

An AI-powered Location-based Recommendation System for Personalized Urban Mobility

Devika Yesodharan Devaprabha, Jia Wang
School of Computing and Mathematical Sciences
University of Greenwich
London, United Kingdom

Abstract—The increasing availability of heterogeneous spatial and contextual data sources presents both an opportunity and a challenge for location-based decision support systems. This paper presents an AI-powered recommendation system designed to facilitate personalized urban mobility while enhancing users' sense of place. The proposed system integrates multiple open data sources through a unified fusion pipeline that aligns various data to support location-based recommendations. Data sources include UK police crime statistics, Ticketmaster event listings, OpenStreetMap points of interest (POIs), online news articles, and routing data. The proposed system comprises four components: a dynamic scoring module for safety and popularity assessment, a conversational AI interface supporting natural language location queries, a personalized recommendation engine informed by user interaction behaviors, and a journey planner that provides user with ranked routes based on safety. The result confirms real-time responsiveness across all four components and contextually meaningful outputs across diverse query types. The key contributions of this work are: (i) a scalable, timely spatial data fusion pipeline for multi-source urban data; and (ii) an AI-augmented decision support framework that promotes more informed and safer urban mobility.

Keywords—Location; Natural language processing; Machine Learning; Conversational AI; Safety; Popularity; Routing;

I. INTRODUCTION

Understanding the safety, events, and contextual characteristics of an unfamiliar urban area remains a significant challenge for many users. Consider a traveler arriving in an unfamiliar city for the first time. Sarah is visiting Glasgow for a weekend and like many solo female travelers, her travel plan is shaped not only by what to do but by where it is safe to do it. To adequately plan a secure and engaging travel plan, she must determine which neighborhoods are safe for solo travel while simultaneously identifying localized events of interest. Currently, this process requires an ad-hoc investigation utilizing travel blogs, social media applications, commercial mapping services, and varied review platforms. For Sarah, she identifies a hostel in the city center, cross-references several review sites and would usually make decisions based on fragmented information based on a partial understanding of the place to visit.

Sarah's experience is not exceptional. Tourists, residents, and daily commuters share a common challenge: achieving comprehensive situational awareness of a specific location requires cross-referencing multiple, disconnected applications platforms simultaneously. Crime statistics where they are publicly available in the UK (such as the crime statistics from the Office for National Statistics¹), are seldom available in the tools for everyday navigation nor for tourism. Event listings for festivals, exhibitions, and local activities are dispersed across services such as Ticketmaster². News coverage of recent incidents resides in separate databases with no spatial interface. Tourist attraction information is distributed across mapping services, reviewing platforms such as Tripadvisor³, and social media platforms such as Instagram⁴. This fragmentation imposes a significant cognitive burden on users, requiring them to manually aggregate disparate sources, mentally reconcile conflicting or incomplete information, and nevertheless remain unaware of critical contextual factors relating to general location safety, ongoing events, or recent local incidents.

The core problem is not the data availability but how multiple datasets can be brought together, spatially contextualized and presented in a form that help enhance a user's understanding of place. This paper describes a system designed to address that gap. The proposed system is a web-based platform that aggregates six types of data sources and applies machine learning-based safety and popularity scoring to each location. Users can interact with the system through four complementary interfaces: a visual map with toggleable data layers, a conversational AI assistant that responds to natural language queries, a personalized recommendation engine that adapts to user behaviors over time, and a day trip planner that generates walking routes optimized around safety and local interest.

This system is designed with two primary user groups in mind. Travelers require a rapid and trustworthy overview of unfamiliar areas, and newly settled residents seek an evolving understanding of their surroundings as they become part of the community. By integrating open data sources into a single spatial framework, the system aims to support more informed, context aware, and safety conscious mobility across diverse user needs.

