IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

SafarSaathi: An AI-Driven Intelligent Travel Assistance Platform with Verified Recommendations and Local Guardian Emergency Support

Sandhya Raju Bachewar[1], Om Deepak Korhale[2], Tanisha Londhe[3], Meenakshi Thalor[4]
Students, Dept. of Information Technology[1-3]
Professor, Dept. of Information Technology[4]
AISSMS Institute of Information Technology, Pune, India
sandhyabacchewar50@gmail.com, omkorhale3@gmail.com
tanisha.londhe@aissmsioit.org, meenakshi.thalor@aissmsioit.org

Abstract : *Solo and women travellers in India commonly rely on five or more disconnected applications for itinerary planning, budgeting, accommodation search, local discovery, and emergency assistance. This paper presents SafarSaathi, an AI-driven travel assistance platform that unifies these functions in a single, safety-first web application. The prototype uses a React and Vite frontend, a Node.js and Express backend, PostgreSQL with PostGIS for spatial queries, Socket.IO for real-time emergency communication, and Gemini function calling for AI orchestration. A deterministic verification layer cross-checks AI-generated places, stays, and prices against curated database records and labels each item as verified, partially verified, or unverified. A Local Guardian network ranks nearby guardians with a weighted multi-factor score and escalates automatically when a guardian declines or times out. The system further provides a budget-constrained trip planner, a safety-scored stay finder with a female-only dormitory filter, and a planned Reel-to-Itinerary pipeline based on keyframe extraction and multimodal vision. We report the design, architecture, and implementation status at approximately 60% completion, together with preliminary functional observations on a synthetic three-city dataset. Quantitative evaluation is planned and is not reported here.*

Keywords: travel assistance, recommender systems, generative AI, function calling, PostGIS, local guardian, emergency SOS, women's safety, multimodal AI.

I. INTRODUCTION

Travel planning for solo travellers, and for women travellers in particular, typically spans several independent applications for itinerary generation, expense tracking, accommodation search, navigation, local discovery, and emergency communication. Switching between these tools increases cognitive load during transit, and none of them is designed to treat personal safety as a primary requirement.

Recent research in tourism recommendation has produced effective techniques for individual tasks, including collaborative filtering for route and point-of-interest (POI) ranking, group constraint satisfaction, real-time recommendation pipelines, and community-driven cultural guidance [1]–[12]. However, the works reviewed for this project each optimise one recommendation task and do not combine verified accommodation, budget-aware planning, community-backed safety, emergency response, and visual travel-reel input in a single assistant.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

Large language models offer a natural conversational interface for travel planning, but when they are used without grounding in curated data they can produce unverified venues or prices [13]. In a safety-sensitive setting, such errors are not merely inconvenient. This motivates an architecture in which the language model selects controlled application tools rather than generating factual claims directly, and in which every returned entity is checked against curated records before it is presented to the user.

This paper presents SafarSaathi, a travel assistance platform built around that principle. The main contributions, as documented in the project design, are:

a hybrid AI orchestration design in which Gemini function calling dispatches six controlled backend tools;

a deterministic three-tier verification layer that labels each relevant output as `VERIFIED`, `PARTIALLY_VERIFIED`, or `UNVERIFIED`;

a Local Guardian network with PostGIS-based spatial filtering, weighted multi-factor ranking, and real-time SOS escalation over Socket.IO;

a budget-constrained trip planner and a safety-scored stay finder with female-only dormitory filtering; and

a planned Reel-to-Itinerary pipeline that converts short travel videos into verified, budgeted itineraries.

The work is at an intermediate stage. At the second project review (5 October 2026) the prototype was approximately 60% complete. Accordingly, this paper reports the design and implementation status and preliminary functional observations, and it clearly distinguishes implemented components from those still in progress. The remainder of the paper is organised as follows. Section II reviews related work, Section III states the problem, Section IV describes the methodology, Section V presents the system architecture, Section VI describes the implementation, Section VII reports results and the evaluation plan, and Section VIII concludes.

