A Hybrid Intelligent System for Personalized Recommendations in Offline Retail

Article

Accepted Version

A. De Paola, P. Ferraro, S. Imperiale, G. Lo Re

Proceedings of the 2025 IEEE International Conference on Smart Computing (SMARTCOMP)

DOI: https://www.doi.org/10.1109/SMARTCOMP65954.2025.00102

Publisher: IEEE

A Hybrid Intelligent System for Personalized Recommendations in Offline Retail

Alessandra De Paola
Department of Engineering
University of Palermo
Palermo, Italy
alessandra.depaola@unipa.it

Pierluca Ferraro
Department of Engineering
University of Palermo
Palermo, Italy
pierluca.ferraro@unipa.it

Sergio Imperiale
Department of Engineering
University of Palermo
Palermo, Italy
sergio.imperiale@unipa.it

Giuseppe Lo Re
Department of Engineering
University of Palermo
Palermo, Italy
giuseppe.lore@unipa.it

***Abstract*—In offline retail settings, there are two major challenges: improving the customer experience through personalized product recommendations and optimizing inventory management through accurate sales forecasting. Conventional recommendation systems assist customers in selecting goods based on personal preferences, similar user behavior, and popularity trends, while forecasting systems help managers predict future sales and optimize inventory levels. However, existing approaches face limitations in offline retail environments due to the scarcity of explicit feedback and the complexity of in-store interactions. To address these limitations, this paper introduces a hybrid intelligent system that combines multiple recommendation paradigms with predictive modeling techniques. By leveraging Recurrent Neural Networks and data-driven statistical models, the system improves both recommendation accuracy and demand forecasting reliability compared to traditional approaches. The effectiveness of the proposed system has been thoroughly evaluated using standard metrics such as Mean Reciprocal Rank at K (MRR@K) and Hit Rate at K (HR@K). Experimental results confirm the effectiveness of the proposed approach in balancing personalization and accuracy, offering significant benefits in offline retail environments.**

***Index Terms*—Recommender Systems, Predictive Analytics Systems, Retail stores**

I. Introduction and Related Work

Recommender systems are intelligent tools that predict and suggest items, content, or services based on a user's past preferences or the behavior of similar users. These systems are an important paradigm in many applications, from video streaming platforms to e-commerce and social media, where they personalize the user experience by identifying relevant items to recommend. Such suggestions reflect the highly varied preferences of users, which can be captured either explicitly, such as through ratings on a defined scale, or implicitly, such as by tracking interactions with content or products [1].

In the domain of non-online retail environments, such as supermarkets, recommender systems take on a unique role, acting as intelligent assistants that aid decision making for both customers and store managers, and these applications highlight how recommender systems can enhance the shopping experience while addressing the specific logistical challenges of offline retail. Despite their success, due to some data characteristics, RS may encounter significant difficulties, and the quality of their recommendations is strictly limited by the density of interactions in the collected data, which is often very sparse [2]. One of the most challenging issues in the design of recommender systems is the cold-start problem, which arises when new users interact with the system without an existing profile, making it difficult to generate personalized suggestions [3]. Common methods to mitigate these challenges include collaborative filtering, which leverages user similarity, content-based filtering, which focuses on item attributes, and hybrid approaches that combine multiple strategies. The development of recommendation systems for offline retail environments has been widely explored in the literature, with several studies addressing the unique challenges of this domain, such as the difficulty in collecting customer interaction data and conducting real-world in-store experiments [4].

In parallel to recommendations, forecasting systems also provide essential tools for store managers by predicting future trends based on historical data. Techniques such as Recurrent Neural Networks (RNNs) are particularly suited for these tasks due to their ability to model sequential data, recognize patterns over time, and predict future sales trends [5]. In addition, while RNNs leverage their dynamic hidden states to store and process past signals due to their ability to identify patterns and model the complex temporal dependencies [6], statistical models provide complementary capabilities by identifying trends and correlations within data, offering valuable insights that support data-driven decision-making [7].