II. RELATED WORK

A. Location-based Recommendation Systems

Location-based recommendation systems must address two fundamental challenges: personalization and scalability. Zhang et al. [1] demonstrate that geographical influence on user behavior is inherently individual, proposing a kernel density estimation framework that personalizes spatial modelling per user, outperforming universal distribution models such as power-law and multi-center Gaussian approaches. Liu et al. [2] extend contextual modelling beyond spatial proximity by applying Skip-gram representation learning to user visit sequences, capturing richer trajectory patterns than Markov-based predecessors. Yang et al. [3] further enrich preference modelling by integrating sentiment analysis of user-generated venue reviews alongside check-in data, demonstrating that qualitative signals meaningfully improve recommendation accuracy. Very few reviewed systems incorporate up-to-date contextual signals, safety considerations, or truly integrated multi-source data fusion, which are the gaps that represent a clear agenda for next-generation recommendation systems capable of serving diverse user needs in dynamic environments.

B. Conversational Interfaces for Geospatial Queries

The use of natural language interfaces for location-based queries has grown rapidly with the availability of large language models (LLM). Google's search integration with Maps allows natural language queries like "coffee shops open now near me," and Apple Maps has integrated Siri for similar tasks. Research in this area has focused on intent detection [4], entity recognition for place names and categories [5], and dialogue management for multi-turn location conversations [6].

Where these systems fall short for our use case is that they are designed for well-known, widely indexed places. Asking "is the Northern Quarter in Manchester safe at night?" does not return a useful answer from any existing conversational interface, the system would need to pull live crime data, assess sentiment from recent news, and synthesize an answer grounded in those sources. The chatbot is designed specifically for this kind of query, grounding its responses in the same real-time data that drives the map.

III. METHODOLOGY

A. Data Sources and Data Retrieval

UK crime statistics through the Police API⁵ provide street-level crime records for any location in England, Wales, and Scotland. Ticketmaster is used to retrieve upcoming public events near the searched location. The proposed system queries events search endpoint using the queried location coordinates and a 10 km radius, filtering for events starting within the next 30 days. Each event record includes the event name, category, start date, and whether it is free or paid. The volume and variety of events are a meaningful signal of how active an area is. High event density correlates with higher footfall, which generally improves perceived safety. OpenStreetMap⁶ (OSM) data are retrieved using Overpass API⁷ provides POIs (points of interest) including cafés, museums, parks, galleries, historic sites, and transport hubs. The system issues an Overpass QL query to retrieve all OSM nodes tagged as amenity, tourism, or historic within a 1 km radius of the user's selected location. The returned records include each POI together with its category and geographic coordinates. These data form the basis for two components of

the application: the POI layer displayed on the interactive map and the pool of candidate locations considered by the journey planning module. NewsAPI⁸ retrieves recent news articles mentioning the searched location from aggregated news article dataset across UK and international sources. Each retrieved article contains four key fields: the headline, a brief description, the publication date, and the source outlet. The system then processes these articles through a natural language processing (NLP) pipeline to extract sentiment indicators. Each article's text is analyzed using VADER (Valence Aware Dictionary and sEntiment Reasoner) [7], a sentiment analysis tool from the NLTK library that is well-suited to short news text. VADER assigns each article a sentiment score ranging from -1 (strongly negative) to 1 (strongly positive) to reflect whether the news coverage of an area is broadly positive or concerning. These scores are aggregated across all retrieved articles to produce a location-based sentiment signal, which is subsequently used as a feature in the safety scoring.

Two other APIs supporting the system's georeferencing and routing capabilities are also adopted. Nominatim⁹ is used for geocoding. OpenRouteService¹⁰ (ORS) provides routing functionality. ORS generates optimal routes between POIs and supports multiple travel modes including walking and cycling.

B. Safety and Popularity Scoring

For each location query, the system generates a safety score and a popularity score based on 13 features grouped into four categories: crime, events, news, and POIs (see Table I).

