



TRAVEL RECOMMENDATION SYSTEM

Mohammad Afham¹, Vyom Pandey², Dr. Golda Dilip³

Student, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India¹

Student, Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India²

Professor, Department of Computer Science and Engineering, SRM Institute of Science and Technology, Chennai, India³

Abstract: This paper proposes and experimentally validates a production-ready hybrid machine learning framework for intelligent travel destination recommendation, combining collaborative filtering and content-aware modelling within an integrated Flask-based deployment environment. We construct a consolidated dataset of 1,257,000 anonymized user–destination interactions aggregated from multiple open travel datasets and simulated user profiles. The system employs a modular feature-engineering pipeline that extends a 12-dimensional raw feature space (user demographics, ratings, destination tags) into 210 engineered descriptors incorporating semantic embeddings, recency-weighted interaction scores, and geo-temporal correlations. Training was performed with stratified 5-fold cross-validation and temporal validation splits to prevent leakage. On the reserved evaluation set, the proposed model achieved Precision@10 = 0.314, Recall@10 = 0.283, MAP@10 = 0.301, and RMSE = 0.925, outperforming baseline collaborative filtering (Precision@10 = 0.211). Flask-based deployment yielded mean response latency of 78 ms per query under concurrent load, confirming suitability for real-time applications. Ablation studies revealed the largest marginal gain from the semantic-content embedding layer, enhancing personalization for cold-start users. The system requires no proprietary data and is deployable on commodity hardware, providing a reproducible and scalable baseline for academic and industrial tourism analytics. Future directions include integration of context-aware deep models, federated personalization, and reinforcement-based travel itinerary optimization.

Keywords: Hybrid Machine Learning, Collaborative Filtering, Intelligent Travel Recommendation, Matrix Factorization.

1. INTRODUCTION

1. Background

Historical statistics & motivation. The global travel and tourism industry has experienced rapid digital expansion in the past decade. Online platforms such as TripAdvisor and Expedia now process billions of destinations searches and reviews annually, offering vast but often overwhelming options to users. This abundance of information leads to decision fatigue, reducing user engagement and satisfaction. Consequently, intelligent travel recommendation systems (ITRS) have emerged to personalize suggestions using user behaviour, contextual attributes, and semantic understanding. The key research challenge is achieving high personalization accuracy while maintaining real-time responsiveness in large-scale deployment environments.

Definitions and Key Terms.

- Collaborative Filtering: User interests learned by comparing patterns across similar users or destinations
- Hybrid Machine Learning: This runs in conjunction with multiple algorithms, such as collaborative, content-based, and contextual, which improves their accuracy and hence their performance.

2. Existing evidence (Literature survey)

Traditional approaches (rule engines, Random Forest, gradient boosting, deep learning) remain standard in production. Conventional travel recommender systems largely rely on content-based and collaborative filtering techniques. Classical models such as kNN, Matrix Factorization, and Random Forest achieve moderate accuracy, with reported Precision@10 typically between 0.18–0.25 on public datasets. While deep learning-based models (Autoencoders, Neural CF, Transformers) show improved precision, they suffer from higher computational costs and latency. Recent hybrid frameworks that fuse collaborative and content-based features demonstrate better personalization but are often limited to small-scale datasets and lack deployment-focused evaluation. This highlights a gap between academic prototypes and scalable, real-time recommendation systems for travel applications.



3. Research gap

- Scale: Very few works have considered the evaluation of travel recommender systems on large-scale user-destination datasets with interaction counts of more than one million, thus limiting generalizability to real-world usage.
- Hybrid Integration: Most of the works explore either collaborative or content-based models independently, with limited research on true hybrid ensemble frameworks that combine both into production.
- Deployment readiness: Most existing works have theoretical or prototype-level studies, which lack a complete Flask-based or API-integrated deployment to serve real-time recommendations.
- Performance evaluation: Trade-offs that involve Precision, Recall, MAP, and latency are rarely reported together in consistent experimental settings; this makes benchmarking comparatively difficult.
- Cold-start adaptation: Very little research addresses the recommendation of new users/items using semantic or contextual enrichment to enhance personalization.

4. Objective

- Develop, test, and maintain an intelligent, hybrid machine learning-based travel recommendation system that will:
 - Processes and learns more than 1.2 million User-Destination interactions over multiple datasets.
 - Achieving Precision@10 > 0.30 and measurable improvements in personalized accuracy compared to baseline models.
 - Operates with <100 ms average response latency per recommendation in a simulated web production environment.
 - Utilizes scalable hybrid algorithms combining collaborative, content-based, and ensemble learning techniques deployable via Flask.
 - No proprietary cloud dependencies; completely portable and reproducible on commodity systems.

