strategies organizations have adopted utilizing these theories. Starting with a mechanism for capturing predictive customer data and using that data to provide enhanced products, the studies move from the frontline to operations and competition, covering customers' trading behavior, product bundling, and investments in customer service. Throughout the examples, the reviews are backed up with the results of interventions to show that these mechanisms are more than just theoretical. They can have real-world impacts, most notably customer satisfaction and retention.

Organizations are able to use thousands of pieces of data on public news articles, sentiment, and message boards to design calls which are a-priori more likely to generate a trading profit for investors. Capturing social media data on millions of investor messages daily, combining this with a history of messages on the same stocks and an innovative exploitation of search patterns to generate indicators which suggest which way a stock will move in the next month. The business has switched from a focus on patents towards maintenance, consulting, and seminars for using their AI products. Their software features fuzzy- and cognitive-logic-based algorithms designed in co-operation with customers, supplying leads and probability indicators in marketing and human resources for instance.

**5. Ethical and Regulatory Considerations**

Organizations are eager to leverage the predictive capabilities of AI to optimize and customize offerings for current and prospective customers. Notwithstanding this attractive value proposition, there are several things organizations must consider when implementing AI in financial services. Various national laws and regulations require independent courts or an equivalent body to be the final arbiter. Human oversight refers to final decision-making; it doesn't mean that each step in the decision-making process needs to be made by a human. As the legal and ethical concepts and frameworks were devised for human intelligence, it is typically not ready to deal with AI. As businesses become more aware of the potential risks associated with biased AI models, more attention is paid to strategies and tools to minimize them. Fairness audits, adversarial debiasing, and algorithmic impact assessments are quickly becoming buzzwords in industry discussions.

The growing awareness of fair treatment of consumers and the importance of algorithmic accountability and transparency is spreading concern. Some fear that these concerns will lead to more restrictive data protection regimes that will impose high technological barriers to entry. These concerns are leading organizations to err on the side of caution and hold back on innovation. Against this backdrop, our research aims to shed light on the ethical and legal issues surrounding the use of AI in financial services, as well as the main concerns and priorities of regulators at both the EU and national levels. The objective is to encourage dialogue and help stakeholders identify the most sensible legislative and regulatory mechanisms to expand AI usage, while at the same time ensuring ethical and responsible molding of AI models and operational use. This approach is viewed as a careful compromise between strict, inflexible instructions and an overly broad regulatory framework that would lack both guidance and legal certainty.

### 5.1. Transparency and Fairness in AI Algorithms

5.1. Transparency and Fairness in AI Algorithms

Transparency. The use of models is somewhat abstracted from real-world decision-making, meaning understanding how financial models work will likely require non-trivial technical and financial expertise. Failing to do so can limit the extent to which related decision-making can take place and be made accountable. Lack of transparency can also limit the potential impact of recommendations, as user engagement and perceived usefulness may be affected. The use of cutting-edge technology and AI in particular could factor into users' concerns about algorithmic transparency and accountability. Addressing these issues proactively has the potential to prevent potential issues further down the line. The mismatch between expert and non-expert understanding is in part a subject of explainable AI research and financial literacy building efforts.

Fairness. Ethical concerns related to decision-makers and decision-making processes may inform consumer and general public opinions about the ethical use of AI in decision-making or finance more generally. In practice, widespread adoption of such systems may be limited by the lack of user trust, especially if potential harms of specific decision-making practices or the model itself are highlighted. Data and technology used in finance to inform decision-making behave similarly to such biases and discrimination risks, meaning concern about fairness, transparency, and accountability in turn should apply. Approaches to making AI decision-making systems fairer are more established and include extensive literature on de-biasing datasets and model training data. Techniques such as these can help to ensure that the underlying patterns between classes are truly discriminatory, not a spurious effect. Other fairness-promoting strategies include legal compliance and subsequent algorithm auditing. Algorithm audits can be performed by independent third-party organizations, but the lack of real data beyond examples in research means this is not a common practice. As such, different auditing procedures should exist to hold financial companies accountable for their fair AI use. Protecting trade secrets is important but can be balanced against public interest needs in this context.

Special considerations should apply in the case of vulnerable individuals, such as minors, low-income consumers, and seniors. These groups are more likely to have limited exposure to other forms of advice and are more likely to trust the recommendations. They are less likely to have the financial knowledge to act on that advice. For this reason, a lower bar for safeguarding these groups under an ethical use of AI is proposed. Ethical considerations aside, AI should also be underpinned by human rights principles. Banking, as a human right, is only recently being discussed. Vulnerabilities such as profiling and credit exclusion, or targeted advertising and price discrimination are core points in the discussion. Non-discrimination, social inclusion, digital education, client empowerment, and redress are core points in the discussion.

