# NeuroRoute: An Agentic AI-Based Personalized Navigation System for Sensory-Friendly Travel

Siddhi Agnihotri[1], Madhvee Singh[2], Surabhi Tyagi[3], Prachi Dhanawat[4]
[1,2,3]*Usha Mittal Institute of Technology, SNDT Women's University, Mumbai 400076, India*

***Abstract*—Travelling independently can be difficult for autistic individuals, not only because of the effort of planning a journey but also because of the sensory overstimulation encountered along the way. Crowded spaces, traffic noise, bright lighting, and construction activity can trigger anxiety and sensory overload, discouraging independent travel; similar difficulties affect elderly travellers and others with sensory processing sensitivities. Conventional navigation applications such as Google Maps, however, optimize routes primarily for travel time and distance and do not account for sensory comfort. To address this gap, this paper presents NeuroRoute, an Agentic AI-based personalized navigation system that recommends routes based on sensory comfort in addition to travel efficiency. The system computes a Sensory Score for each road segment using a rule-based weighted model that considers traffic density, crowd density, noise, brightness, construction activity, and weather conditions, with weights adapted to the user's selected profile (Autistic, Elderly, or General User) and individual preferences. The application is implemented using a React Native (Expo) frontend, a Node.js/Express backend, and MongoDB for data persistence, with routing and mapping powered by OpenStreetMap, OSRM, and the Overpass API. The system further supports locating nearby quiet spaces during sensory overload and lays the architectural groundwork for a Machine Learning-based comfort prediction model and an Agentic AI decision engine. A feasibility study conducted with autistic children at a partner NGO validated the core assumptions guiding the system's design. A comparative evaluation against Google Maps over 30 paired trials (10 routes, three times of day) showed that the recommended route lowered the Sensory Score by about 7% relative to NeuroRoute's alternatives, at an average cost of 19 minutes in travel time that was concentrated on long trips.**

***Index Terms*—Agentic AI, Sensory-Friendly Navigation, Autism Spectrum Disorder, Accessible Technology, Route Planning, Open-StreetMap**

## I. INTRODUCTION

Autism Spectrum Disorder (ASD) is a neurodevelopmental condition that affects how people perceive the world and interact with others. Being a spectrum condition, it impacts people in different ways, but two features are common: differences in social interaction and atypical sensory sensitivity.

These features shape the travel experience, which involves unpredictable environmental stimuli. Moving through a city means passing crowded streets, traffic noise, bright surroundings, and construction activity, and these conditions can trigger anxiety and sensory overload. Travelling also demands constant decisions about which way to go and how to cope with changing surroundings. As a result, some autistic people avoid travelling or depend on others to accompany them, even though independent travel can improve quality of life. Similar barriers affect elderly travellers and other people with sensory or cognitive sensitivities. Crucially, this sensitivity is highly individual: a route that is convenient for one person may be distressing for another, since sensory aversions are idiosyncratic rather than shared across the population [1].

Technology could help overcome some of these obstacles. GPS-enabled smartphones and digital mapping services have made navigation available to many users, yet conventional systems such as Google Maps select routes primarily according to distance, estimated travel time, and traffic conditions. These criteria measure efficiency, but they say little about the sensory and cognitive comfort of the journey, and they leave users with no way to express their own sensory needs. There is therefore a need for a personalized navigation system that considers not only travel efficiency but also the sensory characteristics of the surrounding environment and the individual preferences and aversions of the user. Such a system

should evaluate alternative routes based on sensory conditions, help users find quieter or more suitable places, and support adaptive route selection as conditions change, while keeping the interaction simple enough not to add to the user's cognitive load.

This need has motivated a growing body of research into sensory- and cognition-aware spatial support. Cena et al. [1] proposed PIUMA, a personalized, crowdsourced interactive map that recommends "Safe" Points of Interest (SPOIs) to autistic users through a cascade content-based and collaborative-filtering recommender, factoring in a user's preferences and aversions for noise, crowding, temperature, brightness, and spaciousness. Building on this, Brighenti et al. [2] extended PIUMA using a human-centered, participatory design process involving adults with Level 1 autism spectrum disorder across interviews, prototyping, usability testing, and a field trial. Their results confirmed that sensory information strongly influences whether autistic users choose to visit a place, that map-based exploration is preferred over text menus, and that heavy customization of layout, language, and landmarks is needed to reduce cognitive load, though users tended to consume rather than contribute crowdsourced reviews. Separately, Garc´ıa-Catala et al. [3] developed SmartRoutes, an´ Android application that uses Bluetooth Low Energy beacon trilateration for indoor localization and orientation of users with cognitive disabilities such as Down syndrome, supporting both routine indoor wayfinding and emergency evacuation guidance through short, simple audio and text instructions.

While these works demonstrate the value of sensory-aware POI recommendation [1], participatory and user-centered design for autistic adults [2], and accessible indoor localization with emergency support [3], none of them provides a unified system for continuous, real-time, sensory-aware outdoor route navigation. PIUMA is oriented toward recommending static places rather than turn-by-turn travel, its participatory followup study largely confirms these findings without adding realtime routing, and SmartRoutes addresses indoor rather than outdoor mobility.

To address this gap, we propose NeuroRoute, an Agentic AI-based personalized navigation system for sensory-friendly travel. NeuroRoute supports three user profiles (Autistic, Elderly, and General) and lets users configure preferences such as avoiding crowded, noisy, bright, or construction-heavy areas while favoring parks and safer routes. A rule-based Sensory Scoring Framework evaluates each road segment using traffic density, crowd density, noise, brightness, construction activity, and weather, weighted by the selected profile and individual preferences, so that two users with the same source and destination can receive different route recommendations. Routes with lower scores carry a lower estimated sensory burden. The system is built with a React Native (Expo) frontend, a Node.js/Express backend, and MongoDB, using OpenStreetMap, Nominatim, OSRM, and the Overpass API for mapping, geocoding, routing, and nearby-place discovery. It also integrates a Random Forest-based comfort prediction model and a LangGraph-based Agentic AI decision layer for route comparison and explanation. The remainder of this paper describes the system architecture, the sensory scoring methodology, the implementation, and the results of a feasibility study conducted with autistic users.

## II. LITERATURE SURVEY

1. CARES: Human-Centered AI for Inclusive Tourism [1] is a trip-planning application for autistic adults that recommends points of interest based on user likes and dislikes, using content-based filtering combined with a routing optimizer (VROOM + OSRM) to construct a full-day itinerary. The system was evaluated with 14 users via a usability survey, and its key strength lies in keeping the user in control of the itinerary rather than presenting a single fixed route. Its main limitation is scope: it plans trips and stops rather than providing turnby-turn, real-time walking or driving navigation, and it relies solely on OpenStreetMap data, which does not capture how noisy or crowded a place actually is at a given time.
2. PIUMA: Personalized Spatial Support for People withAutism Spectrum Disorder [2] PIUMA recommends" Safe Points of Interest" and safe paths by combining two recommendation techniques: it first filters places by a user's stated sensory preferences (content-based), then refines suggestions using what similar users liked (collaborative filtering). Its sensory data (noise, crowding, brightness) comes from

crowdsourcing real people rating real places. This is powerful but fragile: if not enough people have rated an area, the system has nothing to recommend (a" cold-start" problem), and the evaluation itself used simulated ratings rather than large-scale real data.