TABLE I. SAFETY AND POPULARITY MEASURE

Category	Feature	Description	Score
Crime	total_crimes	Count of crimes	Safety
	crime_density	Crimes per km ²	Safety
	violent_crime_ratio	Proportion of violent crimes	Safety
	recent_crime_ratio	Crimes from last three months	Safety
Events	total_events	Count of upcoming events	Popularity
	free_event_ratio	Proportion of free events	Popularity
	event_diversity	Unique event categories	Popularity
News	total_articles	News article count	Both
	avg_sentiment	Mean sentiment score	Safety
	news_frequency	Articles per day	Popularity
POIs	total_pois	POI count	Popularity
	poi_diversity	Unique POI types	Popularity
	essential_amenity_ratio	Healthcare & transport POIs	Safety

Crime incidents from the last three months are weighted higher than older records. The news sentiment feature is derived from the aspect-based NLP pipeline that classifies news text across four location-relevant aspects as safety, amenities, transport, and environment (more details see Section V (B)). The *essential_amenity_ratio* further contributes to the composite safety score. Public essential amenities such as healthcare facilities (e.g., urgent care centers, walk-in clinics and pharmacies) and transport POIs (e.g., tube and bus stations) are incorporated on the basis that their spatial density promotes pedestrian footfall and relevant to pedestrian travel. Locations with a higher density of public amenities attract greater pedestrian activity, which increases informal surveillance and correlates with lower perceived risk [8].

⁵ <https://data.police.uk/docs/>

⁶ <https://www.openstreetmap.org/>

⁷ <https://overpass-api.de/api>

⁸ <https://newsapi.org/v2>

⁹ <https://nominatim.openstreetmap.org>

¹⁰ <https://openrouteservice.org/dev/>

The popularity score integrates event density, POI count, and news volume. News contributes through article count and publication frequency. News coverage of local events, including festivals, restaurant openings, and community gatherings, contributes to the popularity signal. Media attention is indicative of heightened area activity and public visibility.

Both safety and popularity scores are computed using a stacked ensemble model, which combines predictions from three gradient boosting algorithms (XGBoost, LightGBM, and CatBoost) to achieve optimal performance compared to individual models (see Table III in Section V). The scores are not stored persistently but are instead computed on demand from the most recently available data at the time of each retrieval.

Safety scores are presented to users on a scale of 0 to 10. Each score is accompanied by a short descriptive label to support users who prefer a plain-language interpretation over a numeric value. The score and label are color-coded into three bands summarized in Table II.

TABLE II. SAFETY SCORE THRESHOLD BANDS

Score Range	Color	Label	Indicator
7.0 - 10.0	Green	Generally Safe	Low crime density, low violent crime ratio, neutral or positive sentiment
5.0 - 6.9	Amber	Stay Aware	Average urban crime levels, moderate crime density
0.0 - 4.9	Red	High Caution	Elevated crime density, high violent crime ratio, or negative sentiment

A green score in the range 7.0 to 10.0 with the label "Generally Safe" indicates low crime density. Green scores reflect a low proportion of violent incidents and neutral or positive local news sentiment. Such scores are typical associated with residential or tourist-oriented areas with few recorded incidents. An amber score in the range 5.0 to 6.9 with the label "Stay Aware" reflects average urban crime levels and moderate crime density. The amber band is consistent with active city centers where incidents occur but do not dominate. A red score in the range 0.0 to 4.9 with the label "High Caution" indicates elevated crime density. Red scores reflect a high proportion of violent crime categories or sustained negative news sentiment.

The classification thresholds are empirically derived from safety scores computed across 500 UK locations, collected using the same API pipeline as the deployed system. The sampling from dense city centers to small market towns with a median score in the range of 4.5 to 5.5. Central districts of major cities including Liverpool, Manchester, and Birmingham consequently receive amber or red classifications, whereas suburban and rural locations predominantly fall within the green band. Both the numeric score and the descriptive label are displayed together to ensure the value and its interpretation remain visible simultaneously.

C. Conversational AI Assistant

A conversational AI assistant is proposed and is based on an intent detection component that classifies each user message into one of several categories relevant to location-based queries related to safety, event discovery, news and incidents, general location questions, and trip planning. For

each message, the detected intent determines which data are retrieved and how the response is constructed. Safety queries use the latest crime statistics and local news sentiment for the selected area, whereas trip planning queries call the journey plan component and return a proposed itinerary. When a LLM connection is available the system passes the retrieved data as context to DeepSeek¹¹, which generates a natural language answer.

