

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Experiments with RouteNav, A Wayfinding App for Blind Travelers in a Transit Hub

Permalink

<https://escholarship.org/uc/item/1gv8z9jt>

Authors

Ren, Peng

Lam, Jonathan

Manduchi, Roberto

et al.

Publication Date

2023-10-22

DOI

10.1145/3597638.3608428

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Experiments with RouteNav, A Wayfinding App for Blind Travelers in a Transit Hub

Peng Ren\*  
Jonathan Lam\*  
University of California, Santa Cruz  
Santa Cruz, California, USA  
pren1@ucsc.edu, jokwlam@ucsc.edu

Roberto Manduchi  
University of California, Santa Cruz  
Santa Cruz, California, USA  
manduchi@ucsc.edu

Fatemeh Mirzaei  
University of California, Santa Cruz  
Santa Cruz, California, USA  
fmirzaei@ucsc.edu

## ABSTRACT

RouteNav is an iOS app designed to support wayfinding for blind travelers in an indoor/outdoor transit hub. It doesn't rely on external infrastructure (such as BLE beacons); instead, localization is obtained by fusing spatial information from inertial dead reckoning and GPS (when available) via particle filtering. Routes are expressed as sequences of "tiles", where each tile may contain relevant points of interest. Redundant modalities are used to guide users to switching goalposts within tiles. In this paper, we describe the different components of RouteNav, and report on a user study with seven blind participants, who traversed three challenging routes in a transit hub while receiving input from the app.

## CCS CONCEPTS

• **Human-centered computing** → **Accessibility technologies** ; *Human-centered computing*.

## KEYWORDS

wayfinding, orientation, mobility

### ACM Reference Format:

Peng Ren, Jonathan Lam, Roberto Manduchi, and Fatemeh Mirzaei. 2023. Experiments with RouteNav, A Wayfinding App for Blind Travelers in a Transit Hub. In *The 25th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '23)*, October 22–25, 2023, New York, NY, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3597638.3608428>

## 1 INTRODUCTION

Wayfinding ("the process of navigating through an environment and traveling to places by relatively direct paths" [43]) can be challenging for everyone. The importance of wayfinding is reflected by the widespread diffusion of carefully designed signage [12, 59] in public spaces such as airports and transit hubs. Tactile wayfinding (e.g., tactile paving [9, 45], first introduced in Japan) is becoming increasingly pervasive, especially in Europe and Pacific Asia [51], and there has been some interest in acoustic wayfinding (e.g., soundscape design [27]). However, there is no doubt that visual signs

\*Both authors contributed equally to this research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ASSETS '23, October 22–25, 2023, New York, NY, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0220-4/23/10.

<https://doi.org/10.1145/3597638.3608428>

(which are inaccessible to blind or visually impaired travelers) have a paramount role in the current wayfinding infrastructure. We are particularly interested in navigation and wayfinding in transit hubs, such as large train/subway stations and intermodal transit centers with transfer terminals. Traveling with public transit can generate stress and fatigue, especially when transfers are involved and the traveler is not familiar with the route [13, 17]. Partly because of lack of accessible wayfinding information, use of public transit to reach unfamiliar destinations for people who are blind can be exceedingly difficult, stressful [21], and potentially dangerous.

Many different approaches have been proposed for wayfinding support. Generally speaking, a wayfinding system contains three main components: a map of the environment (though mapless wayfinding has also been proposed [26, 41, 64]); a localization system, which estimates the current position of the user in the map; and a user interface, which provides information to the user about where to move next, along with any appropriate contextual information. Note that the "map" may contain much more than just the floor plan of the environment. For example, for systems that use visual-based localization, a 3-D representation of the surfaces (along with their visual features) may be stored in the map [55]. Systems that use the signal strength (RSSI) from Wi-Fi [10] or Bluetooth Low Energy (BLE) beacons [66], or the characteristic local variations of the magnetic field [34], need to store a map associating spatial locations with RSSI signatures or magnetic field features.

This article describes the development of a wayfinding iOS app specifically designed for blind travelers in a transit hub. Dubbed RouteNav, this system is designed to be a comprehensive travel support app; we focus here on the within-hub wayfinding component. These are the main contributions of this work:

*Indoor-outdoor localization without additional infrastructure.* We introduce an indoor-outdoor localization system based on fusion of spatial data from inertial sensors (processed by a state of the art machine learning algorithm for dead reckoning, RoNIN [36]) and GPS, when available and reliable. Remarkably, RouteNav requires no external infrastructure for localization. While infrastructure such as high-density Bluetooth Low Energy (BLE) beacons can certainly improve localization accuracy, the installation and maintenance cost (which includes a laborious fingerprinting phase [4]) requires buy-in from the agencies managing the spaces. Proposals for infrastructure development targeting small communities (such as blind travelers) may receive enthusiastic initial response, but this does not often translate into sustained financial support [19]. Note that RoNIN does not use visual information from the smartphone's camera, meaning that users do not need to hold the phone upright to obtain a clear view of the scene.

*Tile-based navigation.* We introduce a novel representation of the space and routes to be traversed by means of *tiles*. A tile is a polygonal area of space whose extent is no smaller than the expected localization resolution. The traveler’s location at each time is robustly mapped to one tile, and all information provided to the user is based on the current tile. This form of “space discretization” helps manage localization inaccuracies, at the same time providing users with a simple and intuitive representation of their current location within and progression along the route.

*Redundant modalities user interface.* RouteNav’s interface was carefully designed using redundant modalities. To minimize the effect of poor localization in the directions given to the user, we developed a mechanism of switching goalposts. The direction to the target goalpost is produced via user-solicited mechanisms (through speech or haptic modalities), as well as unsolicited notifications (speech notifications produced when the user approaches or enters a tile, and a sound/vibration backdrop signaling whether or not the user is walking in the correct direction).

*User study.* We conducted a study with 7 blind participants, who tested RouteNav in very realistic and challenging conditions. For this study, we mapped most of the walkable area of the intermodal Palo Alto Transit Center in California. This includes the train station, serving the Palo Alto area, and a bus stop area with 10 slots, served by three different transit agencies. For our tests, we selected three different routes with start and end points located either at the train tracks, or at a bus stop. These routes (total length: 495 m) included crossing the tracks using an underground tunnel (a GPS-denied zone), walking along boarding areas and walkways, crossing crosswalks, walking on a ramp, and finding the opening in a fence. Of note, GPS localization was consistently poor in several areas of the train station. Compared to routes within buildings, which are normally along corridors and hallways, and are physically constrained by walls, the indoor/outdoors routes considered in our study are generally less constrained and often surrounded by open spaces, which makes navigation without sight quite challenging. While the localization module remained mostly unchanged through the course of the study, we iteratively perfected the user interface from feedback, comments and suggestions from our participants.

## 2 RELATED WORK

Accessible wayfinding apps that use GPS for localization have a long history. Back in the early 90s, Golledge and colleagues proposed using GPS for a personal navigation system for blind travelers [31, 32], along with auditory displays mechanisms (using spatialized sound or speech interface) to support guidance [44]. The first commercial accessible navigation system (Arkenstone Strider) appeared in 1994. With the advent of smartphones equipped with screen readers, blind travelers had access to a plethora of wayfinding apps, from the mainstream Apple Maps and Google Maps (which in 2019 implemented a detailed voice guidance option for walking trips) to more specialized apps such as GoodMaps Outdoors, APH Nearby Explorer, and BlindSquare. These systems rely on GPS for positioning, and on geographical information data set for guidance. Unfortunately, GPS localization is often inaccurate, and the GPS signal is too weak to be received indoors. Both Apple Maps and

Google Maps enabled a modality that improves localization accuracy by aiming the camera at a building, whose image is matched against a 3-D point data set (in locations for which such data sets are available).

Wayfinding systems for blind walkers moving in environments where GPS is unreliable can roughly be categorized according to whether or not they require additional infrastructure for localization. To the first category belong infrared beacons (e.g., TalkingSigns [11]), BLE beacons (e.g., NavCog [4, 33, 56] or GuideBeacon [15]), RFID tags (e.g., PERCEPT [29]), Ultra-Wide Band (UWB) beacons (e.g., SUGAR [48]), and visual markers (e.g., NaviLens color barcodes [47]). Systems that do not require infrastructure installation can be further subdivided, based on whether or not they require prior off-line calibration (often called *fingerprinting*), which can be expensive and may need to be repeated when environmental conditions change. These include Wi-Fi beaconing [2] and magnetic navigation [30, 52]. For example, Google Maps and Apple Maps provide indoors localization in selected venues based on signals from Wi-Fi and BLE beacons, with an accuracy of several meters. Note that computer vision systems that work by matching an image with an underlying 3-D map [55, 60], a strategy commercially pursued by GoodMaps, also belong to this category. Systems that require neither additional infrastructure nor fingerprinting may use dead reckoning from inertial sensors [6, 23, 49, 53] or visual/inertia SLAM [58] from cameras or depth sensors, possibly integrated with inertial sensors [18, 28, 42]. Wayfinding systems for blind travelers that specifically address mixed indoor/outdoor situations include CityGuide [14], which uses GPS and BLE beacons located at specific points of interest inside a building and at the building’s entrances and was tested with 6 participants with visual impairment, and ARIANNA [20], which uses computer vision and inertial sensors (no human testing was reported). Visual SLAM can typically provide more accurate localization than dead reckoning from inertial sensors, but it requires continuously unoccluded view of the scene from the camera, which can be challenging in crowded environments, or if the camera is contained in a smartphone (e.g., it won’t work if one keeps the phone in their pocket). Wayfinding systems that do not require a map of the environment, but guide the walker through a path already taken, have also been proposed [26, 64].

The design of an appropriate, accessible user interface is of paramount importance for navigation apps [37, 39, 40, 57, 65]. Multiple interface mechanisms have been considered for communicating directions to walkers (e.g., directions to waypoints, when to turn, or the presence of points of interest nearby): speech [4, 23, 30], vibration [7, 25, 56], sound [24, 46, 54], or immersive sound (e.g., Microsoft Soundscape). Balata et al. [8] proposed a set of navigation instructions tailored to indoor-outdoor transitions. An open standard (Wayfindr) for communicating wayfinding information to blind travelers was approved as ITU-T F.921.

