



The 17th International Conference on Ambient Systems, Networks and Technologies (ANT)  
April 14-16, 2026, Istanbul, Türkiye

## Seamless Outdoor–Indoor Pedestrian Positioning System with GNSS/UWB/IMU Fusion: A Comparison of EKF, FGO, and PF

Jiaqiang Zhang<sup>a,\*</sup>, Xianjia Yu<sup>a</sup>, Sier Ha<sup>a</sup>, Paola Torrico Moron<sup>a</sup>, Sahar Salimpour<sup>a</sup>, Farhad Keramat<sup>a</sup>, Haizhou Zhang<sup>a</sup>, Tomi Westerlund<sup>a</sup>

<sup>a</sup>*Turku Intelligent Embedded and Robotic Systems (TIERS) Lab, University of Turku, Finland*

### Abstract

Accurate and continuous pedestrian positioning across outdoor–indoor environments remains challenging because GNSS, UWB, and inertial PDR are complementary yet individually fragile under signal blockage, multipath, and drift. This paper presents a unified GNSS/UWB/IMU fusion framework for seamless pedestrian localization and provides a controlled comparison of three probabilistic back-ends: an error-state extended Kalman filter, sliding-window factor graph optimization, and a particle filter. The system uses chest-mounted IMU-based PDR as the motion backbone and integrates absolute updates from GNSS outdoors and UWB indoors. To enhance transition robustness and mitigate urban GNSS degradation, we introduce a lightweight map-based feasibility constraint derived from OpenStreetMap building footprints, treating most building interiors as non-navigable while allowing motion inside a designated UWB-instrumented building. The framework is implemented in ROS 2 and runs in real time on a wearable platform, with visualization in Foxglove. We evaluate three scenarios: indoor (UWB+PDR), outdoor (GNSS+PDR), and seamless outdoor–indoor (GNSS+UWB+PDR). Results show that the ESKF provides the most consistent overall performance in our implementation.

© 2026 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer review under the responsibility of the scientific committee of the Program Chairs

**Keywords:** Localization, PDR, IMU, UWB, GNSS, Kalman Filter;

### 1. Introduction

Accurate and continuous pedestrian positioning across outdoor and indoor environments remains challenging due to the complementary but individually limited nature of available sensors. Outdoors, Global Navigation Satellite System (GNSS) provides globally referenced positions but suffers in urban areas from signal blockage, multipath, and dilution of precision near buildings. Indoors, GNSS is largely unavailable; Ultra-Wideband (UWB) can deliver decimeter-level accuracy where anchors are deployed, but offers only local coverage and is sensitive to non-line-of-

\* Corresponding author. Tel.: +358 414775904.

E-mail address: [jiaqiang.zhang@utu.fi](mailto:jiaqiang.zhang@utu.fi)

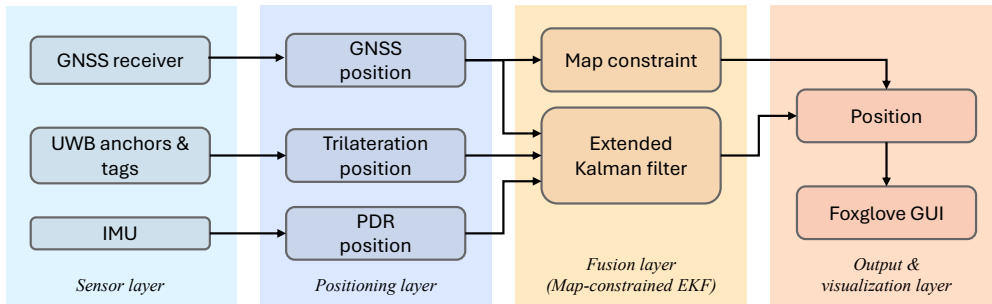


Fig. 1: Proposed evaluation scheme in this study

sight (NLOS) bias. Inertial sensing enables self-contained motion tracking at a high rate, yet pure integration drifts over time. In practice, maintaining seamless accuracy as a user traverses doorways, corridors, and urban canyons requires principled fusion of these modalities together with environmental knowledge that constrains physically plausible motion.

A broad spectrum of probabilistic estimators has been applied to multi-sensor pedestrian positioning. Extended Kalman filtering (EKF), especially in its error-state form, is attractive for real-time embedded systems due to its computational efficiency and structured treatment of high-rate inertial updates. Factor graph optimization (FGO) offers a smoothing perspective that can incorporate heterogeneous constraints over a window, often improving global consistency and robustness to intermittent measurement availability. Particle filtering (PF) provides a flexible nonparametric alternative capable of representing multimodal posteriors and handling strong nonlinearities or non-Gaussian errors that frequently arise in GNSS multipath and UWB NLOS conditions. Yet, for seamless outdoor–indoor pedestrian positioning using the same GNSS/UWB/IMU stack, a carefully controlled comparison of these three estimator families under consistent sensing, map constraints, and evaluation protocols remains limited.