D. Personalised Recommendation System

The system includes a recommendation component that adapts location suggestions to individual users over time. When a user interacts with the map such as clicking on POIs, generating trip plans, or exploring categories, the system records these interactions against a session identifier. Over subsequent visits, the recommendation engine uses this interaction history to adjust how candidate locations are ranked. More frequently engaged categories receive greater weight.

The ranking function combines three signals: a base score from safety and popularity, a category preference score derived from interaction history, and a freshness signal for locations with recent activity such as new events or news coverage. A user who consistently explores museums and independent cafés will see these appeared more prominently in future suggestions, even if other categories score higher on raw popularity alone. This component is most visible in the journey planner, where the selected stops reflect the user's established preferences. The ranking also influences the order in which POIs appear. New users with no interaction history still receive recommendations based on popularity with safety prioritized.

E. Safety-Aware Journey Planner

The system generates journey stops through a three-stage process. First, non-destination POIs (parking spaces, ATMs, waste bins) are filtered from the candidate pool. Second, remaining POIs are scored based on category priority (e.g., museums rank higher than fast food outlets), popularity, and safety. Third, the highest scoring candidates are ordered using a nearest neighbor algorithm, which sequentially selects the closest unvisited stop from the user's current location. Walking times between stops are calculated via ORS. The itinerary is displayed both as a numbered list in a side panel and as a connected route with colored markers on the map. Each stop includes its name, category, safety score, and walking time from the previous location.

IV. IMPLEMENTATION

A. Data Pre-processing and Feature Engineering

Crime data is filtered to remove records outside the 1 km search radius. Duplicate entries which can appear when the same incident is recorded under multiple categories are removed by matching on location coordinates and month. A recency weight is then applied: crimes reported in the last three months receive a weight of 1.0, those between three and six months old receive 0.7, and those older than six months receive 0.4. This ensures the safety score reflects current conditions rather than historical patterns that may no longer apply to the area.

News articles are deduplicated by URL to remove syndicated copies of the same story. Articles with fewer than 50 words are excluded as they contain insufficient text for meaningful analysis. The system retrieves news articles with the location name as the search keyword. To obtain current

¹¹ <https://www.deepseek.com>

data, the query filters results to articles published within the previous 30 days and ranks them by relevance.

POI records from the Overpass API include feature types that are not meaningful visit destinations such as parking spaces, cash machines, post boxes, and waste collection points. These are removed using a category exclusion list. Each remaining POI is assigned a category priority score used by the journey planner: museums and historic sites score highest (9 to 10), caf  s and galleries score mid-range (6 to 8), and fast-food outlets score lowest (2 to 3).

The 13 features in Table II are normalized to a value between 0 and 1 before being passed to the scoring model. This prevents features with large raw values, such as total crime count from dominating features expressed as ratios.

The ensemble scoring model is trained on data collected for 500 UK locations, covering a range of urban densities from central London to small market towns. For each location, the same set of APIs are queried, and the same feature extraction pipeline is applied. Safety and popularity labels are derived from the same data sources used by the live system, ensuring consistency between training and inference. Five-fold cross-validation is used throughout model training and evaluation (Table III).

B. System Architecture and Data Flow

The system development follows a client-server architecture. The backend is built with FastAPI¹², a Python web framework suited to asynchronous data-fetching tasks. The frontend is a React application rendered in the browser, displaying an interactive Leaflet map with several toggleable data layers. The full source code of this project is publicly available at GitHub¹³.

Fig.1 illustrates the end-to-end data flow. When a user searches for a location, the frontend sends the place name to the backend. Nominatim geocodes the place name to georeferenced coordinates. The backend then issues parallel HTTP requests to the UK Police API, Ticketmaster API, the Overpass API, and NewsAPI simultaneously using Python's asynchronous request framework. All the data sources are queried at the same time, which keeps response times short. Results are cached in Redis with source-specific expiry times: crime data is cached for one hour, POI data for 24 hours, news for 15 minutes, and events for 30 minutes. On repeated searches for the same location, cached data is returned immediately with no external API calls. Once data retrieval is complete, all the features (see Table II) are fed into a stacked ensemble model to generate safety and popularity scores. All results, including scores, raw data layers, and an AI generated location summary, are returned to the frontend in a single JSON response, simultaneously populating the map, the side panel, and the summary card in the dashboard as the final product (Fig. 2 in section V).