## 3 THE ROUTENAV SYSTEM

RouteNav was designed to be an integrated travel app. Users can set up “trips” using a web application, where these trips may include multiple means of transit (see Fig. 1, *Setup* panel). Using the RouteNav mobile app, one can select a trip, and receive information about the different legs of the trip and the required transfers. Once

at a transit location (e.g., a train station), RouteNav connects to a cloud server to receive a *tile route* towards the next *boarding area* (Fig. 1, *Using Mobile App* panel).

This contribution focuses on the mobile RouteNav app functionalities when at a transit location. The goal of this app is to help a blind walker reach a certain boarding area, by traversing a sequence of *tiles* as described in Sec. 3.1. As with any wayfinding systems, RouteNav needs to be able to locate and track the walker, provide useful directions, and inform the walker of their current location and progress along the route. The localization module is described in Sec. 3.2, while the information produced and the user interface are described in Sec. 3.3.

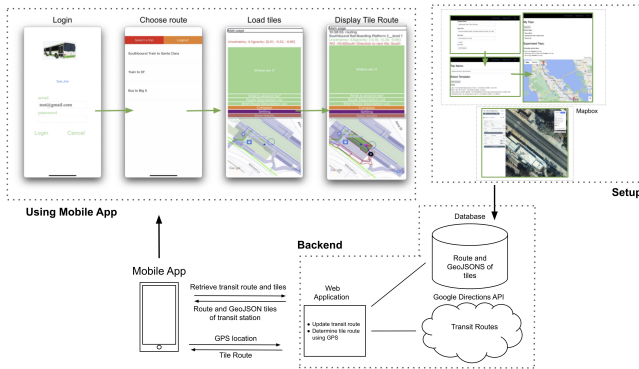


Figure 1: Overview of the RouteNav system.

### 3.1 Mapping

We used MapBox Studio<sup>1</sup> to map relevant areas and features in the location of interest. Note that the Palo Alto Transit Center contains two levels (a ground level and an underground passageway). We maintain separate representations of these two levels, and switch between the two when appropriate.

**3.1.1 Tiles.** Tiles are discrete spatial units in the form of non-overlapping polygons. At each time, the walker’s position is mapped to one tile, or to an *out-of-tile* state (Sec. 3.2.5). A walking route is an ordered sequence of tiles. Tile navigation generalizes the familiar concept of waypoint navigation. We introduced tile navigation for two main reasons:

*Localization inaccuracy mitigation:* Given the expected localization error (possibly up to several meters), it would be inappropriate (and actually, counterproductive) to produce information and directions that involve highly localized features or landmarks (such as the exact location of a turn or of an entrance). By using relatively large spatial units (commensurate with the expected spatial uncertainty), we can cope with poor location accuracy.

*Simplified path representation:* By describing a path as a sequence of tiles, including critical tile features (e.g., whether the tile is a crosswalk, ramp, or underground tunnel) and directions for moving from tile to tile (e.g., bear left at the next tile), we aim to provide a simple and intuitive representation that can be easily communicated

<sup>1</sup><https://www.mapbox.com/mapbox-studio>

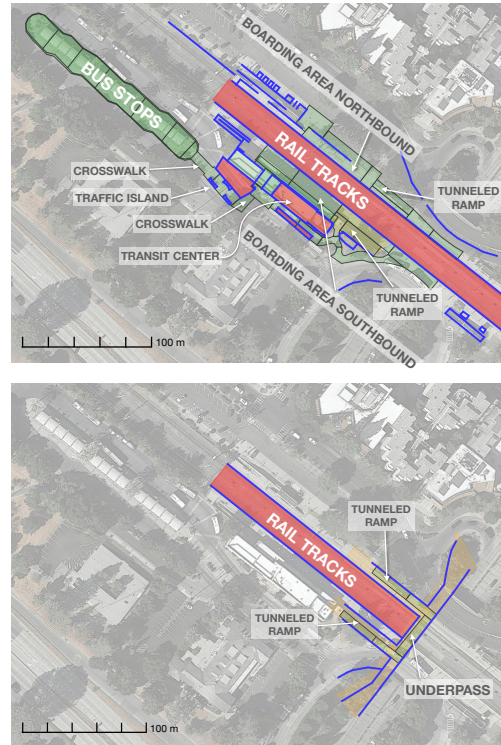


Figure 2: Maps of the Palo Alto Transit Center. Top: Ground level. Bottom: Underground level. Tiles are shown in color, separated by thin black lines. GPS-denied areas are shown with yellow hue. Red areas represent no-go zones. Blue thick lines represent impenetrable wall segments.

and memorized. In addition, keeping track of the tiles traversed is a simple way to represent progress along a route.

The tiles we traced for the walking area of the Palo Alto Transit Center are for the most part of rectangular shape, with length between 10 and 25 m. In choosing a tile’s size, one must trade off different needs: resilience to localization inaccuracy (which calls for larger tiles) against faithful path representation (which would favor smaller sizes, especially for non-rectilinear segments). Tiles are represented as GeoJSON polygons. They have a descriptive name and include a number of fields to record characteristics such as their type (e.g., walkway, crosswalk, ramp, tunnel), their main orientation and length, whether they slope, etc.

The Palo Alto Transit Center includes an underground passageway that overlaps with walkable areas at ground level, and for this reason we also defined a set of underground tiles. The underground level is accessible through two ramps. Rather than assigning tiles in the ramps to either level, we generate spatially overlapping tiles over the ramps. The localization/tracking system is in charge of switching assignment from ground to underground level, and vice-versa.

**3.1.2 Points of Interest and Landmarks.** Tiles may contain *points of interest* (POI), represented as either points or polylines in GeoJSON. POIs are recorded with fields that encode their type, such as obstacle,

bench, or tactile (e.g., for a tactile strip on the pavement), as well as verbal descriptions.

In addition to tiles, we mapped very large polygons (*landmarks*) that are not expected to be traversed by the walker (nearby roads, the train tracks). These landmarks are only listed in the *What's Around Me* functionality (Sec. 3.3.4), in order to offer some level of geographical context.

**3.1.3 Goalposts.** Within each tile, we define two *goalposts*, which are used as reference points to provide directional information. Note that, differently from waypoints (which are locations that the traveler is supposed to reach, or at least pass by), a goalpost is solely used to establish a direction of travel (travelers are not necessarily expected to reach a goalpost).

The goalposts are placed at an appropriate distance to the tile's junctions with the nearby tiles. A walker located in a tile is always directed towards one of the two goalposts in the next tile in the route. Specifically, if a walker is more than 4 meters away from the next tile, the closest goalpost in the next tile is chosen as "target". As the walker gets closer, the target is moved to the further goalpost in the tile (see Fig. 3, right panel).

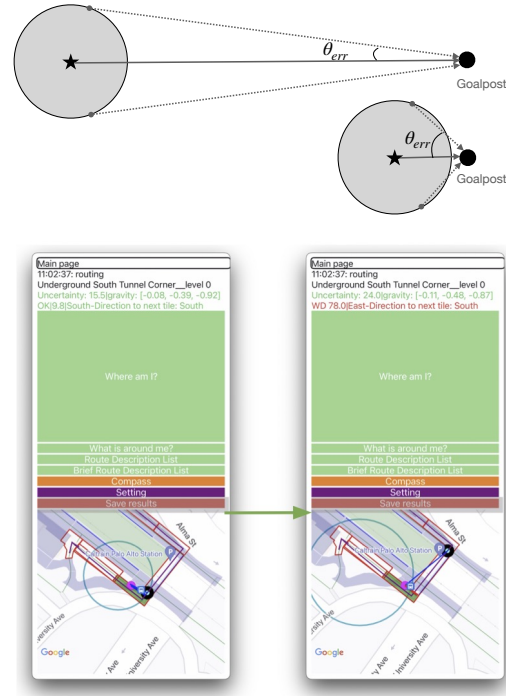
This "switching goalposts" strategy was devised to mitigate the effect of location inaccuracy on the direction provided to the user. As shown in Fig. 3, top panel, if the user is close to a target goalpost, a localization error can induce a very large angular error in the direction to the goalpost. Our strategy ensures that the user is always at a safe distance (larger than the expected positional error) from the target goalpost. Note that one could consider more than two goalposts for long tiles; we will investigate this possibility in future work.

**3.1.4 Walls and No-Go Zones.** We define a number of impenetrable *wall segments* and of *no-go zones*. These features are very useful when used in conjunction with the particle filtering module (Sec. 3.2.3), as they help reduce the risk of computing inconsistent trajectories, and may also help with drift correction for the inertial tracker. No-go zones are areas (represented as polygons) that, while not strictly impenetrable, are not expected to be traversed by a blind traveler. These include buildings close to the public, areas covered with vegetation (bushes), and a few hazardous places. Note that most road surfaces were *not* marked as no-go zones (even though walking there was inadvisable), and occasionally our participants did walk on those roads. No-go zones are shown in red in Fig. 2.

## 3.2 Localization and Tracking

The localization module of RouteNav makes use of the GPS signal (when available and reliable) as well as of the signals produced by the iPhone's inertial sensors (accelerometers and gyros). The two sources of information are fused together using particle filtering, which also leverages prior spatial information.

**3.2.1 GPS.** Using iOS' `CLLocationManager.location`, the GPS location is available in the coordinate property, while its uncertainty radius is in the `horizontalAccuracy` property. Unfortunately, we found that oftentimes the nominal uncertainty radius reported by Core Location is small even when localization is patently wrong (see e.g. Fig. 4). We thus decided not to use GPS information

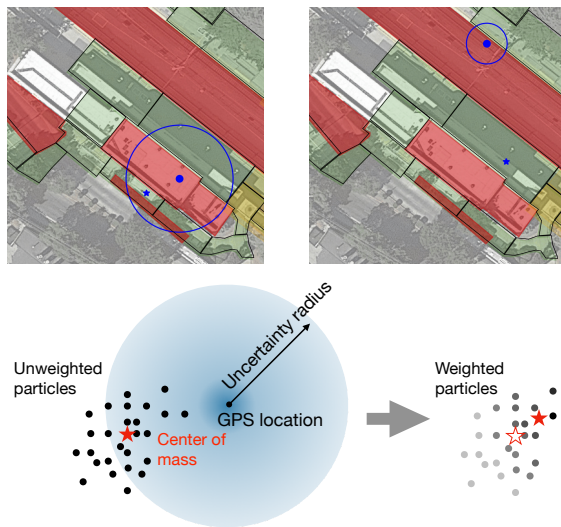


**Figure 3: Top: Schematic representation of the dependence of directional error on distance. In both cases, the true location of the walker is at the star; however, due to localization error, the system may locate the user anywhere in the gray disk, defined by the uncertainty radius. Because of localization error, the direction to the target goalpost (thick dot) as produced by the system (shown by dotted arrows) may be different from the correct direction to the target goalpost, shown by a solid arrow. The angle formed by the two directions is the angular error  $\theta_{err}$ . As the figure suggests, larger angular errors  $\theta_{err}$  should be expected for the same uncertainty radius when the user is closer to the goalpost. Bottom: Screenshots from the RouteNav app. The pink dot represents the walker's location as estimated by the system. The dark circle represents the next selected goalpost. Once the walker enters a tile, the closest goalpost in the next tile is selected as target. As the walker progresses in the tile, the target is moved to the further goalpost.**

when the user location is in a *GPS-denied* zone (shown with yellow hue in Fig. 2). GPS-denied zones are defined by a prior on-site survey.