3. PIUMA Participatory Study (Brighenti, Cena & Keller,2025) [3] This is a follow-up to PIUMA that redesigned the same idea but built it with autistic users from the start (Participatory Design), tested over six weeks with 35 adults. It confirmed that users preferred simple map-based interaction over text instructions, and that showing sensory information changed whether people chose to visit a place. Its main lesson for NeuroRoute is about design, not algorithms: keep the interface simple and visual. Its limitation is that people mostly consumed sensory data rather than contributing it, so the crowdsourcing loop was weak.
4. SmartRoutes: Sensory Navigation System for Indoor Localization and Orientation of Users with Cognitive Disabilities [4]SmartRoutes targets people with cognitive disabilities and solves a different problem: guiding someone indoors (using Bluetooth beacons) and outdoors (using GPS/OpenStreetMap), including an emergency-evacuation mode with short voice instructions. It proves that indoor+outdoor navigation for accessibility needs is achievable, but it needs physical beacon hardware installed in every building, only exists for Android, and update's location slowly (20–30 seconds) too slow for continuous outdoor travel.

Across all four systems, each address only one piece of the problem trip planning, sensory-based recommendation, participatory interface design, or indoor localization but none combines real-time outdoor GPS navigation with continuous, sensory-aware rerouting driven by live conditions such as traffic, noise, crowd density, and weather, personalized to an individual sensory profile. This is precisely the gap NeuroRoute is designed to address.

None of the four prior systems generate a reasoned explanation of why a specific route or place was chosen — they either say nothing (CARES, Google Maps), show a static sensory tag (PIUMA, PIUMA Participatory), or explain navigation steps rather than the route decision (SmartRoutes). NeuroRoute is the only one in this set that produces a factor-grounded, natural-language justification for the route recommendation itself, which is closer to how explainable-AI systems in other

Table I
Neuroroute Vs. Google Maps

| Aspect | Google Maps | NeuroRoute |
|---|---|---|
| Route Optimization Goal | Time and distance only | Sensory comfort and travel efficiency, jointly |
| Sensory Awareness | No dedicated consideration of noise, crowds, or brightness | Rule-based Sensory Score considers traffic, crowds, noise, brightness, construction, and weather |
| Personalization | Traffic and travel-mode preferences; no dedicated sensory user profile | Autistic, elderly, and general user profiles with individual preferences |
| Routing Engine | Proprietary Google routing | OpenStreetMap (OSM) + OSRM (open-source and customizable) |
| Explanation of Route Choice | No dedicated sensory-based explanation of route selection | Natural-language explanations via Gemini, grounded in the Sensory Score |
| Special-Needs Support | No dedicated sensory-processing or cognitive-support routing mode | Designed specifically to support sensory-processing needs |
| Quiet-Space Discovery | No dedicated quiet-space discovery feature | Supported via Overpass API to discover places such as parks, libraries, and cafés |
| Cost Model | Freemium model; API usage may incur charges at scale | Open-source routing stack with no per-request OSRM routing fee; hosting and other service costs may apply |

Table II
Comparison Of Output Explanations Across Reviewed Systems Vs. Neuroroute
domains (e.g., credit scoring) justify a decision rather than just stating it.

| System | How the Output Is Explained to the User | Strength | Weakness |
|---|---|---|---|
| CARES | No explicit explanation layer. Recommendations are presented incrementally, one point of interest (POI) at a time, keeping the reasoning implicit in the step-by-step flow. | Reduces cognitive overload by avoiding too much information at once. | Users are not explicitly told why one place was selected over another. |
| PIUMA | Recommendations are displayed as a ranked list with sensory tags (e.g., "quiet" or "crowded") based on crowdsourced ratings. | Sensory tags provide simple, human-readable justifications. | Relies on the quality of crowdsourced data and does not explain why the algorithm ranked one option above another. |
| PIUMA (Participatory Study) | Uses the same tag-based approach, refined through user feedback about the information they prefer to see. | Highlights users' preference for simple, visual cues rather than raw numbers or charts. | Provides descriptive labels rather than generated explanations of the recommendation process. |
| SmartRoutes | Provides short, spoken, turn-by-turn instructions using text-to-speech (TTS), designed for cognitive accessibility. | Easy to understand and follow while navigating. | Explains *how* to move rather than *why* a particular route was selected over alternatives. |
| Google Maps | Presents a highlighted route and estimated time of arrival (ETA), without a dedicated explanation of sensory-based route selection. | Fast, familiar, and easy to use. | Does not provide an explicit, factor-level explanation of sensory trade-offs between routes. |
| NeuroRoute | Generates plain-language explanations using a large language model (LLM), such as Gemini, grounded in the computed Sensory Score. Explanations identify factors such as noise, crowds, or construction that influence route comfort. | Provides reasoned, factor-level justifications rather than simple labels or instructions. | Explanation quality depends on the availability, configuration, and reliability of the LLM component. |

## III. METHODOLOGY

NeuroRoute is designed as a closed-loop, sensory-aware navigation system that combines rule-based scoring, supervised machine learning, and an agentic decision layer to recommend routes suited to a user's sensory profile. The overall system architecture is shown in Fig. III.1. The user first selects a sensory profile — Autistic, Elderly, or General User and specifies routing preferences such as avoiding crowds, noise, bright lights, or construction. Given a source and destination, the routing layer queries OpenStreetMap data through OSRM/GraphHopper to generate up to three geometrically distinct candidate routes. For each candidate, environmental data across six factors — traffic, crowd density, noise, brightness, construction activity, and weather — is collected and combined with the user's preference weights to compute a rule-based sensory score, after which a trained machine learning model predicts a personalized comfort/stress rating conditioned on the user's profile. An agentic decision engine then selects the single best route and generates a natural-language explanation, after which the route is presented and navigation begins. During navigation, the system continuously monitors live conditions; if the sensory score of the active route degrades beyond an acceptable range, the agent re-evaluates alternatives and suggests a better route. Independently, if the live sensory score at the user's position indicates overload, a second agent recommends nearby quiet places where the user can pause. Post-journey feedback is stored and used to periodically retrain the ML model. Each candidate route is scored using a weighted linear combination of six environmental factors, as given in Equation (1):

$$S = W_1(\text{Traffic}) + W_2(\text{Crowd}) + W_3(\text{Noise}) + W_4(\text{Brightness}) + W_5(\text{Construction}) + W_6(\text{Weather}) \quad (1)$$