The chatbot operates as a separate request path. When a user sends a message, the backend fetches the relevant data for the current location and passes it to the LLM with a structured prompt. The journey planner is triggered as a sub-process both when the chatbot detects a trip planning intent and when the user clicks "Plan Day Trip" in the side panel from the dashboard.




Figure 1. High-level system architecture and data flow.

V. RESULTS AND EVALUATION

A. Scoring Model Performances

Table III presents the performances of the selected machine learning models. The stacked ensemble consistently outperforms each individual model across both safety and popularity scores. The safety R^2 of 0.82 means the model accounts for 82% of the variation in safety scores across locations a meaningful improvement over the best individual model (XGBoost at 0.76). The gain confirms that combining the three models captures patterns that no single algorithm identifies alone. The popularity model performs slightly below safety across all configurations. This is expected since popularity is a broader, noisier signal drawn from multiple features, whereas safety relies mainly on crime statistics.

TABLE III. SCORING MODEL PERFORMANCE COMPARISON

Model	Safety R^2	Safety RMSE	Popularity R^2	Popularity RMSE
XGBoost	0.76	0.142	0.72	0.158
LightGBM	0.74	0.148	0.71	0.161
CatBoost	0.75	0.145	0.73	0.155
Stacked Ensemble	0.82	0.121	0.78	0.139

¹² <https://fastapi.tiangolo.com/>

¹³ <https://github.com/devikayd/Hybrid-AI-Location-App-.git>

B. NLP Pipeline Performance

The *avg_sentiment* feature in Table I is produced by an aspect-based sentiment pipeline that classifies news text across four location-relevant aspects as safety, amenities, transport, and environment. Table IV shows the precision, recall, and F1-score for each aspect evaluated on a held-out test set.

TABLE IV. TABLE I. SENTIMENT CLASSIFICATION ACCURACY

Aspect	Precision	Recall	F1-Score
Safety	0.78	0.72	0.75
Amenities	0.81	0.79	0.80
Transport	0.76	0.74	0.75
Environment	0.73	0.68	0.70
Macro Average	0.77	0.73	0.75

The pipeline performs most reliably on amenities (F1: 0.80), where news language tends to be more specific and consistent. Safety-related language is more varied in tone, which accounts for the slightly lower recall (0.72) the model misses some safety-relevant mentions. Environment scores lowest across all metrics, reflecting the ambiguity of environmental language in local news. Overall, the macro-average F1 of 0.75 is sufficient for the sentiment signal to contribute meaningfully to the safety score.

C. Conversational AI Performance

The chatbot's intent detection component is evaluated on 200 labelled user queries covering five intent types. Table V reports the intent classification performance along with the number of test samples per class.

TABLE V. TABLE II. INTENT CLASSIFICATION PERFORMANCE

Intent Type	Precision	Recall	F1 Score	Sample
safety query	0.91	0.88	0.89	52
event search	0.87	0.92	0.89	41
poi search	0.85	0.83	0.84	38
comparison	0.82	0.79	0.80	29
general info	0.78	0.81	0.79	40
Macro Average	0.85	0.85	0.84	200

Safety queries achieve the highest precision (0.91), which is the most important result for the system's core use case misclassifying a safety question as something else would be the most consequential error. Event search shows the highest recall (0.92), meaning the system rarely misses an event-related query. The *general_info* intent scores lowest (F1: 0.79), which is expected given its broad, catch-all nature. The macro-average F1 of 0.84 across 200 queries confirms the chatbot reliably identifies what users are asking about before retrieving and presenting the relevant data.

D. Evaluation

The system was evaluated through a structured functional demonstration of ten UK locations across multiple spatial scales (London, Greenwich, Birmingham, Leicester, Leeds, Manchester, Borough Market, Liverpool, London Bridge, Blackpool), ranging from dense urban centers and popular tourism locations (such as London Bridge and Borough Market) to smaller towns and local district (such as Blackpool and Greenwich). For each location, data retrieval, score calculation, journey planning, and chatbot accuracy were assessed.