The proposed intelligent hybrid system integrates a recommendation module with a forecasting subsystem to effectively address the complex challenges encountered in non-online retail environments. The hybrid recommendation module employs a multifaceted approach that integrates content-based filtering, user-based collaborative filtering, and popularity-based methods to recommend items to customers. The content-based filtering method identifies item affinities [8], the user-based collaborative filtering method recommends items based on the predilections of comparable customers, and the popularity-based methods accentuate the most efficacious products. These recommendations are delivered to customers through a mobile application, which interacts with several sensory devices in-store that collect raw measurements, such as cameras, RFID readers, and bluetooth beacons [9], to provide personalized offers in real time.

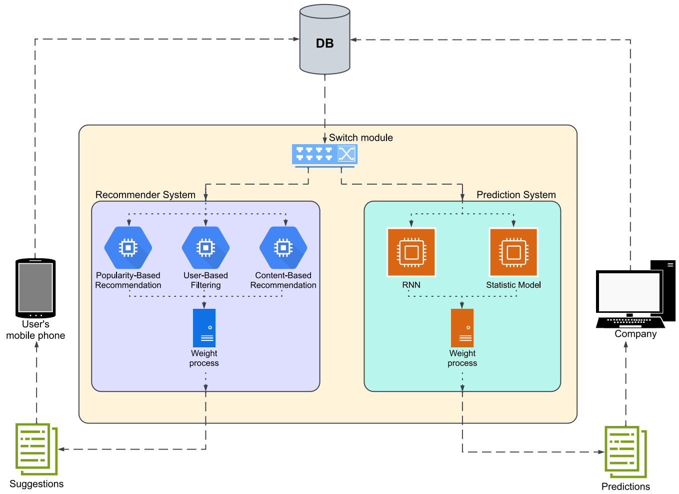

Fig. 1. The architecture of the proposed hybrid system.

The forecasting subsystem provides store managers with predictive analytics tools that leverage RNNs and statistical models. RNNs are employed to analyze temporal sales data, while statistical models uncover macro-level trends, enabling managers to anticipate demand and optimize inventory. The integration of these two modules, as proposed, establishes a bridge between customer personalization and operational efficiency. For example, a forecasting module might identify an upcoming spike in demand for specific products, prompting the recommendation module to prioritize these items in customer offers.

The proposed system is designed to offer a cohesive, adaptive framework suited to the dynamic nature of non-online retail. By leveraging the strengths of recommendation and forecasting technologies, the proposed system offers a comprehensive solution that enhances customer satisfaction while driving business growth.

A series of experiments conducted on a publicly available dataset have demonstrated the efficacy of the proposed system, thereby validating its capacity to address the challenges faced by non-online retail stores. The evaluation process focused on metrics such as MRR@K and HR@K. Initial findings indicate the system's capacity to deliver robust performance and support retail decision-making processes. The remainder of this paper is organized as follows: Section II details the architecture and components of the proposed hybrid system. Section III evaluates the system's performance through experiments, and Section IV concludes with insights and future directions.

II. Proposed Architecture

In this work, we propose a hybrid system that integrates multiple recommendation and prediction modules to enhance decision-making in offline stores. The system has been specifically designed to address the dual goals of improving customer experience through personalized recommendations and optimizing store operations through accurate forecasting. The system's distinct yet complementary modules enable a unified approach to addressing the challenges faced in offline retail settings.

In offline retail environments, the acquisition of direct ratings on individual items is inherently challenging due to the absence of online tracking mechanisms. As discussed in [10], user preferences are often only discernable through indirect feedback methods. This can be illustrated by repeated purchases of the same item or patterns in shopping behavior. These may reflect user preferences that would otherwise be captured by explicit ratings. Addressing this challenge requires the integration of a preprocessing component which, upon receiving raw data from in-store sensors and sales records, applies data fusion and machine learning techniques to transform them into a structured rating system [11]. This system can then be incorporated into recommendation algorithms, facilitating more precise and nuanced insights. In this paper, we assume that this entity is already existing.

