

AI-Based Personalized Travel Planning and Recommendation System: A Unified Framework for Intelligent Trip Planning

Shubham Pandey, Satyam Anand, Sujal Das, Raushan Kumar Tiwari, Vandana Tripathi

Department of Computer Science and Engineering, MGM's College of Engineering and Technology, Noida, India

Published October 2026

Abstract

A growing number of travel planning activities involve managing multiple pieces of often competing data about destinations, budget, weather, attractions' availability, transportation options and personal preferences that are typically distributed across a number of isolated online services. Traditional travel-planning tools tackle these issues one at a time: choosing or navigating, and rarely include personalization, predictive analytics, and multi-constraint reasoning within a single interface. We propose a unified framework for personalized travel planning that integrates personal preferences, context-based destination recommendation, budget- and weather- aware filtering, and route and distance analysis, and AI-assisted itinerary generation. The framework is inspired, and to some extent generalized, from the design of the ongoing academic research prototype TripCraft, an AI-powered agent for trip planning. We build upon the state of the art of recommender-systems theory, context-aware recommendation, tourist trip design (route optimization) literature, and recent large language model-based agents for travel-planning to develop a system-level description, a hybrid AI/ML approach, and a step-by-step algorithmic process for creating personalized itineraries that take preference, budget, weather and route constraints into account. We provide preliminary proof of concept based on a synthetic evaluation on fabricated records using baseline comparisons and constraint-compliance metrics. The findings are only indicative of the possible effectiveness, but do not constitute proof of concept. The expected contribution is the conceptual and methodological foundation for unifying preference-, budget-, weather- and route awareness in intelligent travel planning.

Keywords: artificial intelligence; personalized travel planning; machine learning; recommender systems; itinerary generation; TravelTech

1 Introduction

Digital travel platforms have proliferated over the last two decades, transitioning from static directory listings to dynamic booking, review and mapping platforms, and growing substantially in both the volume and diversity of available travel data — reviews, geotagged points of interest, live weather feeds, pricing and transport schedules [3], [4]. Advances in machine learning and recommender-systems research [1], [2] have made it feasible to model user preferences and produce context-sensitive suggestions across many domains, including tourism [3]. A traveller planning a multi-day trip must, however, reconcile several interacting constraints simultaneously: destination and attraction interest, available budget, expected weather, trip duration, transportation options and the geographic feasibility of the resulting route. Existing tools typically address one or two of these dimensions in isolation [3], [4], leaving travellers to manually reconcile recommendations derived from separate booking, weather, mapping and review applications.

Artificial intelligence provides a natural mechanism for this integration problem: content-based and collaborative filtering can model preferences [1], [2] and context-aware techniques can incorporate situational factors such as weather [8], [9] whereas combinatorial optimization developed for the tourist trip design problem can generate feasible multi-day routes [7]. Large language models have more recently been explored for translating structured recommendations into natural-language itineraries [10], [11], [12], though current evidence suggests that unconstrained language-model planning alone struggles to satisfy the full set of real-world constraints involved in trip planning [11]. There is, however, no widely adopted framework that treats preference, budget, weather and route information as jointly, rather than independently, optimized inputs. This motivates the central research question addressed in this paper:

“How can AI and machine-learning techniques be integrated with travel information, user preferences, budget, weather and route information to generate personalized travel plans?”

Our objective is to propose a unified AI-based framework that integrates heterogeneous travel information and user-specific preferences to support personalized trip planning, investigating how AI/ML techniques can be combined for preference analysis, destination recommendation, budget- and weather-aware planning, route/distance analysis and personalized itinerary generation. The contributions of this paper are: (i) a unified conceptual framework treating preference, budget, weather and route information as jointly considered inputs; (ii) a layered system architecture informed by the TripCraft AI project's technology stack; (iii) a multi-factor AI/ML methodology; (iv) an algorithmic formulation for constraint-aware itinerary generation; and (v) a planned evaluation methodology to guide future experimental validation.