where S is the computed sensory score for a road segment, normalized to a 0–100 scale, with lower values indicating a calmer, more sensory-friendly path. $W_1$ through $W_6$ are the weighting coefficients assigned to each environmental factor (Traffic, Crowd, Noise, Brightness, Construction, and Weather, respectively), such that $\sum_{i=1}^{6} W_i = 1$. These weights are not fixed; they are drawn from a base profile corresponding to the user category (e.g., Autistic User, Elderly User, or General User) and are further adjusted according to the individual's stated sensory preferences (e.g., avoidance of crowds, noise, bright lights, or construction), after which they are re-normalized to preserve the unity constraint. where each factor is measured on a 0–10 scale per road segment, sourced in production from services such as TomTom/Google Roads (traffic), time-of-day and streetlight density (brightness), and OpenWeatherMap (weather). The weight vector W is drawn from a base profile Autistic User, Elderly User, or General User each prioritizing different factors (e.g., the Autistic User profile weights Crowd and Noise most heavily, while the Elderly User profile weights Construction and Traffic more heavily). Where the user has flagged a specific preference, the corresponding weight is boosted by a factor of 1.5 and the full vector is re-normalized to sum to 1, so that preferences shift the relative importance of factors rather than inflating the score itself. To obtain a route-level score, the route geometry is sampled at 12 evenly spaced points, each is scored against the weight vector, and the mean is scaled to a 0–100 range for comparison and ML input. The final route selection is performed by an agent implemented as a four-node LangGraph state graph, structured as: compare routes, apply ML prediction, select the best route, and explain the decision. The candidate routes are first ordered by their rule-based sensory score. A Random Forest classifier, trained offline on the six environmental factors together with a one-hot encoded user profile, then predicts a categorical comfort label Comfortable, Moderate, or Stressful for each route. The predicted label and the sensory score are combined through a deterministic weighted rule, $R_{rank} = (C_{rank} \times 40) + S$, where $C_{rank} \in \{0,1,2\}$, to select the final route; this step is kept free of any language model so that the core decision remains reproducible and auditable. Only the final node invokes a language model (Gemini 3.8 Flash), prompted with the chosen route's sensory score, profile, and preferences to generate a short, plain language justification, with a deterministic fallback used if the model returns an empty response. This division keeps every consequential decision scoring, classification, selection fully deterministic, while confining generative AI strictly to the explanation layer.

A second, narrower agent handles sensory overload during navigation. While the user is navigating, the application samples the sensory score at the user's live GPS position every 20 s, using the same weighted formulation as Equation (1) applied to a single point rather than a route. When this score exceeds an overload threshold τ, the Node.js backend retrieves candidate rest locations from OpenStreetMap through the Overpass API: named points of interest tagged as parks or gardens, libraries, places of worship, viewpoints, cafes, or restaurants within 1.5 km of the user. Requests are raced across multiple public Overpass mirrors and cached for five minutes; if every mirror fails, a small static list of known locations is used, restricted to the same radius so that a distant place is never presented as nearby. The candidates are forwarded to the FastAPI agent service, which is implemented as a two-node LangGraph state graph: rank the candidates, then explain the top choice.

Each candidate place p receives a quietness score

$$Q(p) = B(c_p) + \frac{d_p}{20}, \qquad (2)$$

where $d_p$ is the straight-line (haversine) distance in metres between the user and p, and $B(c_p)$ is a base score set by the place category $c_p$: 10 for parks and gardens, 15 for libraries, 20 for places of worship, 25 for viewpoints, 35 for cafes, and 45 for restaurants. Lower values of Q are preferable. The distance term adds one point per 20 m, so a much closer place of a slightly noisier category can outrank a distant quiet one, reflecting that a user in overload needs relief nearby. The five lowest-scoring candidates are returned in ascending order of Q in Equation (2). As in the route agent, this ranking is deterministic and contains no language model. The base scores are hand-set ordinal assumptions rather than learned values: no labelled dataset of per-venue quietness exists, so training a model would require fabricating labels, and a transparent rule was preferred. The final node calls the language model (Gemini 3.8 Flash) once to produce a one-sentence, plain-language suggestion for the top-ranked place, with a 10 s timeout, a single

retry, and a templated fallback sentence so that a language-model failure never blocks the recommendation. After a suggestion is dismissed, further suggestions are suppressed for three minutes.

To validate the core assumptions behind the system before deeper technical investment, a preliminary feasibility survey was conducted with 13 autistic children and adolescents, with the assistance of parents, caregivers, or NGO staff. The survey used simple questions, pictures, colours, and a threepoint response scale to accommodate differences in age and communication ability across participants.

The key findings of the feasibility survey are shown in Fig. III.3. The survey identified noise as the most common sensory trigger, while strong demand was observed for quietplace recommendations and quieter-route suggestions. Individual differences in brightness sensitivity were also observed, supporting the need for configurable sensory preferences rather than a single fixed profile. All participants reported familiarity with smartphones or tablets, supporting the feasibility of a mobile-first application for this population.

The sensory-trigger responses obtained from the survey are presented in Fig. III.4, which breaks down participant reactions to three common sensory triggers — crowded places, noisy environments, and bright or flashing lights — across a threepoint scale (Happy/Okay, not sure/In-between, Upset/Don't like it). Noisy environments elicited the strongest negative reaction, with 62% of participants reporting distress, followed by crowded places at 46%. Reactions to bright or flashing lights were more evenly distributed, with 31% reporting distress and 23% reporting no distress, reinforcing that sensory sensitivity

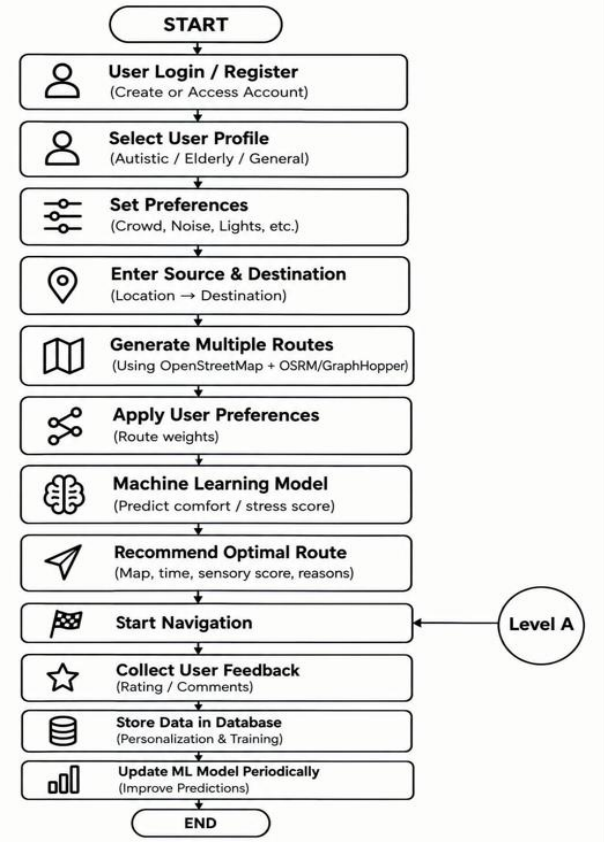



Fig. III.1. System architecture of NeuroRoute.

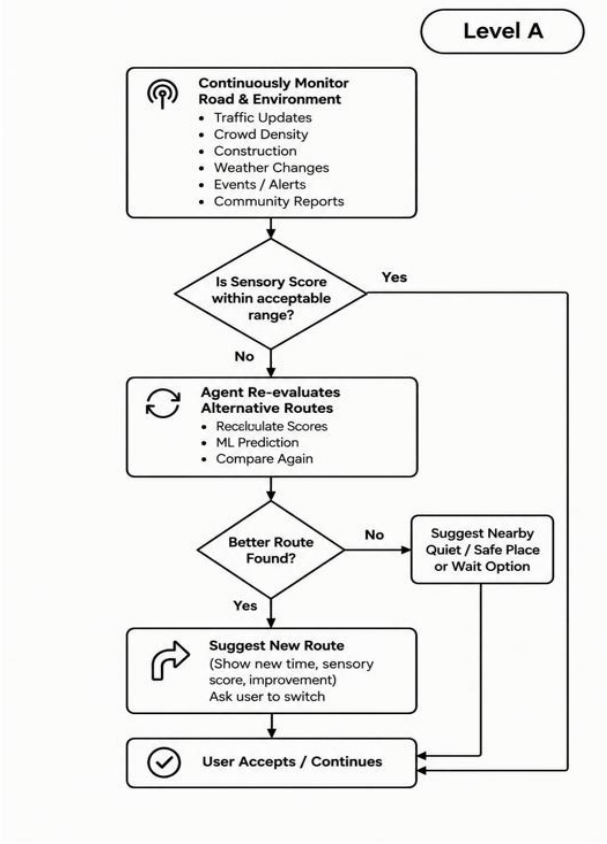



Fig. III.2. Decision flow of NeuroRoute.