The core of the hybrid system lies in its ability to analyze user interactions, sales data, and item attributes to extract meaningful patterns and provide insights. The system is designed to operate in two primary phases: the recommendation phase and the forecasting phase. These phases are elaborated upon in the subsequent sections of this paper. The idea behind the proposed hybrid recommender system is to use three well-known recommender paradigms, analyze the output of each of them, and then combine them to obtain a list of suggested items to be sent to the user's mobile phone, as illustrated in Figure 1.

The algorithm handles five key parameters: the unique identifier of the user u, who receives suggestions, the set of feedback (R_u) given by users, the number of items (n_{items}) to be suggested in the output of the recommendation algorithm, the store's sales history, and all data relating to the items i in the store. The output of each module will be a list of recommended items, which will be combined in a weighted way to create the final proposals.

A. Switch module

The primary function of the *switch module* is to determine whether the incoming request is directed to the *recommendation system* or the *prediction system*.

Furthermore, in the event that the request pertains to the recommender system, the module ascertains the availability of previous user feedback and adapts its operations to activate the appropriate paradigms. In instances where such feedback is present, the module meticulously prepares the data to align with the stipulated criteria of the three recommendation paradigms. Conversely, in the absence of feedback, the module defaults to activating only the popularity-based recommender. In the case of novice users or those with limited interactions, the absence of sufficient data precludes the formation of meaningful patterns or preferences. The switch module, therefore, dynamically adjusts its behavior based on the availability of data, thereby ensuring the system's continued functionality and the provision of useful recommendations, even in the absence of user interaction.

B. Popularity-Based module

This module functions independently of user interaction history, thereby ensuring its reliability as a component capable of generating recommendations based solely on item sales data. The algorithm first analyzes the input sales data and creates two distinct subsets: one that includes the entire sales history and another that focuses only on sales data from the last seven days. This dual approach enables the module to account for both long-term trends and recent fluctuations in item popularity within the store. To achieve this, the system employs a weighted formula that combines the rankings derived from the entire sales history ($P_{History}$) and the most recent sales data ($P_{LastWeek}$). This formula is adapted from the method described in [12], which integrates categorical and quantitative data:

$$O_{pb}(i) = \lambda \cdot P_{History}(i) + (1 - \lambda) \cdot P_{LastWeek}(i), \quad (1)$$

where i represents an individual item, and λ is a parameter that balances the importance of long-term and recent sales data.

C. User-based module

The method starts by establishing a similarity measure, which in the proposed module is defined as the cosine similarity [12].

The output of the cosine similarity computation is an n x n matrix, where rows and columns correspond to users and each cell contains the similarity score between two users. The range of these scores is from -1, indicating completely opposing preferences, to 1, indicating identical preferences. The determination of recommendations is then made by computing a weighted average of similarity scores and individual item ratings. Items that have been evaluated by users with higher similarity scores are assigned a greater weight.

D. Content-based module

Content-based recommender systems have been shown to perform optimally in scenarios where items can be described through rich sets of attributes. These systems depend on a user's own ratings and interactions with other items to identify relevant recommendations. This methodology is especially beneficial for novel items with limited or no reviews, as recommendations are based on item attributes rather than user correlations. The objective of such systems is to match users with items that bear similarities to those previously liked by the user, with the use of item-specific attributes serving as the foundation for comparison. Using data analysis combined with natural language processing, this module is able to enrich unstructured data, extracting pertinent features from textual elements such as product name, category, and description [13]. A pivotal aspect of this module pertains to the *Tf-Idf* (*Term Frequency - Inverse Document Frequency*) function, a numerical metric that quantifies the importance of words within a document, particularly in the context of a larger document collection. The function is further divided into two components: *Tf* and *Idf*, described in [12]. The *Tf-Idf* score combines these two components as follows:

$$(tf\text{-}idf)_{i,j} = tf_{i,j} \times idf_i \quad (2)$$

The proposed algorithm processes categorical information about products, such as the product's name, description, brand, and category. The text undergoes tokenization and analysis through a natural language processing pipeline, including steps such as tokenization, part-of-speech tagging, and dependency parsing. The utilization of this pipeline ensures the systematic conversion of textual data into a structured form, thus enabling detailed linguistic analysis and facilitating subsequent numerical representation.

Subsequent to the aforementioned tokenization, the *Tf-Idf* function transforms the processed descriptions into numerical vectors, thereby creating a matrix wherein each word is assigned a weight based on its importance in the document relative to the entire corpus. To recommend items, the system computes a ranked list of the most similar products for each user's top-ranked items, sorted by similarity score.

E. Recurrent Neural Network module

The *Recurrent Neural Network* (RNN) module for sales forecasting has been developed to capitalize on the temporal dependencies inherent in historical sales data. This module employs *Long Short-Term Memory* (LSTM) [14] layers, selected for their capacity to remedy the vanishing gradient problem, a prevalent challenge in standard RNNs when processing extensive sequences [5]. The employment of LSTMs in this context is particularly advantageous due to their capacity to discern short-term and long-term patterns, a necessity for sales data influenced by seasonal trends and external factors [15].

The model accepts as input a *tensor* representing sequences of historical data, with each sequence comprising multiple time steps. These time steps can contain features such as historical sales data, environmental factors (e.g., temperature readings captured at store entrances via sensors [16]), and the average number of customers entering the store at varying times of the day. The structured nature of these inputs enables the model to discern complex interdependencies among features over time.

Furthermore, layering multiple LSTM layers enhances the model's ability to learn hierarchical representations of the input data. Each layer processes the output of the previous one, thereby enabling the network to extract increasingly abstract features. However, it should be noted that an excessive number of layers can lead to overfitting and increased computational demands. To mitigate this, *dropout regularization* [17] is applied between layers, randomly deactivating neurons during training to improve generalization. The final fully connected layer transforms the output of the LSTM's final time step into the desired number of outputs. In the context of sales forecasting, the output is a single value representing the predicted sales for the subsequent time step. This design ensures that the model remains flexible and can be adapted to other forecasting tasks by modifying the output size or input features.

F. Statistical module

The statistical module has been developed for the purpose of analyzing historical sales data and predicting future sales. It accomplishes this through the implementation of machine learning techniques that are based on temporal and trend-based features. To forecast sales, the module implements gradient boosting, a machine learning technique that repeatedly constructs predictive models by integrating weak learners (typically decision trees) into a robust ensemble [18].

The algorithm starts with an initial model, often based on a simple prediction such as the mean of the target variable. Subsequent models are sequentially added to correct residual errors from previous predictions. Each new tree in this ensemble minimizes the gradient of the loss function, thereby identifying areas that are underperforming and refining the overall model. This iterative process persists until either a predetermined number of trees is constructed or further enhancements in model performance become negligible. This ensures an optimal balance between accuracy and computational efficiency. The final prediction is an ensemble of all trees, with their outputs weighted based on the learning rate.

In the context of retail, Gradient Boosting Decision Tree models have been shown to significantly outperform traditional forecasting methods, producing more accurate sales forecasts [19]. The output of the gradient boosting module can be expressed using the formulation presented in [18], and corresponds to a numerical prediction of sales for the next time period.

G. Weight Process Module

The *Weight Process Module* plays a crucial role in both the *recommendation subsystem* and the *prediction subsystem*, ensuring a balanced integration of multiple sources of information. Its primary function is to aggregate the outputs of the individual modules within each subsystem, weighting their contributions according to specific application requirements and optimization constraints.

