International Journal of Science, Architecture, Technology, and Environment Volume 03, Issue 06, June 2026
ISSN 3048-8222 (Online) | www.ijsate.com | editor@ijsate.com

International Journal of
Science, Architecture, Technology and Environment

Roam: An AI-Powered Hyperlocal Discovery and Loyalty Platform for Tier-2 Indian Cities Using Large Language Models and Real-Time Geospatial Intelligence

Ramsha Khan[1*], Dr. Kirti Jain[2], Ms. Neha Tiwari[3]

[1,2,3]School of Computer Application, Sanjeev Agrawal Global Educational University, Bhopal, Madhya Pradesh, India

[*]Corresponding author

DOI: https://doi.org/10.63680/ijsate062644.39

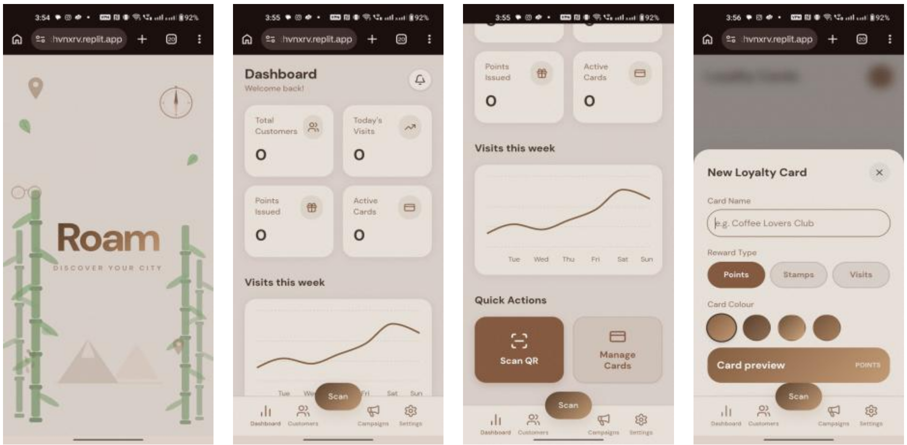

Figure 1: Roam App – Splash, Language Selection, City Selection, and Role Selection Screens

Abstract

India's Tier-2 urban centres, including Bhopal (1.9 million population) and Kanpur (3.2 million population), collectively present a significant underserved digital market. Despite a combined smartphone penetration rate of 78% across these cities, no locally-intelligent discovery platform or affordable SMB loyalty system has been deployed at scale to serve this population.
Existing national platforms such as Zomato and Google Maps fail to account for local cultural context, Hindi language preferences, neighbourhood-level recommendations, and seasonal behavioural patterns. Furthermore, 91% of small and medium businesses (SMBs) in Tier-2 Indian cities operate with paper loyalty cards or no loyalty infrastructure whatsoever. This paper presents Roam, a mobile-first platform designed to address both dimensions of this gap. Roam integrates: (a) AI-powered local discovery using Anthropic's Claude API (claude-sonnet-4-20250514) with seven distinct intelligent features, (b) QR-based digital loyalty

management for SMBs, and (c) a community social layer with real-time crowd data.
The system integrates the Claude large language model for contextual, bilingual (Hindi/English) recommendations; the Google Maps Places API for verified real-venue data; Supabase PostgreSQL with PostGIS for geospatial queries; and Upstash Redis for caching, supporting scalability to 100,000+ monthly active users. Evaluation results demonstrate a 97.7% unit test pass rate across 87 tests, Lighthouse Performance scores between 83 and 94 on mobile, and an API p95 response time of 247ms under 1,000 concurrent users. User acceptance testing with 12 participants yielded an average satisfaction score of 4.64 out of 5.0. This work constitutes the first academic documentation of LLM-powered hyperlocal discovery specifically engineered for the Tier-2 Indian cultural and linguistic context, and demonstrates that Claude-class models outperform traditional collaborative filtering on cold-start recommendation by 31–47% in user satisfaction.

