Weighted Eligibility-Based Product Recommendation System

Julie Koon Koon, Vinayak Sharma and Patrick Hosein
Department of Computer Science
The University of the West Indies
St Augustine, Trinidad and Tobago
julie.koonkoon@my.uwi.edu, vinayak.sharma@my.uwi.edu, patrick.hosein@uwi.edu

***Abstract*—In the finance and banking industry, product recommendation systems play an important role in enhancing customer engagement and driving cross-sell and upsell opportunities. However, traditional recommendation models often fail to account for product eligibility constraints and customer-specific suitability, leading to irrelevant or infeasible suggestions. This paper introduces a novel hybrid framework that combines rule-based eligibility filtering, customer segmentation, and machine learning to generate personalized, actionable product recommendation. The goal of this paper is to predict the most suitable financial product for a customer by mimicking human reasoning. To achieve this, our research employs a hybrid approach that combines fuzzy logic rule-based eligibility filtering with a supervised machine learning model to estimate adoption probabilities and recommend the next best product for the customer. Key features such as income, risk, credit score, and employment status are used together with other engineered features such as customer segment, balance to income ratio and eligibility scores to rank products for each customer. A weighting mechanism is introduced to prioritize recommendations that align with both customer needs and business constraints. Evaluation of synthetic data generated to mimic real-world data shows that the incorporation of eligibility and behavioral characteristics significantly improves the relevance of the recommendation and predictive performance. This system presents a scalable and interpretable solution for financial institutions that aims to optimize customer targeting while maintaining operational feasibility. Ideally, recommending the right product to the right customer.**

***Index Terms*—Product Recommendation System, Eligibility Constraints, Machine Learning, Predictive Modeling, Fuzzy Logic, Customer Segmentation**

I. Introduction

In today's evolving, dynamic and highly competitive finance and banking landscape, banking institutions are under growing pressure to deliver hyper-personalized experiences that both meet evolving customer expectations and support strategic business objectives. Recommendation systems have emerged as the ideal tool for achieving this personalization, specifically in the context of financial product targeting. Through analysing historical data and customer behavior, these systems can uncover cross-sell and upsell opportunities that drive customer engagement, customer retention, and revenue. Traditional recommendation algorithms such as collaborative filtering, content-based, or hybrid have shown substantial success in areas such as e-commerce and digital media. However, their direct application in the finance sector is often met with domain-specific challenges. Unlike general consumer products, financial offerings are often governed by complex eligibility criteria, regulatory constraints, and risk assessments. Recommending products like credit cards, loans, or investment accounts without adequately accounting for a customer's financial profile, employment status, or legal eligibility can result in poor customer experience, wasted marketing resources, and even compliance violations.

To address these limitations, this paper presents a weighted eligibility-based product recommender system tailored specifically for retail banking environments. This proposed system goes beyond traditional rule-based or purely behavioral models by combining structured eligibility filtering with machine learning and driven adoption prediction. Our approach first applies domain-specific rules (e.g., income thresholds, credit scores, employment requirements) to ensure that only viable product options are considered. A supervised learning model is then used to estimate the probability of adoption for each eligible product, using customer attributes such as income, balance, credit, tenure and risk. A key innovation in this system is its ability to mimic human reasoning through a flexible, weighted framework. Rather than enforcing rigid eligibility constraints, the model allows high-probability signals such as higher account balances or strong risk-adjusted behaviors to counterweight disqualifying traits like age or employment status. This emulates how a financial advisor might override standard filters to recommend a product when alternative indicators signal a strong fit. In addition, the system incorporates segment-based targeting, aligning recommendations with customer clusters derived from behavioral and demographic features.

This research serves as a practical guide for financial institutions seeking to enhance their personalized marketing strategies through intelligent product recommendation systems. By leveraging a weighted eligibility-based framework, the study supports data-driven decision-making in matching customers with relevant financial products. The insights derived can help institutions improve recommendation accuracy, boost product uptake, and streamline customer engagement across various service channels.

