

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

PALMS: Plane-based Accessible Indoor Localization Using Mobile Smartphones

Permalink

<https://escholarship.org/uc/item/7bw6797s>

Authors

Cheng, Yunqian

Manduchi, Roberto

Publication Date

2024-10-14

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

PALMS: Plane-based Accessible Indoor Localization Using Mobile Smartphones

Yunqian Cheng

Computer Science and Engineering
University of California, Santa Cruz
Santa Cruz, United States
Email: ychen827@ucsc.edu

Roberto Manduchi

Computer Science and Engineering
University of California, Santa Cruz
Santa Cruz, United States
Email: manduchi@ucsc.edu

Abstract—In this paper, we present PALMS, an innovative indoor global localization and relocalization system for mobile smartphones that utilizes publicly available floor plans. Unlike most vision-based methods that require constant visual input, our system adopts a dynamic form of localization that considers a single instantaneous observation and odometry data. The core contribution of this work is the introduction of a particle filter initialization method that leverages the Certainly Empty Space (CES) constraint along with principal orientation matching. This approach creates a spatial probability distribution of the device’s location, significantly improving localization accuracy and reducing particle filter convergence time. Our experimental evaluations demonstrate that PALMS outperforms traditional methods with uniformly initialized particle filters, providing a more efficient and accessible approach to indoor wayfinding. By eliminating the need for prior environmental fingerprinting, PALMS provides a scalable and practical approach to indoor navigation.

Index Terms—indoor localization, global localization, relocalization, particle filter, plane detection, CES constraint, mobile smartphones

I. INTRODUCTION

Indoor wayfinding, the ability to determine and follow a path in a building towards a desired destination, is an open area of research, with applications including travel (e.g., finding a gate in an airport), health (e.g., finding a doctor’s office), and recreation (e.g., exploring a museum). Wayfinding is particularly challenging for those who are blind or have low vision, as these individuals cannot rely on visual landmarks for orientation, and cannot easily consult a map of the place.

Different types of smartphone apps for assisted wayfinding are available. Some (like Google Maps and Apple Maps) are general-purpose localization and guidance apps that can work indoors (even when GPS signal is unusable) using radio signal from Wi-Fi access points or Bluetooth Low Energy (BLE) beacons. This requires a prior *fingerprinting* phase which, generally speaking, involves collecting signal measurements (received signal strength or RSSI) over a dense grid of spatial locations in the building. Indoor Atlas relies on “magnetic signatures” for mapping and also requires prior fingerprinting of the environment. GoodMaps is a wayfinding app specifically designed for blind users. It enables indoor mapping using images taken from a smartphone camera, which are matched, using computer vision techniques, against a database of visual

features collected in a prior phase (what could be termed *visual fingerprinting*).

A disadvantage of these technologies is that localization is only available in environments that have been fingerprinted, which represents a small portion of the available buildings. A more scalable approach, one that requires no fingerprinting, is offered by systems that attempt to localize and track the walker using dead-reckoning. For example, Crabb et al. [1] experimented with wayfinding using ARKit, an iOS framework for vision-inertial odometry (VIO). A user of this system walks while holding the smartphone in such a way that the camera has a clear view of the environment. A structure-from-motion algorithm, complemented with inertial data, measures the phone’s velocity and orientation with respect to a fixed reference frame. By integrating the velocity vectors through time, the walker’s trajectory can be reconstructed. The unavoidable drift (due to instantaneous bias and noise) can be compensated for via particle filtering, which uses the floor plan of the building to condition the reconstructed trajectory (in particular, by discouraging trajectories going through a wall). A similar approach was taken by Ren et al. [2], [3] but relying on inertial sensors only, using step-based pedestrian dead reckoning (PDR) or machine learning-based odometry (RoNIN).

A main problem with dead-reckoning odometry is that the starting location and orientation needs to be known, requiring some form of initial localization. In addition, occasional relocalization may be necessary because of sporadic system fault or poor tracking due to accumulated drift. In some cases, recognizable visual landmarks, (e.g., an EXIT sign) whose location in the building is known, can be used for initial localization or relocalization (in the following, we will use the term *localization* for both cases). However, map annotations of this type are not often available.