All ten locations returned data from at least four data sources (crime statistics, events, POIs and news). Gaps in crime coverage for rural postcodes and reduced Ticketmaster listings outside major cities reflect genuine data availability

rather than system failures. Safety and popularity scores were produced for all ten locations. The journey planner returned complete itineraries for eight locations. The remaining two (Greenwich and Blackpool) returned shorter itineraries of two and three stops due to limited POI availability, which is the expected behavior for smaller and less busy locations.

Chatbot responses were judged contextually accurate for 43 of 50 test cases (consistent with the macro-average F1 of 0.84 in Table V), with the seven less useful responses arising from locations with limited news coverage.

E. The AI-powered Location Application

The final product is a hybrid AI location app. Its dashboard layout includes an interactive map view, an AI summary of the queried location, a day trip planner, and a chatbot (see Fig. 2).

The search bar at the top accepts location names, postcodes, or addresses, and displays up to four recent searches for quick navigation. The default location is London, so the map opens with London data on first load. Once location is selected, the map centers on it and fetches data from all sources automatically. The interactive Leaflet map supports zooming and panning from city to street level. Layer toggles control data visibility with distinct marker colors: red for crimes, blue for events, green for POIs, and orange for news. Clicking any marker displays a popup with the item's title, category, and date.

When "Day Trip Planner" is selected, a blue dashed route connects the stops with numbered markers colored by safety level (red for high, amber for moderate, green for low). The route disappears when no trip is planned.

The side panel displays safety and popularity scores, aggregated counts for crimes, events, news, and POIs, an AI generated location summary, a step-by-step walking itinerary, and a conversational chatbot. Results are organized by category and can be expanded or collapsed individually.

Fig. 3 demonstrates the safety-aware journey planner feature using Liverpool as an example. The map displays a blue dashed route connecting three numbered stops in the Lime Street area in Liverpool. Each stop is marked with a numbered circular marker colored by safety level. The Day Trip Planner panel lists each stop with its name, category, safety score, and estimated walking time from the previous stop. The three stops, James Morris, BIIRD, and Future Sound, are all categorized as Music, with a total estimated walking time of 21 minutes. All the three stops display a safety score of 5.0/10 with the amber label "Stay Aware". Such scoring is consistent with the moderate crime level of central Liverpool (50 recorded incidents). The AI summary panel shows over 50 upcoming events in the area.




Figure 2. The AI-powered location-recommendation system including the top search bar, the side panels displaying safety and popularity with AI summaries as well as routing results.




Figure 3. An example of safety-aware journey planner. The blue waking route connects three POIs (categorized as music) with safety score and AI summary of the area shown in the side panels. The green circle marks the starting point and the purple circle marks the destination. Colors of the stops along the route indicate safety bands.

F. Comparative Evaluation against Existing Applications

This section evaluates the proposed system against three widely used platforms representing the closest existing tools across navigation, tourism recommendation, and urban mobility: Google Maps, Tripadvisor, and Citymapper¹⁴. For each platform, specific gaps are identified and the corresponding contributions of the current system are quantified.

Google Maps integrates real-time traffic, transit feeds, and user-contributed reviews through Google Places API, with

routing optimized for travel time. It does not integrate official crime statistics with route recommendations and area summaries. There is no safety score, no crime density feature, and no violent crime signal surfaced to users. Whilst the latest versions include conversational search powered by a large language model (Gemini), safety responses are generated from the model's pre-trained knowledge rather than from current published crime records or news coverage. A query such as "Is Peckham safe to visit at night?" receives a generic response with no grounding in the most recently available crime statistics or news sentiment. The current system addresses this

¹⁴ <https://citymapper.com>

gap directly. Searching Peckham returns a Safety Score of 4.5/10 (High Caution) derived from 100 recorded crimes in the past three months and 95 local news articles processed through the VADER sentiment pipeline. This provides users a well-justified answer of safety which is data grounded.