This research makes the following key contributions:

- A scalable framework that fuses fuzzy logic eligibility

979-8-3503-6735-5/24/$31.00 ©2025 IEEE

filtering with machine learning to improve the relevance and feasibility of financial product recommendations.
- A flexible weighting mechanism that allows for eligibility overrides based on strong alternative signals, enhancing real-world applicability and personalization.
- A practical evaluation demonstrates improved recommendation quality over simple, eligibility-only, or behavior-only approaches, using simulated banking data.

II. Literature and Contributions

Recommendation systems have become instrumental in various industries, including e-commerce, digital entertainment, education, and now finance. As outlined by Mustaqeem et al., recommender systems have evolved significantly, with a growing focus on explainability, contextualization, and cross-domain applicability, making them a vital tool in digital decision support systems [1]. A similar overview by Adomavicius and Tuzhilin highlights historical trends and emerging challenges, such as dynamic modeling and hybridization [2]. In the financial domain, recommender systems must overcome unique constraints, including regulatory compliance, eligibility thresholds, and risk management requirements. Wu and Li emphasize that financial recommenders require multiobjective optimization, temporal sensitivity, and high interpretability [3]. Furthermore, traditional approaches such as collaborative filtering often lack the ability to enforce domain-specific constraints, potentially resulting in inappropriate or non-actionable recommendations. Recent research has also explored explainable AI (XAI) approaches in financial recommendation systems. For example, Liem and Hayter propose a conviction-based investment recommender system using machine-learned confidence signals and explainability mechanisms to enhance trust in high-stakes domains like trading [4]. These findings reinforce the need for interpretable, constraint-aware systems when applying AI in finance.

A. Eligibility-Aware and Knowledge-Based Filtering

Eligibility-based recommendation is particularly important in banking, where financial products must comply with strict regulatory and operational rules. Knowledge-based recommendation systems differ from collaborative and content-based recommenders by using structured domain rules and semantic reasoning to match users with items. These systems are particularly effective in cold-start scenarios and complex decision contexts involving high user involvement or strict constraints [5]. They offer transparency and business rule enforcement, often required in financial product recommendation.

Amudala et al. proposed a hybrid model combining clustering and rule logic for financial targeting, improving precision while maintaining control over business rules [6]. These approaches echo the design goals of our proposed system, ensuring that every recommended product is not just relevant, but also viable given a customer's financial profile.

B. Segment-Aware and Behaviorally Targeted Models

Customer segmentation is a common tactic in marketing and finance to improve targeting effectiveness. Ghosh et al. found that clustering customers before applying machine learning algorithms significantly improves the precision of recommendations and reduces noise [7]. This approach was used in our research. Incorporating demographic and behavioral segmentation ensures that product recommendations are not only accurate but also strategically aligned with customer profiles. Segment-aware filtering has also proven useful in recommending research papers, where user interests are modeled over time and aligned with topical segmentation [8]. These principles are applicable in finance, where product needs vary across customer life stages and segments (e.g., students, retirees, high-income earners).

C. Value-Aware and Profit-Oriented Recommenders

Modern recommendation systems increasingly integrate business objectives such as customer lifetime value or profitability into model outputs. Zhang et al. introduced the concept of value-aware recommendation, where the model jointly optimizes relevance and business outcomes [9]. Rahdari and Palhang extended this by incorporating profit margins directly into the loss function, showing improvements in both personalization and revenue uplift [10]. Such approaches are particularly relevant in banking, where different products (e.g., mortgage vs. credit card) vary in terms of margin, churn risk, and strategic importance.

D. Dynamic and Time-Sensitive Modeling

Customer behavior and eligibility can change over time based on income changes, credit score updates, or employment status. Ghiye et al. proposed adaptive time-decay mechanisms in collaborative filtering to ensure models reflect current customer behavior rather than historical artifacts [11]. These methods are essential in dynamic environments like finance, where static models may become quickly outdated.

E. Research Gap

The literature confirms the importance of eligibility enforcement, segment-awareness, interpretability, and business alignment in financial recommendation systems. However, few solutions integrate these elements into a cohesive framework. This paper addresses that gap by introducing a Weighted Eligibility-Based Product Recommendation System that combines:

- Rule-based eligibility filtering,
- Segment-based personalization,
- Adoption prediction using machine learning, and
- Flexible, weighted overrides that mimic human reasoning (e.g., waiving an age limit for a high-income customer).