In this work, we will only consider prior information about the environment in the form of a floor plan of the building. Floor plans are often available, especially for public places, where they are frequently posted online. They normally represent structural elements (e.g. walls) that, unlike furniture or other objects, are expected to remain stable in time. Floor plan-based localization is highly desirable as it removes the need for environment fingerprinting or detailed map annotations. This

task is similar to that of a visitor of a building trying to locate themselves within a posted map when the “You are here” marking is missing. One may, for example, search the floor plan for a location where the presence of nearby features (e.g., doors, wall corners, hallway size) and their spatial relation, as indicated in the map, conform with the visible scene at their location. Our proposed PALMS (Plane-based Accessible Indoor Localization Using Mobile Smartphones) system uses 3-D data, acquired by an iPhone LiDAR, to detect specific features of the visible wall geometry which are then matched against the wall structure shown in the floor plan.

Regardless of the methodology considered, it is reasonable to expect that, in a building with self-repeating structural elements (e.g., a long corridor with uniformly spaced doors), for each location there may be multiple other locations from which the visual scene looks similar. To mitigate this ambiguity, we consider a dynamic form of localization: users take an image (or, in our case, a 3-D scan) of the environment from their location; then, they start walking in any available direction, following a certain path (e.g. random). Over a (hopefully short) period of time, the system determines their location and orientation, by combining information from the initial data collection along with dead-reckoning odometry information from the sensor data, conditioned by a particle filter. Intuitively, while multiple locations in the floor plan may be consistent with the visual data recorded at a given location, the same path (from dead-reckoning) is likely to be unfeasible (because of physical constraints) if started from an incorrect location. Inspired by [4], [5], we implement this intuition by defining a prior spatial probability distribution for the user’s location, which in our case is derived from a measure of how well the observed walls match with local characteristics of the floor plan. Then, the customary particle evolution process is started, driven by the sensor data recorded as the user is walking. Particles that hit a wall are resampled from the remaining particles, which typically leads to a rapid coalescing of the particle into one cluster. The distribution of particles is tracked until convergence is declared, at which time the user’s location is assumed to be reliably determined.

Here are the main contributions of this work:

- We present a novel method (PALMS) to assess the probability that a 3-D scan, captured at a location using an iPhone equipped with LiDAR and processed to extract planar walls, was taken at a certain location in a floor plan for one of N possible orientations of the LiDAR’s reference frame (where in our case, $N = 4$). This probability is computed on the basis of the wall patches extracted from the LiDAR scan, which are then matched against the walls marked in the floor plan. A probability distribution of locations across the floor plan (*heatmap*) is quickly computed through by means of a linear convolution for each possible orientation of the LiDAR’s reference frame.
- We propose a method for dynamic localization that uses a particle filter initialized from the heatmap computed at each of the N considered orientations
- We introduce a 2-step criterion that can be used to declare

“convergence”, the point in time after which the current user’s location can be assumed to be correctly identified. Our experimental data shows that the proposed PALMS-based initialization leads to reduced convergence time and increased localization accuracy after convergence compared to a uniform initial probability distribution. In particular, the average RMSE for PALMS was found to be 6.7 times less than for *Uniform*, while the proportions of locations with error less than 1 meter was found to be 4.9 times higher for PALMS than for *Uniform*.

II. RELATED WORK

Smartphone-based pedestrian tracking and localization systems in indoor, GPS-denied environments have received considerable attention recently. Many researchers have proposed schemes to estimate the user’s location based on the Received Signal Strength Indicator (RSSI) values of wireless technologies, including WiFi [4], Bluetooth Low Energy (BLE) beacons, Radio Frequency Identification (RFID), Ultra-Wideband (UWB), and magnetic field fingerprinting [3]. W-RGB-D [4] uses a heatmap of WiFi signal strength to initialize a particle filter for quick convergence in ambiguous environments. We follow a similar approach as described in section III-B. However, RSSI methods’ requirement for prior calibration or “fingerprinting” before use restricts the scalability of these methods. Similarly, image-retrieval methods [6] require a pre-built image database, while 3-D structure-based methods [7] require pre-built SfM models.

In contrast to fingerprinting-based approaches, some prior works perform indoor localization without prior visitation, leveraging only a known floor plan. Among these works, LiDAR-based methods are particularly notable. Such methods utilize 3-D-to-2-D point cloud matching [8], geometrical feature matching [9], usually coupled with particle filters [10] or graph optimizations [11]. Image-based methods [1], [7], [12]–[20] aim to perform localization from single or multiple images by utilizing image features [13], semantic information (doors [12], [14], [15], windows [15], texts [16], and exit signs [1]), or scene geometries (vertical lines and vertical planes [13]). [17] samples 3-D points from a raised 2-D floor plan and performs point cloud matching with the RGB-reconstructed 3-D scene.

