

GUIDENAV: User-Informed Development of a Vision-Only Robotic Navigation Assistant For Blind Travelers

Hochul Hwang
hochulhwang@umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Soowan Yang
tntn6771@dgist.ac.kr
DGIST
Daegu, Republic of Korea

Jahir Sadik Monon
jmonon@umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Nicholas A Giudice
nicholas.giudice@maine.edu
University of Maine
Orono, Maine, USA

Sunghoon Ivan Lee
silee@cs.umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Joydeep Biswas
joydeepb@cs.utexas.edu
University of Texas at Austin
Austin, Texas, USA

Donghyun Kim
donghyunkim@umass.edu
University of Massachusetts Amherst
Amherst, Massachusetts, USA




Figure 1: We developed a vision-only route following system considering technical requirements for practical guide robots. Our guide dog robot achieved kilometer-scale guiding with blind and low-vision (BLV) users. Robust path following performance and positive feedback from BLV users mark an important milestone toward the real-world deployment of guide dog robots.

Abstract

While commendable progress has been made in user-centric research on mobile assistive systems for blind and low-vision (BLV) individuals, references that directly inform robot navigation design remain rare. To bridge this gap, we conducted a comprehensive human study involving interviews with 26 guide dog handlers, four white cane users, nine guide dog trainers, and one O&M trainer, along with 15+ hours of observing guide dog-assisted walking. After de-identification, we open-sourced the dataset to promote human-centered development and informed decision-making for assistive systems for BLV people. Building on insights from this

formative study, we developed GUIDENAV, a vision-only, teach-and-repeat navigation system. Inspired by how guide dogs are trained and assist their handlers, GUIDENAV autonomously repeats a path demonstrated by a sighted person using a robot. Specifically, the system constructs a topological representation of the taught route, integrates visual place recognition with temporal filtering, and employs a relative pose estimator to compute navigation actions—all without relying on costly, heavy, power-hungry sensors such as LiDAR. In field tests, GUIDENAV consistently achieved kilometer-scale route following across five outdoor environments, maintaining reliability despite noticeable scene variations between teach and repeat runs. A user study with 3 guide dog handlers and 1 guide dog trainer further confirmed the system's feasibility, marking (to our knowledge) the first demonstration of a quadruped mobile system retrieving a path in a manner comparable to guide dogs.

CCS Concepts

• **Human-centered computing** → **Field studies**; **Accessibility technologies**; • **Computer systems organization** → **Robotic autonomy**; • **Computing methodologies** → **Vision for robotics**.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '26, Edinburgh, UK

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>

Keywords

Assistive Robotics, Guide Dog Robot, Navigation

ACM Reference Format:

Hochul Hwang, Soowan Yang, Jahir Sadik Monon, Nicholas A Giudice, Sunghoon Ivan Lee, Joydeep Biswas, and Donghyun Kim. 2026. GUIDENav: User-Informed Development of a Vision-Only Robotic Navigation Assistant For Blind Travelers. In *Proceedings of International Conference on Human-Robot Interaction (HRI '26)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

What should be the first step in developing a mobile navigation assistant system for blind and low-vision (BLV) individuals? The answer would be to identify the most meaningful and important needs of BLV travelers. While commendable strides have been made in BLV user-centered development of navigation systems [6, 26, 31, 33, 39, 40, 56, 63, 68], defining a navigation system that is truly practical for BLV individuals remains challenging. Building a navigation assistant system for BLV travelers is not simply a matter of replicating autonomous driving systems—like a delivery robot using dense map-based navigation moving from any point A to B on a prebuilt map [41]. Prior studies on BLV navigation assistance showed that the level of navigation assistance can vary widely, from mobility-only support (i.e., obstacle avoidance) to full autonomy (i.e., governing navigation direction) [56]. Other works [40, 68] investigated exploratory navigation setups that allow BLV travelers to visit new places, such as exhibitions, with an additional scene description assistance. Studies on interactive navigation of guide dogs and handlers have also highlighted the importance of trust and a sense of agency as prerequisites for introducing autonomous navigation to BLV people [31].

Such prior research provides valuable insights into BLV navigation relating to robotic guides; however, many detailed technical questions remain unanswered: What is the proper sensor suite—can LiDAR be an appropriate option? What conditions (e.g., terrain types, weather, lighting) must the robot handle? How much distance should the system support for daily navigation? And is it valid to assume that users have prior knowledge of the places they navigate? Answering such questions requires multi-dimensional considerations. For example, evaluating the suitability of LiDAR depends on power consumption, onboard computing capacity, operating conditions, and cost. These factors, in turn, are shaped by BLV users' daily navigation patterns: how far they typically walk, whether they commute in rain, fog, or snow, and the acceptable robot weight (which constrains battery and computer size).

These questions may never be fully resolved until a system is deployed to end users. Despite extensive research on spatial cognition and navigation in psychology, cognitive science [20], and HRI, resources capturing the lived practices of BLV travelers remain scarce. This constrains the translation of theoretical insights into practical, human-centered navigation systems. By examining BLV users' daily activities, commuting habits, common destinations, and the shortcomings of current mobility aids, development errors can be significantly reduced. Unfortunately, open-source data capturing this information is minimal, creating a major obstacle for human-centered research on mobile navigation systems.

To bridge this gap, we conducted a comprehensive formative study involving interviews with 26 guide dog handlers, four cane users, nine guide dog trainers, and one O&M trainer, along with 15+ hours of observing how guide dog handlers navigate with their dogs. We specifically examined how BLV individuals interact with guide dogs, which informed the design of our pragmatic navigation assistant. The dataset, termed GUIDEDATA, was carefully anonymized, organized for ease of use, and open-sourced via a project webpage to promote human-centered research and inform design decisions in the development of assistive systems for BLV people.

Building on insights from our formative study, we aimed to define an ideal navigation system to support the independent mobility of BLV individuals. We selected a quadruped robot platform to ensure high mobility across diverse terrains, and decided not to rely on GPS or LiDAR. LiDAR is costly in terms of price, computing power, and weight, which shortens operating time and increases overall system cost. It also demonstrates limited robustness under adverse weather conditions (e.g., rain, fog, snow) and when dealing with transparent objects (e.g., glass doors, shop windows, plastic barriers) or reflective surfaces (e.g., water puddles, metallic doors, shiny tiled floors), encountered in everyday environments. In addition, LiDAR-based navigation algorithms typically require dense, prebuilt 3D maps, which demand significant computation and intensive memory for embedded systems [54]. GPS, meanwhile, is unreliable in urban environments and indoors due to signal degradation near tall buildings as well as its inherent top-down measurement limitations.

Instead of global, metric map-based navigation methods [5, 18, 38], we propose a navigation framework, titled GUIDENav, following a Visual Teach-and-Repeat (VT&R) system [23]. In this setup, an operator demonstrates routes by driving the robot using a remote controller, after which the robot autonomously repeats the demonstrated path for a BLV user. This framework aligns closely with how guide dogs are trained and work with their handlers when they are familiarized with the environment. Moreover, daily activities typically involve routine routes, and our navigation framework specifically targets these common scenarios as a starting point. During the teach phase, the robot records a stream of images and constructs a topological map (instead of a dense metric map) by saving keyframe images along the path. Using the image-based subgoals (i.e., keyframe images) enables vision-only navigation using a single RGB camera on a standalone system (see Fig. 1).

While promising, VT&R involves a fundamental challenge: vulnerability to scene variation. Because the algorithm identifies the robot's location and selects subgoals by comparing observed scenes to saved keyframes, environmental changes (e.g., lighting, object relocation) can drastically reduce retrieval accuracy, often causing the robot to lose its path [23]. To mitigate this, we adopt the strategy used in [7, 66]: leveraging visual place recognition (VPR) embeddings [1, 48] instead of raw image pixel matching [13]. VPR descriptors improve robustness to scene variation by incorporating semantic understanding learned from internet-scale datasets. In addition, we integrate Reloc3r [15], a deep neural network for camera pose estimation, to compute the robot's position and orientation relative to target keyframes. These estimates are then used to compute low-level control commands (i.e., linear and angular velocities). We found that Reloc3r outperforms dense feature-matching

approaches [19, 43, 65], which often suffer from degraded accuracy when depth information is unreliable (common in outdoor environments) or fail under scene variations.