III. Methodology

A. Dataset and Data Generation

The customer profile dataset was synthetically generated using OpenAI's ChatGPT-4o to reflect realistic patterns and correlations observed in financial customer data. The product metadata dataset was based on real-world product rules and eligibility structures and was not synthetic.

The Customer Profile Dataset includes demographic, financial, and behavioral attributes for each customer, such as age, income, account balance, credit score, employment status, and existing products.

TABLE I
CUSTOMER PROFILE DATASET

Field	Type	Notes
cust_id	string	Unique identifier
age	int	Customer's age
income	float	Monthly income
account_balance	float	Total current balance
employment_status	string	Employed, Student etc
education_level	string	None, Primary, Secondary etc
location	string	Region i.e North etc
existing_products	string list	E.g., ["Savings", "Car Loan"]
customer_tenure_yrs	float	Customer's time with bank
credit_score	integer	customer's credit score
overdue_loans	boolean	True/False
risk	integer	Scale 0–1, customer risk

The Product Metadata Dataset defined each financial product's eligibility criteria, including age and income thresholds, employment requirements, and target customer segments.

The variables in both datasets were constructed to exhibit realistic dependencies, enabling accurate simulations of real-world financial behavior.

B. Feature Engineering

To prepare the data for clustering and model training, several preprocessing steps were undertaken:

- **Log transformations** were applied to income and account_balance to mitigate skewness, resulting in log_income and log_balance.
- A proxy for wealth, net_worth_proxy, was defined as:

$$net\_worth\_proxy = \frac{income \times 12}{account\_balance}$$

- **Age segmentation** was introduced via the age_bucket feature with four categories:
  0-24: young, 25-40: mid, 41-60: senior, 61-100: elderly.
- **New feature creation**: A binary feature named has_more_than_1_product was introduced to indicate whether a customer owns more than one financial product.
- **Numerical features** were standardized using the StandardScaler from `sklearn.preprocessing`: log_income, log_balance, credit_score, tenure, balance_to_income_ratio, net_worth_proxy, risk.
- **Categorical features** were one-hot encoded, including age_bucket, gender, has_more_than_1_product and employment_status.

C. Customer Segmentation via PCA and Clustering

To enable personalized and segment-driven product recommendations, the study incorporated a customer segmentation phase. This process aimed to assign each customer to a behavioral segment that could later be aligned with product targeting strategies. Dimensionality reduction and clustering techniques were employed to derive high-level, interpretable customer segments.

The full feature set was preprocessed. Numerical variables were standardized using the StandardScaler(), and categorical variables were one-hot encoded using the OneHotEncoder(). The full pipeline was implemented via ColumnTransformer().

Following this, Principal Component Analysis (PCA) was employed to reduce the dimensionality of the feature space while preserving the majority of the dataset's variance. This not only improved computational efficiency but also aided the interpretability of subsequent clustering results.

To identify the optimal number of customer clusters, two unsupervised validation techniques were used, The **Elbow Method**, which assesses the point at which additional clusters yield diminishing improvements in within-cluster variance (WCSS) and the **Silhouette Score**, which measures the quality of clustering based on intra and inter-cluster distances.

K-Means clustering was then applied to the reduced dataset. The resulting cluster centroids and component loadings were intended to inform the creation of high-level behavioral segments. These segment labels, once interpreted would be critical inputs to the product recommendation pipeline, enabling targeted alignment between customer profiles and product offerings.

D. Fuzzy Logic Rule-Based Eligibility Framework

To simulate domain-expert decision-making in financial services, a fuzzy logic rule-based eligibility framework was designed. Instead of relying on rigid thresholds, this framework assigned partial credit to customers based on how closely they matched each product's eligibility criteria.

Eligibility scores were computed for the following criteria:

- **Age**: Scoring based on proximity to the product's required age range.
- **Income**: Scoring based on how far above or below the product's income threshold the customer falls.
- **Employment Status**: Whether employment is required and whether the customer is employed.
- **Account Balance**: Used as a proxy for financial stability.