While these methods generally rely on a flow of information over multiple sequential sensor measurements, some approaches [14], [18] try to tackle indoor localization by considering only an instantaneous observation. More recently, LaLaLoc [19] and LaLaLoc++ [20] used deep neural networks to find “embeddings” for panorama images to match with the “embeddings” of 2-D locations on the floor plans. As pointed out by most of these papers, the ambiguity introduced by layout-similarity cannot be completely removed when localizing with respect to layout alone, making it hard to localize using a single observation.

In this work, we consider a dynamic form of finger-printing-free localization that utilizes a publicly available floor plan,

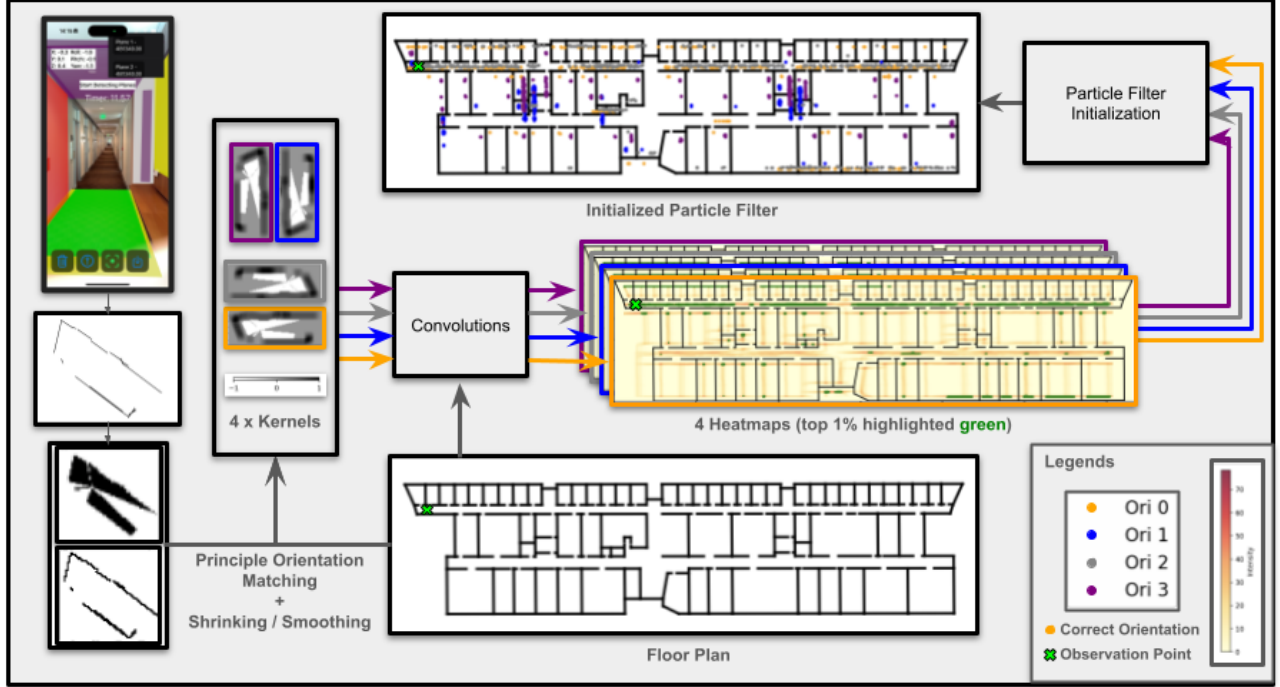


Fig. 1. Architecture of PALMS’ particle filter initialization pipeline. The ARKit-based application (top left) detects vertical planar patches (walls) in a 360° LiDAR scan. We used their 3-D poses to create 4 convolution kernels through principal orientation matching, shrinking, or smoothing. Each kernel is rotated to match a possible orientation. We then convolve the rasterized floor plan with each of the kernels to create 4 heatmaps $H_{0\sim3}$. We extract the top 1% locations from each heatmap and randomly sample 4 groups of particles, each with p particles sharing the same initial drift from the corresponding orientation. These particles will then be placed in one particle filter. Colored edges and dots indicate different orientation groups. Orange represents the correct orientation, and the green cross indicates the ground truth observation point.

an instantaneous observation, and a particle filter updated by odometry data (see section III).

III. METHOD

A. Plane Detection and Principal Orientation Matching

We detect vertical planar patches (walls) in a 360° LiDAR scan and obtain their 3-D poses using a custom ARKit-based iOS app. These patches are then projected onto the ground plane, forming a set of 2-D line segments (in metric units) $\mathcal{L}_{\text{obs}} = \{l_{\text{obs},1}, l_{\text{obs},2}, \dots, l_{\text{obs},n}\}$. We also represent the walls in the given floor plan as a set of 2-D line segments $\mathcal{L}_{\text{fp}} = \{l_{\text{fp},1}, l_{\text{fp},2}, \dots, l_{\text{fp},m}\}$, also expressed in metric units.