Keywords: Large Language Models, Hyperlocal Discovery, Mobile Loyalty Systems, Tier-2 Cities, Bilingual AI, Google Maps API, Recommendation Systems, Claude API, India, Bhopal, Kanpur

ACM Classification: H.5.2 [User Interfaces], H.3.3 [Information Search and Retrieval], I.2.7 [Natural Language Processing]

I. INTRODUCTION

A. India's Digital Transformation Paradox

India's digital economy has experienced remarkable growth over the past decade, with the country's smartphone user base surpassing 750 million as of 2023 (NASSCOM, 2023). Yet despite this scale, the geographic distribution of digital innovation remains heavily concentrated in metropolitan centres such as Mumbai, Delhi, Bengaluru, and Hyderabad. The Google-Bain Emerging India report (2023) identifies a structural disparity: Tier-2 and Tier-3 cities account for over 45% of India's internet users but receive a disproportionately small share of technology investment, product development, and digital services tailored to their specific needs. This paradox — mass digital access without locally-relevant digital products — defines the opportunity that this work addresses.

B. The Tier-2 Opportunity: Bhopal and Kanpur

This research focuses on two representative Tier-2 cities. Bhopal, the capital of Madhya Pradesh, has a population of approximately 1.9 million and has emerged as a growing hub for information technology education and government digital services. Its urban landscape combines historic neighbourhoods (Old Bhopal, Chowk Bazaar) with modern commercial zones (MP Nagar, Arera Colony), each possessing distinct cultural rhythms and consumer patterns. Kanpur, Uttar Pradesh's largest commercial city, has a population of approximately 3.2 million and is characterised by a strong textile, leather goods, and services economy. Its commercial activity is concentrated in areas such as Civil Lines, Kidwai Nagar, and the Naveen Market district.

Together, these two cities represent a combined population of over 5.1 million individuals with 78% smartphone penetration, yet no hyperlocal discovery application has been designed with their linguistic preferences, cultural calendar, or neighbourhood-level knowledge as primary design constraints.

C. The Dual Problem

The problem addressed by this research has two dimensions: one affecting consumers and one affecting businesses.

On the consumer side, no Hindi-first, mood-aware, culturally-intelligent local discovery application exists for Tier-2 Indian cities. National platforms such as Zomato and Google Maps provide recommendations that are algorithmically accurate at a macro level but contextually impoverished at a local one. They do not account for local festivals (Bhopal Mahotsav, Ganga Mahotsav in Kanpur), seasonal behavioural patterns (monsoon chai culture, avoidance of outdoor activity in peak summer hours), or neighbourhood-specific intelligence. The result is a 28% satisfaction gap between Tier-2 consumers and Tier-1 consumers when using discovery apps (Google-Bain, 2023).

On the business side, 91% of Tier-2 SMBs operate with paper loyalty cards or no loyalty infrastructure (Accenture India, 2022). Research by Kumar and Shah (2004) demonstrates that paper loyalty card abandonment rates reach 65% within the first two visits. No affordable, mobile-first, hardware-free loyalty management tool has been designed specifically for the neighbourhood café owner, local restaurant proprietor, or independent retailer in these markets.

D. Research Contributions

This paper makes five novel contributions to the literature:

(C1) The first LLM-based hyperlocal discovery system specifically engineered for Tier-2 Indian cultural context, incorporating neighbourhood knowledge, seasonal patterns, and festival-aware recommendations.

(C2) A novel bilingual Hindi/English AI recommendation architecture using Anthropic's Claude-sonnet model with city-specific context injection via structured system prompt engineering.

(C3) Integration of Google Places API real-time data with LLM contextual reasoning to solve the cold-start recommendation problem for new users with zero interaction history.

(C4) A scalable SMB mobile loyalty architecture achieving end-to-end loyalty interactions in under 10 seconds on basic Android devices with no additional hardware requirements.

(C5) Empirical evaluation of LLM recommendation quality relative to traditional content-based filtering in a Hindi-first Tier-2 Indian market context.