Overall, the feasibility study indicates that NeuroRoute is practical to develop as a student-level navigation application. The selected technologies are suitable for implementing the proposed system, while real-time environmental-data availability remains an important limitation of the current implementation.

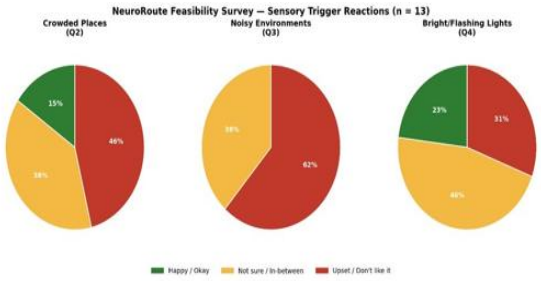



Fig. III.3. NeuroRoute feasibility survey — key findings (n = 13).

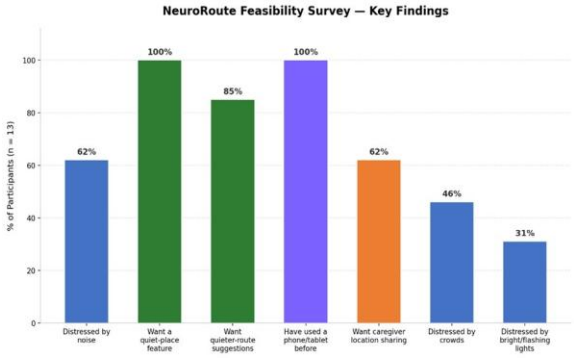



Fig. III.4. NeuroRoute feasibility survey — sensory trigger reactions (n =13).

## IV. RESULTS AND ANALYSIS

This section presents the results obtained from the current implementation of NeuroRoute. Four screens are used as evidence: the home/profile screen, the route-selection screen showing the recommended route, and two navigation screens, one showing the quiet-place card displayed on sensory overload and one showing an adaptive rerouting prompt. Each result is explained in simple terms, followed by an analysis of what it demonstrates about the system's working. The section then reports a comparative evaluation of NeuroRoute against Google Maps over 30 paired trials (Subsection IV-F).

### A. Home Screen and Area Comfort Output

When the user" Surabhi" (Autistic User profile) opens the app, the home screen shown in Fig. IV.1 presents two things before any route is even requested:

- Current area comfort, shown as" Calm — low crowd, low noise." This is a live, one-line summary of the user's present surroundings, calculated from the same environmental factors used for route scoring.
- A simple destination search box with a" Find a comfortable route" button, plus a note that the active preferences are" Avoiding crowds, noise."

Analysis: This screen shows that sensory information is not only used at the routing stage — it is surfaced immediately, before the user even asks for directions. This matches the requirement where sensory preferences are meant to influence the experience continuously, not just at the moment of route comparison. The" Nearby quiet places: No quiet places found nearby" message also confirms that the quiet-place discovery feature (Overpass API) is functionally wired in, even though no result was found for this particular location — an honest empty-state, not a broken one.

### B. Route Generation and Sensory Scoring Output

For the same source-destination pair, the system generated three candidate driving routes through OSRM, shown

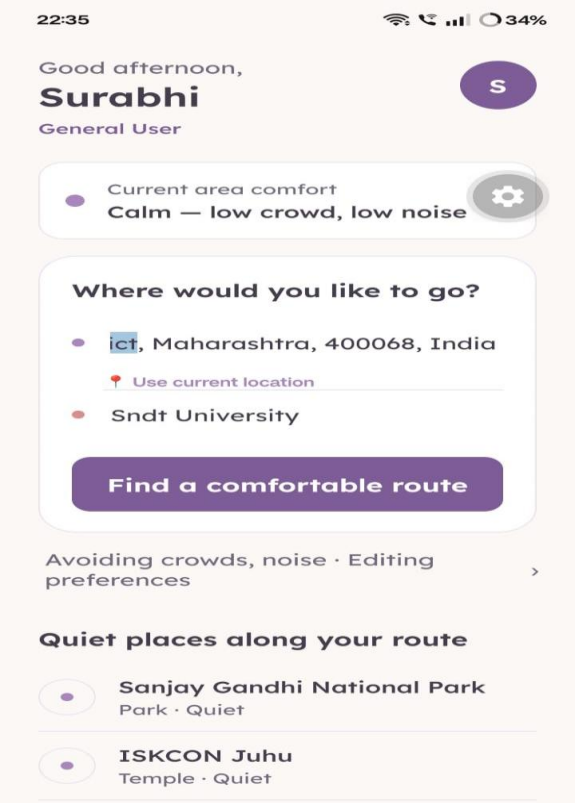



Fig. IV.1. Home screen showing user greeting, current area comfort, and route search.in Fig. IV.2, and each one was passed through the rulebased sensory-scoring engine. The results are summarized in Table III.

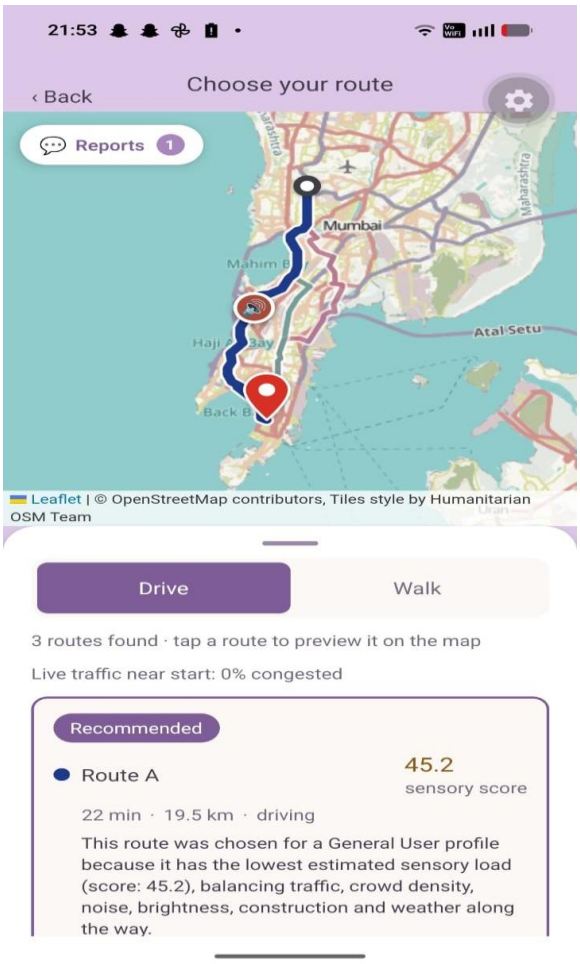



Fig. IV.2. Route options screen with map and Recommended Route A.