Intuitively, to localize the *observation point* (the center of projection of the LiDAR scan), we could try place \mathcal{L}_{obs} over \mathcal{L}_{fp} while adjusting \mathcal{L}_{obs} ’s orientation and location, as if placing down a piece of a jigsaw puzzle looking for a match. We limit the set of orientations by observing that, in most buildings, only a discrete set of wall orientations is observed. In our experiments, we considered buildings with walls intersecting at 90°, resulting in 2 possible orientations for any wall. Note that our algorithm can be easily extended to more complex buildings (e.g., walls intersecting at multiple of 45°).

Our first task is to determine an angle θ such that, by rotating all observed line segments \mathcal{L}_{obs} by θ , any one such segment is aligned with one or more line segments in the floor plan. Note that if θ satisfies this property, all angles $\theta + k \cdot 90^\circ$ for $0 \leq k < 3$ will also satisfy it, and thus need to be considered. We determine θ by first finding the two orthogonal principal orientations of the observed line segments as well as of the floor plan segments. θ is taken to be the smallest angle that pairwise aligns the principal orientations found in the two domains.

B. Heatmaps Creation

Consider the observed line segments \mathcal{L}_{obs} , rotated by $\theta + k \cdot 90^\circ$. Our goal is to define, for any point in the floor plan, the likelihood that this point represents the observation point. Ideally, at the correct location and orientation, the observed and floor plan line segments should overlap. By rasterizing these segments into bitmaps and multiplying them pixel by pixel, we obtain a score that, after normalization, defines a probability distribution for the location. Testing multiple locations is akin to computing a convolution on the rasterized floor plan, with the convolution kernel (termed RW for recorded walls) representing the walls observed in the LiDAR scan (see Fig. 1)

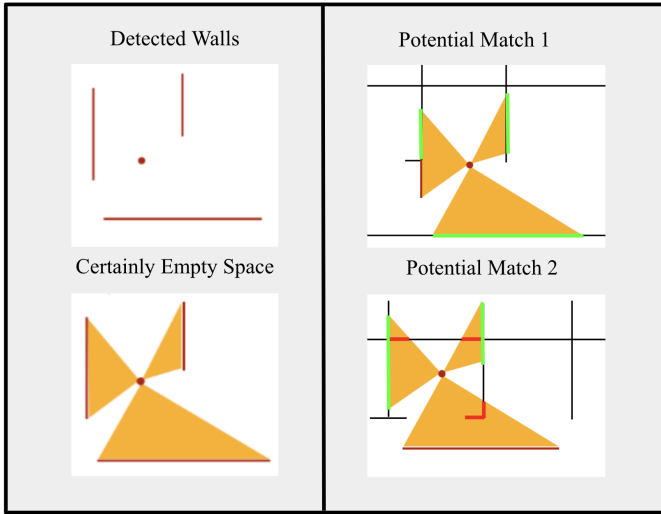


Fig. 2. Certainly Empty Space (CES). CES enforces visibility constraints to eliminate physically impossible matches. The figure shows two scenarios: in match 1 (top right), the constraint is met as no floor plan segments fall within the CES (orange). In match 2 (bottom right), several segments intersect with the CES (bright red), making it a less viable match.

While intuitively appealing, this simple approach suffers from two main drawbacks. First, both the floor plan and the pose of the walls extracted from the LiDAR scan may be affected by errors. As a result, the observed line segments may not perfectly overlap with the floor plan segments even when \mathcal{L}_{obs} is centered at the correct observation point. To account for this, we smooth the observed line segments bitmap with a Gaussian kernel, which basically “widens” the line segments to account for possible localization errors. Second, this criterion does not leverage visibility constraints. Suppose that an observed line segment is matched to a floor plan segment and that another floor plan segment is located between the location being inspected and this matched segment (see e.g. Fig. 2). This unmatched floor plan segment would not be penalized by our simple “goodness of fit” measure, yet its presence is physically impossible, as it would occlude view of the matched wall segment.