Field tests of our robot demonstrate that our system can achieve kilometer-scale, long-range autonomous navigation, surpassing prior guide robot demonstrations that relied on more complex sensing and mapping [6, 68]. GUIDENAV requires no metric maps, GPS, LiDAR, or depth sensing, keeping the platform lightweight and low-cost. By chaining short, visually guided segments, GUIDENAV avoids drift accumulation and adapts to environmental changes, maintaining robustness even with modest computing resources—an important factor for extended operating time.

In user studies, GUIDENAV successfully guided BLV individuals along multiple outdoor routes, including one measuring 2.02 km round trip. To our knowledge, this represents the first demonstration of kilometer-scale outdoor autonomous guiding for BLV participants using a camera-only system. Participants described the experience as easy to use, while professional trainers noted both its reliability and differences from guide dogs. Workload and usability assessments confirmed low cognitive demand and high satisfaction for most users, while also pointing to areas for refinement and potential applications. In summary, our contributions are threefold:

- (1) A large open-source dataset containing rich information about guide dog–handler interactions, real-world navigation challenges, and daily activities of BLV travelers in navigational contexts.
- (2) A lightweight, visual teach-and-repeat framework that eliminates reliance on GPS, LiDAR, and depth sensing.
- (3) We present the first real-world validation of kilometer-scale, vision-only autonomous navigation with BLV participants on their daily routes, measuring turning success rates, interventions, and collisions, and demonstrating robust long-range performance under diverse outdoor conditions.

2 Related Works

2.1 Existing Technological Solutions

Un-actuated systems have been designed in various form factors, such as smart canes [9, 21, 60, 63, 69, 70, 74, 75, 84] and wearable or handheld devices [2, 12, 30, 34, 45, 55, 56, 59, 64, 85]. While these devices are lightweight and compact, they typically offer only limited navigation assistance—for example, notifying users of nearby obstacles or local paths, but no goal-directed navigation. More critically, the use of acoustic or vibrotactile feedback (rather than physical pulling or pushing) often causes distraction and mental fatigue, raising serious safety concerns, as BLV individuals rely heavily on environmental sounds for localization and situational awareness [4].

In contrast, *actuated systems* can assist BLV individuals in a manner more similar to sighted guides or animal guide dogs. Such systems can enable smooth collision avoidance, higher walking speeds, and shorter travel distances [6, 10, 51]. Since the mid-1970s, wheeled mobile robots have been the most widely studied form factor for actuated systems [24, 26, 35, 37, 40, 42, 49, 50, 52, 67, 68, 72, 73]. However, wheeled systems face a critical limitation: they cannot reliably handle uneven terrains such as stairs or curbs [68], making them impractical for real-world outdoor navigation where assistance is most needed.

Recent advances in quadruped robots have opened new opportunities for mobility assistance. Unlike wheeled systems, quadrupeds provide natural and efficient locomotion across uneven terrains and heterogeneous environments [8, 27, 53, 81]. However, existing studies often give limited attention to users’ needs, largely due to a lack of understanding of how guide dogs and their handlers work. For instance, [8, 81] explored the use of a soft leash to pull a person, whereas both our formative study and prior literature [26, 31, 32] emphasize that handlers require a rigid harness handle to perceive immediate feedback on the guide dog’s motion for safety. Similarly, some studies positioned the handler behind the robot during navigation [6, 27], whereas in practice guide-dog handlers walk beside and one step behind the dog’s front feet, enabling them to assess the environment when the dog stops for safety [31].

In summary, quadruped robots hold substantial potential compared to other assistive platforms, yet current guide-dog robot research lacks a comprehensive user-centered perspective. Our objective is to develop a pragmatic guide-dog robot that matches—and ultimately surpasses—the abilities of animal guide dogs, providing BLV individuals with a reliable option for independent navigation.

2.2 Vision-based Navigation and VT&R

While prior VT&R studies demonstrated kilometer-scale path following [23], they often relied on precise metric localization or expensive sensors. For example, many VT&R variants used LiDAR to register the robot’s pose [23]. More recently, there has been a major departure from such approaches, with new methods aiming to be lightweight and vision-centric. Fischer et al. [13] proposed a bio-inspired scheme that leverages odometry and minimal visual correction (sparse cues) to navigate reliably with low computational cost. However, methods employing naive template matching (e.g., Normalized Cross Correlation [44]) are fragile to visually similar patterns and often produce false matches under appearance ambiguities.

FFI-VTR [76] advanced this line of work by constructing a topological keyframe graph of feature descriptors, minimizing pixel-level feature distance between matches without requiring global localization. PlaceNav [66] further reframed VT&R subgoal selection as an image-retrieval task: it uses VPR embeddings (CosPlace [1]) to pick the next topological subgoal and applies Bayesian filtering for temporal consistency [82]. While PlaceNav shares some similarity with our GuideNav in subgoal selection, GuideNav also employs a trained network for pose estimation, which enhances robustness to scene and lighting variations and enables reliable long-range navigation.

3 Formative Study Dataset

Existing qualitative research spans guide dog ownership [57, 79], interaction [11, 29, 86], navigation and communication [25, 47, 63, 80]; in parallel, developers of quadruped and other mobility assistive systems report insights from participatory design [6, 39], experiments [14, 36, 46, 77], and observations [17, 32, 35]. Few prior studies have released in-depth demonstration videos [46] or human-robot interaction data [6, 36]; however, there remains a significant gap in large-scale, open-source qualitative resources, especially

those documenting handler–guide dog observations or detailed trainer/handler interviews [8, 27, 81].

To address this gap and inform the human-centered design of GUIDENAV, we conducted a formative study documenting interviews, navigation, video observations, the matching process, and blindfolded walking sessions. Our de-identified and curated dataset, termed GUIDEDATA, comprising 39 interview transcripts, 31 images, and 126 videos with a total duration of over 15 hours, has been open-sourced via our project webpage.

3.1 Procedure

The dataset includes detailed documentation for each data type and is composed of five key elements:

- **Interviews with guide dog handlers:** Semi-structured Q&A sessions on how they use guide dogs in daily life (transcripts).
- **Interviews with guide dog trainers:** Semi-structured Q&A sessions focusing on how guide dogs are trained and deployed to BLV individuals (transcripts).
- **Observation sessions of guide dog–assisted navigation:** Videos capturing how handlers and their guide dogs navigate through environments they walk in regularly (videos).
- **Observation sessions of matching training:** Videos documenting how an experienced handler and a new guide dog undergo the initial familiarization process (videos).
- **Blindfolded walking sessions by the authors:** Under a guide-dog trainer’s supervision, we conducted blindfolded walks with trained guide dogs to experience handler–dog interaction and the navigational cues needed for safe travel. The video recordings document dog-provided support (e.g., locating sidewalks) versus decisions requiring human judgment (e.g., assessing street-crossing safety), yielding rich evidence on how animal guide dogs deliver navigation assistance.

3.2 Participants

GUIDEDATA includes semi-structured interviews with 11 male and 21 female BLV individuals, and observation sessions involving seven guide dog users and one cane user. Participants had an average age of 64.63 years and roughly 24.21 years of experience using navigation aids, with the majority having experience with multiple guide dogs. The dataset also contains interviews, and observation sessions with nine professional guide dog trainers and one O&M specialist.

To ensure the dataset is informative while protecting privacy, we report key demographic attributes such as gender, age, and experience, while anonymizing participant identities. Further demographic details are available in Supplementary Material 1.