II. LITERATURE REVIEW

Twelve recent works on tourism and group recommender systems were reviewed. Table I summarises their techniques, reported outcomes, and limitations as recorded in the project review.

TABLE I: Summary of Reviewed Literature

Ref.	Technique	Reported outcome	Limitation recorded in review
[1]	MapReduce-based nMR-DBSCAN clustering of trajectory groups with TF-IDF and cosine-similarity semantic strategies	Combines clustering with semantic and distance information for large-scale personalised route recommendation	Not stated in the paper; route-focused rather than a complete assistance platform
[2]	Pearson-similarity collaborative filtering with approximate constraint satisfaction for group packages	Top-M group travel packages; validated with 210 respondents	Constraints entered manually; cold start for new destinations unresolved
[3]	Survey of group recommender techniques (collaborative, content, context, semantic, community, tagged, hybrid, demographic)	Taxonomy of techniques and evaluation methods	Survey only; no new model validated
[4]	Domain thesaurus with Naïve Bayes review mining, Apriori rules, and Fuzzy C-Means group clustering	Precision 0.875 and recall 0.625 for the review classifier	False-positive rate of 0.4; moderate reliability

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 4, November 2026 Impact Factor: 8.2

Ref.	Technique	Reported outcome	Limitation recorded in review
[5]	Slope-One collaborative filtering with cosine similarity for tourism service robots	93.50% average accuracy versus 90.79% for the comparison baseline	Only the route recommendation function is evaluated
[6]	Regression, Random Forest, and CNN modules for destinations, food, activities, and transport (TourBuddy)	Single mobile app across four domains; overall accuracy above 85%	Destination model about 79.3%; tied to the Sri Lankan context
[7]	Matrix-factorisation collaborative filtering with a Kafka real-time pipeline and microservices	85.3% accuracy, 82.7% precision, 4.7/5 satisfaction, 120 ms response, 99.8% stability	Recommendation only; no safety, budgeting, or emergency services
[8]	Collaborative filtering over crowdsourced check-ins with a heuristic travel route planning algorithm	Hit rate 20% higher than a hottest-first baseline	Reliability of noisy crowdsourced check-ins
[9]	Visual-enhanced probabilistic matrix factorisation using CNN features from geo-tagged photos	F1 improved by about 7–9% and RMSE by over 30% relative to PersTour	Requires sufficient geo-tagged photos; higher training cost
[10]	Community-collaborative cultural tourism platform with Google Maps API navigation	88% satisfaction among 67 trial users	Small sample; long-term community impact not validated
[11]	Facial-attribute detection and travel-group-type prediction with a Bayesian model over more than 10M photos	Precision at rank one improved by 4–9%; group-type gains up to 23%	About 80% facial detection accuracy; noise from strangers in photos
[12]	Fuzzy C-Means user profiling with a hybrid collaborative recommender	Unified framework for individual and group recommendation	Conceptual framework; no experimental validation reported

A. Recommendation and Planning Techniques

Collaborative filtering is the dominant technique in the reviewed work. Slope-One filtering was applied to route recommendation for tourism service robots [5], matrix factorisation was combined with a streaming pipeline to evaluate a deployed recommender [7], and crowdsourced check-ins were used to rank POIs before a heuristic route planner assembled multi-POI packages [8]. Group-oriented approaches address the aggregation of individual preferences, either through approximate constraint satisfaction [2], review mining and fuzzy clustering [4], or hybrid user profiling [12]; a survey of group recommender techniques is given in [3]. Clustering of trajectories with semantic strategies supports personalised routes [1]. These works motivate the constraint-based planning and POI ranking used in SafarSaathi.

B. Multi-Service and Community-Based Systems

TourBuddy [6] integrates destination, food, activity, and transport recommendation in one mobile application and is the closest architectural analogue to SafarSaathi's single-application design. Community collaboration for cultural tourism was validated with trial users in [10], which motivates the guardian-contributed tips in SafarSaathi. The use of visual

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

content and people attributes from photographs [9], [11] is relevant to the planned Reel-to-Itinerary feature, although those works operate on photographs rather than video.