2 Literature review

Traditional and personalized recommender systems: Early travel-recommendation systems relied on directory-style listings or simple similarity retrieval. Borràs et al. [3] survey intelligent tourism recommender systems and show that tourism recommendation has been studied from the viewpoint of several services, including destination and activity recommendation. Chaudhari and Thakkar [4] review travel recommender systems across aspects such as transportation, trip duration, accommodation, weather and events, and point out the wide range of objectives and contexts these systems address. At a more general level, Adomavicius and Tuzhilin [1] classify recommendation into content-based, collaborative and hybrid approaches and describe context-aware extensions, while Burke [2] surveys hybrid recommenders and explains how combining different techniques can offset the weaknesses each has on its own.

Machine learning and context-aware recommendation: Liu et al. [5] propose a personalized travel-package model that captures both tourist interests and package characteristics, including seasonal and compositional effects. The authors report better recommendation performance than traditional methods on the dataset they evaluated. Yoon and Choi [9] present R2Tour, a real-time context-aware tourism recommender that uses temperature, precipitation and visitor demographics. Their work shows that situational signals allow recommendations to adapt to changing external conditions instead of relying on static historical patterns alone.

Weather- and budget-aware planning; route optimization: Braunhofer et al. [8] describe STS, a mobile recommender that takes weather forecasts at candidate points of interest into account, and report higher choice satisfaction than a weather-agnostic baseline in a live user study. Ge et al. [6] propose cost-aware latent-factor models that represent tourist interest and travel cost together,

including a variant with a Gaussian prior to model cost uncertainty. Gavalas et al. [7] survey exact, heuristic and metaheuristic approaches to the tourist trip design problem, an extension of the orienteering problem that deals with sequencing visits under time-window and travel-budget constraints. This work gives an established algorithmic basis for the routing and scheduling part of travel planning.

AI-based itinerary generation and generative AI: Chen et al. [10] introduce TravelAgent, an LLM-based travel planning system with four modules: tool usage, recommendation, planning and memory. It aims to produce reasonable, comprehensive and personalized itineraries in dynamic travel scenarios, includes explicit modelling of travel constraints, and is evaluated by the authors with human and simulated users. Xie et al. [11] present TravelPlanner, which is a benchmark rather than a planning system. It contains 1,225 planning intents with reference plans. In their evaluation, current language agents struggled to gather information and to handle several constraints at once, and GPT-4 reached a success rate of only 0.6% in the reported setting. This suggests that language-model agents alone are not yet reliable at satisfying realistic multi-constraint requirements. Chaudhuri et al. [12] introduce the TripCraft benchmark, which adds real-world constraints such as public transit schedules and events, and propose continuous metrics for temporal and spatial itinerary coherence rather than pass/fail checks. Taken together, these benchmarks make constraint-aware itinerary quality a clear evaluation target, and [11] shows that LLM-only approaches do not yet handle it reliably. This motivates the hybrid design considered in this paper, where structured AI/ML components handle feasibility and a generative component is reserved for natural-language presentation.

Figure 1 shows how the works reviewed above are distributed across the four focus areas of this review, based on the reference list.




Figure 1. Number of reviewed works per literature-review focus area.

3 Problem statement

While existing travel platforms provide valuable services, such as booking, navigation, and destination insights, to help plan trips, users are often left comparing different options across various dimensions (e.g., price, weather, transportation, and attractions) in different apps before making a final decision. In the research literature, previous work has primarily focused on specific aspects of trip planning, such as destination ranking [3]–[5], context-aware re-ranking [8], [9], route planning [7], or cost-aware recommendation [6]. By contrast, the present work focuses on the higher-level challenge of jointly considering preference, transportation, price, and weather constraints within a single planning framework. Recently, large language model (LLM)-based approaches to itinerary planning have gained attention due to their ability to generate natural language descriptions of plans [10]–[12]; while the TravelAgent approach [10] includes some elements of constraint modeling, the TravelPlanner benchmark [11] highlights limitations in the capability of language models to handle multiple constraints effectively. This creates an opportunity to explore a mixed approach, which combines the constraint-aware planning and recommendation capabilities of AI/ML with the natural language presentation and personalization features of LLMs, as demonstrated by the TripCraft AI system, the development of which serves as the motivation for the present work.