In the *recommendation subsystem*, this module weights the outputs of the *popularity-based*, *user-based*, and *content-based* recommendation approaches. This ensures that the final recommendations effectively balance long-term popularity trends, collaborative filtering insights, and item-specific attributes. The weighted combination is computed as follows:

$$O_{rs}(i) = \alpha \cdot O_{pb}(i) + \beta \cdot O_{ub}(i) + \gamma \cdot O_{cb}(i), \tag{3}$$

where $O_{rs}(i)$ represents the final recommendation score for item i, while $O_{pb}(i)$, $O_{ub}(i)$, and $O_{cb}(i)$ denote the scores derived from the popularity-based, user-based, and content-based recommendation modules, respectively. The parameters α, β, and γ are weighting coefficients that satisfy the constraint $\alpha + \beta + \gamma = 1$. These coefficients allow the system to dynamically adjust the influence of each module based on user preferences, contextual needs, or business requirements, ensuring a robust and adaptable recommendation strategy. This weighted approach is intended to mitigate the limitations of each individual technique, such as the cold-start problem in user-based systems and the limited scope of content-based systems.

Similarly, in the *prediction subsystem*, the module combines the forecasts generated by the *Recurrent Neural Network* (RNN) and the *Statistical Model* to enhance predictive accuracy and stability. The final sales forecast is computed as:

$$O_{ps} = \delta \cdot O_{RNN} + \zeta \cdot O_{stat}, \tag{4}$$

where O_{ps} is the final prediction, O_{RNN} is the forecast produced by the RNN module, and O_{stat} is the output of the statistical model. The coefficients δ and ζ determine the relative contributions of the two models and are optimized during the training phase to minimize validation error, subject to the constraint $\delta + \zeta = 1$.

The RNN is capable of capturing sequential dependencies and long-term trends in data, while the statistical model demonstrates proficiency in identifying structured patterns and external influences, such as seasonal variations or market trends. The integration of these two distinct paradigms into a unified system serves to mitigate the risk of overfitting to a single type of pattern and enhance its generalization capabilities.

Both the RNN and statistical modules produce a scalar output representing the predicted sales for the upcoming time step. This architectural design enables a straightforward integration of the two outputs through weighted linear aggregation. The resulting prediction, O_{ps}, preserves the temporal resolution of the input data while capturing a balanced combination of sequential modeling and statistically derived patterns. Performing this fusion at the output level ensures that the system remains modular, interpretable, and easily adaptable to a variety of forecasting scenarios.

This hybrid approach provides a cohesive framework that aligns recommendation strategies with sales forecasts. The system's integration of predictive analytics and personalized recommendations enables informed decision-making processes in inventory management, dynamic pricing, and customer engagement strategies.

III. Experimental Results

The evaluation of the proposed system involves two primary challenges: the acquisition of appropriate training data and the assessment of the accuracy of the system. The first challenge is to identify datasets that adequately represent the offline retail environment. This task is further complicated by the inherent difficulty of collecting real-world sales data. The second challenge is to define appropriate evaluation metrics and ensure the reliability of the results. Given these challenges, the evaluation focuses primarily on the recommendation subsystem, as the unavailability of publicly available sales data precludes a rigorous evaluation of the prediction subsystem.

Selecting a dataset that accurately represents a physical retail environment posed a significant challenge. In this paper, we opted for the *UCSD Amazon Dataset* from the University

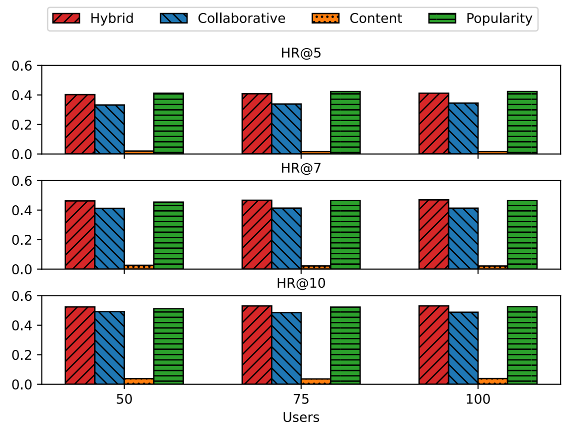

Fig. 2. Comparison of *HR@K* across different models with varying number of users.

