The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM)
ISSN: 2602-3199

The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM), 2026

Volume 40, Pages 98-104

ICBAST 2026: International Conference on Basic Sciences and Technology

Travel DNA: Persona-Driven Destination Prediction with Sentiment-Enhanced Hybrid Modeling

Busra Gedikoglu
Konya Technical University

Sait Ali Uymaz
Konya Technical University

Abstract: This study develops a sentiment analysis-based personalized travel destination prediction system using 12.57 million user reviews from the online accommodation platform Airbnb. The proposed system operates on a dataset covering 20 European cities through a five-stage pipeline: (i) data extraction and preprocessing, (ii) sentiment analysis with the multilingual transformer-based XLM-RoBERTa language model, (iii) traveler profiling with the K-Means clustering algorithm, (iv) trajectory mining with Markov chains, and (v) a three-signal hybrid prediction engine. The proposed hybrid architecture combines three signals: a Markov signal that models inter-city transition probabilities, a history signal reflecting users' revisit tendency, and a sentiment-city affinity signal that matches users' sentiment profiles with city characteristics. These signals are combined with weights optimized per traveler persona. In experiments conducted on 10,436 users (20.9% coverage rate) who reviewed during the test period and had travel history in the training period, the correct destination was found in the first recommendation (Hit@1) at 24.42%, in the top three (Hit@3) at 40.27%, in the top five (Hit@5) at 50.80%, and the Mean Reciprocal Rank (MRR) measuring ranking quality was 33.56%. Ablation study measuring each component's contribution reveals that the hybrid model outperforms all individual signals across every metric and that the proposed sentiment-city affinity signal achieves a 76% relative improvement in Hit@1 over the popularity-based baseline.

Keywords: Destination prediction, Sentiment analysis, Hybrid recommendation system, Persona-based recommendation, Markov chain

Introduction

With the growth of the digital sharing economy, online accommodation platforms such as Airbnb constitute a unique data source for understanding travel behavior. The millions of user reviews accumulated on these platforms contain travelers' emotional experiences, expectations, and satisfaction levels beyond numerical ratings. Bridges and Vásquez (2018) demonstrated that Airbnb reviews contain meaningful information despite the vast majority being positive. This rich textual data enables the development of sentiment-driven personalized prediction models beyond traditional recommendation systems.

Adomavicius and Tuzhilin (2005) categorized recommender systems into two primary paradigms: collaborative filtering (CF), which leverages preferences of similar users, and content-based filtering, which matches item features. However, the data sparsity problem in the travel domain significantly limits the effectiveness of CF methods. Meng et al. (2023) emphasizes the importance of evaluating context-aware approaches that incorporate additional information, such as time, location, and usage scenarios, into the model to overcome this limitation. Furthermore, traditional approaches fail to adequately model users' emotional experiences and sequential dependencies in travel trajectories. Sequential recommendation systems attempt to fill this gap by modeling temporal dependencies in users' action sequences. Quadrana et al. (2018) comprehensively reviewed methods in this area. Markov chains are widely used in travel trajectory prediction due to their low computational cost and interpretability advantages (Xu et al., 2022). However, these models also disregard traveler psychology and emotional compatibility.

In this study, to address the aforementioned gaps, a system called Travel DNA is proposed. The system extracts travelers' emotional profiles through multilingual sentiment analysis, identifies five distinct traveler personas via a clustering algorithm, and presents a three-signal hybrid prediction model with optimized weights for each persona. The main contributions of this study can be summarized as follows:

(1) A large-scale multilingual sentiment analysis study on 12.57 million Airbnb reviews covering 20 European cities,
(2) A sentiment-based profiling method that classifies traveler behavior patterns under five distinct personas,
(3) Proposal of a signal designed in this study called Personalized Sentiment-City Affinity, demonstrating a 76% improvement in the Hit@1 metric over the popularity-based baseline,
(4) Development of a three-signal hybrid prediction model with persona-specific weight optimization and validation through a comprehensive ablation study.

Related Work

Travel Recommendation Systems

Travel recommendation systems aim to predict future destination preferences based on users' past behavior. Collaborative filtering methods discover preferences of similar users through matrix factorization techniques (ALS, SVD) and k-nearest-neighbor algorithms (Adomavicius & Tuzhilin, 2005). However, CF methods suffer from degraded performance due to the data sparsity problem in the travel domain. Meng et al. (2023) provided a comprehensive review of context-aware recommender systems. Content-based approaches perform matching using destination features (price, location, facility type); however, they remain insufficient in capturing the emotional dimension of user experience.