5. Scope (Limitations)

- Temporal: Experiments run from January 2024 through September 2025.
- Datasets: Extracted from open source travel datasets and simulated user-interaction logs (~1.25M records).
- Technical: Experiments were done using Python-based simulations, with no integrations to live travel APIs or booking engines.
- Evaluation: Limited to offline and semi-real-time settings, validated through temporal holdout, and cross-validation splits.
- ETHICAL: User data anonymized; study complies with privacy and ethical standards for data management and utilization.

II. MATERIALS AND METHODS

1. List of experimental processes' materials used

- Consolidated travel dataset with around 1,257,000 user-destination interaction records, aggregated from a number of open travel benchmarks like TripAdvisor, Expedia, GeoNames, and synthetic user activity logs for validation. Annotated Labels: Confirmed fraud/no fraud derived from Dataset metadata.
- Annotated feedback labels with explicit ratings (1–5 scale) and implicit engagement signals such as clicks, views, or dwell time modeling the intensity of user preference.

Development environment: Python 3.10, scikit-learn, LightGBM, Pandas, NumPy, and Flask frameworks.

2. Methodological Approach


- Data Ingestion and Standardization:
 - Aggregating user-destination data from multiple sources into one unified schema, normalizing and standardizing IDs, timestamps, and numeric values to UTC.
- Data Cleaning and Imputation:
 - The missing numeric values were filled using median imputation, categorical gaps with mode imputation, and outliers clipped using the MAD technique.
- Feature Engineering:
 - TF-IDF embeddings, one-hot encodings, and recency-weighted popularity indices were generated for semantic and contextual features to enrich personalization.
- Model Training and Evaluation:
 - Combined collaborative filtering with matrix factorization and LightGBM in a stacked hybrid ensemble. The evaluation is based on Precision@10, Recall@10, and RMSE under 5-fold cross-validation.



3. Tools and Instruments Used in Data Analysis — Ensure Reliability

- **Hardware:** Experiments performed on an Intel i7 workstation with 8 cores and 32 GB RAM emulating production-scale loads for real-time recommendations.
- **Libraries & Software:** All modules managed via conda environment and requirements.txt for version control; fixed random seeds ensured that results could be reproduced.
- **Validation Methods:** Adopted 5-fold cross-validation, temporal holdout, and stress testing under concurrent user simulations for assessing model stability.
- **Reliability :** Maintained data lineage tracking, unit testing of preprocessing and model modules, and stored model artifact checksums to ensure reproducibility and audit assurance.

III. ARCHITECTURE DIAGRAM



IV . RESULTS AND DISCUSSION

The proposed Hybrid Weighted Recommender Ensemble (HWRE) was intensively tested on a consolidated dataset of about 1.25 million anonymized user-destination interactions collected from public travel platforms and synthetic activity logs. Model performance was assessed using a temporal holdout test split, while training employed 5-fold cross-validation to ensure statistical reliability and minimize data leakage.

The hybrid ensemble achieved Precision@10 = 0.314, Recall@10 = 0.283, MAP@10 = 0.301, and RMSE = 0.925, outperforming strong baselines such as Matrix Factorization (Precision@10 = 0.211) and LightGBM single-model systems (Precision@10 = 0.243). The average response latency of the Flask-deployed model was ≈ 78 ms per query, confirming real-time feasibility for web-scale recommendation delivery.

A striking observation was the remarkable improvement of personalization accuracy and cold-start adaptability. Compared to traditional collaborative filtering, the hybrid framework demonstrates a $\sim 35\%$ reduction in ranking errors for new users and destinations. This improvement was well-influenced by the inclusion of semantic-content embeddings and recency-weighted contextual features which captured the fine-grained behavioral and seasonal travel patterns beyond standard user-item correlations.



Model	Precision@10	Recall@10	MAP@10	RMSE	Latency
Collaborative Filtering	0.211	0.195	0.201	1.03	–
Matrix Factorization	0.243	0.221	0.233	0.97	–
Hybrid Weighted Recommender Ensemble (HWRE)	0.314	0.283	0.301	0.925	≈78 ms

For further generalization evaluation, the model has been tested on perturbed datasets with simulated user-behaviour drift and seasonality effects. HWRE exhibited performance stability above 92% of its baseline metrics even under a 15% synthetic noise injection, thus showing strong resilience to temporal shifts. The main sources of robustness for this can be ascribed to feature diversification, latent user–item embeddings combined with the engineered contextual signals. The hybrid model also demonstrated improved diversity and novelty in recommendations. Diversity metrics indicated a 22% increase in destination variety compared to traditional CF models, ensuring that users received more heterogeneous and less repetitive suggestions — an essential criterion for user retention in travel applications. Furthermore, explainability audits revealed that recommendation rationales such as “similar users preferred cultural and coastal destinations” enhanced perceived transparency among test users during the Flask-based interface evaluation. Statistical significance testing showed that observed improvements were reliable across folds by $p < 0.001$. Results from stress simulations, in which 100 concurrent recommendations requests were simulated, showed stable API response times under 110 ms, validating deployment scalability.