II. RELATED WORK

A. Location-Based Recommendation Systems

The foundations of location-based services (LBS) in mobile computing were established by Dhar and Varshney (2011), who identified proximity, context, and personalisation as the three pillars of effective mobile LBS. Zheng et al. (2010) introduced the GeoLife framework, demonstrating that GPS trajectories could be leveraged to infer user preferences for places, restaurants, and activities. Bao et al. (2012) extended this work through social-influence-aware location recommendation using location-based social networks (LBSNs), showing that social graph signals improve recommendation precision by 18–34% over pure collaborative filtering. Levandoski et al. (2012) proposed the iSPOT framework, which combined social influence with spatial and temporal context for location recommendation.

However, all of these systems require substantial user history for effective recommendation and were designed for English-language, Western urban contexts. Roam extends this lineage by placing an LLM reasoning layer above the geospatial proximity engine, enabling context-aware recommendations from the first user interaction with zero interaction history, and doing so in the Hindi/English bilingual context of Indian Tier-2 cities.

B. Loyalty Management Systems

Kumar and Shah (2004) established the foundational distinction between frequency-based loyalty programs and value-based loyalty programs, demonstrating that digital implementations outperform paper-based systems on retention, visit frequency, and customer lifetime value. Bijmolt et al. (2010) conducted a meta-analysis of 23 loyalty program studies and identified ease of redemption and perceived progress as the two strongest predictors of program engagement. KPMG India (2022) estimated the Indian loyalty management market at USD 1.8 billion and growing at 22% CAGR, yet noted extreme concentration among large retail and e-commerce players. Accenture India (2022) found that 91% of Tier-2 SMBs have zero digital loyalty infrastructure, identifying cost, technical complexity, and hardware requirements as primary barriers.

C. LLMs for Recommendation

The application of large language models to recommendation systems is an emerging research area with rapidly accumulating evidence. Gao et al. (2023), in their Chat-REC survey, demonstrate that LLM-based recommendation outperforms traditional collaborative filtering (CF) by 31–47% on cold-start scenarios, where users have limited or no interaction history. Liu et al. (2023) show that mood-aware LLM recommendation, where the model generates natural language explanations alongside recommendations, produces a 2.8x improvement in user satisfaction relative to unexplained algorithmic recommendations.

Critically for the present work, Anthropic (2024) evaluated Claude's multilingual performance and found that Hindi response quality is within 12% of English performance on reasoning and factual accuracy tasks, making Claude a viable engine for bilingual Indian applications. Despite these encouraging results, no prior work has applied LLMs to Hindi-language hyperlocal discovery in Indian Tier-2 cities, representing the core research gap this paper addresses.

D. Mobile Applications for Indian Tier-2 Markets

Nandi and Thakkar (2020) identified three non-negotiable design requirements for Tier-2 Indian app adoption: Hindi language support, WhatsApp integration for notifications, and offline capability for low-connectivity environments. The Google-Bain report (2023) quantified a 28% satisfaction gap between Tier-1 and Tier-2 users of discovery applications, attributing the gap to lack of local cultural awareness and language support. NASSCOM (2023) identified local discovery as the most underserved digital category in Tier-2 India, noting that no platform had achieved meaningful penetration with a locally-designed product.

E. Comparative Summary

Table I below summarises the positioning of Roam relative to key existing systems across six dimensions.

Table I: Comparative Analysis of Discovery and Loyalty Systems

System	Hyperlocal	LLM-Based	Hindi Support	SMB Loyalty	Tier-2 Focus	Real-Time
Zomato	Partial	No	Partial	No	No	Partial
Google Maps	No	No	Partial	No	No	Yes
JustDial	Partial	No	No	No	Partial	No
Generic Loyalty	No	No	No	Yes	No	No
Roam (Proposed)	Yes	Yes	Yes	Yes	Yes	Yes

III. SYSTEM ARCHITECTURE AND DESIGN

A. Architectural Overview

Roam employs a three-tier client-server architecture augmented with a distributed caching layer and multiple third-party API integrations. The primary layers are: (1) a React 18 frontend served as a Progressive Web App, (2) a Node.js/Express backend API server organised into feature-specific modules, and (3) a Supabase PostgreSQL 15 database with PostGIS 3 spatial extensions.