Sentiment Analysis and Tourism

Advances in natural language processing, particularly transformer-based pre-trained language models, have enabled significant progress in sentiment analysis. The BERT model proposed by Devlin et al. (2019) established a new performance standard in text classification tasks. The XLM-RoBERTa model developed by Conneau et al. (2020) has been trained on over 100 languages, providing a strong foundation for multilingual sentiment analysis. In the tourism domain, sentiment analysis has been employed to improve demand forecasting performance with granular sentiment scores (Li et al., 2023) and to support deep learning approaches in personalized itinerary recommendation (Halder et al., 2024). However, the direct use of sentiment profiles for personalized destination prediction has been explored to a limited extent. This study distinguishes itself from existing work by designing a personalization signal that combines users' personal sentiment profiles with a city similarity matrix.

Sequential Recommendation Systems

Sequential recommendation systems model temporal dependencies in users' action sequences. Quadrana et al. (2018) comprehensively classified methods in this area. Markov chains are widely preferred in travel trajectory prediction due to their interpretability and computational efficiency (Xu et al., 2022). Rendle et al. (2010) achieved successful results in the next-basket prediction problem with factorized personalized Markov chains. Deep learning-based models have emerged prominently in recent years: GRU4Rec, proposed by Hidasi et al. (2016) in session-based recommendations, and SASRec, proposed by Kang and McAuley (2018), achieved high performance in sequential recommendations with an attention mechanism. In this study, the hierarchical Markov approach was preferred to maintain interpretability and produce stable predictions under limited data conditions.

Method

The proposed Travel DNA system consists of a five-stage pipeline: (i) data collection and preprocessing, (ii) multilingual sentiment analysis, (iii) traveler profiling, (iv) trajectory mining, and (v) hybrid prediction model. Each stage is described in detail in this section.

Dataset

The dataset consists of user reviews obtained from the Inside Airbnb (2025) project, covering 20 major European cities. Cities covered: Amsterdam, Barcelona, Berlin, Brussels, Budapest, Copenhagen, Dublin, Edinburgh, Florence, Lisbon, London, Madrid, Milan, Munich, Paris, Porto, Prague, Rome, Venice, and Vienna. A total of 12,570,926 reviews encompass 9,181,198 unique user profiles. Table 1 presents the basic statistics of the dataset.

Table 1. Dataset statistics

Feature	Value
Total Reviews	12,570,926
Unique Users	9,181,198
Number of Cities	20
Training Set (75%)	9,428,194
Validation Set (5%)	628,546
Test Set (20%)	2,514,186

Temporal split was used as the evaluation methodology: data was divided chronologically into 75% training, 5% validation, and 20% test sets. This approach prevents data leakage from future data into training (Quadrana et al., 2018). During prediction, only each user's history from the training period (last visited city) is used as input, and the first review in the test period is considered as the target.

Multilingual Sentiment Analysis

Sentiment analysis on review texts was performed using a model fine-tuned for sentiment classification on top of the XLM-RoBERTa base model developed by Conneau et al. (2020). XLM-RoBERTa is a multilingual transformer model pre-trained on over 100 languages, capable of processing reviews in European languages such as English, French, German, Spanish, and Portuguese without requiring translation. The model produces three categorical outputs—negative, neutral, or positive—and a confidence score for each review. These outputs were converted into a weighted sentiment score in the [-1, +1] range by multiplying the category value (-1, 0, +1) by the confidence score. The resulting score captures the emotional direction and intensity of the user's experience in a particular city.

Traveler Profiling

A five-dimensional profile vector was constructed for each traveler: mean sentiment score, sentiment variability (standard deviation), mean price segment, number of unique cities visited, and logarithmic visit frequency. Price segments were defined in four categories based on room types in Airbnb listing data: shared room (0), private room (1), entire home/apartment (2), and hotel room (3). Users were divided into five distinct traveler personas using the K-Means clustering algorithm (Sinaga & Yang, 2020). The optimal number of clusters was determined by the elbow method. Table 2 shows persona definitions and their distributions in the test set. Persona names are interpretive labels assigned based on the behavioral characteristics of clusters and optimized signal weights.