Table III
Route Comparison Output for A General User Profile

| Route | Sensory Score | Time | Distance | Label shown to user |
|---|---|---|---|---|
| Route A | 46.6 (lowest) | 21 min | 23 km | "Recommended" — chosen for a General User profile |
| Route B | 50.2 | 33 min | 30.9 km | "A moderate option — some sensory factors are higher here." |
| Route. C | 51.5 (highest) | 42 min | 32.3 km | Alternative, not recommended |

The map also shows a live-traffic reading near the start point ("32% congested"), confirming that a real-time traffic factor was included in the scoring, not just static map data.

C. Analysis of the Route Selection Logic

This result shows that Route A was selected as the recommended route for the General User profile because it achieved the lowest sensory score of 46.6. In this example it also happened to have the shortest travel time (21 minutes), whereas Route B, with a moderate sensory score of 50.2, takes 33 minutes, and Route C, with the highest sensory score of 51.5, takes 42 minutes. The recommendation, however, follows the Sensory Score (together with the machine-learning comfort label described in Section III), not the travel time: the comparative trials in Subsection IV-F include many cases in which the recommended route is not the fastest. Routes B and C are presented as alternative options.

In simple terms, this shows that NeuroRoute does not simply pick the quickest path the way a conventional map app would. It makes a trade-off: it accepts a longer travel time in exchange for a calmer, less overwhelming journey. This is exactly the behaviour the project set out to achieve treating sensory comfort as being just as important as speed, not as an afterthought.

The three scores (46.6, 50.2, 51.5) are also close to each other rather than wildly spread apart. This is expected, since all three routes pass through broadly the same city area and are subject to similar general conditions (weather, general traffic patterns for that time of day). The scoring engine is still able to detect and rank the smaller differences in the specific roads and environments each route passes through, which shows that the weighted formula is sensitive enough to separate routes even when the overall conditions are similar.

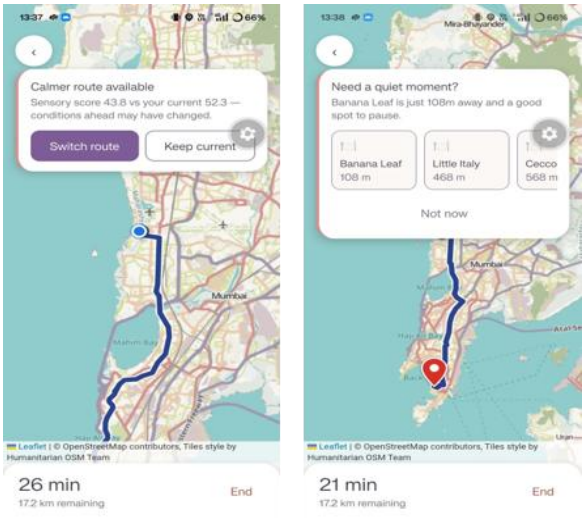
The fact that a clear "Recommended" label was correctly attached to the route with the numerically lowest score, across all three candidates, confirms that the comparison and ranking logic is working correctly end-to-end: OSRM route generation → sensory scoring → comparison → labelling → display.

All three route scores (46.6, 50.2 and 51.5) are below 65. This is why no quiet-place prompt appeared on the routeselection screen. The prompt only appears during navigation, when the score crosses the threshold (see Subsection IV-D).

D. Navigation Support: Quiet Places on Sensory Overload and Adaptive Rerouting

NeuroRoute checks the sensory conditions while the user is travelling. The score is not fixed for the whole trip: it changes as the user moves through areas with different traffic, crowd and noise levels. Two behaviours observed during one navigation session (screenshots taken one minute apart) are shown in Fig. IV.3.

1) Quiet-place card on sensory overload: While the user is navigating, the application samples the Sensory Score at the user's live GPS position every 20 s, using the weighted formula of Equation (1) applied to a single point instead of a whole route. When this score exceeds the overload threshold ($\tau = 65$), the system treats the situation as high sensory load and, without the user having to ask, displays the "Need a quiet moment?" card on top of the map (Fig. IV.3(b)). The candidate places are retrieved through the Overpass API and ranked by the quietness score Q(p) of Equation (2), and the language model then writes a one-sentence suggestion for the top-ranked place. In Fig. IV.3(b) the card reads "Banana Leaf is just 108m away and a good spot to pause" and lists three candidate places with their distances. Table IV gives their quietness scores calculated from Equation (2).




(a) (b)

Fig. IV.3. Navigation screens. (a) Adaptive rerouting prompt: "Calmer route available" with Sensory Score 43.8 against the current 52.3. (b) "Need a quiet moment?" card with nearby quiet places, shown when the live Sensory Score exceeds the overload threshold.

Table IV

Quiet-Place Candidates Shown on the Card and Their Quietness Scores $Q(P) = B(C_p) + D_p/20$ (Lower Is Better) [a]Name truncated on screen.

| Rank | Place | Category (B) | $d_p$ (m) | Q(p) |
|---|---|---|---|---|
| 1 | Banana Leaf | Restaurant (45) | 108 | 50.4 |
| 2 | Little Italy | Restaurant (45) | 468 | 68.4 |
| 3 | Cecco...[a] | Restaurant (45) | 568 | 73.4 |

Each card shows the name and distance of the place, so the user can choose the closest or the most suitable one, and a "Not now" option lets the user dismiss the card and continue (further suggestions are then suppressed for three minutes). The navigation bar at the bottom stays visible with the remaining time (21 min), the remaining distance (17.2 km) and an "End" button, so the journey is not interrupted.

Analysis: The card is automatic and driven by the score: it appears only when the live Sensory Score at the user's position crosses the overload threshold, and not at the planning stage. This matters for sensory-sensitive users, who may find it hard to notice or ask for help when overwhelmed. The order of the cards matches the ascending order of Q(p) in Table IV. Because all three places belong to the same category, they differ only through the distance term, which adds one point per 20 m, so the nearest place is ranked first. The formula also lets a much closer place of a slightly noisier category outrank a distant quiet one, reflecting that a user in overload needs relief nearby. The prompt is short, offers a choice instead of forcing a detour, and matches the project goal of reducing stress, not only avoiding it in advance.

2) Adaptive rerouting: The same monitoring supports reevaluation of the active route. In Fig. IV.3(a) the application shows "Calmer route available": an alternative route has a Sensory Score of 43.8 against 52.3 for the route currently followed, a reduction of 8.5 points (about 16%), with the note that conditions ahead may have changed. Two buttons, "Switch route" and "Keep current", leave the decision with the user, so the system suggests but never forces a change, and the bottom bar again shows the remaining time and distance (26 min, 17.2 km). This screen demonstrates that the route score is recomputed during the journey and compared with the alternatives, as described in the methodology, rather than being fixed at the moment the route is chosen.

E. Testing Performed

Because the current sensory-scoring engine is a rule-based weighted model, its correctness cannot be measured using classification metrics like accuracy, precision, or recall. Instead, its correctness was verified through functional and outputconsistency testing, that is, by checking whether the system behaves exactly as the defined rules say it should. The test cases and results are listed in Table V.