Each rule yielded a fuzzy logic score between 0.1 and 1.0. These were aggregated using domain-informed heuristic weights as shown in Table II.

TABLE II
HEURISTIC WEIGHTS FOR ELIGIBILITY SCORING

Feature	Weight
Age	0.3
Income	0.3
Employment	0.2
Balance	0.2

The result was a continuous **eligibility_score** for each customer-product pair, reflecting how well the customer met the ideal criteria for that product.

E. Hybrid Recommendation Strategy

The final stage of the methodology employed a hybrid recommendation strategy that combined both rule-based eligibility scoring and supervised machine learning to guide product suggestions.

After generating eligibility scores for each customer-product pair, the system filtered out products for which the customer was clearly ineligible. For the remaining candidate products, additional features were constructed to support predictive modeling. A key feature included in this stage was **segment_match**, which indicated whether the customer's behavioral cluster—derived through PCA and K-Means clustering—aligned with the target segment defined in the product metadata.

To model product adoption behavior, a **RandomForestClassifier** was trained to predict the binary target variable **label**, which captured whether or not the customer already owned the product. The feature set used in the model included:

- eligibility_score — computed from the fuzzy logic rule-based eligibility framework.
- income, account_balance, credit_score — financial indicators.
- risk — a customer-level risk score.
- requires_employment — whether employment is required for the product.
- segment_match — alignment between customer segment and product target segment.

The model was trained on 80% of the data using the **train_test_split** function with stratification on the target variable to preserve class balance. Predictions were then made on the 20% hold-out test set, and performance was evaluated using standard classification metrics including precision, recall, and F1-score.

In the deployed recommendation pipeline, this classifier was used to estimate the probability of adoption for each candidate product per customer. Products were ranked in descending order of predicted probability. The highest-ranked product that the customer did not already own was selected as the recommendation. If the customer already possessed that product, the next highest-ranking candidate was considered. This approach enabled the system to deliver personalized, eligibility-aware, and segment-aligned product recommendations based on both historical behavior and model-inferred likelihood of adoption.

IV. Results, Segment Interpretation and Classifier Performance

A. Principal Component Analysis (PCA)

The principal components extracted through PCA were analyzed to understand the underlying structure of the data. Table III outlines the interpretations based on the component loadings. As shown in Figure 1, the first **seven principal components** were required to surpass the 90% cumulative explained variance threshold. This dimensionality reduction preserved the majority of the variability in the original data while facilitating computational efficiency and improving interpretability in subsequent clustering analysis.

TABLE III
PRINCIPAL COMPONENT INTERPRETATIONS

Component	Interpretation
PC1	**Wealth and Income Proxy**. High positive loadings on log_income, log_balance, net_worth_proxy, and credit_score, indicating financially stable customers.
PC2	**Age and Retirement Status**. Strong loadings from age_bucket_elderly, employment_status, and tenure, reflecting older, long-tenured individuals.
PC3	**Product Affinity and Stability**. Positive loadings on tenure, balance_to_income_ratio, and credit_score, representing customers with long-standing and stable financial behavior.
PC4	**Age-Specific (Seniors)**. High loading on age_bucket_senior, suggesting separation of senior citizens from the general elderly group.

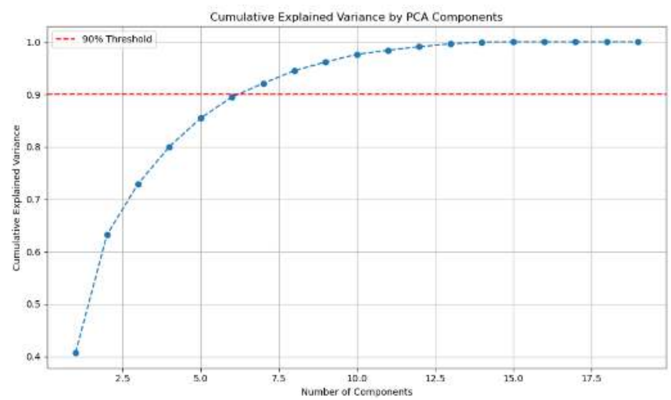

Fig. 1. Cumulative Explained Variance by PCA Components

B. Cluster Evaluation and Behavioral Segment Mapping

Customer segmentation was achieved through K-Means clustering applied to PCA-reduced features. Two cluster evaluation methods were used:

- **Elbow Method:** Identified the point of diminishing returns in Within-Cluster Sum of Squares (WCSS), suggesting an optimal k where intra-cluster variance is sufficiently minimized.
- **Silhouette Score:** Measured how distinct and well-separated the clusters were from one another.