4 Comparative analysis of existing approaches

To make the research gap identified above concrete, Figure 2 compares the approach families reviewed in the Literature Review on four qualitative capability dimensions: personalization, budget-awareness, weather-awareness and route optimization. Scores are based on the design described in each cited source. Full (1) means the approach treats the dimension as an explicit, central part of its design; partial (0.5) means it is handled only indirectly or in a limited way; and absent (0) means the source does not address it. The comparison is qualitative and reflects our reading of prior work, not measured experimental performance. The score for the proposed framework reflects design intent only, since it has not yet been implemented.




Figure 2. Qualitative capability-coverage comparison across reviewed approach families, culminating in the proposed unified framework.

As Figure 2 illustrates, none of the approach families reviewed here covers all four dimensions at once. Hybrid and context-aware recommender systems [1], [2], [9] score well on personalization but not on budget or route optimization; cost-aware [6] and weather-aware [8] systems each add one further dimension; route-optimization research [7] covers routing but gives personalization only limited attention; and current LLM-based approaches [10]–[12] cover personalization well and the other dimensions partially, although benchmark evidence [11] indicates that they do not satisfy constraints reliably. The proposed framework is designed to cover all four dimensions jointly, which is the gap this paper addresses.

5 Research questions

The research questions that will be addressed in this study based on the research gap are as follows:

RQ1: How can user preferences, budget limitations, weather conditions, and route feasibility be combined in a travel itinerary generation framework?

RQ2: How does a recommendation approach that incorporates multiple constraints during recommendation generation compare with conventional recommendation approaches in terms of recommendation quality and itinerary feasibility?

RQ3: What is the influence of individual constraints, including budget limitations, weather conditions, and route feasibility, on itinerary quality?

6 Research objectives

An integrated approach to personalizing travel plans with multiple user and environmental constraints during itinerary generation will be developed by this study.

Specifically, this study will achieve the following objectives:

1. Designing a personalized travel recommendation framework that incorporates user preferences and trip requirements.
2. Incorporating the user's budget constraints into the recommendation and itinerary generation framework.
3. Considering the impact of weather conditions during activity and destination recommendation.
4. Including route feasibility and distance in the itinerary building process.
5. Comparing the recommendation approach based on the incorporation of constraints during recommendation generation against conventional recommendation approaches using appropriate evaluation metrics.
6. Evaluating the impacts of individual constraints on the itinerary quality through ablation analysis.

7 Proposed system

Our proposed AI-based travel-planning framework comprises the following modules, formalizing — at a research level — the design intent of the TripCraft AI project, which is itself currently at the requirement-analysis and system-design stage and has not yet been implemented or experimentally validated:

- Preference and Constraint Processing Module: normalizes raw user input (destination type, interests, style, budget, duration, dates) into a structured preference-and-constraint representation.
- Travel Data Integration Module: retrieves destination, attraction, weather, route and accommodation information from external sources (Weather, Maps, Places and Hotel APIs, generically).
- Recommendation Module: ranks candidate destinations/attractions using content-based, collaborative or hybrid, context-aware techniques [1], [2], [9].
- Weather- and Route-Analysis Module: adjusts candidate suitability by forecast conditions [8] and assesses geographic feasibility using tourist-trip-design techniques [7].
- Budget Analysis Module: filters and allocates candidate cost against the stated budget, following cost-aware latent-factor and category-wise allocation strategies [6], [10].
- Itinerary Generation Module: assembles ranked, feasible candidates into a day-wise plan and, optionally, renders it in natural language [10], [12].




Figure 3. AI-based personalized travel planning workflow.

8 System architecture

The proposed conceptual architecture organizes the framework into eight layers (Figure 4), presented explicitly as a proposed rather than implemented design. Information flows upward from the Travel Data Integration Layer and downward from the User Interaction Layer, converging in the AI/ML Recommendation and Optimization/Constraint layers, where preference, budget, weather and route signals are jointly evaluated before an itinerary reaches the Output and Personalization Layer. Where relevant, each layer is illustrated with the technology choice documented in the TripCraft AI project's own system analysis (e.g., a React.js/Tailwind front end, a FastAPI back end, MongoDB/SQL storage and a TensorFlow/Scikit-learn/PyTorch AI/ML layer), without claiming that these components have been built or tested.