of California San Diego [20], which includes product metadata and user reviews. We used only the *Home and Kitchen* sub-dataset because of the similar characteristics of the items in our considered environment. The metadata dataset was used for the content-based filtering module, and the ratings dataset allowed the evaluation of collaborative filtering and popularity-based recommendation methods by providing user-item interactions and timestamps. For the popularity-based approach, we assumed that the number of ratings for a given item is indicative of its overall popularity. However, this assumption has limitations. An ideal approach would involve a more precise evaluation, including an analysis of temporal dynamics. This would involve using timestamp data to give more weight to interactions closer to the time of recommendation. However, given that a timestamped review does not necessarily indicate when a user actually purchased an article, this approach was considered impractical for our study.

To assess the quality of recommendations, we employed standard evaluation metrics designed to measure the system's ability to distinguish *relevant* items from *irrelevant* ones [21]. Specifically, we considered an item to be relevant to a user if its rating exceeded the user's average rating. In the context of offline retail recommendation systems, *MRR@K* and *HR@K* are particularly well-suited metrics, since they evaluate the relevance and ranking of recommended items (presented to users as a *top-K* list), providing insights into the effectiveness of the system.

Hit Rate@K (HR@K) measures the fraction of users for whom at least one relevant item appears among the top-K recommended results. This binary metric assesses whether at least one of the top-K items in the system's predicted ranking matches the user's actual preferences, without considering the exact position within the list [22]. In the context of retail recommendation systems, HR@K provides a coarse-grained indicator of the system's ability to surface relevant products within a concise shortlist. As such, it is a key metric for assessing the practical effectiveness of personalized recommendations. However, HR@K provides a limited perspective of recommendation quality, as it treats all correct predictions equally, regardless of their position in the ranking, and does not reward the presence of multiple relevant items. This lack of granularity may mask important differences in system performance, particularly in scenarios where the position of relevant items within the list strongly influences user behavior.

Mean Reciprocal Rank@K (MRR@K) addresses this limitation by accounting for the position of the first relevant item in the top-K list. Specifically, it computes the reciprocal of the rank at which the first relevant item appears, then averages this value across all users. In doing so, MRR@K encourages models not only to include relevant items, but also to rank them as highly as possible [22]. However, MRR@K is sensitive to the definition of relevance. In our evaluation, we adopt a strict user-centric relevance threshold based on whether an item is rated above the user's personal average. This definition increases the difficulty of achieving high MRR scores, making it a challenging but insightful metric for assessing ranking effectiveness in personalized recommendation settings.

Before discussing system performance, we briefly outline the hyperparameter settings used in our experiments. The recommendation window was evaluated at $K \in \{5, 7, 10\}$ for each recommender system. The dataset was split into an 80%-20% ratio, where 80% of the data was allocated to training and the remaining 20% to testing. To ensure statistical reliability, each evaluation was conducted over 50 runs, with different subsets of users selected randomly. We considered varying numbers of users, specifically $U \in \{50, 75, 100\}$, to analyze the system's behavior under different user sample sizes. We performed all evaluations using the Recbole [23] metrics library.

Figures 2 and 3 present the experimental results, averaged across all runs, illustrating the performance of the *Collaborative*, *Content-based*, *Popularity-based*, and *Hybrid* modules.