This is a form of black-box functional testing: the internal formula was not inspected directly in this test round, but the input (source, destination, profile) and the output (routes, scores, labels) were checked against the expected rule ("lower score should always be ranked first"). All observed outputs matched the expected rule with no exceptions across the tested scenario.

Table V Functional Test Cases and Observed Results

| Test Case | Expected Behaviour | Observed Result | Status |
|---|---|---|---|
| Generate multiple candidate routes for one source-destination pair | 2–3 distinct routes returned by OSRM | routes returned (A, B, C) | Pass |
| Compute sensory score for each route | Each route gets a distinct numeric score | Scores of 54.6, 57.4, 58.5 computed | Pass |
| Mark the lowest-scoring route as recommended | Lowest score = "Recommended" label | Route A (54.6, lowest) labelled Recommended | Pass |
| Show live traffic condition near the start point | Traffic percentage displayed | "20% congested" displayed | Pass |
| Display area comfort on home screen without a route request | Comfort label shown using currentlocation data | "Calm — low crowd, low noise" shown | Pass |
| Handle a location with no nearby quiet places | Show empty-state message instead of an error | "No quiet places found nearby" shown | Pass |
| Score stays below 65 during navigation | No quiet-place prompt shown | No prompt shown | Pass |
| Score goes above 65 during navigation | Quiet-place card appears on the map | Card appeared with nearby places | Pass |
| Places listed with distances | Each place shows type and distance | Three restaurants shown with distances in metres, ordered by quietness score | Pass |
| User taps "Not now" | Card closes, navigation continues | Card closed, route unchanged | Pass |
| A lower-scoring route becomes available during navigation | Prompt offers "Switch route" or "Keep current" | "Calmer route available" (43.8 vs. 52.3) shown with both options | Pass |

F. Comparative Evaluation against Google Maps

1) Test Design: To assess how NeuroRoute's recommendations compare with those of a conventional navigation application, the system was tested against Google Maps (driving mode) on ten origin–destination pairs in the Mumbai region, from short urban trips of about 2 km to cross-city trips of about 50 km (Table VI). Every pair was tested at three times of day on 12 October 2026: morning (about 07:00–07:30), afternoon (about 14:10–14:45) and evening (about 22:20–22:56). This gave 30 paired trials, documented by 80 screenshots taken within a few minutes of each other in the two applications. All NeuroRoute trials used the General User profile.

For each trial we recorded the time, distance, Sensory Score and estimated traffic delay of each of the three NeuroRoute routes (A, B, C), the live congestion near the start point, and the time and distance of the default (fastest) Google Maps route. In all 30 trials the route marked "Recommended" by NeuroRoute was Route A and was also the route with the lowest Sensory Score, so Route A is taken as the NeuroRoute result in every comparison. Paired differences were tested with the paired t-test and the Wilcoxon signed-rank test, and

effect size was expressed as Cohen's $d_z$. The full transcribed dataset is available as a spreadsheet.[1]

2) Sample Trip: Fig. IV.4 shows one of the 30 trials (Goregaon East to Churchgate, 07:15). NeuroRoute returned Route A (68 min, 28.6 km, Sensory Score 47.5), Route B (52 min, 28.6 km, 47.8) and Route C (87 min, 38.9 km, 56.1), each with an estimated traffic delay and, for the recommended route, a plain-language reason for the choice. Route A was recommended because its score was the lowest, although only by 0.3 points over Route B, which was 16 min faster over the same distance. Google Maps returned a single fastest route of 39 min over 32 km with a toll of INR 100 and no sensory information, so NeuroRoute's recommendation was 29 min

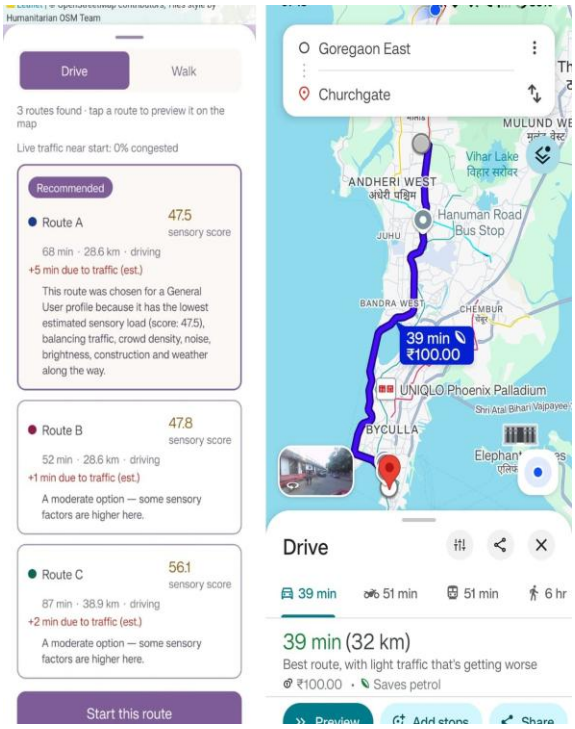



Fig. IV.4. Sample trial (Goregaon East to Churchgate, 07:15). Left: NeuroRoute route options with time, distance, Sensory Score and traffic delay; Route A is recommended. Right: Google Maps fastest route (39 min, 32 km).

(74%) slower. The trial illustrates the general pattern reported below: NeuroRoute selects by Sensory Score and accepts a time cost, and this trade-off is visible only in NeuroRoute because Google Maps exposes no comparable measure.

3) Travel Time: Table VI lists the travel time of every trial and Fig. IV.5 shows the mean per route. NeuroRoute's recommended route took longer than the Google Maps route in 25 of the 30 trials (Table VII). The mean time was 57.5 min for NeuroRoute and 38.2 min for Google Maps, a mean difference of 19.3 min (95% CI 10.3–28.3 min; paired $t(29) = 4.39$, $p = 0.0001$; Wilcoxon $p < 0.001$; $d_z = 0.80$, a large effect). The two sets of times are strongly correlated (log–log

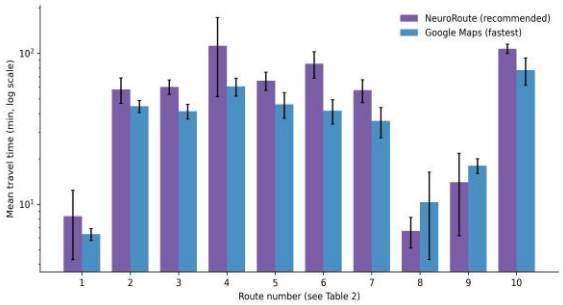



Fig. IV.5. Mean travel time per route, averaged over the three time slots (error bars: standard deviation; log scale). Route numbers as in Table VI. $r = 0.92$, Fig. IV.7), so NeuroRoute orders trips by duration much as Google Maps does, with an added time offset. The result is robust to the largest outlier (Dahisar to Siddhivinayak, morning, 175 min): without it the mean difference is 16.1 min (Wilcoxon $p < 0.001$).