Figure 4. Layered architecture of the proposed framework.

9 Methodology

9.1 Data sources and collection

The framework requires access to several categories of travel information: destination and attraction metadata, geographic coordinates, route/distance data, weather forecasts, accommodation and transportation information, estimated activity costs, and user preference/feedback data collected through the interaction layer. These correspond generically to the Weather, Maps, Places and Hotel API categories identified in the TripCraft AI project's system analysis; no specific provider responses or performance figures are reported here, as such integration has not yet been completed.

9.2 Preference and recommendation modelling

User preferences (destination type, activities, budget range, duration, travel style) are represented as a structured feature vector combining categorical and numerical attributes, consistent with standard content-based and hybrid recommender-system practice [1], [2]. Candidate destinations/attractions are ranked using a weighted combination of factors:

$$\text{Score} = w_1 \cdot \text{Preference} + w_2 \cdot \text{Budget} + w_3 \cdot \text{Weather} + w_4 \cdot \text{Route} - w_5 \cdot \text{Penalty} \quad (1)$$

where weights would be tuned empirically once evaluation data becomes available; this scoring function is a proposed methodology, not a validated result. Given the cold-start conditions typical of a new platform, a hybrid, context-aware configuration [1], [8], [9] is proposed as a reasonable starting point.

9.3 Budget, weather and route optimization

Budget constraints act both as a hard filter and a soft ranking factor, following the cost-aware latent-factor approach of Ge et al. [6] and the category-wise allocation strategy of Chen et al. [10] (accommodation, transportation, attractions, contingency reserve). Weather suitability is computed per candidate and forecast date, down-weighting outdoor attractions and up-weighting indoor alternatives under unfavourable forecasts, following Braunhofer et al. [8]. Route feasibility draws on tourist-trip-design heuristics [7] — shortest-path computation, graph-based sequencing, and meta-heuristic local search — to order visits within the available daily time budget; no specific algorithm is claimed to be implemented at this stage.

9.4 Itinerary generation algorithm

Algorithm 1 summarizes the proposed constraint-aware procedure for converting ranked candidates into a day-wise itinerary.

Algorithm 1. Constraint-aware itinerary generation

```

Input: preferences P, budget B, duration D, location L, dates T
1: collect P, B, D, T, L
2: candidates ← retrieve(L, T)
3: candidates ← remove hard-constraint violations
4: for c in candidates:
5:   score(c) ← w1.pref(c,P) + w2.budget(c,B) + w3.weather(c,T) + w4.route(c,L)
6: rank candidates by score (desc.)
7: for each day in D:
8:   select top feasible candidates within remaining budget / time
9:   sequence selection via route heuristic
10:  update remaining budget, time
11: assemble day-wise itinerary
12: (optional) render itinerary in natural language via LLM component
Output: personalized, budget- and weather-aware itinerary

```

9.5 Planned evaluation approach

9.5.1 Experimental setup

A preliminary controlled evaluation was run using the project's supplied synthetic dataset: 100 user profiles, 120 attraction records, 300 destination-date weather records, and 660 directed route records. These records are constructed placeholders, not observed travel data. Each user was evaluated with three approaches: a popularity baseline ranking attractions by popularity rating; a preference baseline ranking category matches with popularity; and a proposed greedy method that applies budget, weather, opening-hour, daily schedule, and route-transition checks while scoring candidate attractions. The evaluation considered up to three attractions per day, with the day beginning at 09:00 and ending by 17:00. Outdoor conditions were treated as suitable when the dataset's outdoor_suitability_score was at least 0.45 and precipitation probability was no more than 75%.

9.5.2 Results

Across 100 synthetic user profiles, the popularity baseline achieved 100.0% budget compliance, the preference baseline 100.0%, and the proposed method 100.0%. Mean attraction cost was INR 1677.20, INR 1713.20, and INR 1175.20, respectively. The mean number of selected attractions was

9.75, 9.75, and 7.48, respectively. Weather suitability among selected attractions was 74.8% for popularity, 81.1% for preference, and 100.0% for the proposed method.