The results indicate that both the hybrid and popularity-based models achieve consistently strong performance across all evaluation metrics, demonstrating their effectiveness in generating relevant recommendations while maintaining a balance between personalization and accuracy. In particular, *MRR* values remain competitive across all settings, suggesting that while not all suggested items are always relevant, the most relevant one tends to be positioned at the top of the recommendation list. Both the hybrid and popularity-based methods consistently outperform the collaborative and content-based approaches, confirming their robustness in recommendation tasks. The strong performance of the popularity-based method can be attributed to the nature of the dataset, where user preferences align with highly rated and frequently reviewed items, making popularity a strong predictor of relevance. However, in offline retail environments settings, customer preferences are often more heterogeneous and influenced by contextual factors such as location, promotions, and seasonality. In such cases, the hybrid approach presents a more flexible and effective solution, by dynamically adapting to individual user behavior while leveraging global trends. These findings suggest that

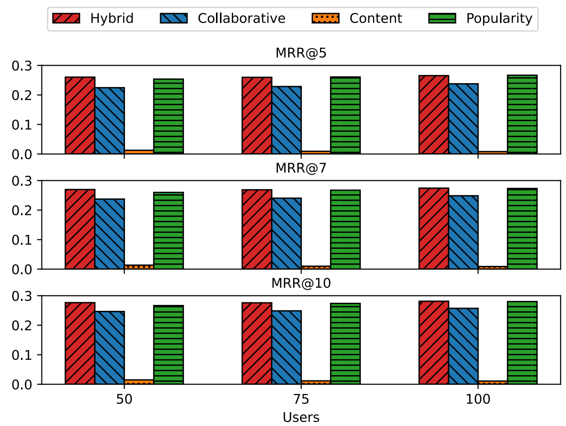

Fig. 3. Comparison of *MRR@K* across different models with varying number of users.

the hybrid system is particularly advantageous for *known users*, where it leverages historical interactions to enhance personalization, whereas the popularity-based filter remains a strong baseline due to its independence from user history. Finally, the ability of the hybrid approach to dynamically adjust weighting parameters further enhances its adaptability, making it well-suited for non-online retail environments where customer behavior and market dynamics can change rapidly.

IV. Conclusions

In this work, we have proposed an intelligent hybrid system designed to enhance decision-making processes in offline retail environments. By integrating recommendation and prediction models, the system aims to optimize both customer experience and store management. The experimental evaluation demonstrated the effectiveness of the proposed recommendation system, showing its adaptability across different scenarios [24] and its ability to maintain robust performance as the number of users increases. Future work will focus on assessing the prediction module, using real-world data collected in a retail store, allowing us to validate the forecasting capabilities of the system. We plan to explore further optimizations to the weighting mechanisms within the hybrid recommendation model, ensuring that the integration of different recommendation paradigms maximizes both accuracy and diversity.

Acknowledgments

This research is partially funded by the "Smart Venues" Project (POC Sicilia 2014/2020).

References

[1] V. Agate, F. Concone, S. Gaglio, and A. Giammanco, "A hybrid recommender system for cultural heritage promotion," in *2021 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2021, pp. 203–208.

[2] Z. Batmaz, A. Yurekli, A. Bilge, and C. Kaleli, "A review on deep learning for recommender systems: challenges and remedies." *Artif. Intell. Rev.*, vol. 52, no. 1, 2019.

[3] A. De Paola, S. Gaglio, A. Giammanco, G. Lo Re, and M. Morana, "A multi-agent system for itinerary suggestion in smart environments," *CAAI Transactions on Intelligence Technology*, vol. 6, no. 4, pp. 377–393, 2021.

[4] P. Li, G. Zhang, L. Chao, and Z. Xie, "Personalized recommendation system for offline shopping," in *2018 International Conference on Audio, Language and Image Processing (ICALIP)*, 2018, pp. 445–449.

[5] Y. Bengio, P. Simard, and P. Frasconi, "Learning long-term dependencies with gradient descent is difficult," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.

[6] V. Varadharajan, N. Smith, D. Kalla, F. Samaah, K. Polimetla, and G. R. Kumar, "Stock closing price and trend prediction with lstm-rnn," *Journal of Artificial Intelligence and Big Data*, vol. 4, 2024.