In order to enforce visibility constraints, we introduce the notion of *certainly empty space* (CES) (see Fig. 2), which represents the areas where, based on the observation from a certain location, no wall in the floor plan should be found. Specifically, if a wall (line segment) was observed from a certain location, no other floor plan line segment should be found in the triangle subtended by the observed line segment at that location. The union of such triangles for all visible line segment forms the CES area. By centering the CES space on different locations in the floor plan, we can find which locations break the visibility constraint (because they contain walls within the CES area). This is implemented by defining a second kernel (CES), with values of 1 at the CES area and 0 otherwise, and convolving this with the rasterized floor plan. To account for possible wall localization errors, we first slightly “shrink” (rescale) the CES kernel, then combine the

two kernels into one, equal to $RW - \alpha \cdot CES$, where we set $\alpha = 1$ in our experiments.

In practice, we first rotate the observed line segments \mathcal{L}_{obs} by angle θ (Sec. III-A), and compute the kernel $RW - \alpha \cdot CES$. Convolution of the rasterized floor plan with this kernel forms a heatmap H_0 . We then rotate the kernel by multiples of 90° to compute additional heatmaps H_1, H_2, H_3 . The heatmap H_k represents the probability distribution of the observation point across the floor plan, given that the reference frame for the LiDAR scan was at an angle of $\theta + k \cdot 90^\circ$ with the floor plan’s reference frame.

C. Particle Filter Initialization from Heatmaps

In our experiments (Sec. IV), the “ground truth” observation point P_{gt} consistently appeared among the locations with top 1% of all heatmap values. We thus simplify all heatmaps by binarizing them, using a threshold equal to the 99 percentile of all values in all maps (see Fig. 1). Each location marked as 1 in a binarized heatmap H_k represents a candidate observation point for the k -th orientation. Note that the proportion of 1’s may vary across heatmaps for different orientations.

We use these binary heatmaps to initialize a particle filter tracking the motion of the user as they walk away from the observation point as driven by sensor data. Following [2], each particle tracks a state formed by location and angular drift. The filter is initialized as follows. For each binary heatmap, we generate a number of particles sampled at uniform spatial density within the areas where the heatmap is equal to 1. We then place all generated particles on the floor plan, where particles generated from the H_k heatmap are assigned a label (k) and are given an initial orientation drift of $\theta + k \cdot 90^\circ$. During particle filter updating, particles that hit a wall are resampled from the remaining particles, which typically leads to a rapid coalescing of the particles into one cluster. Note that when a particle with label k is resampled, the resampled particle maintains the same label, location (with noise added), and angular drift (with noise added).

D. Convergence Criterion

It can thus be expected that the particle filter will converge to the correct user’s location after a while (global localization problem [21]). For example, [1] shows an empirical characterization of the time it takes for the location error to be less than 1 meter. However, in practical scenarios, the system does not have access to this error metric. Therefore, a measurable criterion is necessary to reliably declare “convergence”, allowing the user to assume that their location is being tracked accurately from that point onward. We propose a 2-step convergence criterion based on the particles’ distribution.

In Step 1, we check if one label dominates the current particle set (at least 80%), since particles with the correct initial orientation are less likely to hit a wall and more likely to survive. We mark this point as t_1 .

In Step 2, we cluster particles with the dominant orientation using the mean shift algorithm and check if the dominant cluster contains at least 50% of these particles. At this point

t_2 , we consider the particle filter fully converged, and the algorithm outputs the mean position of the dominant cluster as the predicted user location P_{pred} (See Fig. 3). Like [2], we run the mean-shift algorithm to update the dominant cluster every 20 time steps after t_2 to keep track of the user’s location.

IV. EXPERIMENTS

A. Data Collection

We collected data from 14 paths (2 per 7 starting points) in 3 buildings, averaging 147m in length. Each dataset includes LiDAR-detected vertical planes (using an iPhone 14 Pro) and sensor data recorded in the paths. We scanned the surroundings by rotating the phone 360° and walked away from the observation point, holding the phone with the camera facing forward. Odometry data, collected using iOS’ ARKit, drives the particle filter initialized with our PALMS algorithm. We used this odometry data and another particle filter initialized at the correct location and orientation to generate ”ground truth” location data. The dataset will be publicly available upon acceptance.

B. Global Localization with PALMS

In this experiment, we evaluate the advantages of PALMS initialization using 3-D scan data. We compare it against two methods without this data. In the first method (*Uniform*), particles are uniformly sampled across the floor plan with random initial angular drift between 0° and 360°. The second method (*Uniform + Ori*) uniformly samples particle locations and assigns initial angular drifts $\theta + k \cdot 90^\circ$ for each of the k groups ($0 \leq k < 3$). θ is calculated using the principal orientations of \mathcal{L}_{fp} and the first few tracking points. For *Uniform*, convergence Step 1 is determined by setting a threshold on particle dispersion, while Step 2 monitors for dominant clusters like PALMS. All three methods use 2000 particles and run 100 trials per path, with random re-initialization for each trial.