3.3 Data Postprocessing

To protect privacy, all video recordings are post-processed to blur participants’ faces. For observation videos and images, we applied an automatic face anonymization algorithm [16, 83] to obscure participants and visible license plates. All frames were carefully reviewed to correct any failures of the automated process using a video editing tool. For interview transcripts, we removed all personally identifiable information. The resulting dataset is available on Kaggle under a CC0 license for unrestricted public use. Furthermore,

we ensured that any author-identifiable details were anonymized in the dataset, webpage, and Kaggle repository. To our knowledge, this is the first comprehensive dataset of interviews and observational recordings that captures rich information on how guide dogs provide assistance in practice. This dataset establishes a strong foundation for defining the specifications of robotic guide dogs and guiding the development of assistive navigation systems.

4 Technical Requirements for Guide Robots

Designing a practical guide robot for blind and low-vision travelers is uniquely challenging. Relatively low employment rates demand cost-sensitive solutions, yet the robot must reliably handle diverse real-world environments encountered in daily travel. Grounded in our GUIDEDATA, we highlight the critical hardware and navigation requirements that define the specifications of a deployable guide robot. While many dimensions are relevant, we focus on hardware and navigation as foundational; other aspects such as communication, personalization, or user interfaces [31] remain also important but are deferred for future work.

4.1 Hardware Constraints

4.1.1 Runtime and responsiveness. The system must operate for at least two hours [31] at normal walking speeds (up to at least 1.4 m/s) [3], with all perception and control onboard to avoid network dependence. We propose that end-to-end (perception to control) latency must remain under 100 ms so that safety-critical commands (e.g., “stop,”) are executed promptly and trusted by the handler. Memory must support storage of large-scale maps.

4.1.2 Form factor and weight. The robot must pass through narrow doors, fit under seats, and be light enough for one person to lift when necessary [78], yet heavy and stable enough to safely halt an adult in emergencies. Legged or wheel–legged locomotion is preferred to negotiate curbs, stairs, and uneven terrain that are encountered in daily travels.

4.1.3 Minimal Sensor Suite: Vision-Only Approach. A forward-facing RGB camera is essential for capturing rich visual cues. In contrast, LiDAR – though common in guide robot research – adds cost, power draw, and vulnerability to weather (e.g., drizzle, fog, snow) that corrupt point-cloud returns which can be problematic as handlers may get caught in rain and snow (GH03, GH12). Handlers reported accidents caused by transparent hazards, breaking one’s own leg after falling from ice, or glass doors, which LiDAR often fails to detect. Similarly, GPS proved unreliable in dense urban areas where several participants reported not using GPS-based apps due to frequent signal loss. These findings underscore the rationale for a vision-only approach, as the most practical and scalable option for daily navigation.

4.2 Navigation Requirements for Guide Robots

Navigation for BLV travelers imposes requirements distinct from conventional mobile robotics. A guide robot must compactly encode long routes across varied terrains, adhere to orientation principles emphasized in mobility training, and reliably negotiate the diverse obstacles encountered in sidewalk environments. These requirements motivate lightweight, robust approaches such as Visual Teach




Figure 2: GUIDENAV teach-and-repeat framework. In the *teach* phase, an expert performs remote control of the robot along the desired route. The robot records a sequence of images, and the topomap generator constructs a topological map by extracting keyframe images. In the *repeat* phase, the robot performs visual place recognition by matching current image features, f_t , extracted by the feature extractor Φ , with the topomap using cosine similarity, and then selects a subgoal. The pose estimator, $\mathcal{R}$, computes the relative pose between the current and subgoal images. GUIDENAV achieved kilometer-scale autonomous navigation under varying scene conditions. The system was also used by three guide dog handlers and one guide dog trainer, who provided positive feedback on its potential to support mobility assistance for BLV individuals.

and Repeat (VT&R) [23], where routes are compactly stored and repeated under changing conditions.

4.2.1 Mapping and Perception. Guide robots must represent kilometer-scale routes spanning both structured indoor corridors and unstructured outdoor paths (e.g., cracked sidewalks, curbs, stairs). Compact map representations (e.g., visual descriptors) are preferable to dense 3D point clouds for efficiency. Perception must remain reliable across lighting, weather, and seasonal changes, including transitions from daylight to low-light operation.

4.2.2 Domain Knowledge for BLV Travel. Navigation must integrate orientation and mobility (O&M) practices. For instance, before street crossings the robot should stop at curb edges within one stride, allowing the traveler to probe with their foot and decide when to proceed. O&M training emphasizes maintaining straight-line travel and executing turns only at right angles on command. In areas without sidewalks, teams often “shoreline” the left road edge to preserve orientation – patterns the robot should replicate.

4.2.3 Obstacle Avoidance. Trainers specify that safe clearance requires at least 0.9 m laterally and 1.8 m vertically. Beyond standard obstacles, guide robots must detect overhanging (branches, signs, wires) and transparent hazards (glass panes, puddles, ice) that often challenge vision systems. Handler interviews also highlighted everyday obstructions – manholes, construction equipment, and parked cars – that block pedestrian paths and must be negotiated.

5 Proposed GUIDENAV Architecture

Unlike monolithic end-to-end policies, our approach integrates modular components—topological mapping, visual place recognition, relative pose regression, and control—to achieve kilometer-scale outdoor route following on embedded hardware.

5.1 Problem Formulation

We define the guide dog route following problem as enabling a robot to safely repeat expert-taught pedestrian routes using only onboard sensing and compute. Formally, the robot state lies in $SE(2)$, $x_t = (p_t, \psi_t)$, where $p_t \in \mathbb{R}^2$ is position and ψ_t is orientation. At each timestep, the robot observes I_t through a camera.

During the **teach phase**, an expert demonstrates a safe route $\Gamma = [I_1, \dots, I_N]$. From the demonstration, we construct a sparse topological map $\mathcal{M} = \{(I_i, f_i)\}_{i=1}^N$, where $I_i \in \mathbb{R}^{H \times W \times 3}$ are keyframes and f_i are visual embeddings extracted from a place recognition encoder. The keyframes are extracted using our topological map generator along with a manual filtering process detailed in Section 5.2.

During the **repeat phase**, the robot estimates the most likely subgoal $I_{g_i} \in \mathcal{M}$ and estimates the relative transformation $\xi_{t \rightarrow g_i} \in se(2)$. This relative pose is converted into velocity commands (v_t, w_t) under a nonlinear feedback controller, producing safe closed-loop motion along the expert-taught route.

GUIDENAV functions more like a global planner, autonomously following expert-taught subgoals without GPS or dense metric maps. By design, the system does not include a dedicated local planner, but its modular architecture allows one to be seamlessly integrated for detours or dynamic avoidance when required. This balance of simplicity and modularity ensures both efficiency on embedded hardware and adaptability for future extensions.

5.2 Topological map generator

Each collected frame I_t is encoded using the DINOv3 foundation model [62] $\mathcal{E}$ into an ℓ_2 -normalized global descriptor $z_t = \mathcal{E}(I_t)$. We select keyframes with an *adaptive selector* that jointly enforces (1) minimum temporal spacing, (2) appearance diversity over a recent buffer via cosine similarity, and (3) adaptive similarity thresholds that adjust based on trajectory progress. This embedding-driven approach yields compact yet discriminative topological maps, retaining only a small fraction of frames while improving robustness under viewpoint/illumination changes and reducing memory footprint (e.g., ~24 MB for 1 km travel). A human operator may filter low-quality keyframes (e.g., motion-blurred), though this step could be automated in future work.