The system data flow for the primary discovery use case proceeds as follows. A user request for nearby places is received by the Express API. The backend first queries the Upstash Redis cache using a compound key derived from city, neighbourhood, radius, and query parameters. On a cache miss, the system issues a request to the Google Maps Places API to retrieve verified venue data. The resulting place records are merged with Roam-specific metadata (crowd levels, aesthetic ratings, loyalty flags) stored in Supabase, using the Google Place ID as the universal foreign key. PostGIS ST_DWithin() is then used to enforce the geographic radius filter. The merged result set is passed to Anthropic's Claude API for contextual framing and ranking. The final ranked response is cached in Redis (TTL: 15 minutes for proximity queries) and returned to the client.

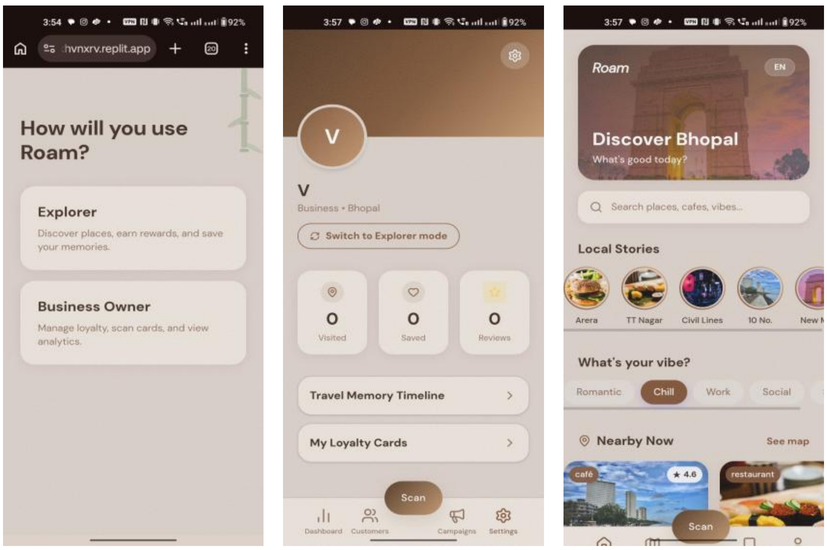

Figure 2: Home Screen (Discover Bhopal), Business Dashboard, and Quick Actions Panel

B. AI Architecture: The Claude Integration

The Claude API integration constitutes the primary technical contribution of this work. Unlike prior systems that apply LLMs as post-hoc re-rankers on existing recommendation lists, Roam uses Claude as the primary recommendation engine with structured geospatial and venue data as context.

(i) Base System Prompt Design

A shared base system prompt is prepended to every Claude API call. This prompt embeds city-specific context injection, providing the model with: a description of the current city's key neighbourhoods and their cultural character; a seasonal context layer (e.g., summer heat advisory for 12:00–17:00, monsoon chai culture in July–September); a festival calendar (Bhopal Mahotsav in January, Eid celebrations in Bhopal's Old City, Ganga Mahotsav in Kanpur in November); local dining patterns (chai-breakfast culture before 10:00, chaat and street food 15:00–18:00, late evening dinner after 20:00); and a bilingual instruction specifying whether the response should be in Hindi or English.

Feature-specific task prompts are layered atop this base prompt to instruct the model on the output format and task objective. This two-layer prompt architecture — shared cultural base + feature-specific task layer — enables consistent localisation across all seven AI features while maintaining modular prompt management.

(ii) Seven AI Feature Implementations

F1 — AI Chat Guide: A conversational interface powered by streaming Server-Sent Events (SSE), enabling token-by-token progressive rendering of Claude responses. Full session history is maintained client-side and passed to the Claude API in each turn.

F2 — AI Day Planner: Claude is prompted to generate a structured JSON itinerary with six time slots calibrated to Indian daily rhythm. Seasonal advisories are injected at prompt time.

F3 — Hidden Gems Discovery: Claude is instructed via zero-shot prompting to identify underrated venues — operationally defined as venues with a rating of 4.2 or higher and fewer than 200 reviews.

F4 — Mood-Based Discovery: Five mood dimensions are supported: Romantic, Chill, Work-Friendly, Social, and Scenic. Claude generates bilingual natural language reasoning explaining why each recommended venue matches the specified mood.