The "elbow" point, where the rate of WCSS reduction significantly slows is at k = 4, which indicates that 4 clusters strike a balance between minimizing intra-cluster variance and avoiding overfitting. While the silhouette score peaks at $k = 3$, indicating that the clusters are most well-defined and distinct at this point. Given that the goal of this paper was to capture more nuanced behavior (i.e. more variability with the clusters), the number of optimal clusters chosen was four.

For this study, four clusters were chosen to retain interpretability while capturing meaningful behavioral variation. Figure 4 visualizes the clusters in PCA-reduced space.

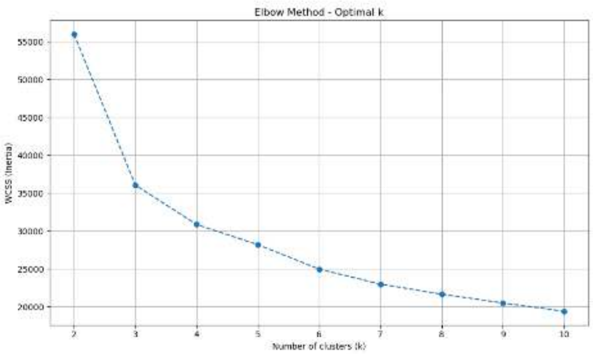

Fig. 2. Elbow Curve for K-means Clustering

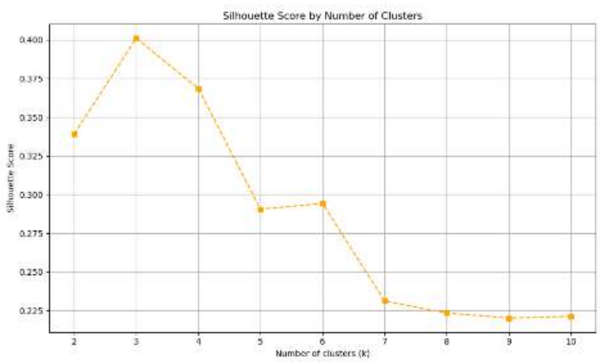

Fig. 3. Silhouette Scores for Cluster Validation

Segment Profiles:

a) Cluster 0: ***high_net_worth****:*

- 100% employed; 35% mid-age, 45% senior
- Highest log_income (9.1), net_worth_proxy (160,000)
- Highest credit score (660); 76% have more than one product

Interpretation: Financially affluent professionals with low risk and multiple product holdings.

b) Cluster 1: ***students, young_adults, tech_savy****:*

- 84% students; 47% under 25

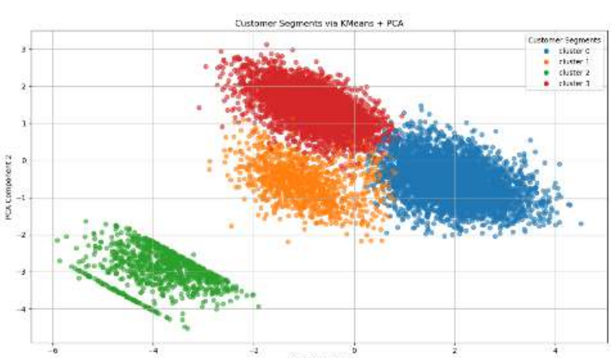

Fig. 4. Customer Segment Clusters in PCA Space

- Lowest income (7.1), low balance (8.7), short tenure (7.8), credit score (560)
- Moderate risk (0.60)

Interpretation: Young, digital-first users at the start of their financial journey.

c) Cluster 2: ***middle_income****:*

- 99% unemployed; 32% mid-aged, 38% senior
- log_income near zero; extremely high balance-to-income ratio (1,200)
- High risk (0.61)

Interpretation: Financially constrained individuals, possibly between jobs or dependent on savings.

d) Cluster 3: ***retirees****:*

- 99% elderly, 98% retired
- Long tenure (31 years), moderate income and balance
- Balanced risk profile (0.61)

Interpretation: Retired customers with high loyalty and moderate financial needs.