5.3 Visual place recognition with temporal consistency

Each incoming image I_t is encoded by CosPlace [1] into a normalized descriptor $f_t = \Phi(I_t; \theta) \in \mathbb{R}^{512}$. Cosine similarity enables efficient retrieval of candidate subgoals. To prevent erratic frame switching, we enforce temporal consistency through a belief update

Table 1: User Study Participant Demographics

ID	Age	Gender	Vision Level	Experience*	Role
H01	63	F	Totally blind	35	User
H02	66	F	Legally blind	9	User
H03	67	M	Totally blind	15	User
T01	37	F		10	Evaluator
T02	35	F		4	User
T03	28	F		7	Evaluator

* Indicates the number of total years the participant trained guide dogs.

over topological states, following the approach of Suomela et al. [66]. Based on observation, this yields smoother subgoal selection and mitigates the kidnapping problem [71].

5.4 Direct relative pose estimation

Given a current frame I_t and retrieved subgoal I_{g_i} , we estimate their relative pose using ReLoc3r [15] as our pose regression network $\mathcal{R}$:

$$\mathcal{R}(I_t, I_{g_i}; \Theta) = T_{t \rightarrow g_i} \in SE(3), \quad (1)$$

with $T_{t \rightarrow g_i} = \begin{bmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix}$. The transform is projected onto the ground plane to yield

$$\xi = [\Delta x, \Delta y, \Delta \psi]^\top \in \mathbb{R}^3 \quad (2)$$

a compact 2D displacement and yaw representation directly consumed by the controller. Rather than relying on Perspective-n-Point (PnP) [22], which requires dense feature matching and is vulnerable to low texture and depth errors, using a direct regression approach provides robust pose estimates suitable for long-range navigation.

5.5 Goal Position to Velocity Commands

We model the robot with unicycle kinematics, where the relative pose $\xi = [\Delta x, \Delta y, \Delta \psi]^\top$ is expressed in polar form:

$$\rho = \sqrt{\Delta x^2 + \Delta y^2}, \quad \alpha = \arctan 2(\Delta y, \Delta x), \quad \beta = \psi - \alpha, \quad (3)$$

with ρ the goal distance, α the heading error, and β the orientation misalignment. A nonlinear feedback law regulates these terms:

$$v_{\text{raw}} = v_{\text{max}} \left(1 - e^{-k_\rho \rho} \right), \quad (4)$$

$$w_{\text{raw}} = k_\alpha \alpha + k_\beta \beta, \quad (5)$$

where $(k_\rho, k_\alpha, k_\beta)$ are control gains. For stability, we further apply three heuristics: reduced speed near the goal, velocity damping under large heading errors, and amplified orientation correction in final alignment. Finally, coordinated scaling enforces velocity limits, yielding smooth, bounded trajectories consistent with unicycle dynamics and reliable teach-and-repeat navigation.

6 Evaluation of GUIDENAV Components

We evaluated GUIDENAV in both indoor and outdoor environments to assess the robustness of its key components. As shown in Fig. 3, the topological map maintained reliable node prediction under lighting changes and partial occlusions.

6.1 Topological Map Node Prediction

We first examined the effect of different keyframe selection strategies on route following. Specifically, we compared a simple distance-based heuristic with topological maps constructed from visual embeddings. Learning-based embeddings consistently improved repeat reliability over the distance heuristic, with DINOv3 providing the most robust localization across varied lighting and viewpoint conditions.

6.2 Relative Pose Estimation

We next compared direct relative camera pose regression with a feature-matching + PnP pipeline. While dense matching can yield accurate correspondences, it frequently failed under textureless surfaces and extremely bright conditions. In contrast, direct regression achieved higher stability, lower failure rates, and smoother control signals, making it better suited for long-horizon outdoor navigation.

7 Full System Evaluation with Stakeholders

We evaluated our full system in realistic contexts through two complementary studies. First, we conducted a field study with three BLV individuals, all experienced guide dog handlers. Second, we replicated an official guide dog evaluation by testing the robot in a standard training environment: a blindfolded guide dog trainer served as the user, while another trainer acted as evaluator. Our objectives were twofold: (1) to assess the effectiveness and acceptance of the guide dog robot among BLV users, and (2) to benchmark performance against the criteria applied to animal guide dogs for real-world deployment.

7.1 Participants

We recruited three experienced guide dog handlers, each capable of independently navigating their local communities without a sighted guide. All relied primarily on a guide dog as their mobility aid. Eligibility criteria were (1) visual acuity at or below the legal blindness threshold defined in the Social Security Act §1614 [61], and (2) at least six months of experience working with a guide dog. Demographic details are shown in Table 1. H01 and H03 were totally blind (with H03 retaining light perception), while H02 had residual vision (20/2200 in one eye, lower in the other). In addition, three professional guide dog trainers were recruited, each with a minimum of four years' experience. T01 observed the robot guiding a handler and provided evaluation as an expert bystander. T02 and T03 jointly completed the standard guide dog evaluation course, with T02 blindfolded in the role of user and T03 serving as evaluator.

7.2 Guide Dog Robot

Our guide dog robot is a compact, stand-alone quadrupedal platform designed for both indoor and outdoor route following. Our system is built on the Unitree Go2 robot equipped with a rigid harness handle that allows a rigid connection with the user. A single Intel RealSense D435i camera is used for GUIDENAV. Relative pose estimation is performed using Reloc3r-512 [15] with mixed-precision inference to achieve a high accuracy and inference speed. GUIDENAV runs at 5 Hz on a NVIDIA Jetson AGX Orin.




Figure 3: GUIDENAV outputs during guidance of H02, showing topological map node prediction and relative pose estimation. The system correctly predicts the node despite drastic lighting changes and unseen objects, and maintains accurate pose estimation.

Table 2: Robot Solo and Handler-Guide Dog Robot Team Navigation Results

Trajectory ID	Environment		Performance		
	# TSR	Distance (m)	# Collisions	# Interventions	Time
Robot Solo (FidelcoFront → Storage)	5/6	130 m	0	2	4 min 8 s
Robot Solo (FidelcoFront → Parking)	6/7	160 m	0	1	4 min 2 s
Robot Solo (LongWalkLoop)	17/17	1670 m	0	1 ^{†‡}	44 min 58 s
Robot Solo (ResidentialLoop)	13/14	430 m	1	3	31 min 12 s
Robot Solo (ClassroomLoop-Indoor)	4/6	30 m	0	3*	2 min 42 s
Total	45/50 (90%)	2420 m	1	10	87 min 2 s
H01+Robot (Lot → Library)	4/4	390 m	0	0	12 min 14 s
H01+Robot (Library → Lot)	3/3	390 m	0	0 [†]	12 min 54 s
H02+Robot (Home → BusStop)	7/7	70 m	0	0	3 min 10 s
H02+Robot (Dining → FineArts)	3/3	670 m	0	0	13 min 43 s
H02+Robot (FineArts → Dining)	2/2	670 m	0	1 [†]	18 min 6 s
H03+Robot (Home → Church)	8/8	1010 m	0	1 [‡]	21 min 5 s
H03+Robot (Church → Home)	8/8	1010 m	0	3 [‡]	24 min 3 s
T02+Robot (Guide dog evaluation loop)	11/11	730 m	1	1 [‡]	18 min 57 s
Total	46/46 (100%)	4940 m	1	6	124 min 12 s

[†] Technical malfunction (camera disconnection/overheating) - excluded from intervention count; [‡] Single guidance intervention required for unexpected obstacle avoidance; ^{*} Single physical intervention required to remove entangled cables from robot leg; TSR stands for Turning Success Rate

7.3 Procedure

The experiment consisted of four phases:

- (1) **Guide dog demonstrations:** We observed that all handlers (H01–H03) walked familiar community routes with their guide dogs. H01 walked from the parking lot drop-off to her workplace (Lot ↔ Library). H02 began walking from home to the bus stop (Home → BusStop) but switched to another route (Dining ↔ FineArts) to avoid drawing neighbors' attention. H03 walked a residential loop from home to the end of the sidewalk near a church which was over a kilometer one-way (Home ↔ Church). Additionally, trainer T02 walked a standard evaluation route blindfolded under the supervision of evaluator T03.
- (2) **GUIDENAV Teach (Data collection with robot):** Based on these observations, a researcher teleoperated the robot along each route, with minor adjustments for efficiency and safety (e.g., walking on the edge of a road instead of the center, by-passing unnecessary stops such as poles). The camera mounted on the robot recorded video to construct the topological map.
- (3) **GUIDENAV Repeat (Autonomous guiding trials):** GUIDENAV then autonomously guided the users. Prior to each trial, handlers adjusted the harness handle (length and angle) and practiced walking with the robot (all handlers except for the trainer had prior walking experience with similar prototypes). On a single command, the robot started and navigated to the destination. Users were notified just before the start and after the robot stopped when it reached the goal. One researcher

logged data using a laptop for debug purposes (does not affect GUIDENAV performance), and another video-recorded the walk.