**3.2.2 Inertial Localization.** We use RoNIN, a state-of-the-art algorithm for inertial dead-reckoning whose code was made available by the authors [36]. RoNIN uses a deep network to process data from the phone's accelerometers and gyros, producing velocity vectors that can then be integrated in time to obtain the walker's position. We used the ResNet-18 architecture [35] as considered in [36].

In our iOS implementation, RoNIN produces velocity vectors at 25 Hz. Remarkably, the velocity vectors produced by RoNIN are largely independent of the actual attitude of the phone, meaning that they represent the user’s velocity direction irrespectively of how one is holding the phone. This is a very useful feature of the algorithm: users can keep their phone wherever they want, and even move it (e.g., from tucked in one’s pocket to hand-held) without affecting positioning. RoNIN was shown [49] to also function for blind individuals walking with a long cane or a guide dog, whose gait may be expected to have different characteristics than for sighted walkers [38]. However, we found that RoNIN suffers from two main problems. First, it is liable to orientation drift accumulated in the phone’s attitude estimation produced by Core Motion. Second, we found that both magnitude and direction of the velocity vectors computed by RoNIN may be incorrect for some users (see also [49]).




**Figure 4: Top panel: GPS errors.** The star represents the true location of the walker, while the dot represents the location estimated by GPS, surrounded by a circle with radius equal to the uncertainty radius. Note that while in the example shown to the left, the uncertainty radius (14.5 m) is consistent with the actual error (10.5 m), in the example to the right the uncertainty radius is small (5.4 m) in spite of the very large error (36.5 m). **Bottom panel: An illustration of the strategy used for fusion of localization from GPS and inertial data (Sec. 3.2.4).** Each particle in the cloud is assigned a weight that depends on its distance to the GPS location. As a result, the center of mass of the particles moves closer to the GPS location.

**3.2.3 Particle Filtering.** Particle filtering is a form of Bayes filtering that is particularly adept at exploiting prior spatial information in the form of impenetrable walls [61]. Our implementation of particle filtering is similar to that described in [49]. The posterior probability of the user’s location at time  $t$ , given all inertial measurements before  $t$ , is represented by a set of  $N = 750$  “particles”. In addition to the user’s location, the system tracks a *drift* angle that represents

the current angular bias of the phone’s attitude estimation. Each particle is thus characterized by its location ( $p_i(t)$ , where  $i$  is the particle’s index), a drift angle  $d_i(t)$ , as well as a weight  $w_i(t)$ . The particles are propagated at each time instant  $t$  based on the velocity vector  $v(t)$  produced by RoNIN. For each particle,  $v(t)$  is first rotated by the drift angle  $d_i(t)$ , resulting in a new vector  $\bar{v}_i(t)$ . Gaussian noise is added to both magnitude and orientation of  $\bar{v}_i(t)$ . Then, the particle’s position is moved to  $p_i(t) + \bar{v}_i(t) \cdot \Delta t$  (where  $\Delta t = 1/25$  s in our implementation). In addition, Gaussian noise is added to the value of  $d_i(t)$ . If the particle’s motion would take it across a wall, or to a location in a no-go zone, the particle’s weight is set to 0, and the particle is resampled at the next iteration. All particles are resampled when the Effective Sample Size [22] (which measures the weight diversity across particles) is less than half the number of particles with weight larger than 0. The user’s location  $p(t)$  at each time is produced as the dominant mode of the weighted distribution of particles, computed using the mean shift algorithm [16, 49].

Note that the inertial tracker uses an arbitrarily oriented reference system, with Z-axis pointing in the direction of gravity, whereas GPS coordinates are in the WGS-84 (geodetic) system. We transform data from both sources into the UTM (metric) system. While WGS-84 to UTM coordinate conversion is straightforward, conversion from the inertial tracker frame to UTM requires an initial calibration, which we perform as follows. Each route starts from a specific location with known coordinates. The walker walks for a few steps in a known direction. Based on the computed trajectory, a rotation matrix and translation vector are found that map the tracker’s reference frame to the UTM system.



**Figure 5: Examples of particle distributions.** Red areas represent no-go zones. Blue thick lines represent impenetrable wall segments. The black star represents the barycenter of the particles in the dominant cluster. Note that the example in the rightmost panel has two clusters of particle (one shown in pale blue, one in dark red), with the barycenter of the second cluster shown by an orange star.

**3.2.4 Fusion of GPS and Inertial Localization.** To combine localization information from GPS and RoNIN, we use the fusion algorithm originally proposed in [50]. The GPS location uncertainty is modeled by an isotropic bivariate normal distribution, with mean at the GPS location and standard deviation  $\sigma$  equal to twice the uncertainty radius, as produced by Core Location. The current weight  $w_i(t)$  of each particle, located at  $p_i(t)$ , is multiplied by the value of this normal distribution at  $p_i(t)$ . Intuitively, particles that are closer to the GPS location are given higher weight (see Fig. 4). When resampling is performed, these particles are more likely to be resampled than those located further away. This causes the mode of

the particle’s distribution to move closer to the GPS location. As expected, the effect of GPS is larger when the uncertainty is low (as particles close to the GPS location get higher relative weight). When the GPS uncertainty radius is larger than 10 m, or when the current estimated location is in a GPS-denied zone, we simply ignore the effect of GPS.

This fusion strategy is critical for RouteNav. GPS (when available) helps correct drift accumulated by RoNIN, and, in turn, RoNIN mitigates sporadic GPS errors. Particle propagation proceeds with data from RoNIN even when reliable GPS data appears or disappears, thus ensuring seamless handling of indoor-outdoor transitions.

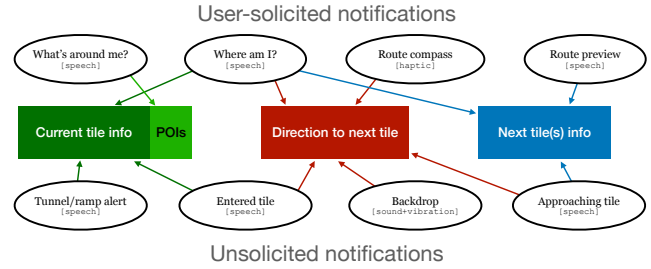
**3.2.5 Tile Assignment.** The estimated user location  $p(t)$  at each time is assigned to one tile, or to the out-of-tile state. To reduce the risk of undesired back-and-forth switches between tiles due to noisy localization, we use a simple hysteresis mechanism. We first compute the distances  $\{D_T\}$  between  $p(t)$  and each tile polygon (where  $T$  is the tile index). Given the current tile assignment  $T(t)$ , we re-assign  $p(t)$  to a tile other than  $T(t)$  only if  $D_{T(t)} > 2$  m, in which case  $p(t)$  is assigned to its closest tile (note that this is not necessarily the next tile in the route). If, however,  $D_{T(t)} > 7$  m, the location is assigned to the *out-of-tile* state. Subsequent locations are still assigned to out-of-tile, until the location is less than 4.5 m to its closest tile. At that point, the location is assigned to its closest tile, and the tile route is recomputed if necessary. Note that with this mechanism, a location  $p(t)$  is assigned to a tile even if it is outside of it, at a distance no larger than 7 m. This is not dissimilar from conventional navigation mechanisms that map a vehicle location to the nearest road.

**3.2.6 Walking Direction.** We determine the current walking direction by joining the current location  $p(t)$  with  $p(t - t_D)$ , where  $t_D$  is chosen such that the path traversed from  $p(t)$  to  $p(t - t_D)$  has length of 3 meters. Information about the angle between the current walking direction and the next goalpost is used in the *tile-entered* notification as well as in the *Where Am I* message (Sec. 3.3.1). This value is also used to determine the type of sound/vibration backdrop.

**3.2.7 Phone Azimuth.** While the measured walking direction is independent of the phone’s orientation, the *Route Compass* functionality (Sec. 3.3.1) is based on the current azimuth angle of the phone, which is obtained from the `magneticHeading` attribute of the `CLLocation` class (part of `Core Location`). The azimuth represents the direction of the projection onto the horizontal plane of the Y-axis of the phone.

**3.2.8 Underground Ramp Detection.** In order to cross the rail tracks in the Palo Alto Transit Center, one needs to take an underground passageway, which can be entered through two tunneled downward ramps. As mentioned earlier, these ramps have spatially overlapping tiles at ground and underground level. When the walker enters one such ramp from ground level, their position is initially assigned to a ground level tile. As the walker progresses through the ramp, the level is switched to underground through one of two possible mechanisms:

**Elevation Measurement.** The iPhone’s altimeter can be used for an approximate elevation measurement, which can be thresholded to determine whether or not the user is in the ramp. Unfortunately,



**Figure 6: An overview of the RouteNav interface. The colored rectangles represent different categories of information provided: direction to the next tile (more precisely, to the target goalpost in the next tile); information about the current tile, possibly including points of interest (POIs); and information about the next tile or tiles in the route. The different types of notifications (solicited or unsolicited) produced by the system, which convey information in one or more categories, are shown with round shapes, along with the associated interface modality (speech, sound, vibration, haptic tap). Solicited notifications are generated when the user double taps the corresponding button, except for the Route Compass notification, which is generated when the user stops walking and holds the phone horizontal.**

we found that the altitude data from the phone is affected by a bias that may change from trial to trial. Note that, unlike the case of determining which floor in a building one is at, e.g. when exiting an elevator [4], we need to be able to measure small differences in elevation from the ground level, to confirm that the walker did in fact enter or exit the tunneled ramp.