[7] D. Swami, A. D. Shah, and S. K. Ray, "Predicting future sales of retail products using machine learning," *arXiv preprint arXiv:2008.07779*, 2020.

[8] P. Ferraro and G. Lo Re, "Designing ontology-driven recommender systems for tourism," *Advances onto the Internet of Things: How Ontologies Make the Internet of Things Meaningful*, pp. 339–352, 2014.

[9] A. De Paola, G. Lo Re, S. Gaglio, and M. Ortolani, "An ambient intelligence architecture for extracting knowledge from distributed sensors," *ACM International Conference Proceeding Series*, vol. 403, pp. 104–109, 2009.

[10] A. De Paola, P. Ferraro, G. Lo Re, M. Morana, and M. Ortolani, "A fog-based hybrid intelligent system for energy saving in smart buildings," *Journal of Ambient Intelligence and Humanized Computing*, vol. 11, pp. 2793–2807, 2020.

[11] V. Agate, P. Ferraro, and S. Gaglio, "A cognitive architecture for ambient intelligence systems," in *International Workshop on Artificial Intelligence and Cognition (AIC 2018)*, Palermo, Italy, July 2018.

[12] C. C. Aggarwal, *Data Mining: The Textbook*, 1st ed. Springer Cham, 2015, springer International Publishing AG, part of Springer Nature 2015. eBook Packages: Computer Science, Computer Science (R0). [Online]. Available: https://doi.org/10.1007/978-3-319-14142-8

[13] V. Agate, F. Concone, and P. Ferraro, "Wip: Smart services for an augmented campus," in *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*, 2018, pp. 276–278.

[14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 11 1997.

[15] V. Agate, S. Drago, P. Ferraro, and G. Lo Re, "Anomaly detection for reoccurring concept drift in smart environments," in *18th International Conference on Mobility, Sensing and Networking (MSN)*. IEEE, 2022, pp. 113–120.

[16] G. Lo Re, M. Morana, and M. Ortolani, "Improving user experience via motion sensors in an ambient intelligence scenario," *PECCS 2013 - Proceedings of the 3rd International Conference on Pervasive Embedded Computing and Communication Systems*, pp. 29—34, 2013.

[17] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 56, pp. 1929–1958, 2014.

[18] J. H. Friedman, "Greedy function approximation: A gradient boosting machine." *The Annals of Statistics*, vol. 29, no. 5, pp. 1189 – 1232, 2001.

[19] Y. Shi, "Sales prediction optimization via gradient boosting decision tree," *Highlights in Science, Engineering and Technology*, vol. 44, pp. 121–128, 04 2023.

[20] R. He and J. McAuley, "Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering," in *Proceedings of the 25th International Conference on World Wide Web*, ser. WWW '16. International World Wide Web Conferences Steering Committee, 2016, p. 507–517.

[21] A. Gunawardana and G. Shani, "A survey of accuracy evaluation metrics of recommendation tasks," *J. Mach. Learn. Res.*, vol. 10, p. 2935–2962, Dec. 2009.

[22] A. Jadon and A. Patil, "A comprehensive survey of evaluation techniques for recommendation systems," in *Computation of Artificial Intelligence and Machine Learning*, 2025, pp. 281–304.

[23] W. X. Zhao, J. Chen, P. Wang, Q. Gu, and J. Wen, "Revisiting alternative experimental settings for evaluating top-n item recommendation algorithms," in *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*, 2020.

[24] A. De Paola, P. Ferraro, S. Gaglio, G. Lo Re, M. Morana, M. Ortolani, and D. Peri, "A context-aware system for ambient assisted living," in *Ubiquitous Computing and Ambient Intelligence: 11th International Conference, UCAmI 2017, Philadelphia, PA, USA*. Springer, 2017, pp. 426–438.