- (4) **Post-Trial Interviews and Surveys:** After each walk and a brief break, participants completed a semi-structured interview and workload/usability surveys. Interviews were audio-recorded, transcribed using a speech-to-text service [58]. We note that large language model-based generative AI tools were used only to improve the clarity and style of writing.

7.4 Findings

All participants successfully reached the goal in their familiar daily routes within their local community with the guidance of the guide dog robot. The distance and the complexity of the daily routes varied as detailed in Table 2. For example, H03 successfully walked with the robot for over 2 km (1 km one-way) where he usually goes for a walk with his guide dog. H02 in the other hand, walked for a short distance (70 m) in her residency area, however, there were multiple angled and sharp turns to complete the path. The data collection time and GUIDENAV runtime was held on a different day and time, and various changes in the environment (e.g., parked car, a crowd of pedestrians passing by - jogging, riding bicycles, light changes) made VT&R more challenging due to appearance change. In total, six interventions were required for guiding the user over 4.9 km.

7.4.1 Satisfying and Easy Outdoor Travel in Familiar Routes. GuideNav enabled all BLV handlers to successfully complete their familiar community routes, ranging from 70 m to over 2 km round trip, with minimal interventions. Participants consistently expressed surprise that the robot was fully autonomous, even through routes with multiple sharp turns. As H01 remarked, *“Really it was a hundred percent itself? Even around plant planters, that was all on its own. Wow. It was a success.”* Similarly, H02 reflected after a repeat trial, *“So it really did know where it was going and we did do the route, wow.”* Walking with the robot was described as straightforward and low effort. H03 summarized, *“It seems like it’s got some potential for it, for sure. It wasn’t hard or anything. It was pretty easy.”* Both handlers and trainers highlighted the robot’s ability to negotiate uneven terrain. During observation of H03 walking with his guide dog, we noted surface tree roots. And later in the teach phase, we teleoperated the robot to safely maneuver over them, and in the repeat phase, the robot autonomously guided the user across the same spot. This led H03 to remark it was *“pretty seamless,”* and T01 to observe, *“Yeah, it actually avoided [them].”* Collectively, participants framed the experience as *“easy to walk with”* and emphasized its significance beyond laboratory trials. As H03 put it, *“I honestly think that either yours or this will be the future of blindness, the mobility of blindness.”* These reactions highlight not only the robot’s technical robustness but also its perceived potential as a practical mobility aid.

7.4.2 Comparison with Guide Dogs: Trainers’ Perspectives. Professional trainers evaluated our guide dog robot against standard guide dog assessment criteria, noting both promising capabilities and differences. The blindfolded trainer (T02) described the experience as sturdy and enjoyable, while the evaluator (T03) concluded, *“She looked comfortable. I thought it was very safe in its movements. I didn’t feel like she was going to run into anything.”* An experienced

Table 3: Perceived workload (raw 0–20 NASA–TLX scores)

Component	H01	H02	H03	T02	Median
Mental demand	4	5 → 4	8	10	6.25
Physical demand	1	1 → 2	5	10	3.25
Temporal demand	2	0 → 0	5	5	3.50
Performance	1	0 → 5	5	4	3.25
Effort	2	5 → 2	7	16	5.25
Frustration	1	10 → 4	1	3	2.00
Total	11	21 → 17	31	48	25.00

trainer (T01), who had previously seen the robot under manual control, remarked on its progress and usability: *“Compared to two years ago, the technology has come a long way as far as the actual usability of it in real life. I was impressed by its ability to navigate [...] the concept of putting in a route and following it specifically was very impressive.”*

At the same time, T02 emphasized that the robot’s movement remained fundamentally different from that of a guide dog, describing it as more abrupt and less fluid: *“To put it very basically, the movement of the robot is just very robotic. It doesn’t have the organic flow that a dog does. Even when you’re approaching a curve with a dog, you can feel through the harness that the dog is slowing. With the robot, you get fewer of those body cues because the robot doesn’t have them.”*

7.4.3 Workload and Usability Evaluation. NASA–TLX instrument [28] results showed a clear divide between BLV handlers and professional trainers (Table 3). All three BLV handlers (H01–H03) reported very low workload (median TLX scores below 6 for mental, physical, and frustration), describing the robot’s guidance as requiring minimal cognitive or physical effort. In contrast, the blindfolded trainer (T02) rated workload much higher (48/100), reflecting both a stricter professional benchmark (guide dog expectations) and first-time use of the robot. One handler (H02) showed a large reduction in workload across sessions, with mental demand, effort, and frustration dropping substantially in her second trial, aligning with contextual factors – her first route was disrupted by neighborhood distractions, although it did not affect GUIDENAV’s performance.

System Usability Scale (SUS) scores also varied. Two handlers (H01, H03) rated the system extremely high (92.5 and 95.0), describing it as intuitive and effortless. H02 rated it much lower (32.5 → 40.0), reflecting initial difficulties and possibly reduced comfort due to her residual vision when told the robot would guide autonomously. The trainer (T02) rated the robot at 52.5, notably lower than her benchmark for a guide dog (80.0), underscoring a usability gap. Overall, BLV users experienced the system as easy and non-frustrating (low TLX, high SUS), while trainers applied stricter standards. These results suggest that GUIDENAV can achieve very high usability for some users, but evaluations must account for adaptation, personal context, and professional expectations.

8 Discussions and Conclusion

Our study shows that GUIDENAV can already deliver usable guidance in familiar community routes, yet several design refinements

remain critical. Both handlers and trainers emphasized the need for smoother motion cues and explicit landmark indication (e.g., signaling curbs, crossings, or turns) to better support orientation practices taught in mobility training. Beyond these improvements, trainers envisioned diverse applications: as a lower-maintenance alternative for those unable or unwilling to care for a guide dog; as a training tool for instructors to assess mobility skills, balance, and pace before matching a person with a guide dog; as a temporary substitute when users are between guide dogs (getting a new one); and as a flexible solution for pre-planned routes in settings such as hospitals, conferences, or community centers. Because the system requires only a single camera, it could be embodied in multiple robot forms and deployed at scale, broadening access for both cane users and guide dog users. Together, these findings highlight not only the feasibility of vision-only teach-and-repeat navigation, but also its potential to fill important gaps in assistive mobility, opening new directions for human-centered guide robot design.