**User as a Sensor.** Most persons who are blind are able to recognize (from auditory clues) whether they are in a narrow tunnel. Based on this consideration, we implemented a mechanism that produces an alert message (Sec. 3.3.3) asking users to confirm whether or not they entered a tunnel. This alert is triggered when the walker is localized in an underground tunnel ramp tile, or if they have been localized in a nearby tile for a certain period of time. Based on the user’s response, the tile/level assignment is adjusted. This is an instance of the “user as a sensor” paradigm, first introduced in [23]. We use a similar mechanism to confirm whether or not the walker is on a ramp branching out and extending parallel to a walkway in another section of the station.

### 3.3 User Interface

The RouteNav iOS app was designed to be operated with VoiceOver. Its main screen contains a list of buttons, visible in Fig. 7, which can be scrolled through using right/left swipes. Note that one button (*Where Am I*) has a much larger screen area than the others; this was suggested by a participant (P1) to make it easier to locate, since this functionality is used much more frequently than the others. A map showing the current user’s location in the tile route is shown in the lower part of the screen, for use by sighted walkers and for debugging. The main functionalities of the RouteNav app are described below (see also Fig. 6).

**3.3.1 Direction to Target Goalpost.** We offer a redundant variety of modalities for communicating the direction towards the target goalpost:

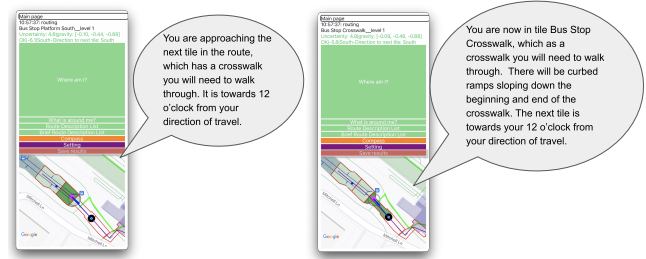
**Where Am I:** The *Where Am I* button produces a direction to the next goalpost in speech form. Users can select (from the Setting button) whether to use clockface directions (e.g., 2 o'clock), “simple” directions (e.g., front-right), or degrees (with respect to the frontal direction). Importantly, these directions refer to the *past walking direction* of the user (Sec. 3.2.6). Users were informed that if they stop and then rotate their body, they still need to remember their past walking direction to make sense of the direction from the app. In addition to the direction to the next goalpost, *Where Am I* provides the following information: (1) the name of the current tile; (2) the name of the next tile in the route; (3) the distance to the next tile; (4) what direction to take at the end of the current tile; (5) and how many tiles are left in the route till destination. The distance to the next tile (3) is expressed only as more/less than 10 meters. Using higher resolution would not be advisable given possible localization errors. After (2), important information about the next tile is given if applicable, such as whether the tile is a ramp, tunnel, or curb. The *Where Am I* button can be used at any time in the route.

**Route Compass:** The *Route Compass* is a simple pointing mechanism that is automatically activated when the user is static (not walking) and holds the phone approximately level to the ground. The angle between the phone’s Y (longest) axis and the direction to the next goalpost (Sec. 3.2.7) is measured as the user slowly rotates the phone left and right. When this angle crosses  $0^\circ$ , an haptic tap (a rapid and slight vibration) is generated. This mechanism is similar to the “wand” modality considered by Azenkot et al. [7].

**Sound/Vibration Backdrop:** When a user is walking and the angle between the estimated direction and the direction to the next goalpost is within  $\pm 45^\circ$ , a short whooshing sound (`mail-sent.caf` in iOS) is generated, along with a short vibration, every other second. If the angular difference becomes larger, the sound is changed to a more unpleasant one (`nfc-scan-failure.caf`), and a short vibration is repeated twice every other second. If the user stops, both sound and vibration stop. Note that another possible interface modality to convey direction to the target goalpost is via 3-D audio (e.g., Microsoft Soundscape). We will consider this and other sonification modalities (e.g. [3]) in future work.

**Approaching/Entered at Tile:** These unsolicited notifications are produced (in speech form) when the walker is approaching the next tile in the route, as well as when the walker’s position is mapped to a new tile. Examples of these notifications are given in Fig. 7. Each notification is only produced once per tile, and is preceded by a chime sound to get the walker’s attention.

**3.3.2 Route Preview.** The *Route Description List* button generates a list with the sequence of tiles from the current one to destination. Items in the list can be accessed with standard right/left VoiceOver swipes. Each tile in the sequence is listed with its name, its length, any important characteristic (e.g., if it slopes down, if it is a tunnel or a crosswalk) as well as the recommended walking direction through it. This direction is expressed in an allocentric frame (cardinal directions, e.g. “walk South”). Finally, the walking direction



**Figure 7: Examples of speech notifications generated when the walker is approaching a new tile in the route (left), and when the walker enters new tile (right).**

recommended once at the end of the tile to enter the next tile is provided in egocentric form (e.g., left, straight, or right).

**3.3.3 Tunnel/Ramp Alert.** When the walker is located in the proximity of either tunneled ramp (shown as S1(N) or S1(S) in Tab. 3), or of the walkway ramp shown as (S4), an alert message is generated, starting with a chime, followed by the words “Alert. RouteNav thinks that you may be in a tunnel (ramp).” Then, a list with two buttons (Yes or No) is presented, for the user to select from.

**3.3.4 Points of Interest and Landmarks.** By double tapping the *What’s Around Me* button, an audio description is produced with the POIs in the current tile, ordered along the expected traversal direction of the tile, as well as distant landmarks.

## 4 EXPERIMENT

### 4.1 Population

We recruited 7 participants (2 female, 5 male) for our study, with ages ranging from 59 to 74. All participants were blind, with at most some residual light perception. The participants’ characteristics are summarized in Tab. 1. P4, P5 and P6 used a dog guide, while the others used a long cane for mobility. All participants had substantial experience with independent travel, except for P1, who said that he normally travels with his spouse. P2 was an accomplished independent traveler in his young age, though now only travels accompanied by his sighted partner. None of the participants had experience traveling independently in the Palo Alto Transit Station. P3, P4 and P5 had been there before, accompanied by sighted guides, but said they had no relevant spatial memory of the layout of the station.

All participants were iPhone users and were proficient with VoiceOver (although P1 often experienced difficulties with double-tapping during the test). Asked about what (if any) travel apps they used, P3, P4 and P7 said they sometimes used Nearby Explorer from APH, and P3, P5, P6 and P7 said they used GoodMaps. P4 also used Microsoft Soundscape.

### 4.2 Modality

**4.2.1 Routes.** We identified three representative routes for our tests (see Fig. 8):

**Northbound to Southbound tracks (NB2SB).** This route (130 m long) starts from a point on the Northbound train platform, and



ID	Age	Gender	Onset of blindness	Independent traveler?	Mobility aid
P1	74	M	Since teen age	Rarely travels alone	Long cane
P2	74	M	Since teen age	Yes (when younger)	Long cane
P3	60	M	Since birth	Yes	Long cane
P4	65	F	Since birth	Yes	Dog guide
P5	69	M	Since 3 years of age	Yes	Dog guide
P6	59	F	Since teen age	Yes	Dog guide
P7	70	M	Since birth	Yes	Long cane


**Table 1: Participants' characteristics.**

ends at a point (boarding area) of the Southbound track through the underground passageway. An underground tunnelled ramp entrance can be found after walking in the SE direction for about 15 m. The ramp (24 m long) has a gentle slope downwards, and ends at a T junction. The route turns right to SW, and continues on a level underground passageway that borders a very trafficked (and noisy) highway underpass. After 30 m, there is a T junction; the route turns right (NW) on an upward tunneled ramp that parallels the first one. Upon exiting the ramp (at ground level), the route continues straight for about 33 m on the Southbound track platform, bordering the Transit Center building. Walkers need to pay attention not to trip on light posts, benches and other obstacles, and not to walk into the tracks.


**Southbound tracks to bus stop (SB2Bus).** This route (145 m long) starts from a point in the Southbound train tracks, and leads to a point in the bus stop area. After walking SE for about 16 m, one needs to find a narrow passage between the SE end of the Transit Center building and the underground tunnelled ramp entrance. Upon turning right onto this passage, the walker may follow a paved pathway that curves around the SE end of the building (flanked by dirt patches). This connects to a raised passageway, flanked to the left by a bushed area and to the right by the Transit Center building. This passageway continues straight for 36 m; however, after 27 m one needs to find to the left the entrance to a narrow downward ramp, parallel to the passageway, that leads to the street level. At the base of the ramp, walkers need to turn left onto a crosswalk and traverse it. The crosswalk ends on an elongated traffic island; one should turn right on this island, then traverse it walking NW. This traffic island contains several concrete planters and a tactile strip. After 25 m, the island ends onto a fence, with a narrow entrance onto a crosswalk. Walkers need to find the entrance through the fence, cross the crosswalk in the NW direction, then continue walking on the bus stop island until they reach the boarding area (about 10 m after the crosswalk).

**Bus stop to Northbound tracks (Bus2NB).** This is the longest (and arguably most complex) route (225 m). It follows (in reverse) the SB2Bus track from the bus stop to the entrance of the underground tunnelled ramp located near the Southbound tracks. From there, one needs to enter the tunnel, and follow (in reverse) NB2SB to a boarding point in the Northbound tracks.