C. Research Gap

Table II compares the reviewed literature with SafarSaathi on selected capabilities, using the assessment recorded in the project review and the implementation status at the second review. The reviewed studies primarily address specific recommendation, planning, personalisation, or optimisation problems in isolation. To the best of the project team's review, none of them interposes a deterministic verification layer between model output and the user, integrates real-time emergency dispatch with spatial guardian matching, or accepts video input. SafarSaathi proposes the integration of planning, budgeting, accommodation, safety, local guardian support, local discovery, and reel-based itinerary generation within one assistant.

TABLE II: Reviewed Literature versus SafarSaathi

Capability	Reviewed work	SafarSaathi status
Route / POI planning	Present	Implemented
Budget planning	Limited	Implemented
Accommodation recommendation	Study-specific	Implemented
Safety, guardian, and SOS	Not focused	Implemented (KYC mock)
Function-calling AI orchestration	Not central	Implemented
Retrieval-augmented guidance	Not central	In progress
Visual (reel) input	Partial (photos)	In progress
Translation assistance	Not central	Future scope
Collaborative filtering	Present	Future

III. PROBLEM STATEMENT

Solo and women travellers in India depend on multiple disconnected applications for planning, accommodation, safety, and discovery, and no single verified, safety-first assistant is available to them [13]. The project identifies four practical problems:

1) *Fragmented travel stack.* Separate planners, budget tools, booking sites, maps, and safety applications force frequent context switching.

2) *Disconnected safety support.* Existing safety applications focus on matching with people and are not integrated with trip duration, route context, or budget-constrained accommodation [13].

3) *Unverified AI output.* Generative itinerary tools may present venues, stays, or prices that cannot be traced to a trusted record, exposing travellers to unverified stays.

4) *Local and field constraints.* Travellers lack trusted hyper-local knowledge and face practical barriers such as unfamiliar scripts on menus and signboards.

The core engineering problem is therefore to design a single AI-assisted platform that produces realistic, budget-constrained travel plans, grounds safety-sensitive outputs in curated spatial data, and enables rapid local emergency response.

A. Objectives

The project defines six objectives: (i) a conversational travel companion whose answers are grounded in verified records and labelled with trust badges; (ii) day-wise trip plans with category-wise budget splits within a fixed budget; (iii) safety-scored stay recommendation using spatial filtering, including female-only dormitories; (iv) a Local Guardian network with one-tap real-time SOS, multi-factor ranking, and automatic escalation; (v) proactive, city- and date-aware recommendations grounded in guardian tips through retrieval-augmented generation (RAG); and (vi) conversion of short travel reels into verified, budgeted itineraries.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

B. Scope and Constraints

The prototype covers three Indian cities (Pune, Jaipur, and Goa) and a synthetic seed dataset. Target users are solo travellers, women travellers, first-time backpackers, and interstate visitors. Real identity (KYC) verification, payment gateway integration, international destinations, and full multilingual translation of menus and signboards are outside the current phase. Authentication and KYC are mock implementations for demonstration.

IV. PROPOSED METHODOLOGY

A. Development Methodology

The project follows an iterative, phase-wise software engineering process comprising requirement analysis, literature survey, system design, implementation in seven phases (P0–P6), multi-level testing, and formal reviews. This allows the core spatial, AI orchestration, and SOS platform to be built and verified before the more experimental RAG and multimodal components.

B. System Workflow

A user session proceeds as follows. (1) The user selects a traveller persona and an active city. (2) The user submits a natural-language or voice query, a trip and budget form, a stay filter, a reel upload, or a one-tap SOS trigger. (3) The React client sends the request to an Express REST endpoint or to the Socket.IO event server. (4) The service layer executes PostGIS queries and, where applicable, invokes a Gemini tool call. (5) The verification layer checks returned entities against curated records and attaches trust badges. (6) Structured, badge-labelled cards are returned to the interface, while SOS events are synchronised live between traveller and guardian views.