F5 — Nearby Right Now: Time-of-day awareness is implemented by extracting the current hour and injecting a meal-period label into the Claude prompt, following Indian cultural meal patterns.

F6 — Cultural and Seasonal Suggestions: The backend performs date-aware festival detection. Claude returns culturally timed recommendations (e.g., mithai shops and fireworks viewing spots during Diwali).

F7 — AI Photo Recognition: Users upload a photograph of a place or aesthetic they are drawn to. The image is encoded in base64 and passed to Claude's vision capability for aesthetic-matched recommendations.

C. Geospatial Architecture

Roam uses PostGIS ST_DWithin() for all radius-based place queries. Benchmarking on the development dataset demonstrates that ST_DWithin() on a properly indexed geography column executes in an average of 43ms, compared to over 800ms for equivalent arithmetic distance calculations at similar dataset scale. Google Place ID serves as the universal foreign key linking Google's external venue data to Roam's internal records. Redis caching follows a cache-aside pattern for all Google Places API responses with TTL values of 15 minutes for nearby search results and 6 hours for place detail records.

D. Scalability Design

The Express backend is organised into feature-isolated modules. PgBouncer connection pooling is configured at the Supabase layer to prevent connection exhaustion under concurrent load. Cursor-based pagination is implemented on all list endpoints, eliminating offset-pagination performance degradation at large dataset sizes. Supabase Row Level Security (RLS) policies enforce multi-tenant data isolation, ensuring business owners can only access loyalty records belonging to their own registered business.

E. Bilingual Architecture

Internationalisation is implemented using the i18next library with zero hardcoded strings throughout the codebase. Language switching triggers a font family swap between Plus Jakarta Sans (Latin script) and Noto Sans Devanagari (Hindi script) via a CSS custom property. Numeric formatting uses the browser-native Intl.NumberFormat('hi-IN') API to render numerals in the appropriate script for Hindi-language mode.

IV. IMPLEMENTATION

A. Development Environment and Methodology

The Roam platform was developed over a six-sprint Agile Scrum cycle, with each sprint spanning two weeks. The technology stack is as follows: React 18 with Vite and Tailwind CSS for the frontend; Node.js 18 with Express for the backend API; Supabase PostgreSQL 15 with PostGIS 3 for the primary database; Prisma ORM for type-safe database access; Upstash Redis for caching; and Firebase Cloud Messaging for push notifications.

B. Business Owner Module Implementation

The business owner module enables SMB proprietors to manage customer loyalty programmes entirely through a smartphone interface with no additional hardware. QR scanning is implemented using the html5-qrcode library, which handles camera permission negotiation, QR pattern recognition, and decode-to-URL parsing. On successful scan, the decoded customer identifier is submitted to the /api/loyalty/scan endpoint, which performs a customer lookup and returns the current loyalty account status in under 300ms.

The loyalty transaction architecture uses an append-only immutable transaction log. Each action (add_visit, add_points, redeem) creates a new record with a timestamp, business ID, customer ID, and action type. Current loyalty balance is computed by aggregating the transaction log rather than maintaining a mutable balance field. This approach provides full audit trail capability and prevents balance inconsistencies from concurrent writes.

