Intelligent Tourism Itinerary Generation Through Natural Language Processing and Hybrid Recommendation Systems

Saqib Shouqi
School of Computing,
Informatics Institute of Technology,
Sri Lanka
saqibshouqi@gmail.com

Sahan Priyanayana
School of Computing,
Informatics Institute of Technology,
Sri Lanka
sahan.p@iit.ac.lk

S.C. Thennakoon
School of Computing,
Informatics Institute of Technology,
Sri Lanka
chirasthi.20232005@iit.ac.lk

Viruna Vidaswin
School of Computing,
Informatics Institute of Technology,
Sri Lanka
viruna.20231903@iit.ac.lk

Methupa Ninduwara
School of Computing,
Informatics Institute of Technology,
Sri Lanka
methupa.20232225@iit.ac.lk

Keywords—Natural Language Processing, Recommender Systems, Tourism Analytics, Machine Learning, Intelligent Systems

I. Introduction

Traditional travel planning systems rely on rigid form-based interfaces with predefined dropdown menus, limiting user's ability to express nuanced preferences. This constraint often results in generic itineraries that fail to capture individual travel styles, budgets, or interests. To address this gap, we developed TravelMate AI, a commercial itinerary builder that introduces two key innovations: natural language processing (NLP) for interpreting free-text trip descriptions and a hybrid recommendation system combining machine learning with large language models (LLMs). The novelty lies in enabling users to describe their travel preferences conversationally (e.g., "Cultural tour of Kandy with family, medium budget") while maintaining structured outputs suitable for commercial deployment. By eliminating the need for form-based inputs, the system democratizes travel planning for non-technical users while retaining the precision of AI-driven recommendations.

Furthermore, unlike existing commercial solutions, which either rely on structured inputs (TripAdvisor) or quiz-based preference gathering (MindTrip), TravelMate AI processes unstructured text directly, enabling richer and more flexible itinerary customization. This adaptability allows users to input diverse and complex preferences without rigid constraints, making travel planning more personalized and user-friendly.

II. Literature review

Recent advancements in AI-driven tourism systems emphasize constraint-based algorithms and collaborative filtering. Souffriau et al. (2009) pioneered metaheuristics for itinerary optimization, while Dai et al. (2016) integrated genetic algorithms for route planning. Refanidis et al. (2014) introduced dynamic itinerary recommendations responsive to user preferences and geographical constraints. Zhiwen et al. (2014) leveraged collective footprint data for group-based personalization. Commercial platforms such as Layla AI, TripAdvisor, and MindTrip utilize conversational interfaces and review-based recommendations but remain limited by structured inputs, quizzes, or external user content, thus restricting detailed personalization. Chen et al. (2017) further highlighted big data's potential for enhancing personalized recommendations through large-scale datasets.

TravelMate AI overcomes these limitations via three innovations: first, NLP techniques directly extract nuanced preferences from free-form user text; second, location-specific Random Forest models ensure recommendations tailored specifically to Sri Lanka's tourism landscape; and third, large language models (LLMs) dynamically enforce geographical and time constraints during itinerary synthesis. The methodology (Section III) and experimental validation (Section IV) detail these contributions further.

III. Methodology

The proposed system architecture integrates three distinct components designed to work collaboratively: NLP-Driven Preference Extraction, Hybrid Recommendation Engine, and LLM-Based Itinerary Synthesis and Scheduling. Each component is discussed individually below.

1. NLP-Driven Preference Extraction

A spaCy-based pipeline processes user descriptions using fuzzy matching and semantic similarity analysis. Valid locations are identified from a curated list of 23 Sri Lankan destinations, while Word2Vec embeddings map free-text interests (e.g., "historic ruins") to predefined categories (e.g., "culture & heritage"). Companion types and budget levels are extracted through pattern matching, enabling the system to handle varied phrasings like "traveling with my spouse" or "luxury trip."