C. AI Orchestration via Function Calling

The AI assistant does not answer factual queries from free-form generation. Instead, it proceeds in five steps: (1) the query, active city, and traveller persona are received, as text or as voice captured with the Web Speech API; (2) the Gemini model [14] selects one of six registered tools through structured function calling; (3) the selected tool is executed against the Express service layer and the PostgreSQL/PostGIS database; (4) the returned entity names, locations, and prices are verified; and (5) a structured, badge-labelled response is returned to the client. The six tools are `plan_trip`, `calculate_budget`, `find_stays`, `get_local_recommendations`, `trigger_sos`, and `extract_itinerary_from_reel`. When no external API key is configured, a deterministic mock engine (`mockGemini`) returns repeatable responses so that every flow remains demonstrable offline.

D. Verification Layer

Each database-backed item in an AI response is compared with the curated records and assigned one of three trust badges, as defined in Table III.

TABLE III: Verification Badges and Assignment Rules

Badge	Rule
`VERIFIED`	Exact match on name, location, and price in the curated database
`PARTIALLY_VERIFIED`	Entity found in the database, but price or other details differ
`UNVERIFIED`	No match in the curated database

E. Guardian Matching

When an SOS is triggered, available guardians within a configurable radius are retrieved with the PostGIS function `ST_DWithin`, and the geodesic distance to each is measured with `ST_DistanceSphere` [15]. After the distance is normalised, each candidate g is scored as

$$S(g) = 0.45\,(1 - d_{\text{norm}}) + 0.25\,r + 0.15\,\rho + 0.15\,a, \qquad (1)$$

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/IJARSCT-40954

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

where d_{norm} is the normalised distance, r is the guardian rating, ρ is the response rate, and a is the attribute match between traveller and guardian. Candidates are ranked by $S(g)$ and the highest-ranked guardian is alerted over Socket.IO. If the guardian declines or the alert times out, the next-ranked candidate is alerted. The weights are fixed in the current design, and the sensitivity of the ranking to them is planned as part of the evaluation.

F. Trip and Budget Planning

Given a city, trip duration, total budget B, and traveller persona, the planner generates a day-wise itinerary and partitions the budget across four categories. The allocation must satisfy

$$B_{\text{stay}} + B_{\text{food}} + B_{\text{transport}} + B_{\text{activities}} \leq B. \qquad (2)$$

Selections are adjusted if the total exceeds the ceiling, and each itinerary item and budget line item carries a verification badge. The persona influences stay selection, for example through a minimum safety score and the female-only dormitory filter.

G. Reel-to-Itinerary Extraction

The planned multimodal pipeline accepts MP4 or WebM videos of up to 50 MB, extracts keyframes with FFmpeg, identifies landmarks and food spots with Gemini Vision [14], matches recognised entities against the PostGIS places table, and produces a budgeted itinerary in which each item passes through the verification layer. At the second review, the upload interface and demonstration reel selector were complete, whereas keyframe extraction and vision analysis were in progress.

H. Retrieval-Augmented Local Recommendations

Proactive recommendations are planned to combine city and date information with vector search over 50 guardian-contributed tips stored with embeddings, so that suggestions such as festivals, timing, and crowd guidance are grounded in community knowledge. This component was in progress at the second review.

I. Dataset

To obtain repeatable offline testing without external data dependencies, the prototype uses a synthetic seed dataset stored as JSON and CSV and loaded into the relational schema by a reproducible seed script. It covers three cities (Pune, Jaipur, and Goa) and contains 15 guardians, 40 accommodations, 60 places and attractions, 20 festivals and events, and 50 guardian tips. Recorded fields include city, category, coordinates, price, rating, safety score, popularity, and preferences.

V. SYSTEM ARCHITECTURE

SafarSaathi is organised in five layers, shown in Fig. 1, with the implementation status of each module at the second review indicated by colour.