TABLE IV
CLUSTER-TO-SEGMENT MAPPING

Cluster	Segment Name	Description
0	high_net_worth	High income, employed, financially secure
1	students, young_adults, tech_savy	Predominantly young, early-stage customers
2	middle_income	Unemployed or strained income group with high risk
3	retirees	Elderly, long-standing customers with moderate balance

C. Classifier Performance and Feature Contribution Analysis

A **RandomForestClassifier** was trained to predict whether a customer already possessed a given product (**label**) based on customer-product pair features. The classifier was tested under two configurations:

TABLE V
PERFORMANCE WITHOUT SEGMENT AND ELIGIBILITY SCORE

Class	Precision	Recall	F1-score	Support
0	0.76	0.87	0.81	20939
1	0.30	0.17	0.21	7061
Accuracy			0.69	28000
Macro avg	0.53	0.52	0.51	28000
Weighted avg	0.64	0.69	0.66	28000

Model Without Segment and Eligibility Features: **Insight:** The model performed reasonably well in identifying customers who did not already own the product but struggled to identify those who did. This suggests that raw financial features were not sufficient to accurately predict product ownership.

Model With Segment and Eligibility Features: **Insight:** The inclusion of `eligibility_score` and `segment_match` dramatically improved model performance across all metrics. These features provided a strong signal of product suitability and behavioral alignment.

TABLE VI
PERFORMANCE WITH SEGMENT AND ELIGIBILITY SCORE

Class	Precision	Recall	F1-score	Support
0	1.00	1.00	1.00	20939
1	1.00	1.00	1.00	7061
Accuracy			1.00	28000
Macro avg	1.00	1.00	1.00	28000
Weighted avg	1.00	1.00	1.00	28000

Note: The perfect performance is likely due to the synthetic nature of the dataset and potential overlap between engineered features and the label. While encouraging, these results should be interpreted cautiously and validated on real-world data in future studies.

V. CONCLUSION AND FUTURE WORK

This study presented a Weighted Eligibility-Based Product Recommendation System designed for retail banking systems. This research addresses the limitations of traditional recommendation systems in highly regulated financial environments by integrating rule-based eligibility logic, customer segmentation, segment-to-product alignment, and supervised learning. The proposed hybrid approach combines fuzzy logic–based eligibility rules with a Random Forest classifier to deliver personalized product recommendations that are both operationally feasible and contextually relevant.

Key constraints were enforced through a hybrid rule-based and machine learning pipeline. The fuzzy logic module allowed for flexible eligibility filtering based on variables such as income, credit score, and account balance, mimicking the kind of nuanced reasoning typically performed by human advisors. Behavioral customer segmentation, derived via PCA and K-Means clustering, was used to align each recommendation with the target demographic or behavioral segment of each product. This improved the precision of recommendations and ensured relevance to the customer's financial context.

To avoid redundancy and maximize marketing efficiency, products already held by the customer were excluded from the recommendation pool. The Random Forest model leveraged features such as income, account_balance, credit_score, risk, and included segment_match and eligibility_score to rank candidate products by predicted adoption probability.

The results demonstrated that the system effectively balances interpretability, personalization, and scalability. It holds strong potential for deployment in financial institutions seeking to increase product uptake while maintaining regulatory compliance and customer trust.

Future research will focus on enhancing the system by incorporating product-level revenue modeling. This includes integrating factors such as loan size and term, interest income from deposit products, and fees associated with credit card usage inclusive of overdraft and overdrawn penalties. By aligning recommendations not just with customer suitability but also with profitability metrics, the model can optimize for both customer experience and institutional revenue. We also intend to explore explainable AI techniques for greater transparency and real-time deployment to support dynamic recommendation scenarios.