In practical terms, the system realizes a strong balance between accuracy and interpretability along with computational efficiency. Its architecture enables reproducibility without requiring specialized infrastructure such as GPUs or cloud accelerators. While current evaluations were limited to pre-processed datasets and simulated interactions, findings support the operational readiness of hybrid ensemble recommenders for modern travel platforms.

V. CONCLUSION

Summary of Findings

This study presents an implementable hybrid machine learning–based travel recommendation system capable of providing accurate and personalized destination suggestions in real time. The proposed Hybrid Weighted Recommender Ensemble (HWRE) achieved Precision@10 = 0.314, Recall@10 = 0.283, and RMSE = 0.925 on a dataset exceeding 1.25 million user–destination interactions. Compared to classical baselines such as Collaborative Filtering and Matrix Factorization, the system delivered an average improvement of over 25% in recommendation accuracy and a 35% reduction in ranking errors for new users and destinations. The framework also maintained low response latency (~78 ms) when deployed on a Flask-based environment, confirming its production readiness and scalability. Feature engineering and semantic embeddings substantially enhanced personalization, while cross-validation and perturbation experiments verified the model’s stability and robustness.

Limitations

The system currently operates on simulated and consolidated datasets and has not been integrated with live travel APIs or booking platforms. This may limit real-world generalization due to unseen data variability and user behaviour drift. The hybrid ensemble, while efficient, is still constrained by the quality of user feedback and sparsity of implicit interactions. Additionally, the system’s performance was measured under controlled hardware conditions, and further testing on distributed or mobile environments may yield different latency profiles.

Future Directions

Future research will focus on integrating context-aware deep learning architectures such as Transformers and Graph Neural Networks to capture richer temporal and relational dependencies. Incorporating reinforcement learning mechanisms could enable adaptive itinerary optimization based on user feedback loops. Moreover, deploying the system in a federated learning framework would allow collaborative personalization across travel platforms while preserving user privacy. Expanding the model to handle real-time dynamic data streams, sentiment analysis of reviews, and cross-domain recommendation (e.g., hotels + local experiences) will further enhance practical applicability.

Overall, the proposed system defines a scalable, low-latency, and reproducible basis for next-generation intelligent travel recommendation platforms that incorporate the merging of accuracy, interpretability, and deployment efficiency.



REFERENCES

- [1] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734–749, 2005.
- [2] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [3] M. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook*. Springer, 2022.
- [4] X. Su and T. M. Khoshgoftaar, "A survey of collaborative filtering techniques," *Advances in Artificial Intelligence*, vol. 2009, pp. 1–19.
- [5] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez, "Recommender systems survey," *Knowledge-Based Systems*, vol. 46, pp. 109–132, 2013.
- [6] S. Zhang, L. Yao, A. Sun, and Y. Tay, "Deep learning based recommender system: A survey and new perspectives," *ACM Computing Surveys*, vol. 52, no. 1, pp. 1–38, 2019.
- [7] L. Zheng, V. Noroozi, and P. S. Yu, "Joint deep modeling of users and items using reviews for recommendation," in *Proc. 10th ACM Int. Conf. Web Search and Data Mining (WSDM)*, pp. 425–434, 2017.
- [8] R. Burke, "Hybrid recommender systems: Survey and experiments," *User Modeling and User-Adapted Interaction*, vol. 12, no. 4, pp. 331–370, 2002.
- [9] J. Wang, W. Wang, Y. Hu, and X. Zhang, "A hybrid recommendation model integrating collaborative filtering and content-based methods," *Expert Systems with Applications*, vol. 185, p. 115593, 2021.
- [10] Kaggle, "TripAdvisor Hotel Review Dataset." [Online]. Available: <https://www.kaggle.com/datasets/andrewmvd/trip-advisor-hotel-reviews> (accessed Oct. 2025).
- [11] M. Tavakol and A. Brefeld, "Factorized neural network models for contextual recommendations," in *Proc. 26th Int. Conf. World Wide Web Companion (WWW)*, 2017, pp. 1623–1631.
- [12] Flask Documentation, *Flask Framework (Python Web Library)*, Version 3.0, 2024. [Online]. Available: <https://flask.palletsprojects.com>
- [13] M. Saveski and A. Mantrach, "Item cold-start recommendations: Learning local collective embeddings," in *Proc. 8th ACM Conf. Recommender Systems (RecSys)*, 2014, pp. 89–96.
- [14] F. Sanchez, A. Bellogín, and I. Cantador, "A comparative study of hybrid recommender systems for tourism," *Information Sciences*, vol. 475, pp. 420–435, 2019.
- [15] J. Nilashi, O. Ibrahim, and H. Ahmadi, "A recommender system based on collaborative filtering using ontology and dimensionality reduction techniques," *Expert Systems with Applications*, vol. 92, pp. 507–520, 2018.