Table 1. Results of the preliminary evaluation on 100 synthetic user profiles.

Method	Budget compliance (%)	Mean cost (INR)	Mean attractions	Weather suitability (%)
Popularity	100.0	1677.20	9.75	74.8
Preference	100.0	1713.20	9.75	81.1
Proposed	100.0	1175.20	7.48	100.0

9.5.3 Interpretation and limitations

The results indicate that the proposed rule-based procedure can enforce the budget and weather rules encoded in this implementation on the constructed cases. They do not establish better real-world recommendations: the dataset is synthetic, the baselines are simplified, no independent ground-truth itinerary or user satisfaction labels are available, and the proposed method's hard filters directly determine some of its compliance metrics. Route distance is calculated from recorded within-day transitions and is not a comparison against an optimal route. The weather score is also based on a simplified threshold, not meteorological validation. These results should therefore be described as preliminary software-level validation rather than evidence of real-world effectiveness.

9.5.4 Reproducibility

The four input CSV files and the evaluation script should be retained with the paper submission. The reported metrics were calculated over all 100 user profiles for each method. Before making stronger claims, the experiment should be rerun after checking data consistency (including whether attraction coordinates match the stated destinations), adding explicit adverse-weather test cases, and defining an independent reference itinerary or relevance labels for recommendation-quality metrics.

9.6 Deployment setup

Consistent with the TripCraft AI project's planning-phase system analysis, the framework is intended to be deployed as a modular web application: a React.js/Tailwind front end, a containerized FastAPI backend service, MongoDB/SQL storage, and a separate AI/ML service layer hosting the recommendation, optimization and (optional) language-model components, communicating over encrypted (HTTPS) channels. Role-based access and minimal data retention are intended to protect user preference and trip data. As with the rest of this paper, this deployment description is a plan rather than a completed implementation.

10 Dataset and data collection plan

For the initial evaluation of the framework, a synthetic dataset was constructed. It consists of four CSV files: users.csv (100 user profiles), attractions.csv (120 attraction records), weather.csv (300 weather records) and routes.csv (1,320 directed route records, corresponding to 660 unique attraction pairs), covering ten destinations. User profiles represent traveller preferences and constraints through the preferred destination, interests, budget in Indian rupees, trip duration, travel style and preferred activity. Each attraction record specifies its category, estimated cost, visit duration, geographic coordinates, indoor or outdoor classification, opening hours and popularity rating. Weather records give date-specific temperature, precipitation probability, general condition

and an outdoor suitability score for each destination, and route records give the distance and estimated travel time between attractions within a destination.

This structure is intended to support the planned evaluation of the proposed constraint-aware approach against the baseline methods described in the evaluation section. User interests can be compared with attraction categories to assess preference matching; user budgets with attraction costs to assess budget compliance; weather records with the indoor or outdoor classification of scheduled attractions to assess weather suitability; and route distances and travel times, together with opening hours and visit durations, to assess route feasibility. These evaluations are planned and have not yet been performed, and no experimental results are reported in this paper. The dataset is synthetic, does not represent real tourist behaviour or observed travel conditions, and is intended only for initial framework evaluation; validation on real-world data remains future work.

Before experimentation, the data will be preprocessed by handling missing or inconsistent values, normalizing numerical attributes, and converting categorical preferences into suitable machine-readable representations. The dataset will be divided into appropriate development and evaluation subsets where required. Any changes to the dataset introduced during preprocessing will be reported together with the experimental results.

11 Expected system behaviour

Once implemented, the framework is expected to accept user preferences, budget and dates; retrieve and jointly evaluate destination, weather, route and cost information; and return a day-wise itinerary that is feasible with respect to budget and route constraints and adapted to forecast weather, together with an optional natural-language rendering and a feedback channel for future personalization. This description reflects the intended behaviour of the proposed design and is not a report of measured system performance, which is deferred to the planned evaluation described above.