**4.2.2 Procedure.** Each participant after P1 was provided a few days in advance of the study with an audio description of the app and of the experiment, narrated by an experimenter, and asked to listen to it beforehand. Tests were conducted on weekends, to avoid the commuter crowd. Each participant was tested individually. P2 and P3, as well as P5 and P6, were tested on the same day, one after the other. On test dates, participants (except for P5 and P6) were



Northbound to Southbound tracks (NB2SB) – P5



Southbound tracks to bus stop (SB2Bus) – P3



Bus stop to Northbound tracks (Bus2NB) – P1

**Figure 8: Sample paths reconstructed for participants traversing the three considered routes. Each route starts at the location marked by a star. Note that the NB2SB and Bus2NB routes include an underground connection. White line: the most direct end-to-end path. Line-connected dots: GPS measurements. The color of the dots represents the uncertainty radius, from blue (< 5 m) to red (>12 m). Gray line: path reconstructed from inertial sensors using RoNIN. Light blue line: Path reconstruction from fusion of GPS and RoNIN via particle filtering. Sample pictures of the participants traversing the routes are shown, along with screenshots of the app and graphical representations of the direction to the next goalpost.**

greeted at the parking lot to the North of the transit center. From there, the experimenters walked with them to the starting point of NB2SB. Once there, verbal and written consent (as approved by our IRB) was obtained, after which participants were given a demo of the RouteNav app. They were asked to experiment with the different buttons, and to try the *Route Compass* functionality. Participants were asked to adjust the speed of VoiceOver speech as desired. They were provided with bone conduction Bluetooth headphones (Shokz OpenRun) connected to the smartphone (an iPhone 12) running RouteNav. These headphones do not occlude the ear canal and are comfortable to wear. Participants were then accompanied to the designated starting point for the first route. The experimenter started the app, selected the destination, and walked with the participant along the designated direction for a few steps. This was necessary for the initial calibration described in Sec. 3.2.3. After that, the experimenter handed over the iPhone

to the participant, who was asked to reach the route destination following the phone’s input.

During the route, participants were accompanied (at a safe distance) by two experimenters. We did our best to be minimally invasive (so as not to involuntarily bias the participants’ judgment) while at the same time ensuring the participants’ safety while walking. In three situations did we intervene by providing more information than that produced by the app: (1) when a participant was at imminent risk, such as when they were about to trip or bump into a hazard, when a cyclist was passing by, or when a person sitting on the floor along the route would not move; (2) when, based on the experimenter’s judgment, the app was giving incorrect directions for an extended period of time, in which case the experimenter suggested walking in a certain direction, or asked the participant to walk with them for a short while, until the app was again able to give useful information (these situations are documented in Sec. 4.3.2); (3) when a participant appeared to have missed or neglected some important audio information produced the phone (in particular, the *Tunnel/Ramp Alert*, which needs to be attended to). In this case, the experimenter would bring this information to the participant’s attention. This is also documented in Sec. 4.3.2. Occasionally, the experimenter would remind a participant to check the Route Compass or to double tap the *Where Am I* button, in order to find the direction to the next tile. This typically happened at the beginning of the study, until the participant became more familiar with the app.

At the end of the first route, participants were asked to rest for a few minutes, after which the next trial started with the same modality. All participants except for P5 and P6 walked along the three routes in this order: NB2SB; SB2Bus; Bus2NB. P5 and P6 arrived together on a Southbound train (though they performed the study individually, one after the other). We decided to change the order of routes for them so that they would not get a “preview” of the route by walking to the beginning of NB2SB on the Northbound track. The order of routes for them was: SB2Bus; Bus2NB; NB2SB.


At the end of the third route, we asked all participants after P1 to answer a short questionnaire, described in Sec. 4.4.

**4.2.3 Interface Design Evolution.** While the localization/tracking module remained basically unchanged throughout the experiment, feedback from the participants was incorporated in a series of iterative improvements, as briefly described in the following.

P1 suggested that we should incorporate a continuous feedback mechanism, confirming that one is walking in the correct direction. This prompted us to add the sound/vibration backdrop feature. The choice of the “wrong direction” sound was further modified after feedback from P2. Sound/vibration backdrop is now one of the most appreciated system features, based on the participants’ comments. P1 also recommended that the *Where Am I* button be given more “real estate” on the screen to make it easier to find. We noticed that he would often miss the unsolicited speech notifications produced by the app, which prompted us to add a preceding chime sound.

Important additions to the information produced by the system were implemented incrementally from feedback from P1, P2, and P3. P1 recommended that relevant information about the next tile (whether the tile is a crosswalk, tunnel or ramp) be incorporated in the *tile-approaching* and *tile-entered* unsolicited notifications

(originally, this was only available on request from *Where Am I*). This simple modification turned out to be very effective, because it informs users about what to expect as they move on. P2 recommended including information about detectable curb changes, which prompted us to add a “curbed ramps sloping down” notification text as seen in Fig. 7 (though P7 later suggested that we change this wording to simply mention the presence of a curb cut). P2 also recommended notifying the user about whether a ramp being approached slopes up or down (this information is now provided by the app). P3 suggested that the notification produced when approaching the walkway ramp (S4 in Tab. 3) should include which side (left or right) one should expect to find the entrance to the ramp. This proved to be another very helpful feature.



**Figure 9: Path reconstructed for the Bus2NB route for P1 (left) and P6 (right). (See caption of Fig. 8 for details.) Note that the system failed to track P6 in the underground connection.**

## 4.3 Observations

**4.3.1 General Observations.** All participants were able to complete all routes, except for two localization/tracking failure cases described below. Occasionally, an experimenter’s intervention was necessary in a few challenging spots, as detailed in the next section. Although the study was conducted on weekends, numerous people were standing, walking or biking in the station, and several unhoused persons were sitting and laying on the ground. The environment was rather noisy, due to the presence of diesel-powered incoming trains, as well as to intense car traffic in the highway underpass near the underground passageway. In spite of this, participants were able to hear the speech generated by the app through the bone-conducting headphones. Traversal times are reported in Tab. 2.

The system experienced major failures, in terms of user localization/tracking, in two cases. For P7, at the end of the NB2SB route, the location estimated for the participant moved from the Southbound platform (where the participant was located) across the rails to the Northbound platform. This was caused by a large, sustained GPS error coupled with low uncertainty radius (see Fig. 4) that eventually pushed all particles (after resampling) to the wrong side of the tracks. Since this happened when the participant had just arrived at the terminal tile, we simply terminated the trial at that point. For P6 in the Bus2NB route, while in the underground passageway, the participant’s path was incorrectly tracked (due to RoNIN undershooting), and the system was no longer able to recover the path (see Fig. 9, right). Since directions from the system were no longer reliable, we decided to walk with the participant

through the rest of the route until destination (Northbound platform).

All participants except for P4 were able to use the Route Compass (P4 told us that she could not feel the haptic tap). In fact, we noticed that the Route Compass was generally used more often than the *Where Am I* button. Although participants were informed that they could keep the phone in their pocket, if they wanted to, all walked holding the phone in their free hand. This was most likely due to the fact that the two most often used interface functionalities (*Where Am I* and Route Compass) required either tapping on the phone’s screen, or holding the phone horizontally. In future work, we will explore the possibility of using a smartwatch for both functionalities, thus relieving the user from having to hold the phone in one hand.

Only P1 made occasional use of the *What’s around me* button. To make sure that participants were at least aware of this functionality, we asked them to use it at least once, when on the traffic island at the SB2Bus route. The traffic island and bus stops contained a tactile strip on the ground, oriented along the direction of the route, and the description of the tile produced by the app did mention it. To our surprise, only P1 and P2 followed the tactile strips, and even then only for a very short extent.


Using a dog guide vs. using a long cane clearly made a difference in the way participants moved along the route. In general, blind travelers walk faster with a dog guide than when using a cane. We noticed that participants with dogs were able to walk with less veering, and the dog helped find the opening in the fence (S7(S) and S7(N) in Tab. 3). At the same time, the dog users had to attend to multiple tasks: maintaining their awareness of the environment, making sense of the information from the app, and giving commands to the dog. That was particularly taxing for P4, whose dog was relatively undisciplined.

Indoor–outdoor transitions were handled smoothly by the system. We noted, however, that when indoors (in the GPS-denied tunneled ramps and underground passageway), the velocity vectors produced by RoNIN had incorrect magnitude for a few participants. As shown in Fig. 9 (thick grey line), this resulted in overshooting (P1) or undershooting (P6).

	P1	P2	P3	P4	P5	P6	P7	Median
NB2SB (130 m)	19	20	9	9	4	8	11	9
SB2Bus (145 m)	27	20	13	22	11	21	10	20
Bus2NB (225 m)	47	28	13	13	6	–	10	13

**Table 2: Traversal times (in minutes) for all participants and all routes. As explained in Sec. 4.3.1, traversal of Bus2NB was interrupted for P6 due to a failure of the localization system.**

**4.3.2 Challenging Spots.** From observation of the participants navigating the routes, we identified a number of “challenging spots”, which are described below and shown in Tab. 3. The same table summarizes the performance of our participants in these spots. When participants negotiated a challenging spot successfully by themselves, we marked this as  $\checkmark$ . We used the symbol  $\circ$  to mark a



	P1	P2	P3	P4	P5	P6	P7
S1	$\checkmark$	$\circ^N, \bullet^S$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
S2	$\circ^N, \bullet^S$	$\checkmark$	$\checkmark$	$\bullet^S, \bullet^N$	$\checkmark$	$\bullet^S, \bullet^N$	$\bullet^S, \bullet^N$
S3	$\checkmark$	$\checkmark$	$\checkmark$	$\bullet$	$\bullet$	$\bullet$	$\checkmark$
S4	$\circ$	$\bullet$	$\circ$	$\circ$	$\checkmark$	$\bullet$	$\checkmark$
S5	$\bullet$	$\bullet$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\bullet$
S6	$\checkmark$	$\bullet$	$\checkmark$	$\bullet$	$\checkmark$	$\circ$	$\checkmark$
S7	$(\checkmark)^N$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$(\checkmark)^N$

**Table 3: Participants’ performances in the challenging spots shown in the images above (with the route overimposed as an orange line), and described in detail in Sec. 4.3.2.  $\checkmark$ : Spot successfully negotiated, possibly by finding an alternate route to bypass the spot (cases marked as  $(\checkmark)$ ).  $\circ$ : The participant asked questions about the information provided by the app, or they were prompted by the experimenter to pay attention to something that was communicated by the app.  $\bullet$ : The experimenter gave spatial directions to the participant, or walked with them away from the challenging spot.**

case in which a participant asked questions about the information provided by the app, or if they were prompted by the experimenter to pay attention to something that was communicated by the app. Conceivably, better training or experience with the system could remove the need for input from the experimenter in these cases. In some cases, the experimenter gave spatial directions to the participant, or walked with them away from the challenging spot, because information from the system was unreliable. These “unrecoverable” situations are marked as  $\bullet$ .