Measurements included: time (in seconds) and distance (in meters) until convergence is declared; RMSE after convergence (in meters); and proportion of predicted locations (after convergence) at less than 1 meter from the ground truth ($\%Error < 1m$). Results in Tab. I (averaged over all trials) show that PALMS significantly outperforms the other methods considered. In particular, the average RMSE for PALMS was found to be 6.7 times less than for *Uniform* (5.6 times less than for *Uniform + Ori*), while the proportions of locations with error less than 1 meter post-convergence was found to be 4.9 times higher for PALMS than for *Uniform* (2.8 times higher than for *Uniform + Ori*). Examples of particles distributions for the three cases at different stages of convergence are shown in Fig. 3.

V. DISCUSSION AND CONCLUSIONS

We presented an indoor localization algorithm using an iPhone LiDAR scan to acquire nearby wall geometry, then finding locations in a floor plan consistent with the observed

TABLE I
POST CONVERGENCE METRICS FOR GLOBAL LOCALIZATION

Init. Method	Metrics			
	Time (s) ↓	Dist. (m) ↓	RMSE ↓	%Error < 1m ↑
Uniform	75.94	129.87	31.52	16%
Uniform + Ori	69.18	121.86	26.03	28%
PALMS	59.14	110.28	4.69	78%

wall geometry for possible camera orientations. This information filters out unlikely locations but cannot precisely pinpoint the user due to geometric ambiguities in self-similar layouts. To address this, we generate a spatial probability distribution for each orientation to initialize a particle filter, which is expected to converge to the correct location.

We showed that this strategy reduces time to convergence and substantially increases accuracy after convergence compared to uniform particle initialization (i.e., without initial orientation and the initial information on the user location provided by the LiDAR scan). A main advantage of our approach is that it does not require prior ”fingerprinting” of the environment, as it only uses a floor plan for prior information. It provides a solution to a critical bottleneck of dead-reckoning tracking algorithms (based on visual or inertial data), that is the determination of the initial user location and orientation.

At the same time, there are a number of limitations that need to be considered. Our system uses an iPhone LiDAR, which is only available on ”Pro” models at the time this paper is written. In addition, the maximum range of this sensor is 5 meters, which reduces its utility in wide open spaces. In future work, we will consider extracting planar wall models directly from images to enable detection of walls beyond the LiDAR’s range, although it could be less accurate. While heatmap-based particle initialization has been shown to reduce convergence time, this system still requires the user to walk for a while (up to a minute) before convergence. This could prove impractical, and we will consider strategies like combining other modalities to reduce convergence time in future work. Finally, the dead-reckoning tracker we considered in this work was based on visual-inertial odometry from data recorded by an iPhone held by the user while walking. We will consider applying the same strategy to inertial-only tracking with particle filtering. In this configuration, users can use the phone to take a LiDAR scan, then conveniently put the phone back in their pocket, and be tracked by the phone’s inertial sensors.

VI. ACKNOWLEDGEMENTS

Research reported in this publication was supported by the National Eye Institute of the National Institutes of Health under award number R01EY029260-01. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

REFERENCES

- [1] R. Crabb, S. A. Cheraghi, and J. M. Coughlan, ”A lightweight approach to localization for blind and visually impaired travelers.”