References

- [1] Gabriele Berton, Carlo Masone, and Barbara Caputo. 2022. Rethinking Visual Geo-localization for Large-Scale Applications. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (New Orleans, LA, USA). IEEE, Piscataway, NJ, USA, 4878–4888.
- [2] Shripad S. Bhatlawande, Jayant Mukhopadhyay, and Manjunatha Mahadevappa. 2012. Ultrasonic spectacles and waist-belt for visually impaired and blind person. In *2012 National Conference on Communications (NCC)* (Kalyani, India). IEEE, Piscataway, NJ, USA, 1–4. doi:10.1109/NCC.2012.6176765
- [3] Richard W Bohannon and A Williams Andrews. 2011. Normal walking speed: a descriptive meta-analysis. *Physiotherapy* 97, 3 (2011), 182–189.
- [4] Nicholas A Bradley and Mark D Dunlop. 2002. Investigating context-aware clues to assist navigation for visually impaired people. In *Proceedings of Workshop on Building Bridges: Interdisciplinary Context-Sensitive Computing*, University of Glasgow. University of Glasgow, Glasgow, UK, 6 pages.
- [5] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. 2016. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics* 32, 6 (2016), 1309–1332.
- [6] Shaojun Cai, Ashwin Ram, Zhengtai Gou, Mohd Alqama Wasim Shaikh, Yu-An Chen, Yingjia Wan, Kotaro Hara, Shengdong Zhao, and David Hsu. 2024. Navigating Real-World Challenges: A Quadruped Robot Guiding System for Visually Impaired People in Diverse Environments. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 44, 18 pages. doi:10.1145/3613904.3642227
- [7] Luis G. Camara, Tomáš Pivoňka, Martin Jilek, Carl Gäbert, Karel Košnar, and Libor Přeucil. 2020. Accurate and Robust Teach and Repeat Navigation by Visual Place Recognition: A CNN Approach. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 6018–6024. doi:10.1109/IROS45743.2020.9341764
- [8] Hua Chen, Zejun Hong, Shunpeng Yang, Patrick M Wensing, and Wei Zhang. 2023. Quadruped capturability and push recovery via a switched-systems characterization of dynamic balance. *IEEE Transactions on Robotics* 39, 3 (2023), 2111–2130.
- [9] Tzu-Kuan Chuang, Ni-Ching Lin, Jih-Shi Chen, Chen-Hao Hung, Yi-Wei Huang, Chunhui Teng, Haikun Huang, Lap-Fai Yu, Laura Giarre, and Hsueh-Cheng Wang. 2018. Deep trail-following robotic guide dog in pedestrian environments for people who are blind and visually impaired-learning from virtual and real worlds. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, IEEE, Piscataway, NJ, USA, 5849–5855.
- [10] DD Clark-Carter, AD Heyes, and CI Howarth. 1986. The efficiency and walking speed of visually impaired people. *Ergonomics* 29, 6 (1986), 779–789.
- [11] Peter J Craigon, Pru Hobson-West, Gary C W England, Chantelle Whelan, Emma Lethbridge, and Lucy Asher. 2017. “She’s a dog at the end of the day”: Guide dog owners’ perspectives on the behaviour of their guide dog. *PLoS One* 12, 4 (April 2017), e0176018.
- [12] Dimitrios Dakopoulos and Nikolaos G Bourbakis. 2009. Wearable obstacle avoidance electronic travel aids for blind: a survey. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 40, 1 (2009), 25–35.
- [13] Dominic Dall’Osto, Tobias Fischer, and Michael Milford. 2021. Fast and robust bio-inspired teach and repeat navigation. In *2021 IEEE/RSJ international conference on intelligent robots and systems (iros)*. IEEE, IEEE, Piscataway, NJ, USA, 500–507.
- [14] David DeFazio, Eisuke Hirota, and Shiqi Zhang. 2023. Seeing-Eye Quadruped Navigation with Force Responsive Locomotion Control. In *Proceedings of The 7th Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 229)*, Jie Tan, Marc Toussaint, and Kourosh Darvish (Eds.). PMLR, Proceedings of Machine Learning Research, 2184–2194. <https://proceedings.mlr.press/v229/defazio23a.html>
- [15] Siyan Dong, Shuzhe Wang, Shaohui Liu, Lulu Cai, Qingnan Fan, Juho Kannala, and Yanchao Yang. 2025. Reloc3r: Large-scale training of relative camera pose regression for generalizable, fast, and accurate visual localization. In *Proceedings of the Computer Vision and Pattern Recognition Conference* (Nashville, TN, USA). IEEE, Piscataway, NJ, USA, 16739–16752.
- [16] Martin Drawitsch. 2020. *deface: Video anonymization by face detection*. <https://github.com/ORB-HD/deface> Originally published at <https://github.com/mdraw/deface>.
- [17] Brian L. Due. 2023. Guide dog versus robot dog: assembling visually impaired people with non-human agents and achieving assisted mobility through distributed co-constructed perception. *Mobilities* 18, 1 (2023), 148–166. arXiv:<https://doi.org/10.1080/17450101.2022.2086059> doi:10.1080/17450101.2022.2086059
- [18] Hugh Durrant-Whyte and Tim Bailey. 2006. Simultaneous localization and mapping: part I. *IEEE Robotics & Automation Magazine* 13, 2 (2006), 99–110.
- [19] Johan Edstedt, Qiyu Sun, Georg Bökman, Märten Wadenbäck, and Michael Felsberg. 2024. RoMa: Robust Dense Feature Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Seattle, WA, USA). IEEE, Piscataway, NJ, USA, 19790–19800.
- [20] Russell A. Epstein, E. Zita Patai, Joshua B. Julian, and Hugo J. Spiers. 2017. The cognitive map in humans: spatial navigation and beyond. *Nature Neuroscience* 20, 11 (2017), 1504–1513. doi:10.1038/nn.4656
- [21] José Faria, Sérgio Lopes, Hugo Fernandes, Paulo Martins, and João Barroso. 2010. Electronic white cane for blind people navigation assistance. In *2010 World Automation Congress*. IEEE, IEEE, Piscataway, NJ, USA, 1–7.
- [22] Martin A Fischler and Robert C Bolles. 1981. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Commun. ACM* 24, 6 (1981), 381–395.
- [23] Paul Furgale and Timothy D. Barfoot. 2010. Visual teach and repeat for long-range rover autonomy. *Journal of Field Robotics* 27, 5 (2010), 534–560.
- [24] Georgios Galatas, Christopher McMurrough, Gian Luca Mariottini, and Fillia Makedon. 2011. eyeDog: an assistive-guide robot for the visually impaired. In *Proceedings of the 4th International Conference on Pervasive Technologies Related to Assistive Environments* (Heraklion, Crete, Greece) (PETRA '11). Association for Computing Machinery, New York, NY, USA, Article 58, 8 pages. doi:10.1145/2141622.2141691
- [25] Florence Gaunet and Sandie Besse. 2019. Guide dogs’ navigation after a single journey: A descriptive study of path reproduction, homing, shortcut and detour. *PLoS One* 14, 7 (July 2019), e0219816.
- [26] João Guerreiro, Daisuke Sato, Saki Asakawa, Huixu Dong, Kris M. Kitani, and Chieko Asakawa. 2019. CaBot: Designing and Evaluating an Autonomous Navigation Robot for Blind People. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility* (Pittsburgh, PA, USA) (ASSETS '19). Association for Computing Machinery, New York, NY, USA, 68–82. doi:10.1145/3308561.3353771
- [27] Kaveh Akbari Hamed, Vinay R Kamidi, Wen-Loong Ma, Alexander Leonessa, and Aaron D Ames. 2019. Hierarchical and safe motion control for cooperative locomotion of robotic guide dogs and humans: A hybrid systems approach. *IEEE Robotics and Automation Letters* 5, 1 (2019), 56–63.
- [28] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, Amsterdam; New York, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [29] Marion A. Hersh and Michael A. Johnson. 2010. A robotic guide for blind people. Part 1. A multi-national survey of the attitudes, requirements and preferences of potential end-users. *Appl. Bionics Biomechanics* 7, 4 (Oct. 2010), 277–288. doi:10.1080/11762322.2010.523626
- [30] Noriaki Hirose, Amir Sadeghian, Marynel Vázquez, Patrick Goebel, and Silvio Savarese. 2018. Gonet: A semi-supervised deep learning approach for traversability estimation. In *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, IEEE, Piscataway, NJ, USA, 3044–3051.
- [31] Hochul Hwang, Hee-Tae Jung, Nicholas A Giudice, Joydeep Biswas, Sunghoon Ivan Lee, and Donghyun Kim. 2024. Towards Robotic Companions: Understanding Handler-Guide Dog Interactions for Informed Guide Dog Robot Design. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 596, 20 pages. doi:10.1145/3613904.3642181
- [32] Hochul Hwang, Tim Xia, Ibrahim Keita, Ken Suzuki, Joydeep Biswas, Sunghoon I Lee, and Donghyun Kim. 2023. System configuration and navigation of a guide dog robot: Toward animal guide dog-level guiding work. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, IEEE, Piscataway, NJ, USA,