**S1(N,S): North/South Tunnel Entrance (NB2SB, Bus2NB).** The tunneled ramps have a width of about 2.5 m, with empty walkable space around. Producing reliable directions helping a blind traveler to find the tunnel entrance is challenging for two reasons: (1) a small directional error may lead one to miss the tunnel entry, and start walking to its side instead; (2) even if one correctly enters the tunnel, the particle filter may locate them outside of it. This is because particles are likely to cluster in two different sets (one inside and one outside the tunnel), with unpredictable results. To mitigate the risk of directional errors, walkers were informed in advance (by the unsolicited notifications, or when they double tapped the *Where Am I* button) whether the next tile contained a tunnel (a feature that, as mentioned earlier, was originally suggested by P1). At least two participants (P4 and P6) were able to identify the presence of the expected tunnel entrance through auditory clues. (Note in passing that the dog guide accompanying P6 had a certain reluctance to enter the tunnel.) The risk of the system incorrectly locating a user

outside the tunnel when they were actually inside (and vice-versa) was substantially mitigated by the “user as a sensor” input alert, which asked users to confirm whether they actually were inside a tunnel. As shown in Tab. 3, all participants except for P2 were able to find the entrance of both tunnels.

**S2(N,S): North/South Underground Passageway T-Junctions (NB2SB, Bus2NB).** The NB2SB and Bus2NB routes required participants to use the underground passageway to move from the North to the South side of the station, and vice-versa. When walking in the passageways, participants needed to find the junction to the upward tunneled ramp taking them to ground level, and turn left (S2(N)) or right (S2(S)) onto it. Note that when walking underground, GPS is completely unreliable, hence the localization system only relies on RoNIN and particle filtering. Unfortunately, RoNIN produced incorrect velocity values for some participants, resulting in incorrect localization and thus unreliable directions. This was especially the case for the two participants with dog guides (P4 and P6), who walked very fast with their dogs, as well as for P7. In the case of P6, the system completely failed to track her path in the passageway, as shown in Fig. 9.

**S3: Turning around the Transit Center Building End (SB2Bus).** In this spot, participants found themselves in an open space, with poor GPS reception. Particle filtering was only partially effective, due to the absence of nearby walls (particles can just disperse in open space). This occasionally led to poor localization and consequently incorrect directions, requiring intervention from the experimenter for P4, P5 and P6. Note that in the case of P5, the experimenter only gave a verbal recommendation to keep to the right in order to avoid entering the nearby road.

**S4: Finding the downward ramp from the raised walkway (SB2Bus).** This was the most challenging situation, so much so that only two participants were able to successfully negotiate it without any input from the experimenters. The difficulty here is that the ramp is parallel to the walkway, with a narrow entrance to the left. This is not dissimilar to the case of tunnel entrance (S1), in that the localization system had great difficulty understanding whether the user was walking on the walkway rather than on the ramp (and vice-versa). Here too, the “user as a sensor” input alert message and early notification of the presence of the ramp were helpful. For three participants (P1, P3 and P4), after they overshot the ramp entrance, the experimenter simply reminded them that the notifications from the system included the presence of a ramp to be taken. These participants were then able to find the entrance to the ramp by themselves.

**S5: Entering the traffic island from the bottom of the ramp (SB2Bus).** P1 crossed the crosswalk diagonally, and started walking on the road, prompting the experimenter to direct him to walk over the curb onto the traffic island. P2, while negotiating one of the large concrete planters, was about to step down the curb onto the road, at which point he was stopped by the experimenter. After this occurrence, we added text to the tile descriptions produced by the system that advised participants against walking down the curb (curb cuts were located at the correct entry points of this traffic island).

**S6: Entering the upward ramp from the traffic island (Bus2NB).** Two participants (P2 and P4) missed the entrance to the ramp. This was due to the fact that, once in the vicinity of the ramp entrance,

they were directed towards SE (parallel to the ramp) due to a minor error in localization (the system assumed that they were already walking on the ramp). Hence, the two participants started walking on the road instead, and were asked by the experimenter to walk back and find the ramp entrance.

**S7(N,S): Finding the opening to the fence walking from North (Bus2NB) or from South (SB2Bus).** The NW edge of the traffic island has a fence with a narrow opening. All participants were able to find the opening when walking in the SB2Bus route. Two participants (P1 and P7), when walking in the Bus2SB route, missed the opening and walked on the road, trailing the traffic island on its SW edge. Both participants were able to find a way (following directions from the app) to walk back on the traffic island and proceed. Though not desirable, we found this to be an acceptable strategy, and marked it as (✓).

#### 4.4 Final Questionnaire

The questions asked at the end of the study, along with a summary of the replies, are listed below.

*RouteNav divides a route into “tiles”. Do you think that this helped you in your navigation?* Participants had mixed opinions on the importance of tiles. P6 thought it was a great idea because it “deals with the imprecision of turns”. P7 said “Absolutely. Got to have it. Don’t know any other way you could do it. You’ve got to break the route up. Nobody could ingest the whole route at one time.” P4 said “Immensely. Because I knew I was on route.” On the other hand, P3 and P5 remarked that, while useful to understand one’s progress along the route, tiles did not convey more information than regular waypoints.

*The “What’s around me” button gives information about landmarks in a tile. Do you think that this information is useful?* P2 mentioned that he was interested in knowing information about hazards around himself. P3 found it useful, though he remarked that he would only use it once he was more familiar with the app, as he was still getting used to parsing information from the app while using the cane. P4, P5, P6 and P7 suggested having different verbosity levels for this information. P7 also suggested that more information about perceivable features (such as “listen for the tunnel entrance” or “trail the wall with the cane”) could be added.

*Do you think that the app gave you correct walking directions?* Most participants thought it did “most” or “90% of the times”. P2 thought so, except when the direction was “across a wall” (because the next goalpost was behind a corner). P6 (whom the system failed to track under the tunnel) answered “No”. It appears that several participants expected the system to not always be reliable, without this being a major concern. This is consistent with the findings of the survey reported in [1].

*Do you think that the app knew where you were at all times?* Here too, participants (except for P6) thought that localization was good most of the time, though P2 said he wanted better localization, and P3 would have liked to know exactly when to turn (but understood that this would be difficult in practice.) Likewise, P4 would have liked more precise information on where exactly wall corners were located.

*Did the “correct direction” audio and tactile cues help you stay on the route?* All participants were enthusiastic about the sound/vibration

backdrop (though P6 didn't notice the vibration, but only the sound). Several commented that it gave them reassurance that they were on the correct route. P3 remarked that a few times, due to traffic noise, he could not hear the sound, but could still feel the vibration. P5 and P7 commented that some may find the sound overwhelming, in which case there should be a setting to turn the sound off. A few participants also commented here that they liked the Route Compass (which uses a haptic tap to indicate the direction to the next goalpost).

*Were you able to follow and understand the verbal notifications and audio descriptions? Did they impede you when navigating around the station? What wording or pieces of information were confusing, if any?* All participants said that this information helped and that it was not overwhelming (though some again mentioned the need for a verbosity level.) P7 mentioned that he would not want to have to tap a button all the time, so he welcomed the unsolicited notifications. P5 said that he was not yet "fluent with the language" of the app, but felt that with some more experience he would be able to process all of the information produced. A similar feeling was expressed by P3 and by P4. P4 said that while there was a lot of information, the information was useful, and she would listen to the message from the *Where Am I* button multiple times until she could process it all.

## 5 DISCUSSION

RouteNav was designed to help blind travelers navigate challenging indoor/outdoor scenarios, where GPS can have very poor quality or be unavailable, and walkers need to traverse complex routes that require negotiating ramps, tunnels, T-junctions or even openings through fences. No additional infrastructure that could facilitate localization was considered, because this infrastructure may not have been installed, or may malfunction. RouteNav does not use visual information from the smartphone's camera, which means that users can hold the phone as best suits them (e.g., in their pocket), rather than aiming so as to obtain a clear view of the environment. Instead, we used state-of-the-art inertial tracking algorithms, combined with particle filtering to fuse this data with GPS while leveraging prior spatial knowledge of the location.

RouteNav's localization module, which fuses data from GPS and the phone's inertial sensors, while also using knowledge of the location of walls and other no-go areas, can to some extent correct GPS errors, and enables tracking even in GPS-denied areas such as the underground passageway (see Fig. 8). Still, due to the occasionally poor quality of GPS signal, as well as to the nature of dead-reckoning systems such as RoNIN, localization was inaccurate at certain locations (with two reported cases of complete tracking failure; see Sec. 4.3.1). It is conceivable that more refined algorithms or better sensors could produce better localization. However, we believe that one should never count on accurate localization everywhere. This would be the case even if additional infrastructure were available (e.g., a BLE beacon may run out of power) or when using visual recognition (e.g., a fixture could have been moved after a visual survey, changing the visual aspect of the scene; or, the presence of a crowd could impede clear view).

We believe that mitigation of poor localization accuracy necessarily hinges on a well-designed user interface. More precisely: the

interface (1) should not assume perfect knowledge of the user's location, and (2) should help users make informed decisions about where to move next. It is critical that system designers do not view users of a wayfinding system as "passive agents", simply reacting to the system's directions. Instead, information from the system needs to be mediated by the user's own perception and judgment. A good interface is one that supports this process.

Multiple provisions of RouteNav were designed to deal with poor localization accuracy. The "switching goalpost" strategy (Sec. 3.1.3) helps reduce directional angular errors by ensuring that the target goalpost is always beyond a certain threshold distance. The "user as a sensor" tunnel/ramp alert notifications are very effective at breaking the ambiguity caused by multimodal particle distributions. It is also important to note that users of RouteNav are never given a "turn here" or similar (e.g., "turn in 1 meter") direction: instead, they are informed about the direction to the next goalpost through multiple modalities, solicited or unsolicited, as explained in Sec. 3.3. This choice was informed by two observations. First, large portions of the walkable area were located in relatively open spaces. "Turn here" directions, while appropriate, for example, in buildings characterized by corridor networks, would be less effective here. Second, issuing a "turn here" direction is only meaningful if the user is assumed to be precisely localized. A blind walker following this direction would expect to find an opening; if instead they hit a wall with their cane (because the notification was given too early or too late), they would then have to search left and right to figure out where the opening is. This is amplified by the difficulty experienced by some blind walkers to precisely follow prescribed turns [5]. Importantly, these occurrences may reduce confidence in the navigation system.

While providing directions to the target goalpost is appropriate in most cases, there are situations that require active "collaboration" for successful negotiation. Consider for example the case of Fig. 3, bottom panel. In this case, the user is approaching a junction where they need to turn left. When the user is close enough to the junction, the further goalpost of the next tile is chosen as the target. At this point, the direction to the target goalpost goes through a wall, until the user clears the corner. The participants of our study were informed of these possible situations, and of the fact that they would need to "work with the system" while finding a way to turn the corner by trailing the wall. By and large, the participants were able to deal with this nuisance well (though P2, as mentioned in Sec. 4.4, explicitly complained about these occurrences.)