12 Discussion

If realized as intended, the framework could offer a single point of interaction for the traveller, itineraries checked for budget and route feasibility rather than left for manual verification, and adaptability to changing weather forecasts, with potential applications extending to business travel, group-trip coordination and tourism-board decision support. Several challenges nonetheless warrant caution: data-quality issues (incomplete metadata, inconsistent pricing, stale forecasts) can degrade recommendations regardless of model sophistication; dependence on multiple external APIs introduces operational risk; new users and destinations face a cold-start problem that hybrid recommendation mitigates but does not eliminate; dynamic conditions (closures, price or forecast changes) require re-evaluation rather than one-shot planning; the route-optimization component is combinatorially demanding at scale [7]; explainability of the multi-factor scoring function is important for user trust; and reliance on preference and interaction data makes privacy-preserving handling and clear consent a core design requirement rather than an afterthought.

13 Limitations

The framework has undergone a preliminary software-level evaluation on synthetic records, but it has not been validated with real-world travel data or users. Results should be interpreted as checks of the implemented rules rather than evidence of real-world effectiveness.

14 Conclusion

This paper reports preliminary results from a controlled experiment on synthetic records. Real-world validation, independent relevance labels and user studies remain necessary before making claims about practical effectiveness.

15 Future work

Planned next steps include implementing a working prototype aligned with the TripCraft AI project's roadmap (recommendation, budget-estimation and route-optimization components); collecting or constructing a suitably sized evaluation dataset; conducting controlled user studies of satisfaction and perceived personalization; exploring reinforcement learning for longitudinal preference adaptation; enabling dynamic itinerary adaptation to real-time weather, pricing or availability changes; incorporating multimodal signals (images, reviews); developing explainable recommendation outputs; and extending the conversational component into a full multi-turn travel assistant, consistent with the chatbot module planned for the TripCraft AI project.

References

- [1] Adomavicius G, Tuzhilin A. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans Knowl Data Eng.* 2005;17(6):734–749.
- [2] Burke R. Hybrid recommender systems: Survey and experiments. *User Model User-Adapt Interact.* 2002;12(4):331–370.
- [3] Borràs J, Moreno A, Valls A. Intelligent tourism recommender systems: A survey. *Expert Syst Appl.* 2014;41(16):7370–7389.
- [4] Chaudhari K, Thakkar A. A comprehensive survey on travel recommender systems. *Archives of Computational Methods in Engineering.* 2020;27(5):1545–1571. <https://doi.org/10.1007/s11831-019-09363-7>
- [5] Liu Q, Ge Y, Li Z, Chen E, Xiong H. Personalized travel package recommendation. In: *Proc. IEEE 11th Int. Conf. Data Mining (ICDM).* 2011. p. 407–416.
- [6] Ge Y, Liu Q, Xiong H, Tuzhilin A, Chen J. Cost-aware travel tour recommendation. In: *Proc. 17th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining.* 2011. p. 983–991.
- [7] Gavalas D, Konstantopoulos C, Mastakas K, Pantziou G. A survey on algorithmic approaches for solving tourist trip design problems. *J Heuristics.* 2014;20(3):291–328.
- [8] M. Braunhofer, M. Elahi, M. Ge, F. Ricci, and T. Schievenin, “STS: Design of weather-aware mobile recommender systems in tourism,” in *Proc. First International Workshop on Intelligent User Interfaces: Artificial Intelligence meets Human Computer Interaction (AIHCI)*, Turin, Italy, 2013, CEUR Workshop Proceedings, vol. 1125.
- [9] Yoon J, Choi C. Real-time context-aware recommendation system for tourism. *Sensors.* 2023;23(7):3679.
- [10] Chen A, Ge X, Fu Z, Xiao Y, Chen J. TravelAgent: An AI assistant for personalized travel planning. *arXiv:2409.08069.* 2024.
- [11] Xie J, Zhang K, Chen J, Zhu T, Lou R, Tian Y, et al. TravelPlanner: A benchmark for real-world planning with language agents. *arXiv:2402.01622.* 2024.
- [12] Chaudhuri S, Purkar P, Raghav R, Mallick S, Gupta M, Jana A, et al. TripCraft: A benchmark for spatio-temporally fine grained travel planning. In: *Proc. 63rd Annu. Meeting Assoc. Comput. Linguistics (ACL)*, Vienna, Austria. 2025. p. 17035–17064.