Tripadvisor ranks hotels, restaurants, and attractions through crowdsourced reviews and supports manual itinerary construction. It provides no real-time feed, no crime data, and no detection of recent incident clusters. Ranking is driven purely by aggregate review score, so a highly rated restaurant in a high-crime area could receive identical score to its equivalent in a safe area. Personalization is limited to browsing history without collaborative filtering or interaction-type weighting. The current system addresses these gaps by embedding and prioritizing safety scores directly into location ranking and recommendation. This ensures that low-safety locations are ranked lower regardless of their review scores. Recommendations adapt over time through interaction tracking. A user who repeatedly engages with museums and independent cafés will find those categories weighted more prominently in subsequent suggestions. The system also prioritizes recency. Unlike Tripadvisor's static review rankings, locations with upcoming events or recent news coverage receive greater prominence, ensuring suggestions reflect current activity rather than historical popularity.

Citymapper provides multi-modal transit routing across major UK cities optimizing for travel time and mode efficiency. It is mostly a transit-only platform. It surfaces no POIs or events along a route, assigns no safety scores to walking segments, and provides no natural language interface. Its underlying data model draws exclusively on standardized timetable feeds and real-time departures, with no integration of crime databases, news APIs, event feeds, or sentiment analysis. While Citymapper serves commuters effectively, it is not designed for location recommendation, which is a task that demands a more in-depth, multi-dimensional understanding of place. The proposed system fills the gap between transit

routing and safety-aware itinerary planning. Stops are ranked using a composite of safety score, popularity, category priority, and interaction-derived user preference. Natural language queries such as "Plan a day in Manchester avoiding high-crime areas" are handled directly through the conversational interface.

VI. USAGE SCENARIOS

A. Scenario A: Solo Traveller

Maria plans a weekend in Manchester and searches for "Manchester Northern Quarter". The system geocodes the input to coordinates (53.4831° N, 2.2367° W) and returns a safety score of 45.6%, a popularity score of 30.1%, crime markers for the past three months, 12 upcoming events, and 47 nearby points of interest rendered as toggleable map layers.

Maria asks the chatbot: "Is it safe to walk around here at night?" The system detects a safety query and generates a contextual response using local crime statistics and news sentiment. The system returns a contextual response indicating very low crime levels for the area, with over 50 upcoming events and a high concentration of local cultural activity, suggesting the area is safe to explore at any time of day.

Maria then clicks "Plan Day Trip" from the side panel. The system filters and scores candidate locations by category priority, safety, and popularity, then applies a nearest-neighbor algorithm to order stops and calculates walking times via ORS. The resulting itinerary includes five stops: The Dave Weckl / Tom Kennedy Project (safety: 5.0/10), PIXIES: P40 Celebrating 40 Years (5.0/10), Imogen and the Knife at Manchester Deaf Institute - The Lodge (5.0/10), An Evening with PinkPantheress (5.0/10), and Bluey's Big Play (5.0/10). All stops are categorized as "Stay Aware". The route is displayed on the map with a total walking time of 1 hour 57 minutes as shown in Fig. 4.




Figure 4. Day trip planner for solo traveller Maria in the area of Manchester Northern Quarter.

B. Scenario B: New Resident

James just moved to Bedminster, Bristol, and uses the proposed location recommendation system to monitor his new neighborhood. The system returns a safety score of 50% and a popularity score of 30% for the area, with over 50 upcoming events detected via Ticketmaster. Weekly searches return updated scores derived from the latest crime records (UK crime statistics is updated monthly), current event listings, and news articles from the past seven days, enabling longitudinal tracking of neighborhood conditions.

The system records user interactions (POI clicks, trip generations, category exploration) against session identifiers. The recommendation engine uses this history to weight categories proportionally to engagement frequency. After several sessions focused on cultural and food related categories, James's trip plans in Bedminster now prioritize Tobacco Factory Theatre, independent cafés, and a local market (all rated 5.0/10 as Stay Aware) rather than generic tourist destinations, reflecting the hybrid recommendation engine's category weighting derived from his interaction history.