2. Hybrid Recommendation Engine

https://doi.org/10.31705/ADScAI.2025.36

The system employs location-specific Random Forest classifiers to generate personalized activity recommendations. Each destination utilizes a separate model trained on a domain-expert-validated dataset containing Sri Lankan tourism activities. Categorical features (budget level, companion type) are encoded via LabelEncoder, while user interests are transformed into binary vectors using MultiLabelBinarizer to handle multi-label classifications. The Random Forest architecture (100 estimators, balanced class weights) addresses dataset imbalances and prevents bias toward frequent activities. During inference, the model combines encoded budget/companion features with interest vectors to calculate activity probabilities. These probabilities are sorted to prioritize contextually relevant suggestions (e.g., prioritizing "Sigiriya Rock Fortress tours" for cultural enthusiasts) while filtering invalid options through a location-specific allow-list derived from tourism domain knowledge. This dual approach ensures recommendations align with both statistical patterns and geographical feasibility.

3. LLM-Based Itinerary Synthesis

Groq's Llama-3-70B model generates structured itineraries using dynamically engineered prompts that enforce arrival/departure times and logical location sequences. The LLM outputs JSON data with daily schedules, accommodation transitions, and cost estimates, which a DOCX templating engine converts into user-ready documents.

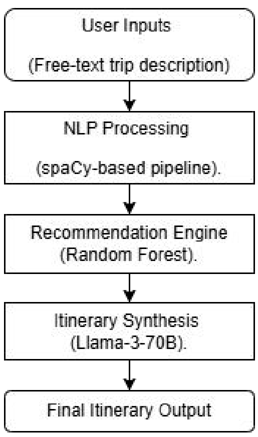

Figure 1: System Architecture of the Itinerary Builder

IV. EXPERIMENTS

The system was deployed with a Sri Lankan tourism agency for real-world validation. Agents reported a 90% reduction in itinerary preparation time, attributing this efficiency to the NLP interface's ability to capture client needs without iterative form adjustments. In a pilot study with 50 users, 82% rated the generated itineraries as "highly personalized," with particular praise for the system's handling of complex requests like "family-friendly hikes near Ella with vegetarian dining options." Qualitative feedback highlighted the value of editable recommendation checkboxes, which let users refine AI suggestions without restarting the planning process.

To evaluate, we measured:

- Random Forest Model Performance: Precision = 85%, Recall = 84%, F1-Score = 84%.
- NLP Entity Extraction Accuracy: 93% when tested on 100 user inputs.

Qualitative Long-Term User Feedback

A two-month follow-up with the initial 50 users showed consistently high satisfaction, particularly due to TravelMate AI's conversational flexibility. Users highlighted that freely expressing evolving preferences notably boosted their continued engagement and trust.

Comparative Analysis of Existing Systems

We compared TravelMate AI with existing solutions: Layla AI, TripAdvisor, MindTrip, and Wonder Plan. Layla AI and Wonder Plan offer conversational interactions but rely on structured suggestions, whereas TripAdvisor and MindTrip use structured questions, quizzes, or external user content, limiting detailed personalization. In contrast, TravelMate AI uniquely leverages advanced NLP to allow travelers to naturally articulate complex preferences, providing more personalized, flexible, and user-centric itineraries.

Limitations

Despite significant advancements in itinerary personalization, TravelMate AI's NLP accuracy can be affected by ambiguous language and subtle cultural nuances, occasionally impacting recommendation precision.

V. CONCLUSION

TravelMate AI demonstrates that NLP and hybrid recommendation systems can bridge the gap between unstructured user inputs and structured travel planning requirements. By combining spaCy-based parsing, Random Forest classifiers, and LLMs, the system supports natural interaction while maintaining commercial-grade output precision. Future work will expand destination coverage and integrate real-time pricing APIs, addressing current limitations in dynamic cost estimation. The commercial adoption by tourism agencies underscores the practicality of this approach in real-world settings.

REFERENCES

[1] Souffriau, W., et al. (2009). Tourist Trip Design Algorithms. Annals of Tourism Research.

[2] Dai, J., et al. (2016). PSR Framework for Travel Route Planning. IEEE Transactions on Knowledge and Data Engineering.

[3] Honnibal, M., & Montani, I. (2017). spaCy: Industrial-strength NLP in Python. Journal of Open Source Software.

[4] Refanidis, I., et al. (2014). myVisitPlannerGR: Personalized itinerary planning system for tourism.

[5] Zhiwen, Y., et al. (2014). Personalized tourist track planning method based on group footprints.

[6] Chen, S., et al. (2017). Personalized travel service big data application system and method.

[7] Ricci, F., et al. (2015). Recommender Systems Handbook. Springer.