This work addresses seamless outdoor–indoor pedestrian positioning by combining GNSS, UWB, and Pedestrian Dead Reckoning (PDR) from a chest-mounted IMU within three probabilistic fusion frameworks augmented by map constraints: an error-state EKF (ESKF), an FGO-based sliding-window smoother, and a PF-based estimator. We target scenarios such as industrial sites or sensitive buildings in urban settings, where users move intermittently between open-sky areas and factory-like interiors. Our key observation is that coarse environmental structure—specifically, building footprints—can be exploited online to (i) reject or down-weight physically impossible GNSS fixes (e.g., spuriously placing a user inside a building without UWB coverage) and (ii) regularize the fused trajectory near transitions, reducing discontinuities and large outliers. By integrating the same map knowledge into all three estimators, we aim to isolate algorithmic differences from environmental priors.

The proposed system is realized on a wearable platform comprising a Raspberry Pi<sup>5</sup>, a GNSS module, an Xsens MTi-300 IMU, and a Decawave DWM1001 UWB tag mounted on a chest strap. PDR is derived from the IMU by detecting steps and estimating per-step displacements and heading increments; these kinematic increments provide drift-resistant relative motion over walking horizons and supply high-rate constraints between absolute updates. UWB ranges to surveyed anchors inside a designated building yield accurate indoor updates, while GNSS provides absolute position outdoors. In the ESKF, we (i) treat PDR increments as control-like motion updates, (ii) loosely couple GNSS position measurements, and (iii) tightly incorporate UWB range observations. In the FGO formulation, PDR increments, GNSS positions, and UWB ranges are expressed as factors over a sliding window, enabling joint refinement of recent states and improved consistency during brief GNSS outages or UWB geometry changes. In the PF, PDR increments define the proposal dynamics, while GNSS and UWB likelihoods re-weight particles; this framework allows the posterior to remain multi-hypothesis in ambiguous transition zones.

To inject environmental knowledge, we query OpenStreetMap via Overpass Turbo to obtain building polygons (Fig. 2a). At runtime, all building interiors are treated as forbidden regions for the fused solution except the single building equipped with UWB anchors, which is designated as free space. The map constraint is implemented consistently across estimators: in the ESKF by rejecting or projecting implausible GNSS updates to the nearest admissible boundary (Fig. 2b); in the FGO by adding feasibility-aware penalties or gating factors that discourage solutions inside forbidden regions; and in the PF by assigning near-zero likelihood to particles that violate the occupancy rule. This

unified treatment prevents physically impossible corrections and reduces multipath-induced excursions along facades, while still permitting motion within the UWB-covered interior.

All data are acquired and synchronized using ROS 2, and we provide real-time GUI visualization in Foxglove to display trajectories and map overlays. The system is quantitatively evaluated in an urban campus with a factory-like indoor lab: outdoors, ground truth is provided by an Xsens MTi-680G with Real Time Kinematic (RTK); indoors, by an optical motion-capture system aligned to the global frame. These complementary references allow continuous accuracy assessment across environment transitions and support a fair comparison of ESKF, FGO, and PF under outdoor, indoor, and outdoor–indoor sequences.

This paper aims to: (i) design a map-constrained fusion pipeline that integrates GNSS, UWB, and IMU-PDR to yield robust, drift-bounded pedestrian trajectories across outdoor–indoor transitions; (ii) implement three representative estimators—ESKF, FGO, and PF—under consistent sensing and constraint assumptions; and (iii) provide an empirical comparison that clarifies when the efficiency of filtering, the global consistency of smoothing, or the flexibility of nonparametric inference offers the greatest benefit for seamless pedestrian positioning.

We make four contributions: (1) a unified GNSS/UWB/IMU-PDR fusion architecture with a consistent building-footprint constraint applied across ESKF, FGO, and PF; (2) a lightweight, real-time use of open-source map polygons to enforce feasibility while permitting motion within a designated UWB-covered building; (3) a reproducible ROS<sub>2</sub>/Foxglove implementation on an embedded Raspberry Pi<sub>5</sub> with commodity sensors and three estimator backends; and (4) a dual-ground-truth evaluation protocol enabling fine-grained assessment of accuracy and continuity during doorway crossings and urban canyon segments. We assume a single permitted building equipped with surveyed UWB anchors, treat all other building interiors as non-navigable, and focus on horizontal positioning for a walking pedestrian; vertical motion and multi-floor transitions are left to future work.