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/IJARSCT-40954

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

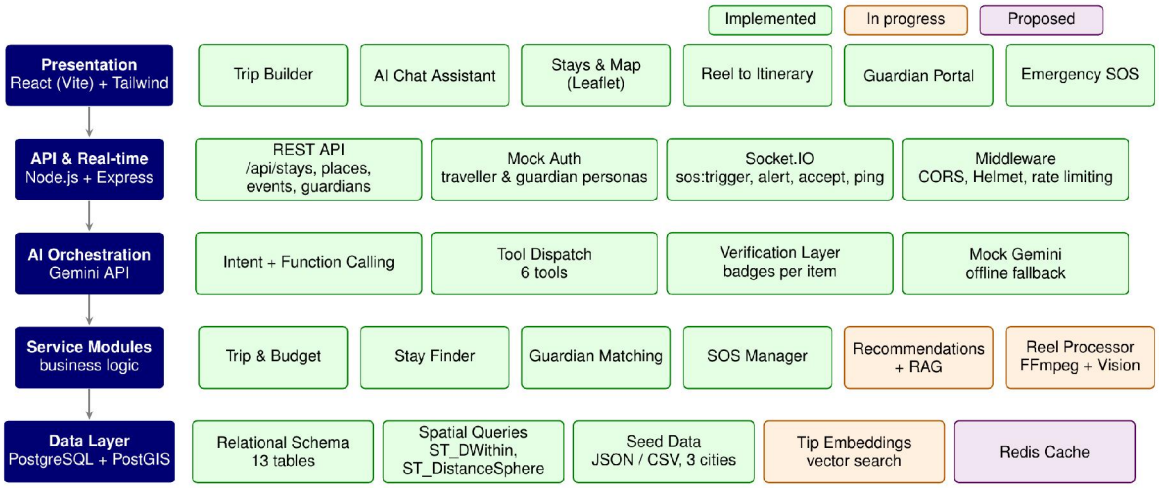


Fig. 1. Five-layer architecture of SafarSaathi, showing the implementation status of each module at the second project review.

A. Layers

1) *Presentation Layer:* The React (Vite) and Tailwind CSS client provides a Trip Builder, an AI Chat Assistant with voice input, a Stays and Map view built on Leaflet, a Reel-to-Itinerary uploader, a Guardian Portal, and a persistent Emergency SOS control [16].

2) *API and Real-Time Layer:* Node.js and Express expose JSON REST endpoints (`/api/stays`, `/api/places`, `/api/events`, `/api/guardians`, `/api/trips`, and `/api/chat`), a persona-based mock authentication service for traveller and guardian roles, and Socket.IO channels for the SOS events `sos:trigger`, `sos:guardian_alert`, `sos:accept`, `sos:status_update`, and `sos:location_ping`. Helmet headers, CORS policies, and rate limiting are applied as middleware.

3) *AI Orchestration Layer:* This layer performs intent detection and function calling over six tools, applies the verification layer, and falls back to the deterministic mock engine when no API key is available.

4) *Service Modules Layer:* Domain logic is encapsulated in six services: Trip and Budget, Stay Finder, Guardian Matching, SOS Manager, Recommendations with RAG, and Reel Processor.

5) *Data Layer:* PostgreSQL with PostGIS stores the 13-table schema. Location-aware entities use `geography(Point, 4326)` columns so that radius filtering and geodesic distance can be computed in the database. Guardian tips are stored with vector embeddings for planned semantic retrieval, and a Redis cache is proposed.

B. Module Design

Table IV lists the inputs, outputs, and status of the principal modules.