**5.2. Compliance with Data Protection Regulations**

When using AI, financial services institutions (FSIs) require vast amounts of customer data gathered from various sources, such as bank statements, transaction data, credit scores, and sometimes even social media. The General Data Protection Regulation, the California Consumer Privacy Act, the Data Protection Act 2018, the Personal Data Protection Act 2012, and others place legal obligations on organizations that collect, handle, and use individuals’ personal data. Depending on their location, the applicable data protection law enshrines data subjects' rights to request information about stored personal data and object to its processing. To date, data breaches in the U.S. cost companies millions per single breach, and a large, centrally managed database may be a target for cyberattacks. Compliance with data protection regulations is now an ethical and a business requirement as they safeguard both the identity as well as traits and actions of our next generation society.

From a customer point of view, the risk of privacy and data breaches is a significant concern when new features go into production that influence their financial status or have access to their transaction data. The trust customers have in current FSIs must not be violated, as once lost, trust must be regained in order to not jeopardize the organizations' reputation. Financial data is especially sensitive due to its inclusion of reputation, power, privacy, and emotional value and should be handled with care. Businesses must be transparent with their use and distribution but also be critical of how business models operate and be upfront with their customers/right-holders in what circumstances surveillance is used to further these business opportunities. Data protection regulations must not be seen as a hurdle for AI development, but seen as a collaboratively developed component of the established technological and social construct between the directorate of enterprise strategy, research, and executive AI development. Compliance needs to be discussed in the same way as markets, risks, ecosystems, data, and AI elegantly fit together as a strategic AI and data cocktail within the enterprise. The board must also risk assess the entire hardware and software AI technology landscape, and showing transparency in this effort can foster trust.

**6. Future Direction**

By studying published research works, reports, and experts' opinions, we identified some of the possible trends to continue enhancing financial product recommendations in the context of future development of AI. First of all, digital services are changing, so the methodology of introducing AI will continuously have to reflect a faster change in customer behavior, habits, and expectations. In addition, self-learning exclusivity might decrease, requiring more experience, faster technology, and complete AI models. Also, the level of human-like understanding through AI might be increasing, bringing natural language and straight-through learning to the next level. Further, distribution might play an important role, which requires an adaptive AI model that increases its performance while getting deeper in financial product knowledge. The techniques of natural language processing might be evolving. The next trend might be the transition to deeper learning techniques. Today's deep learning is not yet a full reflection of human cognitive reasoning capabilities. This might still need future development in terms of a deep-learning paradigm shift. The future of AI in financial product recommendations might also reflect in the relations between the technology providers and financial institutions. This requires collaborations, partnerships, reflection, and further innovation. Ethical as well as regulatory, risk, and compliance issues will continue to shape. Consumer trust in dashboards might come more into focus, not only the analytical trust. Also, apart from upsides, there will be challenges and downsides to face. Therefore, they should be looked at from different angles to reap, strive, and stay ahead.

## 7. Conclusion

Based on the extensive analysis of the previous sections, the transformative potential of AI in personalizing financial services can be concluded. A successful AI journey comprises numerous considerations. One essential aspect is the availability of high-quality data that serves as an input to the AI model. Pursuing system-related AI decision accuracy to match or exceed human error, product-underlying model performance is advised. Based on the evaluation of a wealth of models that could have been applied to the case of financial product recommendations, decision trees, XGBoost, CatBoost, and LightGBM performed very well. Trade-offs to be aware of when choosing among different ML classification models concern cross-validation time, data size, intuition, and explanation of the model, as well as the influence on readily available model implementations. In selecting the appropriate model for financial product recommendations, numerous aspects have to be considered. Ethical considerations need to be part of the robotization of AI, particularly regarding sensitive data.

A second success-related consideration is model accuracy. Random Forest, logistic regression, and K-nearest neighbors underperformed with F1-scores smaller than 40%, the benchmark applied. Further investigation in this research area is needed to control for responsible lending, allowing fair access to financial products while adhering to ethical AI considerations and regulatory compliance. Overall, this aims to inspire further collaboration on the AI topic. In summary, the impact of AI in personalizing financial products is expected to increase. The findings presented may provide guidance to academics and professionals.