Another provision of RouteNav that helps dealing with location inaccuracies is the paradigm of "tile navigation". At each point in time, the user's estimated location is mapped to a discrete tile. Routing decisions (choosing the next tile), directional information (direction to the target goalpost in the next tile), and contextual information (whether the tile is a crosswalk, ramp, or tunnel, along with points of interest), are all based on the tile the user is assumed to be located in. To mitigate occasional localization errors, tile sizes are chosen to be larger than the expected radius of localization uncertainty, and tile assignment is robustified through hysteresis (Sec. 3.2.5). Our current approach to tile definition, which follows the basic criteria described in Sec. 3.1.1, is admittedly ad-hoc. We plan to study more systematic principles for tile definition in future work.

In the current implementation of RouteNav, users must be able to specify their starting location and initial walking direction. This is necessary for initial calibration, as explained in Sec. 3.2.3. In practice, one could use clearly identifiable features for this purpose. For example, users may be asked to find the main entrance to a building, then, once entered, walk straight for a few steps. Other transition features (escalators, elevator) could be leveraged for system initialization (e.g., walk straight for a few steps at the bottom of the escalator; turn right and walk for a few steps when exiting the elevator).

Observations from our study brought to light the importance of contextual information provided to users about the area near their specific location (i.e., the tile they are assigned to), and especially about the areas they will traverse next (the next tiles in the route). After feedback from P1, P2 and P3 (see Sec. 4.2.3), we included verbal information in both solicited and unsolicited notifications pertaining to the presence of a tunnel, ramp, or crosswalk they would encounter, and indications of how to negotiate them (the ramp entrance will be to the left, you will need to walk through the crosswalk or enter the tunnel). The fact that most participants were able to negotiate the most challenging features (see Tab. 3) after hearing these notifications, and in spite of poor localization, shows that participants relied on their own sensorial and cognitive abilities, leveraging the directions from the phone while assessing how much to trust them. This suggests that generation of verbal descriptions of local spatial contexts (beyond the simple features used in RouteNav) could tremendously improve the usability of future wayfinding systems. We believe that this could be an exciting new research direction, which could leverage recent progress on embodied AI for navigation [62, 63], as well as “narrative maps” methodologies developed by experienced Orientation & Mobility professionals.

An important question is whether a system like RouteNav could be used in the “real world”. As shown in Tab. 3, for each participant there was at least one “challenging spot” where intervention from the experimenter was necessary (if only to remind the participant to pay attention to the notifications produced by the app). Would the participants have been able to reach the destination, had the experimenter not been there, ready to help? It’s fair to say that, with more time, at least some of the participants would eventually have succeeded in negotiating those challenging situations (e.g., finding the entrance to a tunnel, correcting their own path if they found themselves entering a road). In other cases (e.g., when the system failed with P6, Fig. 9), help from a sighted bystander would be needed. Still, we believe that RouteNav can offer substantially better guidance than existing systems that do not require installation of special infrastructure such as BLE beacons, or extensive visual surveys to enable vision-based localization. RouteNav’s localization is far more accurate than what is available with Apple or Google APIs, thanks to the particle filter, which leverages spatial priors, and to the fusion with inertial data processed by RoNIN, which enables guidance even in GPS-denied areas. Local contextual information provided with each tile, an important feature of RouteNav, is simply not available in other existing systems. Importantly, the cost and effort of adding this information (specifically, defining tiles, annotated with features and points of interest; mapping impenetrable walls and no-go areas; and mapping GPS-denied zones)

is modest and, in our opinion, does not hamper scalability of this approach. In the case considered for this study, we were able to map tiles, walls and no-go areas from aerial images available in Google Maps, while GPS-denied zones were defined based on a local survey (comparing locations produced by GPS with actual ones). This is substantially less onerous than conducting a complete fingerprinting survey to record BLE or Wi-Fi RSSI fingerprints, or a visual survey to obtain a ground-level accurate 3-D representation of the whole hub. Of course, if other localization systems are available in the same area, RouteNav could be easily adapted to leverage this information. For example, if visual-based localization were available, users navigating the area could keep the phone in their pocket and be tracked by RouteNav, while occasionally pulling out the phone to take a picture of the environment and get a new “fix” from visual localization when necessary.

## 6 CONCLUSIONS

We introduced RouteNav, a wayfinding app designed to be used by blind travelers navigating indoor-outdoor transit hubs. Compared to other similar system, RouteNav does not require the presence of specialized infrastructure (e.g. BLE beacons). RouteNav uses data from the inertial sensors of the smartphone, which is processed using RoNIN, a state of the art dead-reckoning machine learning algorithm. Spatial data is fused with information from GPS (when available and reliable), and the user’s location is tracked using particle filtering. Navigation in RouteNav is structured via routes of tiles with switching goalposts, where descriptions of the tiles (and of the point of interest contained within) are provided to walkers as they progress along the route. The user interface, designed over multiple iterations guided by the input from our participants, provides directions and local information through redundant modalities. User studies were conducted in realistic situations with 7 blind participants, where each participant was tasked with navigating three routes that included taking an underground passageway, walking on ramps, crossing crosswalks, and finding entrances in fences. All participants were able to complete the routes (except for one case of tracking failure), though intervention from the experimenter were often needed in a few “challenging spots” along the routes.

## ACKNOWLEDGMENTS

Research reported in this publication was supported by the National Eye Institute of the National Institutes of Health under award number R01EY029260-01, and by the National Science Foundation under award IIP-1632158. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or of the National Science Foundation. We would like to thank the participants who volunteered for this study.

## REFERENCES

- [1] Ali Abdolrahmani, William Easley, Michele Williams, Stacy Branham, and Amy Hurst. 2017. Embracing errors: Examining how context of use impacts blind individuals’ acceptance of navigation aid errors. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 4158–4169.
- [2] Iyad Abu Doush, Sawsan Alshatnawi, Abdel-Karim Al-Tamimi, Bushra Alhasan, and Safaa Hamasha. 2017. ISAB: integrated indoor navigation system for the blind. *Interacting with Computers* 29, 2 (2017), 181–202.