TABLE IV: Module Design and Status at Second Review

Module	Input	Output	Status
AI chat	Text or voice query, city, persona	Verified answer cards	Implemented
Trip and budget	City, days, budget, persona	Day-wise plan, budget split	Implemented
Stay finder	Price, safety score, female-only filter, location	Filtered stays on list and map	UI done; live map tiles in progress

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

Module	Input	Output	Status
Guardian and SOS	Traveller location, guardian availability	Ranked guardian, SOS state	Core done; KYC in progress
Recommendations and RAG	City, date, guardian tips	Proactive suggestions	In progress
Reel to itinerary	MP4/WebM reel	Verified itinerary	Upload UI done; pipeline in progress

C. Database Design

The relational schema contains 13 tables: `users`, `chat_sessions`, `chat_messages`, `trips`, `itinerary_items`, `budgets`, `places`, `accommodations`, `events`, `guardians`, `kyc_records`, `sos_requests`, and `guardian_tips`. A user owns trips and chat sessions; a trip owns itinerary items and budget lines; itinerary items reference curated places; and a guardian is linked to KYC records, SOS requests, and tips. An SOS request links the traveller and the selected guardian. KYC records are held in a dedicated table.

VI. IMPLEMENTATION

A. Technology Stack

Table V lists the technologies used. Relative to the first review, Google Maps was replaced by Leaflet with OpenStreetMap and CARTO tiles (Google Maps remains an optional toggle), and Redis caching was moved from the implemented to the proposed category.

TABLE V: Technology Stack

Layer	Technologies	Status
Frontend	React (Vite), Tailwind CSS, Lucide React, Leaflet, Web Speech API	Implemented; live map tiles in progress
Backend	Node.js, Express, Socket.IO, Helmet, CORS, rate limiting	Implemented
Database	PostgreSQL, PostGIS, JSON/CSV seed scripts	Implemented
AI	Gemini function calling, mock Gemini engine	Implemented
AI	Gemini Vision, Gemini Embeddings	In progress
Media	FFmpeg keyframe extraction	In progress
Caching	Redis	Proposed
Tooling	Docker Compose, Git/GitHub, Postman, VS Code	Implemented; Postman suite planned

B. Phase-Wise Implementation

Development is organised into phases P0–P6. At the second review, P0 (scaffolding, 13-table schema, seed data, and Docker Compose configuration for PostGIS), P1 (Express REST APIs, security middleware, mock authentication, and spatial endpoints), and P2 (six function tools, mock fallback, verification layer, chat interface with quick-action chips and voice input) were complete. P3 (guardian matching, Socket.IO SOS channels, and traveller and guardian views) and P4 (day-wise itinerary, budget split, stay filters, and map and list views) had their core functionality implemented, with mock KYC and live map tiles in progress. P5 (proactive recommendations and RAG) and P6 (reel processing, documentation) were in progress.

C. User Interface

The interface uses a dark, high-contrast palette (deep slate, safety emerald, warm ochre, and SOS red). The SOS button is persistent on every screen, and the active city, guardian availability, and traveller persona remain visible in the

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

header. Verification badges are displayed on AI output cards, tap targets are at least 44 px, and the layout is responsive for desktop and mobile viewports.

D. SOS Workflow

The SOS lifecycle is modelled as the state chart in Fig. 2. Triggering an SOS starts a 5-second cancellation window to prevent accidental activation. The system then retrieves and ranks nearby guardians using Eq. (1), alerts the top-ranked guardian, and on decline or timeout alerts the next-ranked guardian. After acceptance, status updates and location pings are exchanged until the request is resolved. Traveller location is transmitted only during an active, user-initiated SOS session or an explicit spatial query.

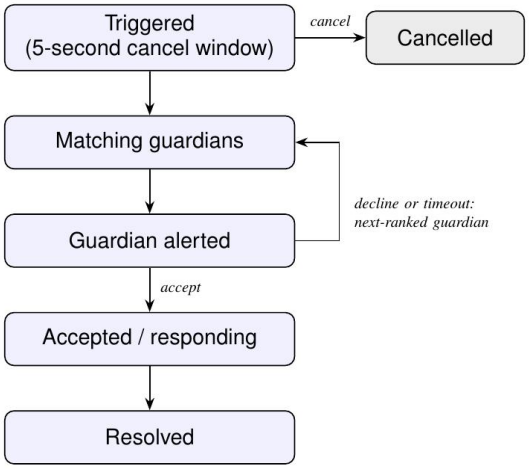


Fig. 2. State chart of an SOS request.