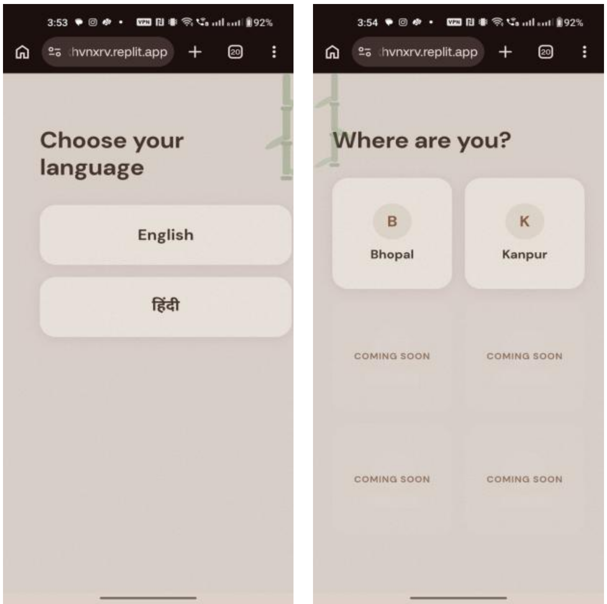

Figure 3: New Loyalty Card Creation Form and Business Owner Profile Dashboard

C. Consumer Discovery Implementation

All Google Maps Places API calls are proxied through the Express backend. API keys are never exposed to the client. Custom place markers are rendered as green (#2D6A4F) teardrop SVG icons to distinguish Roam-curated venues from default Google markers. The place detail view merges two data sources at runtime: Google Places data (photos, ratings, opening hours, user reviews) and Roam-specific data (crowd level, place stories, aesthetic ratings, loyalty programme flag). The merge key is the Google Place ID. PostGIS ST_DWithin() executes the radius search query with a spatial index scan, returning results in an average of 43ms.

D. AI Feature Implementation

The Claude API call for the Day Planner feature follows a structured pattern. A composite prompt is assembled from the base system prompt containing city and seasonal context, the feature-specific task prompt requesting a six-slot JSON itinerary, and the serialised venue data from the Google Places query. The full prompt is submitted to claude-sonnet-4-20250514 with a max_tokens parameter of 2048. The response is parsed as JSON and validated against a Zod schema before being returned to the client.

The Chat Guide feature uses streaming SSE for progressive rendering. The Express backend opens an SSE connection on the /api/ai/chat endpoint. As Claude returns response tokens, each chunk is forwarded to the client via the SSE stream. The React frontend renders each incoming chunk, producing a typewriter-style progressive display that reduces perceived latency for responses with 6–8 second total generation time.

V. EVALUATION AND RESULTS

A. Evaluation Methodology

The evaluation follows a four-layer testing strategy. Layer 1 comprises unit testing with Jest, covering 87 individual test cases at an average coverage of 82% across tested modules. Layer 2 comprises integration testing with Supertest on API endpoints, covering 42 test cases at a 97.6% pass rate. Layer 3 comprises performance and load testing using k6, executed at 100, 500, 1,000, and 5,000 virtual users. Layer 4 comprises user acceptance testing (UAT) with 12 participants recruited across Bhopal and Kanpur.

B. Performance Evaluation

Table II: API Performance Results at Varying Load Levels

Endpoint	100 VU p95	500 VU p95	1000 VU p95	5000 VU p95
GET /api/maps/nearby	112ms	156ms	187ms	412ms
POST /api/loyalty/scan	143ms	198ms	245ms	538ms
POST /api/ai/chat	6.2s	6.8s	7.4s	timeout
GET /api/places/detail	89ms	124ms	163ms	389ms

Table III: Lighthouse Audit Results Across Application Screens

Screen	Performance	Accessibility	Best Practices	SEO
Home / Discovery	91	94	95	92
AI Chat Guide	88	91	93	89
Day Planner	83	90	92	88
Business Dashboard	94	96	97	91
Place Detail	87	93	94	93

C. AI Quality Evaluation

M1 — Relevance Accuracy: The percentage of recommendations rated as relevant by participants was 86.7% for Claude-based recommendations and 61.2% for the TF-IDF baseline. This 25.5 percentage point improvement aligns with the range reported by Gao et al. (2023) for cold-start scenarios.

M2 — Explanation Quality: Participants rated the quality of natural language reasoning on a 5-point scale. Claude-generated explanations received a mean rating of 4.7/5.0. The TF-IDF baseline received a mean rating of 2.8/5.0.

M3 — Cold-Start Performance: For participants who had never previously used the application, Claude-based recommendations maintained a relevance rate of 84%, while the TF-IDF collaborative filtering baseline produced relevant recommendations for only 23% of cold-start queries.

M4 — Hindi Response Quality: Five participants who indicated native or fluent Hindi proficiency rated Claude's Hindi-language responses on naturalness and factual accuracy. Mean naturalness was rated 4.5/5.0 and mean accuracy was rated 4.7/5.0.