- 9778–9784.
- [33] Rie Kamikubo, Seita Kayukawa, Yuka Kaniwa, Allan Wang, Hernisa Kacorri, Hironobu Takagi, and Chieko Asakawa. 2025. Beyond Omakase: Designing Shared Control for Navigation Robots with Blind People. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 671, 17 pages. doi:10.1145/3706598.3714112
 - [34] Robert K Katzschmann, Brandon Araki, and Daniela Rus. 2018. Safe local navigation for visually impaired users with a time-of-flight and haptic feedback device. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 26, 3 (2018), 583–593.
 - [35] Seita Kayukawa, Keita Higuchi, João Guerreiro, Shigeo Morishima, Yoichi Sato, Kris Kitani, and Chieko Asakawa. 2019. BBeep: A Sonic Collision Avoidance System for Blind Travellers and Nearby Pedestrians. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300282
 - [36] J. Taery Kim, Wenhao Yu, Yash Kothari, Bruce Walker, Jie Tan, Greg Turk, and Sehoon Ha. 2023. Transforming a Quadruped into a Guide Robot for the Visually Impaired: Formalizing Wayfinding, Interaction Modeling, and Safety Mechanism. In *Proceedings of The 7th Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 229)*, Jie Tan, Marc Toussaint, and Kourosh Darvish (Eds.). PMLR, Proceedings of Machine Learning Research, 2288–2303. <https://proceedings.mlr.press/v229/kim23c.html>
 - [37] Vladimir Kulyukin, Chaitanya Gharpure, Pradnya Sute, Nathan De Graw, and John Nicholson. 2004. A robotic wayfinding system for the visually impaired. In *Proceedings of the 16th Conference on Innovative Applications of Artificial Intelligence* (San Jose, California) (IAAI'04). AAAI Press, Menlo Park, CA, USA, 864–869.
 - [38] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. 2011. g2o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 3607–3613.
 - [39] 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 *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 41, 16 pages. doi:10.1145/3544548.3580687
 - [40] Masaki Kuribayashi, Kohei Uehara, Allan Wang, Shigeo Morishima, and Chieko Asakawa. 2025. WanderGuide: Indoor Map-less Robotic Guide for Exploration by Blind People. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 676, 21 pages. doi:10.1145/3706598.3713788
 - [41] Joan Lläne. 2020. How Starship Delivery Robots know where they are going. <https://medium.com/starshiptechnologies/how-starship-delivery-robots-know-where-they-are-going-c97d385a1015>.
 - [42] Tachi Lab. n.d. *Guide Dog Robot (MELDOG)*. Tachi Laboratory, The University of Tokyo. <https://tachilab.org/en/projects/meldog.html>
 - [43] Vincent Leroy, Yann Cabon, and Jérôme Revaud. 2024. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*. Springer Nature Switzerland, Cham, 71–91.
 - [44] J. P. Lewis. 2001. Fast Normalized Cross-Correlation. *Industrial Light & Magic* 10 (2001), 7 pages.
 - [45] Bing Li, J. Pablo Muñoz, Xuejian Rong, Jizhong Xiao, Yingli Tian, and Aries Ardit. 2016. ISANA: Wearable Context-Aware Indoor Assistive Navigation with Obstacle Avoidance for the Blind. In *Computer Vision – ECCV 2016 Workshops*, Gang Hua and Hervé Jégou (Eds.). Springer International Publishing, Cham, 448–462.
 - [46] Shuijing Liu, Aamir Hasan, Kaiwen Hong, Runxuan Wang, Peixin Chang, Zachary Mizrahi, Justin Lin, D Livingston McPherson, Wendy A Rogers, and Katherine Driggs-Campbell. 2024. Dragon: A dialogue-based robot for assistive navigation with visual language grounding. *IEEE Robotics and Automation Letters* 9, 4 (2024), 3712–3719.
 - [47] Janice KF Lloyd, Steven La Grow, Kevin J Stafford, and R Claire Budge. 2008. The guide dog as a mobility aid part 2: Perceived changes to travel habits. *Vision Rehabilitation International* 1, 1 (2008), 34–45.
 - [48] Stephanie Lowry, Niko Sünderhauf, Paul Newman, William Maddern, Michael Milford, Duncan Stewart, John J Leonard, David Cox, Peter Corke, and Gordon Wyeth. 2016. Visual place recognition: A survey. *IEEE Transactions on Robotics* 32, 1 (2016), 1–19.
 - [49] Rajesh Kannan Megalingam, Souraj Vishnu, Vishnu Sasikumar, and Sajikumar Sreekumar. 2019. Autonomous Path Guiding Robot for Visually Impaired People. In *Cognitive Informatics and Soft Computing*, Pradeep Kumar Mallick, Valentina Emilia Balas, Akash Kumar Bhoi, and Ahmed F. Zobaa (Eds.). Springer Singapore, Singapore, 257–266.
 - [50] A Allan Melvin, B Prabu, R Nagarajan, Bukhari Illias, A Allan Melvin, B Prabu, R Nagarajan, I Bukhari, and AA Melvin. 2009. ROVI: a robot for visually impaired for collision-free navigation. In *Proc. of the International Conference on Man-Machine Systems (ICOMMS 2009)*. Universiti Malaysia Perlis (UniMAP), Batu Ferringhi, Penang, Malaysia, 3B5–1.
 - [51] Rachel Joy-Taub Miner. 2001. The experience of living with and using a dog guide. *RE: view* 32, 4 (2001), 183.
 - [52] Amal Nanavati, Xiang Zhi Tan, and Aaron Steinfeld. 2018. Coupled Indoor Navigation for People Who Are Blind. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (Chicago, IL, USA) (HRI '18). Association for Computing Machinery, New York, NY, USA, 201–202. doi:10.1145/3173386.3176976
 - [53] NSK Ltd. 2011. NSK Develops a Guide-Dog Style Robot. <https://www.nsk.com/company/news/2011/press111027b.html>. Accessed: 2023-1-17.
 - [54] Chano Park, Peyman Moghadam, Jason L. Williams, Sunwoo Kim, Sridha Sridharan, and Clinton Fookes. 2022. Elasticity Meets Continuous-Time: Map-Centric Dense 3D LiDAR SLAM. *IEEE Transactions on Robotics* 38, 2 (2022), 978–993. doi:10.1109/TRO.2021.3074024
 - [55] Erick Javier Argüello Prada and Lina María Santacruz Forero. 2022. A belt-like assistive device for visually impaired people: Toward a more collaborative approach. *Cogent Engineering* 9, 1 (2022), 2048440. doi:10.1080/23311916.2022.2048440
 - [56] Vinitha Ranganeni, Mike Sinclair, Eyal Ofek, Amos Miller, Jonathan Campbell, Andrey Kolobov, and Edward Cutrell. 2023. Exploring Levels of Control for a Navigation Assistant for Blind Travelers. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) (HRI '23). Association for Computing Machinery, New York, NY, USA, 4–12. doi:10.1145/3568162.3578630
 - [57] K Refson, AJ Jackson, AE Dusoier, and DB Archer. 1999. The health and social status of guide dog owners and other visually impaired adults in Scotland. *Visual Impairment Research* 1, 2 (1999), 95–109.
 - [58] Rev.com. 2024. Rev: Audio & Video Transcription Services. <https://www.rev.com/>. Accessed 2025-09-27.
 - [59] Alberto Rodríguez, J Javier Yebe, Pablo F Alcantarilla, Luis M Bergasa, Javier Almazán, and Andrés Cela. 2012. Assisting the visually impaired: obstacle detection and warning system by acoustic feedback. *Sensors* 12, 12 (2012), 17476–17496.
 - [60] Mohammad Farid Saaid, AM Mohammad, and MSA Megat Ali. 2016. Smart cane with range notification for blind people. In *2016 IEEE International Conference on Automatic Control and Intelligent Systems (I2CACIS)*. IEEE, Piscataway, NJ, USA, 225–229.
 - [61] Social Security. 2010. *Social Security Act §1614*. https://www.ssa.gov/OP_Home/ssact/title16b/1614.htm
 - [62] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Quab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. 2025. Dinov3. *arXiv preprint arXiv:2508.10104* abs/2508.10104 (2025), 67 pages.
 - [63] Patrick Slade, Arjun Tambe, and Mykel J Kochenderfer. 2021. Multimodal sensing and intuitive steering assistance improve navigation and mobility for people with impaired vision. *Science robotics* 6, 59 (2021), eabg6594.
 - [64] Pawel Strumillo, Michal Bujacz, Przemyslaw Baranski, Piotr Skulimowski, Piotr Korbel, Mateusz Owczarek, Krzysztof Tomalczyk, Alin Moldoveanu, and Runar Unnthorsson. 2018. Different Approaches to Aiding Blind Persons in Mobility and Navigation in the “Naviton” and “Sound of Vision” Projects. In *Mobility of Visually Impaired People: Fundamentals and ICT Assistive Technologies*. Springer, Cham, 435–468. doi:10.1007/978-3-319-94595-8_15
 - [65] Jiaming Sun, Zehong Shen, Yuang Wang, Huijun Bao, and Xiaowei Zhou. 2021. LoFTR: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Nashville, TN, USA). IEEE, Piscataway, NJ, USA, 8922–8931.
 - [66] Lauri Suomela, Jussi Kalliola, Harry Edelman, and Joni-Kristian Kämäräinen. 2024. Placenav: Topological navigation through place recognition. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, IEEE, Piscataway, NJ, USA, 5205–5213.
 - [67] Susumu Tachi and Kiyoshi Komoriya. 1985. Guide Dog Robot. *The Robotics Research 2 (The Second International Symposium 1984)* 2 (1985), 333–349.
 - [68] Hironobu Takagi, Kakuya Naito, Daisuke Sato, Masayuki Murata, Seita Kayukawa, and Chieko Asakawa. 2025. Field Trials of Autonomous Navigation Robot for Visually Impaired People. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 678, 8 pages. doi:10.1145/3706599.3706663
 - [69] Hotaka Takizawa, Shotaro Yamaguchi, Mayumi Aoyagi, Nobuo Ezaki, and Shinji Mizuno. 2012. Kinect cane: An assistive system for the visually impaired based on three-dimensional object recognition. In *2012 IEEE/SICE international symposium on system integration (SII)*. IEEE, IEEE, Piscataway, NJ, USA, 740–745.
 - [70] Hotaka Takizawa, Shotaro Yamaguchi, Mayumi Aoyagi, Nobuo Ezaki, and Shinji Mizuno. 2015. Kinect cane: An assistive system for the visually impaired based on the concept of object recognition aid. *Personal and Ubiquitous Computing* 19 (2015), 955–965.