- [3] Dragan Ahmetovic, Federico Avanzini, Adriano Baratè, Cristian Bernareggi, Marco Ciardullo, Gabriele Galimberti, Luca A Ludovico, Sergio Mascetti, and Giorgio Presti. 2023. Sonification of navigation instructions for people with visual impairment. *International Journal of Human-Computer Studies* 177 (2023), 103057.
- [4] Dragan Ahmetovic, Cole Gleason, Chengxiang Ruan, Kris Kitani, Hironobu Takagi, and Chieko Asakawa. 2016. NavCog: a navigational cognitive assistant for the blind. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 90–99.
- [5] Dragan Ahmetovic, Uran Oh, Sergio Mascetti, and Chieko Asakawa. 2018. Turn right: Analysis of rotation errors in turn-by-turn navigation for individuals with visual impairments. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*. 333–339.
- [6] Ilias Apostolopoulos, Navid Fallah, Eelke Folmer, and Kostas E Bekris. 2014. Integrated online localization and navigation for people with visual impairments using smart phones. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3, 4 (2014), 1–28.
- [7] Shiri Azenkot, Richard E Ladner, and Jacob O Wobbrock. 2011. Smartphone haptic feedback for nonvisual wayfinding. In *The proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility*. 281–282.
- [8] Jan Balata, Jakub Berka, and Zdenek Mikovec. 2018. Indoor-outdoor intermodal sidewalk-based navigation instructions for pedestrians with visual impairments. In *International Conference on Computers Helping People with Special Needs*. Springer, 292–301.
- [9] Billie L Bentzen, Janet M Barlow, and Lee S Tabor. 2000. *Detectable warnings: Synthesis of US and international practice*. US Access board Washington, DC.
- [10] Joydeep Biswas and Manuela Veloso. 2010. Wifi localization and navigation for autonomous indoor mobile robots. In *2010 IEEE international conference on robotics and automation*. IEEE, 4379–4384.
- [11] Lesley A Brabyn and John A Brabyn. 1983. An evaluation of "Talking Signs" for the blind. *Human Factors* 25, 1 (1983), 49–53.
- [12] Chris Calori and David Vanden-Eynden. 2015. *Signage and wayfinding design: a complete guide to creating environmental graphic design systems*. John Wiley & Sons.
- [13] Hsuan-Hsuan Chang. 2013. Wayfinding strategies and tourist anxiety in unfamiliar destinations. *Tourism Geographies* 15, 3 (2013), 529–550.
- [14] Seyed Ali Cheraghi, Ali Almadan, and Vinod Namboodiri. 2019. Cityguide: A seamless indoor-outdoor wayfinding system for people with vision impairments. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility*. 542–544.
- [15] Seyed Ali Cheraghi, Vinod Namboodiri, and Laura Walker. 2017. GuideBeacon: Beacon-based indoor wayfinding for the blind, visually impaired, and disoriented. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 121–130.
- [16] Dorin Comaniciu and Peter Meer. 2002. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence* 24, 5 (2002), 603–619.
- [17] Tom Cox, Jonathan Houdmont, and Amanda Griffiths. 2006. Rail passenger crowding, stress, health and safety in Britain. *Transportation Research Part A: Policy and Practice* 40, 3 (2006), 244–258.
- [18] Ryan Crabb, Seyed Ali Cheraghi, and James M Coughlan. 2023. A Lightweight Approach to Localization for Blind and Visually Impaired Travelers. *Sensors* 23, 5 (2023), 2701.
- [19] Bill Crandall and James Marston. 2018. -Development, Evaluation, and Lessons Learned: A Case Study of Talking Signs® Remote Infrared Audible Signage. In *Assistive Technology for Blindness and Low Vision*. CRC Press, 123–150.
- [20] Daniele Croce, Laura Giarre, Federica Pascucci, Ilenia Timirello, Giovanni Ettore Galio, Domenico Garlisi, and Alice Lo Valvo. 2019. An indoor and outdoor navigation system for visually impaired people. *IEEE Access* 7 (2019), 170406–170418.
- [21] Adele Crudden, Jennifer L Cmar, and Michele C McDonnell. 2017. Stress associated with transportation: A survey of persons with visual impairments. *Journal of Visual Impairment & Blindness* 111, 3 (2017), 219–230.
- [22] Arnaud Doucet, Simon Godsill, and Christophe Andrieu. 2000. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and computing* 10, 3 (2000), 197–208.
- [23] Navid Fallah, Ilias Apostolopoulos, Kostas Bekris, and Eelke Folmer. 2012. The user as a sensor: navigating users with visual impairments in indoor spaces using tactile landmarks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 425–432.
- [24] Alexander Fiannaca, Ilias Apostolopoulos, and Eelke Folmer. 2014. Headlock: a wearable navigation aid that helps blind cane users traverse large open spaces. In *Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility*. 19–26.
- [25] German Flores, Sri Kurniawan, Roberto Manduchi, Eric Martinson, Lourdes M Morales, and Emrah Akin Sisbot. 2015. Vibrotactile guidance for wayfinding of blind walkers. *IEEE transactions on haptics* 8, 3 (2015), 306–317.
- [26] German Flores and Roberto Manduchi. 2018. Easy return: an app for indoor backtracking assistance. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [27] Michael D Fowler. 2013. Soundscape as a design strategy for landscape architectural praxis. *Design studies* 34, 1 (2013), 111–128.
- [28] Giovanni Fusco and James M Coughlan. 2020. Indoor localization for visually impaired travelers using computer vision on a smartphone. In *Proceedings of the 17th International Web for All Conference*. 1–11.
- [29] Aura Ganz, James Schafer, Siddhesh Gandhi, Elaine Puleo, Carole Wilson, and Meg Robertson. 2012. PERCEPT indoor navigation system for the blind and visually impaired: architecture and experimentation. *International journal of telemedicine and applications* 2012 (2012).
- [30] Nicholas A Giudice, William E Whalen, Timothy H Riehle, Shane M Anderson, and Stacy A Doore. 2019. Evaluation of an accessible, real-time, and infrastructure-free indoor navigation system by users who are blind in the mall of america. *Journal of Visual Impairment & Blindness* 113, 2 (2019), 140–155.
- [31] Reginald G Golledge, Roberta L Klatzky, Jack M Loomis, Jon Speigle, and Jerome Tietz. 1998. A geographical information system for a GPS based personal guidance system. *International Journal of Geographical Information Science* 12, 7 (1998), 727–749.
- [32] Reginald G Golledge, Jack M Loomis, Roberta L Klatzky, Andreas Flury, and Xiao Li Yang. 1991. Designing a personal guidance system to aid navigation without sight: Progress on the GIS component. *International Journal of Geographical Information System* 5, 4 (1991), 373–395.
- [33] João Guerreiro, Dragan Ahmetovic, Daisuke Sato, Kris Kitani, and Chieko Asakawa. 2019. Airport accessibility and navigation assistance for people with visual impairments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [34] Janne Haverinen and Anssi Kemppainen. 2009. Global indoor self-localization based on the ambient magnetic field. *Robotics and Autonomous Systems* 57, 10 (2009), 1028–1035.
- [35] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [36] Sachini Herath, Hang Yan, and Yasutaka Furukawa. 2020. RoNIN: Robust neural inertial navigation in the wild: Benchmark, evaluations, & new methods. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3146–3152. Code available at <https://github.com/Sachini/ronin>.
- [37] Md Elias Hossain, Khandker M Qaiduzzaman, and Mostafijur Rahman. 2020. Sightless helper: an interactive mobile application for blind assistance and safe navigation. In *Cyber Security and Computer Science: Second EAI International Conference, ICONCS 2020, Dhaka, Bangladesh, February 15-16, 2020, Proceedings 2*. Springer, 581–592.
- [38] Marco Iosa, Augusto Fusco, Giovanni Morone, and Stefano Paolucci. 2012. Effects of visual deprivation on gait dynamic stability. *The Scientific World Journal* 2012 (2012).
- [39] Elmar Krainz, Viktoria Lind, Werner Moser, and Markus Dornhofer. 2016. Accessible way finding on mobile devices for different user groups. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. 799–806.
- [40] Bineeth Kuriakose, Ida Marie Ness, Maja Å skov Tengstedt, Jannicke Merete Svendsen, Terese Bjørseth, Bijay Lal Pradhan, and Raju Shrestha. 2023. Turn Left Turn Right-Delving type and modality of instructions in navigation assistant systems for people with visual impairments. *International Journal of Human-Computer Studies* (2023), 103098.
- [41] Masaki Kuribayashi, Tatsuya Ishihara, Daisuke Sato, Jayakorn Vongkulbhisal, Karnik Ram, Seita Kayukawa, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. 2023. PathFinder: Designing a Map-less Navigation System for Blind People in Unfamiliar Buildings. In *ACM CHI Conference on Human Factors in Computing Systems*.
- [42] Young Hoon Lee and Gérard Medioni. 2016. RGB-D camera based wearable navigation system for the visually impaired. *Computer vision and Image understanding* 149 (2016), 3–20.
- [43] Richard G Long and EW Hill. 1997. Establishing and maintaining orientation for mobility. In *Foundations of orientation and mobility*, BB Blash, WR Wiener, and RL Welsh (Eds.). AFB Press, New York.
- [44] Jack M Loomis, Reginald G Golledge, and Roberta L Klatzky. 1998. Navigation system for the blind: Auditory display modes and guidance. *Presence* 7, 2 (1998), 193–203.
- [45] Jiangyan Lu, Kin Wai Michael Siu, and Ping Xu. 2008. A comparative study of tactile paving design standards in different countries. In *2008 9th International Conference on Computer-Aided Industrial Design and Conceptual Design*. IEEE, 753–758.
- [46] Roberto Manduchi and James M Coughlan. 2014. The last meter: blind visual guidance to a target. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3113–3122.

- [47] Elliot W Martin, Emily Farrar, Evan Magsig, Susan A Shaheen, Ashley Auer, Sarah Hoban, Booz Allen Hamilton, et al. 2020. *Accessible Transportation Technologies Research Initiative (ATTRI) Impact Assessment White Paper*. Technical Report. United States. Department of Transportation. Intelligent Transportation ...
- [48] Alejandro Santos Martínez-Sala, Fernando Losilla, Juan Carlos Sánchez-Aarnoutse, and Joan García-Haro. 2015. Design, implementation and evaluation of an indoor navigation system for visually impaired people. *Sensors* 15, 12 (2015), 32168–32187.
- [49] Peng Ren, Fatemeh Elyasi, and Roberto Manduchi. 2021. Smartphone-based inertial odometry for blind walkers. *Sensors* 21, 12 (2021), 4033.
- [50] Yafei Ren and Xizhen Ke. 2010. Particle filter data fusion enhancements for MEMS-IMU/GPS. *Intelligent Information Management 2* (2010), 417–421.
- [51] Transparency Market Research. 2020. Tactile Paving Market - Global Industry Analysis, Size, Share, Growth, Trends, and Forecast, 2020-2030. Retrieved on 9/4/2022 from <https://www.transparencymarketresearch.com/tactile-paving-market.html>.
- [52] Timothy H Riehle, Shane M Anderson, Patrick A Lichter, Nicholas A Giudice, Suneel I Sheikh, Robert J Knuesel, Daniel T Kollmann, and Daniel S Hedin. 2012. Indoor magnetic navigation for the blind. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 1972–1975.
- [53] Timothy H Riehle, Shane M Anderson, Patrick A Lichter, William E Whalen, and Nicholas A Giudice. 2013. Indoor inertial waypoint navigation for the blind. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 5187–5190.
- [54] David A Ross and Bruce B Blasch. 2000. Wearable interfaces for orientation and wayfinding. In *Proceedings of the fourth international ACM conference on Assistive technologies*. 193–200.
- [55] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. 2019. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12716–12725.
- [56] Daisuke Sato, Uran Oh, Kakuya Naito, Hironobu Takagi, Kris Kitani, and Chieko Asakawa. 2017. Navcog3: An evaluation of a smartphone-based blind indoor navigation assistant with semantic features in a large-scale environment. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*. 270–279.
- [57] Farzaneh Shahini, Vanessa Nasr, and Maryam Zahabi. 2022. A Friendly Indoor Navigation App for People with Disabilities (FIND). In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 66. SAGE Publications Sage CA: Los Angeles, CA, 1922–1926.
- [58] Geraldo Silveira, Ezio Malis, and Patrick Rives. 2008. An efficient direct approach to visual SLAM. *IEEE transactions on robotics* 24, 5 (2008), 969–979.
- [59] Edo Smitshuijzen. 2007. *Signage Design Manual*. Lars Mueller Publishers.
- [60] Hajime Taira, Masatoshi Okutomi, Torsten Sattler, Mircea Cimpoi, Marc Pollefeys, Josef Sivic, Tomas Pajdla, and Akihiko Torii. 2018. InLoc: Indoor visual localization with dense matching and view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7199–7209.
- [61] Sebastian Thrun. 2002. Probabilistic robotics. *Commun. ACM* 45, 3 (2002), 52–57.
- [62] Su Wang, Ceslee Montgomery, Jordi Orbay, Vighnesh Birodkar, Aleksandra Faust, Izzeddin Gur, Natasha Jaques, Austin Waters, Jason Baldrige, and Peter Anderson. 2022. Less is More: Generating Grounded Navigation Instructions from Landmarks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15428–15438.
- [63] Erik Wijmans, Manolis Savva, Irfan Essa, Stefan Lee, Ari S Morcos, and Dhruv Batra. 2023. Emergence of Maps in the Memories of Blind Navigation Agents. *arXiv preprint arXiv:2301.13261* (2023).
- [64] Chris Yoon, Ryan Louie, Jeremy Ryan, MinhKhang Vu, Hyegi Bang, William Derksen, and Paul Ruvolo. 2019. Leveraging augmented reality to create apps for people with visual disabilities: A case study in indoor navigation. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. 210–221.
- [65] Maryam Zahabi, Xi Zheng, Azima Maredia, and Farzaneh Shahini. 2022. Design of Navigation Applications for People with Disabilities: A Review of Literature and Guideline Formulation. *International Journal of Human-Computer Interaction* (2022), 1–23.
- [66] Yuan Zhuang, Jun Yang, You Li, Longning Qi, and Naser El-Sheimy. 2016. Smartphone-based indoor localization with bluetooth low energy beacons. *Sensors* 16, 5 (2016), 596.