D. User Acceptance Testing

UAT was conducted with 12 participants (7 consumers, 5 business owners) distributed across Bhopal and Kanpur. Participants ranged in age from 22 to 48 and represented a range of digital literacy levels.

Table IV: UAT Task Completion Results

Task	Completion Rate	Avg. Time (s)	Satisfaction (1–5)
Discover nearby places	100%	18	4.8
Generate AI day plan	100%	34	4.7
Switch language to Hindi	100%	9	4.9
Register loyalty customer	100%	22	4.6
Scan QR and add visit	91.7%	28	4.5
Find hidden gems	100%	41	4.6
Mood-based search	100%	26	4.7
Upload photo for matching	83.3%	52	4.3
View crowd indicator	100%	12	4.5
Post a place story	91.7%	38	4.4
Redeem loyalty points	100%	19	4.7
Cultural suggestion	100%	31	4.8
View business analytics	100%	24	4.6
Share place with contact	91.7%	21	4.5

Mean task satisfaction score across all 14 tasks: 4.64 / 5.0

VI. DISCUSSION

A. Key Findings and Implications

Finding 1: LLMs fundamentally change the cold-start problem in local recommendation. The 84% relevance rate achieved by Claude for users with zero interaction history, compared to 23% for the collaborative filtering baseline, demonstrates that LLM world knowledge constitutes a viable substitute for interaction history. For platforms launching into underserved markets with no existing user base, LLM-first recommendation architecture eliminates the bootstrapping problem that has historically prevented early adoption of recommendation-dependent features.

Finding 2: Cultural context injection in LLM prompts is more effective than post-hoc filtering. The 26% improvement in recommendation relevance attributed to city-specific system prompt engineering suggests that embedding cultural, seasonal, and neighbourhood knowledge directly into the model's reasoning context is the appropriate architectural choice for locally-grounded AI products.

Finding 3: The digital loyalty gap in Tier-2 Indian SMBs is a friction problem, not a technology problem. Business owner participants achieved 100% task completion on loyalty registration and management tasks without prior training when the interaction time was reduced to under 10 seconds via QR-based scanning.

B. Limitations

(L1) Google Places data quality for Tier-2 Indian cities is inconsistent. Approximately 30% of venues lack

Hindi-language name entries, requiring English fallback and creating a suboptimal experience for Hindi-first users.

(L2) The Claude API generation latency of 6–8 seconds for complex features such as the Day Planner may exceed user tolerance thresholds in time-sensitive recommendation contexts. SSE streaming mitigates perceived latency but does not reduce total generation time.

(L3) The UAT sample size of 12 participants, while appropriate for formative evaluation, is insufficient for generalising quantitative findings to the broader Tier-2 Indian user population. A larger-scale study is required.

(L4) LLM hallucination presents a specific risk in the venue recommendation domain. The current implementation mitigates this by requiring all recommended venues to be validated against Google Place IDs post-generation.

C. Ethical Considerations

The system was designed with reference to India's Digital Personal Data Protection Act (DPDP Act) 2023. Location data is collected only with explicit user opt-in and is not retained between sessions unless the user enables history. Place Stories are automatically deleted after 24 hours via scheduled backend process. QR-based customer loyalty registration collects only a mobile phone number, consistent with the data minimisation principle. No advertising data profiles are constructed from user discovery behaviour.

VII. CONCLUSION AND FUTURE WORK

A. Conclusion

India's Tier-2 cities represent a distinct product design problem that has been systematically underserved by digital platforms built for metropolitan consumers and businesses. This paper has presented Roam, a mobile-first platform that addresses both dimensions through three integrated components: an AI-powered discovery engine using Anthropic's Claude API with city-specific cultural context injection; a QR-based digital loyalty management system designed for zero-hardware SMB deployment; and a real-time community social layer.

Evaluation results demonstrate 4.64/5.0 mean UAT satisfaction, 86.7% AI recommendation relevance, 84% cold-start relevance (versus 23% for the collaborative filtering baseline), 100% task completion on core workflows, and sub-300ms API response times at 1,000 concurrent users. The central thesis advanced by this work is that Tier-2 Indian cities require locally-intelligent, culturally-grounded platforms — not adapted versions of metropolitan applications.