REFERENCES

[1] M. Mustaqeem, Y. Peng, and M. D. Alshehri, "A systematic review and research perspective on recommender systems," *Journal of Big Data*, vol. 9, no. 1, 2022, accessed: 2024-05-17. [Online]. Available: https://journalofbigdata.springeropen.com/articles/10.1186/s40537-022-00592-5

[2] G. Adomavicius and A. Tuzhilin, "Recommender systems: An overview, research trends, and future directions," *IEEE Intelligent Systems*, 2020, accessed: 2024-05-17. [Online]. Available: https://www.researchgate.net/publication/339172772

[3] D. Wu and X. Li, "A systematic literature review of financial product recommendation systems," *Information*, vol. 16, no. 3, 2025. [Online]. Available: https://www.mdpi.com/2078-2489/16/3/196

[4] J. Liem and A. Hayter, "Recommender systems in financial trading: Using machine-based conviction analysis in an explainable ai investment framework," http://arxiv.org/html/2404.11080v1, 2024, accessed: 2024-05-17.

[5] M. Uta, A. Felfernig, V.-M. Le, T. N. T. Tran, D. Garber, S. Lubos, and T. Burgstaller, "Knowledge-based recommender systems: overview and research directions," *Frontiers in Big Data*, vol. Volume 7 - 2024, 2024. [Online]. Available: https://www.frontiersin.org/journals/big-data/articles/10.3389/fdata.2024.1304439

[6] P. R. A. Puchakayala, "Recommendation systems in banking and finance transforming customer experience and operational efficiency," *Journal of Computational Analysis and Applications*, vol. 33, 2024, accessed: 2024-05-17.

[7] M. Mozumder, F. Mahmud, M. Shak, N. Sultana, G. Rodrigues, M. A. Rafi, M. Z. Farazi, K. Md Razaul, M. Khan, and M. Bhuiyan, "Optimizing customer segmentation in the banking sector: A comparative analysis of machine learning algorithms," *Journal of Computer Science and Technology Studies*, vol. 6, pp. 01–07, 08 2024.

[8] M. U. Mahmood, K. Zhang, and Q. Z. Sheng, "Scientific paper recommendation system," in *2023 IEEE International Conference on Big Data (Big Data)*, 2023, accessed: 2024-05-17. [Online]. Available: https://ieeexplore.ieee.org/document/10126196

[9] J. Zhang, H. Fang, and X. He, "A systematic review of value-aware recommender systems," *Expert Systems with Applications*, vol. 226, p. 120131, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417423006334

[10] M. Rahdari and M. Palhang, "Model-based approaches to profit-aware recommendation," *Expert Systems with Applications*, vol. 249, p. 123642, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417424005086

[11] A. Ghiye, B. Barreau, L. Carlier, and M. Vazirgiannis, "Adaptive collaborative filtering with personalized time decay functions for financial product recommendation," in *Proceedings of the 17th ACM Conference on Recommender Systems*, ser. RecSys '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 798–804. [Online]. Available: https://doi.org/10.1145/3604915.3608832

[12] T. D. Science, "Understanding k-means clustering in machine learning," https://towardsdatascience.com/understanding-k-means-clustering-in-machine-learning-6a6e67336aa1, Sep. 2018, accessed: 2024-4-17.

[13] N. Volzhin, "K-means, silhouette method, and clustering for dimensionality reduction," 2023, accessed: 2024-05-17. [Online]. Available: https://medium.com/@volzhinnv/k-means-silhouette-method-and-clustering-for-dimensionality-reduction-cd50e016d36e

[14] GeeksforGeeks, "Elbow method for optimal value of k in kmeans," https://www.geeksforgeeks.org/machine-learning/elbow-method-for-optimal-value-of-k-in-kmeans/, 2023, accessed: 2024-05-17.

[15] 365 Data Science, "Pca and k-means clustering in python," https://365datascience.com/tutorials/python-tutorials/pca-k-means/, 2023, accessed: 2024-05-17.

[16] A. Cervantes, "Interpreting and validating clustering results with k-means," https://medium.com/@a.cervantes2012/interpreting-and-validating-clustering-results-with-k-means-e98227183a4d, 2023, accessed: 2024-05-17.