Table 2. Traveler persona definitions

Persona ID	Persona Name	Number of Users (n)	Key Characteristic
P0	Single-City Traveler	813	Limited history, prediction from sentiment profile
P1	Variable Traveler	1,699	Inter-city transitions, loyalty significant
P2	Mainstream Traveler	7,287	Large mass, sequential transition patterns
P3	Predictable Traveler	578	High loyalty, revisit tendency
P4	Sequential Traveler	59	Sequential transition patterns dominant, highest Markov signal

Hierarchical Markov Chain

Inter-city transition probabilities were modeled with a hierarchical Markov structure. In transition mining, if the interval between two consecutive reviews of the same user was 21 days or shorter, these two cities were considered a connected travel transition; longer intervals were filtered as independent trips. Persona-based first-order transition matrices were constructed only for personas with multi-city visits (P1–P4); since the P0 (Single-City Traveler) group, by definition, has no inter-city transition data, no Markov matrix was generated for this persona. Predictions for P0 users are based on sentiment affinity and global popularity signals. During prediction, the system first consults the second-order Markov chain (last two-city pair), falls back to the persona-based first-order transition matrix when insufficient data is available, and ultimately falls back to the global popularity distribution. This hierarchical fallback mechanism prioritizes the most specific information available while addressing the data sparsity problem.

Hybrid Prediction Model

The proposed hybrid system combines three independent signals with persona-specific weights:

Hierarchical Markov Signal (Mk): Computes the transition probability of the next city based on the user's last visited city/cities. Produces stable predictions under data sparsity conditions through the hierarchical fallback mechanism.

Visit History Signal (Hi): Reflects the user's normalized visit frequency per city. This signal captures travelers' tendency to revisit specific cities (loyalty).

Personalized Sentiment-City Affinity Signal (Se): This is the core contribution of this study. Sentiment scores from the cities the user has visited are multiplied by the inter-city similarity matrix to produce affinity scores for cities not yet visited. The city similarity matrix was computed using cosine similarity from seven features: each city's mean sentiment score, sentiment standard deviation, proportional distribution of four price segments, and logarithmic visitor volume. Features were subjected to z-score normalization with StandardScaler to prevent different scale ranges from affecting the similarity computation.

The final prediction is computed as given in Equation (1):

$$\hat{P}(\text{city}|\text{user}) = w_1^p \cdot \text{Mk(city)} + w_2^p \cdot \text{Hi(city)} + w_3^p \cdot \text{Se(city)} \tag{1}$$

Here, p represents the persona index to which the user belongs, and the constraint $w_1^p+w_2^p+w_3^p=1$ is satisfied. Weights were optimized separately for each persona via grid search on the validation set using the combined metric (0.3×Hit@1 + 0.4×Hit@3 + 0.3×Hit@5). The search space was scanned with a step size of 0.05 in the [0, 1] range for each weight, evaluating a total of 153 combinations that satisfied the constraint $w_1 + w_2 + w_3 = 1$.

Results and Discussion

Evaluation Metrics

System performance was evaluated using metrics widely adopted in the recommender systems literature (Adomavicius & Tuzhilin, 2005): Hit@k (k=1,3,5) measures the rate at which the actual destination appears in the top k elements of the generated recommendation list. NDCG@k (Normalized Discounted Cumulative Gain) evaluates ranking quality with logarithmic discounting. The Mean Reciprocal Rank (MRR) is defined in Equation (2), where |U| denotes the total number of users and rank_i represents the position of the correct destination in the ranked list for user i.

$$\text{MRR} = (1/|U|) \sum 1/\text{rank}_i \tag{2}$$

For example, if the correct prediction is at the first position, it yields 1/1=1; if it is at the third position, it yields 1/3 ≈ 0.33.

Ablation Study

An ablation study was conducted to measure the contribution of the hybrid model and each signal. In the ablation study, system components are disabled one by one to evaluate the impact of each component on performance. For this purpose, six different model configurations were compared. Table 3 presents the results. Table 3 presents the results of six model configurations evaluated on 10,436 test users. The results clearly demonstrate that the proposed hybrid model outperforms all individual signals and baselines across every metric. Particularly noteworthy is that the sentiment-city affinity signal proposed in this study (Hit@1=15.48%) achieves 76% relatively higher performance than the popularity baseline (Hit@1=8.80%) ((15.48-8.80)/8.80×100≈76%). This difference confirms that sentiment-based personalization produces user-specific meaningful signals beyond merely following popularity. The history signal is the strongest individual signal in Hit@1 (24.24%); however, its growth rate slows in Hit@3 (32.13%) and Hit@5 (37.12%), falling behind the sentiment affinity signal (Hit@3=36.74%, Hit@5=46.95%). This indicates that the history signal is strong for the first prediction but insufficient for generating diverse recommendation lists. The sentiment affinity signal offers a more balanced performance profile. The proposed hybrid model achieves the highest performance across all metrics by combining these complementary strengths through a weighted linear combination.