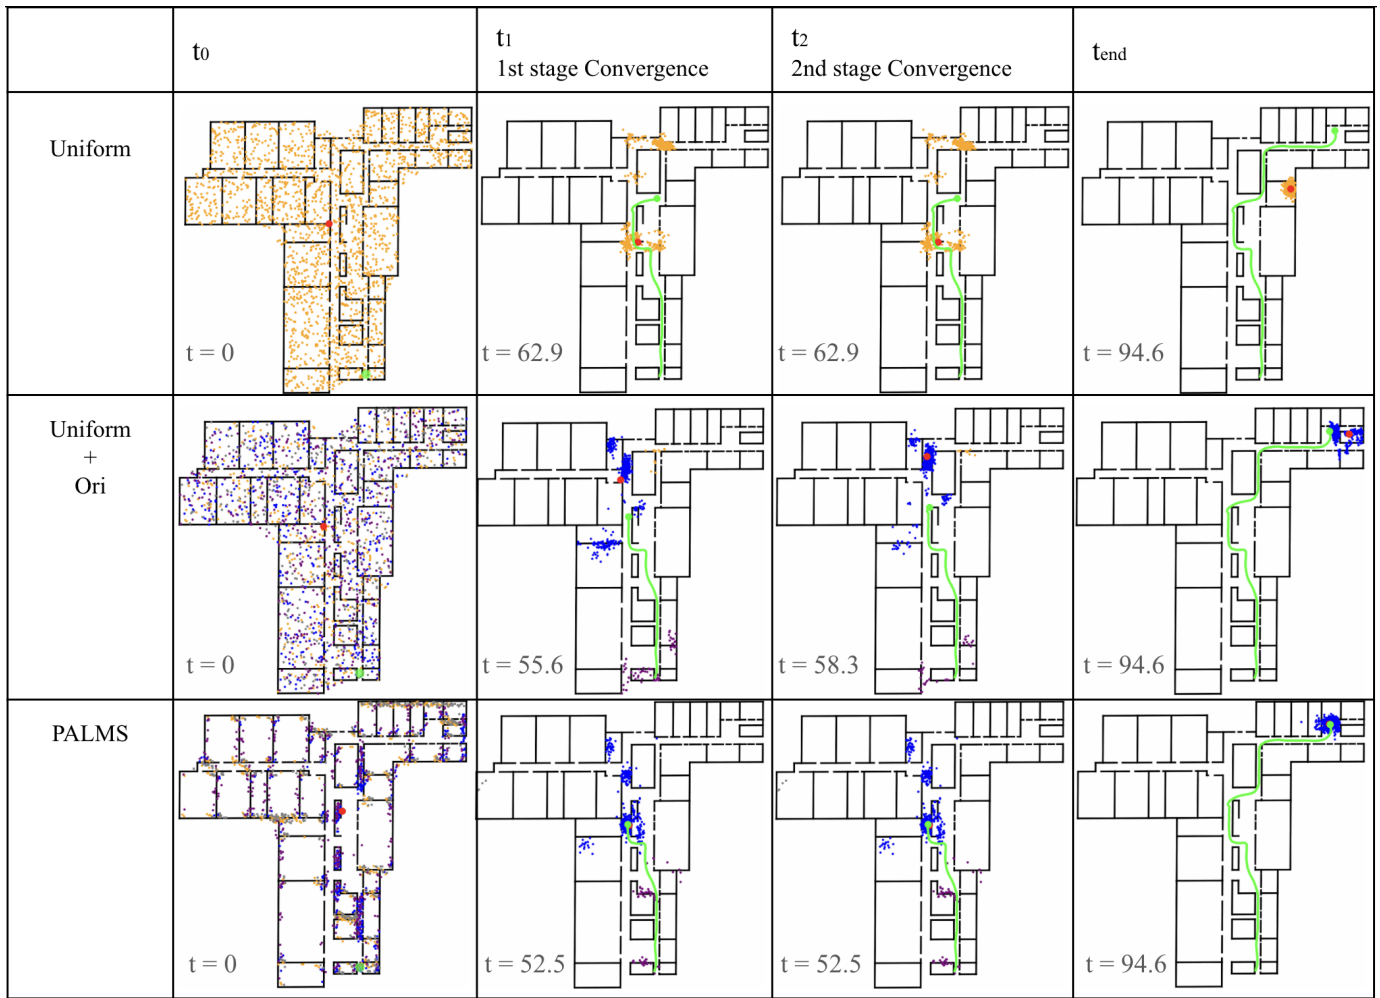


Fig. 3. Visualization of the particle filters under different experiment settings at the time of initialization, at different stages of convergence, and when the path ends. Red dot: mean location of all particles (before t_1), mean position of the dominant group (after t_1), mean position of the dominant cluster (after t_2). Green: ground truth path and location. All three settings share the same starting point, tracking data, ground truth path, and particle count. Particles in different orientation groups are represented by different colors. $t = x$ shows the time in seconds. We can see that PALMS successfully localized the user at t_2 and tracks the user until t_{end} while the other methods failed to do so. Note that for some cases the 1st stage and 2nd stage convergence happen simultaneously.