- [71] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. 2005. *Probabilistic Robotics*. MIT Press, Cambridge, MA.
- [72] Kazuteru Tobita, Katsuyuki Sagayama, Mayuko Mori, and Ayako Tabuchi. 2018. Structure and examination of the guidance robot LIGHBOT for visually impaired and elderly people. *Journal of Robotics and Mechatronics* 30, 1 (2018), 86–92.
- [73] Kazuteru Tobita, Katsuyuki Sagayama, and Hironori Ogawa. 2017. Examination of a guidance robot for visually impaired people. *Journal of Robotics and Mechatronics* 29, 4 (2017), 720–727.
- [74] Iwan Ulrich and Johann Borenstein. 2001. The GuideCane-applying mobile robot technologies to assist the visually impaired. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 31, 2 (2001), 131–136.
- [75] Mohd Helmy Abd Wahab, Amirul A Talib, Herdawatie A Kadir, Ayob Johari, Ahmad Noraziah, Roslina M Sidek, and Ariffin A Mutalib. 2011. Smart cane: Assistive cane for visually-impaired people. *arXiv preprint arXiv:1110.5156* 8, 4 (2011), 7 pages.
- [76] Jikai Wang et al. 2025. FFI-VTR: Lightweight and Robust Visual Teach and Repeat Navigation based on Feature Flow Indicator and Probabilistic Motion Planning. *arXiv preprint arXiv:2507.12800* abs/2507.12800 (2025), 6 pages. <https://arxiv.org/abs/2507.12800>
- [77] Luyao Wang, Qihe Chen, Yan Zhang, Ziang Li, Tingmin Yan, Fan Wang, Guyue Zhou, and Jiangtao Gong. 2023. Can quadruped guide robots be used as guide dogs?. In *IROS* (Detroit, MI, USA). IEEE, Piscataway, NJ, USA, 4094–4100.
- [78] Eric B Weston, Alexander M Aurand, Jonathan S Dufour, Gregory G Knapik, and William S Marras. 2020. One versus two-handed lifting and lowering: lumbar spine loads and recommended one-handed limits protecting the lower back. *Ergonomics* 63, 4 (2020), 505–521.
- [79] Lorraine Whitmarsh. 2005. The benefits of guide dog ownership. *Visual impairment research* 7, 1 (2005), 27–42.
- [80] Michele A. Williams, Amy Hurst, and Shaun K. Kane. 2013. "Pray before you step out": describing personal and situational blind navigation behaviors. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility* (Bellevue, Washington) (*ASSETS '13*). Association for Computing Machinery, New York, NY, USA, Article 28, 8 pages. doi:10.1145/2513383.2513449
- [81] Anxing Xiao, Wenzhe Tong, Lizhi Yang, Jun Zeng, Zhongyu Li, and Koushil Sreenath. 2021. Robotic Guide Dog: Leading a Human with Leash-Guided Hybrid Physical Interaction. In *2021 IEEE International Conference on Robotics and Automation (ICRA)* (Xi'an, China). IEEE Press, Piscataway, NJ, USA, 11470–11476. doi:10.1109/ICRA48506.2021.9561786
- [82] Ming Xu, Niko Sünderhauf, and Michael Milford. 2021. Probabilistic Visual Place Recognition for Hierarchical Localization. *IEEE Robotics and Automation Letters* 6, 2 (2021), 311–318.
- [83] Yuanxuan Xu, Wan Yan, Genke Yang, Jiliang Luo, Tao Li, and Jianan He. 2020. CenterFace: joint face detection and alignment using face as point. *Scientific Programming* 2020, 1 (2020), 7845384.
- [84] Cang Ye, Soonhac Hong, Xiangfei Qian, and Wei Wu. 2016. Co-robotic cane: A new robotic navigation aid for the visually impaired. *IEEE Systems, Man, and Cybernetics Magazine* 2, 2 (2016), 33–42.
- [85] Limin Zeng, Markus Simros, and Gerhard Weber. 2017. Camera-based mobile electronic travel aids support for cognitive mapping of unknown spaces. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services* (Vienna, Austria) (*MobileHCI '17*). Association for Computing Machinery, New York, NY, USA, Article 8, 10 pages. doi:10.1145/3098279.3098563
- [86] Yan Zhang, Ziang Li, Haole Guo, Luyao Wang, Qihe Chen, Wenjie Jiang, Mingming Fan, Guyue Zhou, and Jiangtao Gong. 2023. "I am the follower, also the boss": Exploring Different Levels of Autonomy and Machine Forms of Guiding Robots for the Visually Impaired. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–22.