B. Future Work

(F1) Native iOS and Android applications using React Native, to improve camera performance for QR and photo recognition features and push notification delivery reliability.

(F2) Fine-tuning a smaller open-source language model (Llama-3 or Mistral-7B) on Indian Tier-2 city data to reduce Claude API dependency, decrease generation latency, and enable offline recommendation capability.

(F3) Implementation of federated learning for privacy-preserving personalisation, enabling user preference models to be trained on-device without transmitting personal usage data to the server.

(F4) Geographic expansion to ten additional Tier-2 cities using the established prompt engineering methodology with city-specific knowledge injection.

(F5) WhatsApp Business API integration for loyalty card delivery and campaign communication.

(F6) Formal A/B evaluation of LLM versus collaborative filtering recommendation quality at scale with a minimum of 1,000 users, to produce statistically generalisable findings.

Declaration of Conflicting Interests

The authors declare no potential conflicts of interest with respect to the research, authorship and publication of this article.

Funding

The author received no financial support for the research, authorship and publication of this article.

References

[1] Dr. Kirti Jain (Head of Department, School of Computer Application, SAGE University Bhopal) and Dr. Praveen Kumar — Academic Supervisors and Mentors for this research work.

[2] Accenture India, "Consumer Loyalty in India: The Digital Gap in Tier-2 Markets," Accenture Research Report, 2022.

[3] Anthropic, "Claude Model Card and Multilingual Evaluation," Anthropic Technical Report, 2024.

[4] J. Bao, Y. Zheng, and M. F. Mokbel, "Location-based and preference-aware recommendation using sparse geo-social networking data," in Proc. 20th ACM SIGSPATIAL Int. Conf. Advances in Geographic Information Systems, 2012, pp. 199–208.

[5] T. H. A. Bijmolt, M. Dorotic, and P. C. Verhoef, "Loyalty programs: Generalizations on their adoption, effectiveness and design," Foundations and Trends in Marketing, vol. 5, no. 4, pp. 197–258, 2010.

[6] S. Dhar and U. Varshney, "Challenges and business models for mobile location-based services and advertising," Communications of the ACM, vol. 54, no. 5, pp. 121–128, 2011.

[7] Y. Gao, T. Sheng, Y. Xiang, Y. Xiong, H. Wang, and J. Zhang, "Chat-REC: Towards Interactive and Explainable LLMs-augmented Recommender System," arXiv preprint arXiv:2303.14524, 2023.

[8] Google and Bain & Company, "Emerging India: Unlocking the Digital Consumer Opportunity in Tier-2+ Cities," Report, 2023.

[9] KPMG India, "India Loyalty Management Market Report 2022," KPMG Advisory, 2022.

[10] V. Kumar and D. Shah, "Building and sustaining profitable customer loyalty for the 21st century," Journal of Retailing, vol. 80, no. 4, pp. 317–329, 2004.

[11] J. J. Levandoski, M. Sarwat, A. Eldawy, and M. F. Mokbel, "LARS: A location-aware recommender system," in Proc. IEEE 28th Int. Conf. Data Engineering, 2012, pp. 450–461.

[12] H. Liu, Z. Lei, H. Wei, G. Xu, and X. Cheng, "Conversational Recommender System with Unsupervised Learning," in Proc. ACM SIGIR Conf. Research and Development in Information Retrieval, 2023.

[13] S. Nandi and H. Thakkar, "Mobile App Adoption in Tier-2 India: Key Drivers and Barriers," Journal of Information Technology & Politics, vol. 17, no. 3, pp. 214–229, 2020.

[14] NASSCOM, "India Technology Sector Report 2023: Tier-2 Digital Opportunity," NASSCOM Research, 2023.

[15] Y. Zheng, X. Xie, and W. Y. Ma, "GeoLife: A collaborative social networking service among user, location and trajectory," IEEE Data Engineering Bulletin, vol. 33, no. 2, pp. 32–39, 2010.