- Sensors (Basel, Switzerland)*, vol. 23, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:257321613>
- [2] P. Ren, F. Elyasi, and R. Manduchi, "Smartphone-based inertial odometry for blind walkers," *Sensors (Basel, Switzerland)*, vol. 21, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:235642119>
 - [3] C. H. Tsai, F. Elyasi, P. Ren, and R. Manduchi, "All the way there and back: Inertial-based, phone-in-pocket indoor wayfinding and backtracking apps for blind travelers," *ArXiv*, vol. abs/2401.08021, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:266998972>
 - [4] S. Ito, F. Endres, M. Kuderer, G. D. Tipaldi, C. Stachniss, and W. Burgard, "W-rgb-d: Floor-plan-based indoor global localization using a depth camera and wifi," *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 417–422, 2014. [Online]. Available: <https://api.semanticscholar.org/CorpusID:14117443>
 - [5] R. Maffei, V. A. M. Jorge, V. F. Rey, M. L. Kolberg, and E. P. e Silva, "Fast monte carlo localization using spatial density information," *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6352–6358, 2015. [Online]. Available: <https://api.semanticscholar.org/CorpusID:7200821>
 - [6] K. Dai, T. Xie, K. Wang, Z. Jiang, D. Liu, R. Li, and J. Wang, "Eaainet: An element-wise attention network with global affinity information for accurate indoor visual localization," *IEEE Robotics and Automation Letters*, vol. 8, pp. 3166–3173, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:257754855>
 - [7] Y. Zhou, X. Zheng, R. Chen, H. Xiong, and S. Guo, "Image-based localization aided indoor pedestrian trajectory estimation using smartphones," *Sensors (Basel, Switzerland)*, vol. 18, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:3291004>
 - [8] Y. Watanabe, K. Ramirez-Amaro, B. Ilhan, T. Kinoshita, T. Bock, and G. Cheng, "Robust localization with architectural floor plans and depth camera," *2020 IEEE/SICE International Symposium on System Integration (SII)*, pp. 133–138, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:212645763>
 - [9] X. Wang, R. J. Marcotte, and E. Olson, "Glfip: Global localization from a floor plan," *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1627–1632, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:197634269>
 - [10] Z. Li, K. Kawkeeree, and M. H. Ang, "Localization on indoor topological maps - scam: Scale-compatible adaptive monte-carlo localization," in *Annual Meeting of the IEEE Industry Applications Society*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:256305617>
 - [11] L. Gao and L. Kneip, "Fp-loc: Lightweight and drift-free floor plan-assisted lidar localization," *2022 International Conference on Robotics and Automation (ICRA)*, pp. 4142–4148, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:247187616>

- [12] R. G. Goswami, P. V. Amith, J. Hari, A. Dhaygude, P. Krishnamurthy, J. Rizzo, A. Tzes, and F. Khorrani, "Efficient real-time localization in prior indoor maps using semantic slam," *2023 9th International Conference on Automation, Robotics and Applications (ICARA)*, pp. 299–303, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258869755>
- [13] L. Chen, Y. Zou, Y. Chang, J. Liu, B. Lin, and Z. Zhu, "Multi-level scene modeling and matching for smartphone-based indoor localization," *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 311–316, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:210693118>
- [14] D. Yoo, G. Shan, and B. hee Roh, "A vision-based indoor positioning systems utilizing computer aided design drawing," *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252898727>
- [15] O. A. M. Maldonado, S. Hadfield, N. Pugeault, and R. Bowden, "Sedar - semantic detection and ranging: Humans can localise without lidar, can robots?" *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–8, 2017. [Online]. Available: <https://api.semanticscholar.org/CorpusID:11680434>
- [16] N. Zimmerman, L. Wiesmann, T. Guadagnino, T. Labe, J. Behley, and C. Stachniss, "Robust onboard localization in changing environments exploiting text spotting," *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 917–924, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:247627732>
- [17] H. Chu, D. K. Kim, and T. Chen, "You are here: Mimicking the human thinking process in reading floor-plans," *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 2210–2218, 2015. [Online]. Available: <https://api.semanticscholar.org/CorpusID:15166342>
- [18] H. Hile and G. Borriello, "Positioning and orientation in indoor environments using camera phones," *IEEE Computer Graphics and Applications*, vol. 28, 2008. [Online]. Available: <https://api.semanticscholar.org/CorpusID:13422180>
- [19] H. Howard-Jenkins, J.-R. Ruiz-Sarmiento, and V. A. Prisacariu, "Lalaloc: Latent layout localisation in dynamic, unvisited environments," *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 10 087–10 096, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:233296099>
- [20] H. Howard-Jenkins and V. A. Prisacariu, "Lalaloc++: Global floor plan comprehension for layout localisation in unvisited environments," in *European Conference on Computer Vision*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:253270204>
- [21] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte carlo localization for mobile robots," *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 2, pp. 1322–1328 vol.2, 1999. [Online]. Available: <https://api.semanticscholar.org/CorpusID:9696573>