Table 3. Ablation study results

Model	Hit@1 (%)	Hit@3 (%)	Hit@5 (%)
Random Baseline	5.30	15.89	26.63
Popularity Baseline	8.80	33.77	48.42
Markov Only	8.00	22.77	35.96
History Only	24.24	32.13	37.12
Sentiment Affinity Only	15.48	36.74	46.95
Hybrid Model (Proposed)	24.42	40.27	50.80

Overall Performance

The system achieved Hit@1=24.42%, Hit@3=40.27%, Hit@5=50.80%, NDCG@3=0.3352, NDCG@5=0.3782, MRR=33.56% on the test set. The Hit@5 value indicates that for more than half of the test users, the correct destination was among the top five recommendations. Additionally, multi-step trajectory prediction was evaluated: the system predicts two cities for each user, and the intersection rate (set-precision) between these predictions and the cities actually visited by the user is computed. Average set-precision was measured at 15.43% over 2,000 users.

Persona-Specific Analysis

Grid search optimization produced distinctly different weight distributions for each persona. Table 4 shows persona-based performance and optimized weights, where w_Mk, w_Hi, and w_Se denote the weights assigned to the Markov, visit history, and sentiment-city affinity signals, respectively.

Table 4. Persona-based performance and signal weights

Persona ID	Number of Users (n)	H@1	H@3	H@5	w_Mk	w_Hi	w_Se
P0	813	100	100	100	0.05	0.10	0.85
P1	1,699	23.0	39.6	50.1	0.25	0.55	0.20
P2	7,287	13.9	31.5	43.7	0.40	0.45	0.15
P3	578	54.3	67.3	72.0	0.20	0.65	0.15
P4	59	27.1	54.2	61.0	0.65	0.25	0.10

The weight distributions offer valuable insights into the decision-making processes of different traveler types. For P0, sentiment affinity is the dominant signal (0.85), and these users having only a single-city history result in 100% accuracy. However, the impact of this trivial prediction mechanism on overall metrics should be considered: P0's share in the test set is 7.8% (813 / 10,436), and when P0 is excluded (n=9,623), Hit@1=18.03%, Hit@3=35.23%, Hit@5=46.64% are obtained.In mainstream travelers (P2), the balance between Markov (0.40) and history (0.45) reveals the importance of sequential patterns; in P3, history loyalty (0.65) reflects revisit tendency; and in P4, Markov dominance (0.65) demonstrates the deterministic role of sequential transition patterns. This differentiation demonstrates that persona-specific optimization is superior to uniform weights.

Discussion

Value of personalization: The sentiment-city affinity signal surpassing the popularity baseline by 76% in Hit@1 confirms that users' sentiment profiles carry tangible value for destination preference prediction.

Advantage of the hybrid model: The ablation study demonstrates that no single signal is the best across all metrics. The proposed hybrid model achieves the highest performance across all metrics by combining the complementary strengths of signals through weighted aggregation.

Limitations: The main limitations of the system are: (i) The dataset is limited to Airbnb only. (ii) The scope of 20 European cities is limited. (iii) Predictions cannot be generated for cold-start users; evaluation is restricted to users with training history (10,436 out of 50,000, coverage 20.9%). (iv) The P0 group's 100% accuracy stems from trivial predictions. (v) Comparison with SASRec (Kang & McAuley, 2018) and BERT4Rec (Sun et al., 2019) is left to future work. (vi) The CF signal was removed from the system due to negligible contribution.

Conclusion

In this study, the Travel DNA system was proposed for sentiment analysis-based personalized travel destination prediction using 12.57 million Airbnb reviews. The three-signal hybrid architecture (Markov + History + Sentiment Affinity) outperformed all individual signals and baselines across every evaluation metric (Hit@1=24.42%, Hit@5=50.80%). The proposed sentiment-city affinity signal achieved a 76% improvement in Hit@1 over the popularity-based approach, proving that sentiment analysis is a meaningful source of personalization for destination prediction. Persona-specific weight optimization strengthened the model's interpretability. Future work plans include comparison with deep learning-based sequential models, incorporation of additional signals such as seasonality and price dynamics, and scaling the system to broader city clusters.