When James queries "Has anything happened near East Street recently?", the system detects a news query intent, retrieves articles from the past seven days mentioning East Street, and returns a sentiment-based summary indicating predominantly neutral coverage. The response indicates predominantly neutral coverage focused on a local planning dispute, providing contextual awareness without indicating safety concerns.

VII. CONCLUSIONS

This paper demonstrates an AI-enabled location-based application that gives travelers and new residents in the UK a single place to understand the safety and characters of any location they are interested in. The system integrates publicly available data from multiple sources, applies machine learning scoring to produce safety and popularity signals, and makes those signals accessible through a map interface, a conversational AI assistant, a personalized recommendation engine, and a novel safety-aware journey planner. All four features operate in real time, drawing on the latest available data each time a location is searched.

The journey planner constitutes the primary novel contribution. As demonstrated in the comparison with existing platforms, no existing publicly accessible tool generates walking itineraries for UK locations optimized simultaneously for safety and popularity. Existing solutions prioritize proximity and user ratings but neglect safety considerations for visitors in unfamiliar areas.

The current system has several limitations. First, it is designed for moderate scale deployment and excludes features such as user authentication to maintain project scope. Second, reliance on external APIs introduces dependencies on third

party services, each with rate limits that require careful management. When rate limits are exceeded, functionality is temporarily degraded. Third, data update varies by source. The UK Police API updates monthly, meaning crime data may be up to two months old at retrieval time. Fourth, event data relies exclusively on the Ticketmaster API. Alternative providers such as Eventbrite were evaluated but proved unreliable due to frequent 404 errors and inconsistent endpoint availability. This limits the event coverage to a single source used in the current study. Fifth, the system performs best in areas with comprehensive data coverage, as scores are predictions based on available features and cannot account for factors absent from the data.

Future work includes additional data sources (social media, transportation, environmental sensors), real-time notifications, historical pattern analysis, and comparative location analysis could make the system more useful. If users give permission, live location tracking could send alerts and notifications based on their current location. Implementing enhanced user interaction tracking, such as tracking which items users click in location lists, which URLs they visit, and which location they like, would improve personalized recommendations. A mobile application is planned to extend the system's accessibility, enabling travelers to access safety scores, journey plans, and the conversational assistant directly from their mobile phones while on the move.

REFERENCES

- [1] J.-D. Zhang, C.-Y. Chow, and Y. Li, "iGeoRec: A personalized and efficient geographical location recommendation framework," *IEEE Trans. Serv. Comput.*, vol. 8, no. 5, pp. 701–714, 2015.
- [2] X. Liu, Y. Liu, and X. Li, "Exploring the context of locations for personalized location recommendations," in *Proc. 25th Int. Joint Conf. Artificial Intelligence (IJCAI)*, New York, USA, 2016, pp. 1188–1194.
- [3] D. Yang, D. Zhang, Z. Yu, and Z. Wang, "A sentiment-enhanced personalized location recommendation system," in *Proc. 24th ACM Conf. Hypertext and Social Media (HT)*, Paris, France, 2013, pp. 119–128.
- [4] D. Nadeau and S. Sekine, "A Survey of Named Entity Recognition and Classification," *Linguisticae Investigationes*, vol. 30, no. 1, pp. 3–26, 2007.
- [5] E. Cambria and B. White, "Jumping NLP Curves: A Review of Natural Language Processing Research," *IEEE Computational Intelligence Magazine*, vol. 9, no. 2, pp. 48–57, 2014.
- [6] A. Bordes and J. Weston, "Learning End-to-End Goal-Oriented Dialog," in *Proc. 5th Int. Conf. on Learning Representations (ICLR)*, Toulon, France, 2017.
- [7] C. J. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in *Proc. 8th Int. AAAI Conf. Weblogs and Social Media (ICWSM)*, Ann Arbor, MI, USA, Jun. 2014.
- [8] B. Jiang, C. N. S. Mak, H. Zhong, L. Larsen, and C. J. Webster, "From broken windows to perceived routine activities: Examining impacts of environmental interventions on perceived safety of urban alleys," *Frontiers in Psychology*, vol. 9, p. 2450, 2018.