4) Time of Day, Route Length and Distance: Both systems show the same daily pattern (Fig. IV.6). Google Maps times are highest in the afternoon (43.2 min) and lowest in the evening (32.9 min), and NeuroRoute's own traffic-delay estimate for its recommended route is much larger in the afternoon (mean 14.5 min) than in the morning (4.1 min) or evening (3.0 min). The gap between the two applications is similar in all three slots, and no significant effect of time of day on the percentage difference was found (Kruskal–Wallis $p = 0.50$; the test has low power with 10 trials per slot).

[1] Dataset spreadsheet: https://docs.google.com/spreadsheets/d/1rlFfWOXoxc0dJ3AiEwkgD8daaS34YAp5-QJy-8LXyqQ/edit.

Route length is the main moderator of the time cost (Table VII). On short trips of 10 km or less NeuroRoute was on par with Google Maps (9.7 vs. 11.6 min; slower in 4 of 9 trials), and all five trials in which it was faster occurred here. On long trips it was slower in all 21 of 21 trials, by 28.4 min on average. Distances point the same way: the NeuroRoute route averaged 22.0 km against 19.8 km for Google Maps (28 trials with comparable distances; mean +2.2 km, paired t-test $p = 0.009$, Wilcoxon $p = 0.013$), and only 8 of the 28 pairs were within 10% of each other, consistent with NeuroRoute often choosing a different, longer path.

5) Sensory Load: Google Maps provides no sensory measure, so sensory load can only be analysed within NeuroRoute (Table VIII, Fig. IV.8). The recommended route had the lowest Sensory Score of the three in all 30 trials. Its mean score was 48.7 against 52.5 for the two alternatives, a reduction of 3.8 points (7.3%; paired $t(29) = 9.85$, $p < 0.001$; Wilcoxon $p < 0.001$), with a range of 0.9–10.3 points per trial. Scores were lowest in the evening (47.6) and similar in the morning (49.5) and afternoon (49.0). Two further findings matter. First, the extra minutes relative to Google Maps were unrelated to the size of the sensory benefit (Spearman $\rho = -0.03$, $p = 0.87$), so a longer route did not buy a calmer one. Second, the fastest of NeuroRoute's own three options averaged 44.4 min and matched or beat Google Maps in 10 of 30 trials (within 10% in 13 of 30), which suggests that a time-aware ranking could

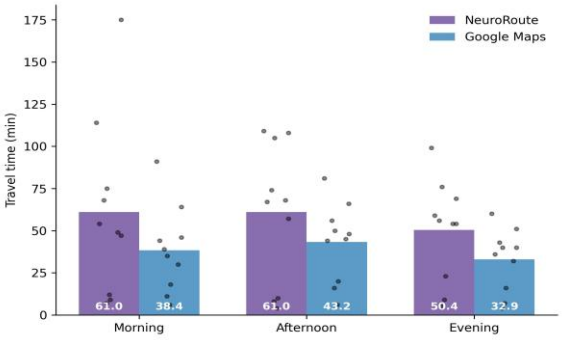



Fig. IV.6. Mean travel time by time of day for NeuroRoute and Google Maps. Dots are individual routes.

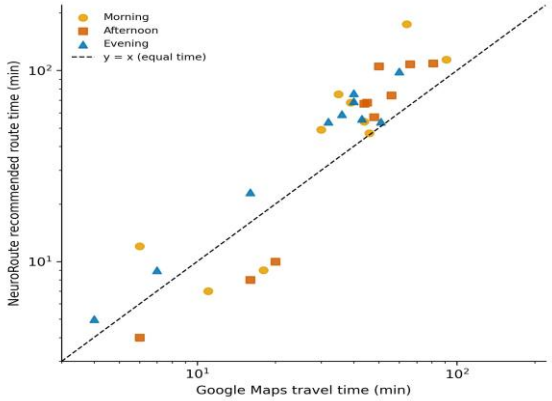



Fig. IV.7. NeuroRoute recommended-route time against Google Maps time, one point per trial (log–log axes; dashed line: equal time; $r = 0.92$). Points above the line are trials where NeuroRoute was slower.
recover much of the gap.

6) Limitations of the Comparison: The comparison covers a single day with one run per route–time pair, and the 30 trials come from 10 routes, so trials from the same route are not fully independent and the p-values should be read as indicative. The two applications sometimes used slightly different start or end labels (e.g., "Dahisar" and "Dahisar East"), so distances are only approximately comparable. In two evening trials (SNDT to Marine Drive, and Vasai Creek to Marine Drive) the Google Maps screenshot showed a slower non-toll route; the fastest toll-route time (32 and 60 min) was used and the distance was excluded. One evening Google Maps reading (Vasai Creek to Dahisar) was taken offline without traffic data. Two NeuroRoute values look anomalous (Dahisar to Siddhivinayak, morning: 175 min for Route A, and 26 min for 29.2 km on Route C); they were kept as displayed. Live congestion near the start read 0% in 24 of 30 trials, so it carried little information. Finally, Google Maps routes were not scored on the Sensory Score scale, so a like-for-like

Table VI

Travel Time (Min) Per Route For Morning / Afternoon / Evening. Nr = Neuroroute Recommended Route, Gm = Google Maps Fastest Route. A Positive Δ Means Neuroroute Is Slower

| # | Route | NR time (min) | GM time (min) | Δ time (%) |
|---|---|---|---|---|

| | | | | |
|---|---|---|---|---|
| 1 | Anand Nagar → Dahisar Stn | 12 / 4 / 9 | 6 / 6 / 7 | +100 / −33 / +29 |
| 2 | Anand Nagar → SNDT | 47 / 57 / 69 | 46 / 48 / 40 | +2 / +19 / +72 |
| 3 | Borivali → SNDT | 54 / 67 / 59 | 44 / 44 / 36 | +23 / +52 / +64 |
| 4 | Dahisar → Siddhivinayak | 175 / 108 / 54 | 64 / 66 / 51 | +173 / +64 / +6 |
| 5 | Goregaon East → Churchgate | 68 / 74 / 56 | 39 / 56 / 43 | +74 / +32 / +30 |
| 6 | SNDT → Gateway of India | 75 / 105 / 76 | 35 / 50 / 40 | +114 / +110 / +90 |
| 7 | SNDT → Marine Drive | 49 / 68 / 54 | 30 / 45 / 32 | +63 / +51 / +69 |
| 8 | SNDT → Santacruz Stn | 7 / 8 / 5 | 11 / 16 / 4 | −36 / −50 / +25 |
| 9 | Vasai Creek → Dahisar | 9 / 10 / 23 | 18 / 20 / 16 | −50 / −50 / +44 |
| 10 | Vasai Creek → Marine Drive | 114 / 109 / 99 | 91 / 81 / 60 | +25 / +35 / +65 |