VII. RESULTS AND EVALUATION

A. Status of Evidence

At the second review the available evidence consists of a working prototype and its software architecture. Automated unit, API, integration, and usability tests had not yet been executed, and no quantitative performance measurements have been collected. This paper therefore does not report accuracy, latency, or user-study figures for SafarSaathi. The values quoted in Section II are results reported by the cited authors and must not be read as results of this system.

B. Preliminary Functional Observations

The following qualitative observations were recorded in the demonstration environment, which uses the synthetic dataset:

all five main modules and the global SOS button are reachable and navigable;
city and persona context is preserved across modules;
the mock AI engine executes every demonstration flow without external API keys; and
the reel upload interface accepts MP4 and WebM files of up to 50 MB, tested with two demonstration reels.

The standard demonstration scenario uses a solo female traveller in Pune who requests a verified 3-day plan within INR 4,500, filters stays by nightly price of at most INR 2,500, safety score of at least 4.5, and female-only dormitories, uploads a demonstration reel, and simulates an SOS alert between traveller and guardian views.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

C. Test Plan

Testing is planned at five levels: unit tests for the scoring formula, badge logic, and tool dispatch; API tests of the spatial endpoints with Postman collections; integration tests of the Socket.IO SOS sequence; UI walkthroughs on desktop and mobile; and usability trials with solo-traveller task scenarios. Table VI lists the ten planned test cases. None of them had been executed at the time of writing.

TABLE VI: Planned Test Cases (Not Yet Executed)

ID	Condition	Expected result
TC-01	Pune, 3 days, INR 4,500, solo-female persona	Plan returned; budget split within INR 4,500; badges present
TC-02	Price ≤ INR 2,500, safety ≥ 4.5, female-only	Only matching stays shown
TC-03	Traveller coordinates and search radius	Available guardians returned in distance order
TC-04	Known distance, rating, response, attribute values	Ranking follows Eq. (1)
TC-05	Top guardian declines or times out	Next-ranked guardian receives the alert
TC-06	Entity with exact name, location, price	`VERIFIED` badge
TC-07	Entity found but price differs	`PARTIALLY_VERIFIED` badge
TC-08	Entity absent from curated database	`UNVERIFIED` badge
TC-09	API key unavailable	Mock response completes the flow
TC-10	MP4/WebM file of at most 50 MB	File accepted; extraction invoked once implemented

D. Evaluation Metrics

Table VII defines the metrics planned for the final evaluation and the targets set in the project synopsis. These are targets, not measured results. The 120 ms latency reported in [7] is cited only as a literature reference point.

TABLE VII: Planned Evaluation Metrics and Synopsis Targets

Metric	Definition	Target (synopsis)
Verified-item share	Share of AI output items with the `VERIFIED` badge	100% for curated entities; non-database items marked `UNVERIFIED`
AI response latency	Query submission to structured response	Under 1,500 ms live; under 150 ms with mock engine
SOS alert-to-accept time	SOS trigger to guardian acceptance	Event propagation under 500 ms; 5 s timeout escalation
Spatial ranking correctness	Adherence to radius filter and ranking formula	100% on TC-03 and TC-04
Budget and filter compliance	Plans within budget; filters respected	Total within INR 4,500; exact filter match
Task success and usability	Completed task scenarios; user rating	At least 90% completion; at least 4.5/5

E. Limitations

The principal limitations of the current prototype are: the dataset is synthetic and limited to three cities; authentication and KYC are mock implementations; RAG grounding and reel processing are incomplete; live map tiles require a provider key; and no quantitative evaluation or user study has yet been performed. Verification is also limited to the coverage of the curated database, so that valid venues absent from it are labelled `UNVERIFIED`. The scaling of the rating and response-rate terms in Eq. (1) and the normalisation of distance are implementation details that remain to be documented.