Recommendations

Based on the findings of this study, future research directions are recommended: (i) Incorporating deep learning-based sequential models such as SASRec and BERT4Rec as additional baselines for performance comparison. (ii) Extending the dataset beyond Airbnb to include reviews from multiple platforms (e.g., Booking.com, TripAdvisor) for cross-platform generalizability. (iii) Integrating temporal features such as seasonality, holiday periods, and price dynamics into the prediction model. (iv) Developing a cold-start solution for users with no prior travel history by leveraging demographic or contextual features. (v) Scaling the system to a broader geographic scope including non-European destinations.

Scientific Ethics Declaration

* The authors declare that the scientific ethical and legal responsibility of this article published in EPSTEM journal belongs to the authors.

* This study uses publicly available data from the Inside Airbnb project. No personally identifiable information was collected or processed. Ethics committee approval was not required as the study involves only publicly shared anonymous review data.

Conflict of Interest

* The authors declare that they have no conflicts of interest.

Funding

* This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Acknowledgements or Notes

* This article was presented as an oral presentation at the International Conference on Basic Sciences and Technology (www.icbast.net) held in Konya/Türkiye on April 02-05, 2026.

References

Adomavicius, G., & Tuzhilin, A. (2005). 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, 17*(6), 734–749.

Bridges, J., & Vásquez, C. (2018). If nearly all Airbnb reviews are positive, does that make them meaningless? *Current Issues in Tourism, 21*(18), 2057–2075.

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., & Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In *Proceedings of ACL* (pp. 8440–8451). arXiv:1911.02116

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT* (pp. 4171–4186).

Halder, S., Lim, K. H., Chan, J., & Zhang, X. (2024). A survey on personalized itinerary recommendation: From optimisation to deep learning. *Applied Soft Computing, 152*, 111200.

Hidasi, B., Karatzoglou, A., Baltrunas, L., & Tikk, D. (2016). Session-based recommendations with recurrent neural networks. In *Proceedings of ICLR*. arXiv:1511.06939

Inside Airbnb. (2025). *Adding data to the debate*. http://insideairbnb.com

Kang, W. C., & McAuley, J. (2018). Self-attentive sequential recommendation. In *Proceedings of IEEE ICDM* (pp. 197–206).

Li, H., Gao, H., & Song, H. (2023). Tourism forecasting with granular sentiment analysis. *Annals of Tourism Research, 103*, 103667

Meng, X., Du, Y., Zhang, Y., & Han, X. (2023). A survey of context-aware recommender systems: From an evaluation perspective. *IEEE Transactions on Knowledge and Data Engineering, 35*(7), 6708–6727.

Quadrana, M., Cremonesi, P., & Jannach, D. (2018). Sequence-aware recommender systems. *ACM Computing Surveys, 51*(4), 1–36.

Rendle, S., Freudenthaler, C., & Schmidt-Thieme, L. (2010). Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of WWW* (pp. 811–820).

Sinaga, K. P., & Yang, M.-S. (2020). Unsupervised K-Means clustering algorithm. *IEEE Access, 8*, 80716–80727.

Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., Ou, W., & Jiang, P. (2019). BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformers. In *Proceedings of CIKM* (pp. 1441–1450).

Xu, Y., Zou, D., Park, S., Li, Q., Zhou, S., & Li, X. (2022). Understanding the movement predictability of international travelers using a nationwide mobile phone dataset collected in South Korea. *Computers, Environment and Urban Systems, 92*, 101753.

Author(s) Information

Busra Gedikoglu
Konya Technical University
Department of Computer Engineering, Konya, Türkiye
ORCID iD: https://orcid.org/0009-0009-1574-8576
*Contact e-mail: *e258229001018@ktun.edu.tr*

Sait Ali Uymaz
Konya Technical University
Department of Computer Engineering, Konya, Türkiye
ORCID iD: https://orcid.org/0000-0003-2748-8483

*Corresponding author(s) contact email

To cite this article:

Gedikoglu, B., & Uymaz, S. A. (2026). Travel DNA: Persona-driven destination prediction with sentiment-enhanced hybrid modeling. *The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM), 40,* 98-104.