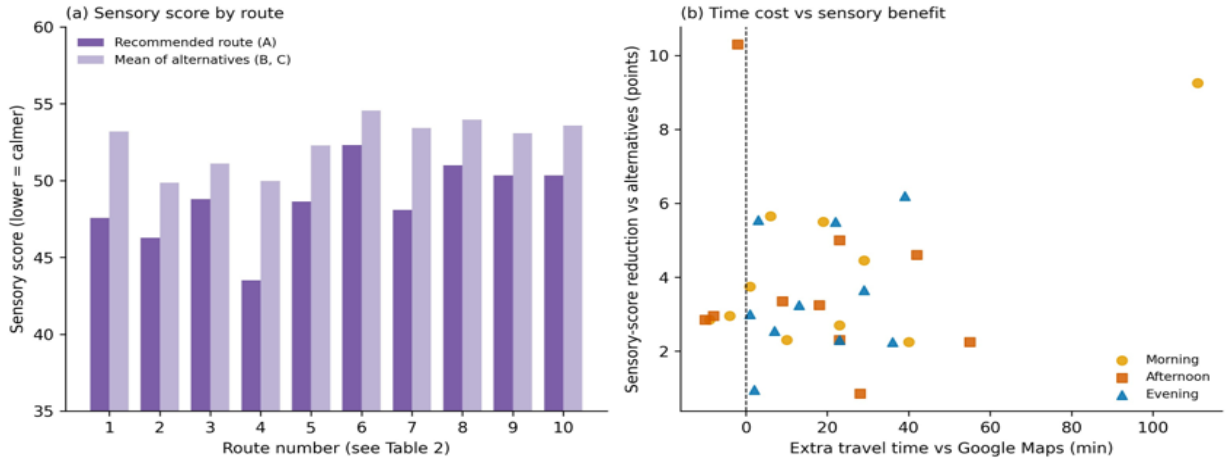



Fig. IV.8. (a) Mean Sensory Score of the recommended route against the mean of the alternatives, by route (numbering as in Table VI). (b) Extra travel time relative to Google Maps against Sensory Score reduction, one point per trial.

Table VII, Table VIII
Paired Comparison of Mean Travel Time (Min) By Time Slot andMean Sensory Score (Lower Is Calmer) And Travel Time Per
Route LengthRoute, Averaged Over the Three Time Slots. Nr-A =

Recommended Route; B, C = Alternatives

sensory comparison remains future work.

rently synthetic rather than sensor-measured. The logic of the scoring and ranking is validated by these results; the real-world precision of the underlying environmental readings will improve once live sensory data sources are integrated.

2) Small-scale manual testing.The test cases above were run manually on individual scenarios rather than automated across a large batch of source-destination pairs. Broader, automated testing across many routes is planned as future work.
3) Fixed threshold. The value of 65 is a rule set by the project team. It is not clinically validated, and different

G. Discussion on Reliability and Current Limitations

A few honest limitations should be stated alongside these results:

1) Synthetic environmental data. Reliable real-time roadlevel data for crowd density, noise, and construction is not always available from external sources, so some of the environmental inputs behind these scores are cur-

users may need different thresholds. Personalised thresholds are planned as future work.

H. Summary of Results

The results demonstrate that the core sensory-aware routing pipeline works as intended: routes are generated, scored, ranked, and the lowest-sensory-score route is correctly and consistently recommended to the user even when it is not the fastest option. This confirms the central claim of the project, that a navigation system can weigh comfort alongside speed, and functional testing across the tested scenarios

showed no mismatches between the defined rule and the displayed output. The system also monitors the sensory score during navigation. When the live score at the user's position exceeds the overload threshold of 65, it automatically offers the nearest quiet places ranked by the quietness score, and when a calmer alternative route becomes available it offers the user a choice to switch, so support is given when the user needs it and not only at the planning stage. The comparison with Google Maps over 30 trials quantifies the trade-off: the recommended route lowered the Sensory Score by about 7% relative to NeuroRoute's alternatives, and on average took 19.3 min longer than the Google Maps route, with no time penalty on short trips but a consistent one on long trips.

## V. CONCLUSION

This paper presented NeuroRoute, a personalized navigation assistance system that incorporates sensory comfort alongside conventional travel efficiency for autistic, elderly, and general users. Building on gaps identified in related sensory-aware and cognition-aware navigation research PIUMA's crowdsourced Safe POI recommender [1], its participatory-design follow-up study [2], and SmartRoutes' indoor localization and emergency guidance system [3] NeuroRoute combines conventional GPS-based route generation with a rule-based Sensory Scoring Framework, profile- and preference-driven personalization, nearby quiet-place discovery, and an architecture for Machine Learning-based comfort prediction and Agentic AI-driven adaptive route selection. The system was implemented using React Native (Expo), Node.js/Express, MongoDB, OpenStreetMap, Nominatim, OSRM, and the Overpass API, with a supporting Python service for the Random Forest comfort model and a LangGraph-based agent workflow. A preliminary feasibility study conducted with 13 autistic children and adolescents at a partner NGO validated the core assumptions behind the system, confirming strong interest in quiet-place discovery and quieter-route recommendations, and reaffirming that sensory sensitivity varies considerably between individuals. A comparative test against Google Maps over 30 paired trials showed that the recommended route lowered the Sensory Score by about 7% relative to NeuroRoute's alternatives, at an average cost of 19 minutes in travel time that was concentrated on long trips. While the current implementation remains at the prototype stage — with realtime environmental sensing, full Agent-to-backend integration, and dynamic rerouting still under development — NeuroRoute establishes a functional foundation for sensory-aware outdoor navigation and demonstrates the technical feasibility of combining mainstream mapping and routing infrastructure with personalized, accessibility-focused route evaluation. Future work will focus on real-time environmental data integration, a time-aware route ranking, complete Agentic AI deployment, improved pedestrian and offline routing, wearable-sensor integration, and larger-scale evaluation with a more diverse user base.

## ACKNOWLEDGMENT

The authors have a great pleasure to express our gratitude to our guide and all those who have contributed and motivated us during our project work.

## REFERENCES

[1] F. Cena, A. Rapp, and C. Mattutino, "Personalized spatial support for people with autism spectrum disorder," in *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization (UMAP '18 Adjunct)*, Singapore, Jul. 2018, pp. 233–238.

[2] S. Brighenti, F. Cena, and R. Keller, "Personalized interactive urban maps for autism: An innovative approach involving persons with level 1 autism spectrum disorder to user-centered design," *Advances in Autism*, vol. 11, no. 1, pp. 52–66, 2025.

[3] M. T. García-Catalá, E. Martín-Barroso, M. C. Rodríguez-Sánchez, M. Delgado-Álvaro, and R. Novak, "Sensory navigation system for indoor localization and orientation of users with cognitive disabilities in daily tasks and emergency situations," *Sensors*, vol. 24, no. 22, Art. no. 7154, 2024.

[4] S. Brighenti, F. Cena, and R. Keller, "PIUMA participatory study," 2025. [Incomplete reference: replace with the actual publication details.]

[5] Expo, “Expo documentation,” [Online]. Available: https://docs.expo.dev/. [Accessed: 2026].
[6] Expo, “Location—Expo documentation,” [Online]. Available: https://docs.expo.dev/versions/latest/sdk/location/. [Accessed: 2026].
[7] Google, “Fused location provider,” [Online]. Available: https://developers.google.com/location-context/fused-location-provider. [Accessed: 2026].
[8] Google, “Google Maps Platform documentation,” [Online]. Available: https://developers.google.com/maps/. [Accessed: 2026].
[9] LangChain, “LangGraph documentation,” [Online]. Available: https://langchain-ai.github.io/langgraph/. [Accessed: 2026].
[10] OpenStreetMap Foundation, “OpenStreetMap,” [Online]. Available: https://www.openstreetmap.org/. [Accessed: 2026].
[11] Project OSRM, “OSRM: Open-Source Routing Machine,” [Online]. Available: https://project-osrm.org/. [Accessed: 2026].