Copyright to IJARSCT
www.ijarsct.co.in
DOI: 10.48175/IJARSCT-40954

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 4, November 2026
Impact Factor: 8.2

VIII. CONCLUSION AND FUTURE WORK

This paper presented SafarSaathi, an AI-driven travel assistance platform that combines budget-constrained planning, safety-scored stay discovery, community-backed local guidance, and real-time guardian SOS support in a single application. Its central design principle is that AI assistance in a safety-sensitive domain should be grounded in curated data and accompanied by explicit verification status, which it realises through function calling over controlled tools and a three-tier verification layer. At the second review the project had progressed from user interface and database design to a working React, Express, and PostGIS prototype with Gemini function calling, verification, and a real-time guardian SOS flow. RAG grounding, reel-to-itinerary extraction, live map tiles, testing, and documentation remain to be completed, and completion of the prototype is planned for December 2026.

Future work includes executing the test plan and reporting the metrics in Table VII, followed by real KYC and background verification for guardians, translation and OCR for menus and signboards, a multilingual interface, payment and booking integration, additional Indian and international destinations, Redis caching, and an installable progressive web application.

REFERENCES

1. Yin, Sun, and Zhao, “Personalized tourism route recommendation system based on dynamic clustering of user groups,” in *Proc. IPEC*, 2021.
2. J. K. Kim, W. C. Kwon, I. Y. Choi, H. Heo, and H. S. Moon, “A group travel recommender system based on group approximate constraint satisfaction,” *IEEE Access*, vol. 12, 2024.
3. Katarya, “A systematic review of group recommender system techniques,” in *Proc. ICISS*, 2017.
4. Valliyammai *et al.*, “An intelligent personalized recommendation for travel group planning based on reviews,” in *Proc. ICoAC*, 2016.
5. G. Wang, T. Dong, J. Chen, Y. Chu, J. Lu, and D. Long, “Collaborative filtering algorithm in the design of intelligent tourism service robots,” in *Proc. ICDSIS*, 2024, doi: 10.1109/ICDSIS61070.2024.10594567.
6. N. A. L. Nethmin *et al.*, “Enhancing tourism experience in Sri Lanka through data driven recommendations,” in *Proc. ICSES*, 2023, doi: 10.1109/ICSES60034.2023.10465348.
7. Z. Zhang, “Evaluation of the effect of tourism recommendation system based on AI,” in *Proc. AICIT*, 2024, doi: 10.1109/AICIT62434.2024.10730422.
8. Yu, Xu, Yang, and Guo, “Personalized travel package with multi-POI recommendation based on crowdsourced user footprints,” *IEEE Trans. Human-Mach. Syst.*, 2016.
9. Zhao *et al.*, “Photo2Trip: Exploiting visual contents in geo-tagged photos for personalized tour recommendation,” *IEEE Trans. Knowl. Data Eng.*, 2021.
10. Thananchana and K. Noinan, “The designing of cultural-based tourism recommendation system with community collaboration,” in *Proc. ECTI DAMT & NCON*, 2022, doi: 10.1109/ECTIDAMTNCON53731.2022.9720393.
11. Chen, Cheng, and Hsu, “Travel recommendation by mining people attributes and travel group types from community-contributed photos,” *IEEE Trans. Multimedia*, 2013.
12. Anjali, Sandhu, and Goyal, “User profiling in travel recommender system using hybridization and collaborative method,” in *Proc. ICCCIS*, 2021.
13. “SafarSaathi: An AI-driven intelligent travel assistance platform,” Project Synopsis, AISSMS Institute of Information Technology, Pune, India, Academic Year 2026–27, unpublished.
14. Google, *Gemini API Documentation: Function Calling and Vision*, 2024.
15. PostGIS Project, *PostGIS Documentation: ST_DWithin and ST_DistanceSphere*, 2024.
16. Official documentation of React, Vite, Tailwind CSS, Express.js, Socket.IO, and Leaflet, 2024.