The rest of the paper is organized as follows. Section 2 provides an overview of related work on GNSS/UWB/IMU fusion and outdoor–indoor transition handling. Section 3 details the system architecture, PDR modeling, measurement models, map-constraint mechanisms, and the three estimator formulations. Section 4 presents results for outdoor, indoor, and outdoor–indoor scenarios and compares the ESKF, FGO, and PF performance. Finally, Section 5 concludes the paper.

## 2. Related Work

Seamless pedestrian localization increasingly fuses absolute outdoor positioning with infrastructure-aided and inertial methods indoors, often augmented by environment awareness or explicit constraints. Multi-sensor fusion is a representative direction: Zhang et al. [1] integrate GNSS/IMU/UWB with map information to mitigate multipath and stabilize trajectories across indoor–outdoor transitions, closely aligned with our use of building polygons as feasibility constraints. A low-cost wearable platform that loosely couples GNSS and UWB with an Inertial Navigation System (INS) backbone demonstrated continuity at indoor–outdoor boundaries in [2]. Robust estimator design for UWB/PDR integration has also been explored: an adaptive EKF with an additional heading constraint improves resilience to UWB non-line-of-sight and reduces PDR drift in [3]. On commodity devices, environment-perception-based switching orchestrates GNSS/INS/BLE for smooth handovers near doorways [4], and BLE-aided particle-filter fusion yields real-time indoor positioning when combined with PDR [5]. A dynamic feasible-region IMU/UWB fusion method restricts the solution space to improve robustness in cluttered interiors [6].

Within inertial-centric pedestrian dead reckoning, recent work prioritizes heading stability and drift mitigation for general-purpose wearables and unconstrained device poses. A multi-source PDR approach fusing accelerometer, gyroscope, geomagnetic, and ancillary cues improves gait detection, step length, and heading estimation indoors [7]; complementary yaw-correction with two low-cost MIMUs further suppresses heading drift and improves downstream position estimates [8]. For smartphone deployments, Jiang et al. implement PDR/GNSS integration using both a Kalman filter and factor-graph optimization—modeling PDR as step-wise factors and GNSS as absolute position factors—and show that factor-graph fusion outperforms KF-based fusion in field trials [9]. A recent overview synthesizes model-based and data-driven pedestrian inertial navigation, highlighting trends in learned PDR, hybrid fusion, and the role of priors and constraints for wearables [10].

Body-mounted wearables that centralize sensors on the torso or head illustrate how PDR priors, visual/radio aiding, and learned models can be combined for robust tracking. A chest-mounted localization vest fuses advanced

adaptive PDR (dynamic step segmentation and adaptive step length) with binocular vision in an EKF, yielding substantial accuracy gains over vision-only baselines across diverse indoor motions [11]. In parallel, ROCIP employs a learned probabilistic motion model with a Rao–Blackwellized particle filter to maintain bounded inertial drift during complex human–environment interactions [12]. A recent fusion system integrating UWB, IMU, and vision further underscores the benefits of multi-sensor opportunism for indoor robustness [13]. Taken together with map-aware GNSS/IMU/UWB fusion [1] and feasible-region IMU/UWB constraints [6], these threads suggest a design space in which wearable PDR provides high-rate kinematic increments, absolute anchors (GNSS/UWB) provide global consistency, and environmental constraints regularize the estimator near building boundaries—an approach we adopt within an ESKF.

### 3. Methodology

#### 3.1. Proposed Pedestrian Positioning System

The proposed framework for this study can be found in Fig. 1. A chest-mounted platform provides step-wise PDR from an Xsens MTi-300 IMU, while absolute position updates are obtained from GNSS outdoors and UWB-based trilateration indoors. All measurements and estimates are expressed in a local ENU frame anchored at the test site.

To compare filtering and smoothing paradigms under the same sensing conditions, we implement three back-ends: a planar error-state EKF, a sliding-window factor graph optimizer, and a particle filter. The three methods share the same PDR front-end and consume the same absolute position topics (GNSS fix and UWB trilateration fix).

The Chest-Mounted PDR front-end runs at an effective rate of 50 Hz and outputs step-triggered horizontal displacement increments in ENU. Steps are detected from the band-pass filtered dynamic acceleration magnitude within the typical walking frequency range. Step length is estimated using a Weinberg-style relation driven by the peak–trough amplitude, where the peak–valley difference is defined as  $\Delta A_k = A_k^{\text{pk}} - A_k^{\text{val}}$ , and the resulting step length is computed as  $\Delta s_k = k (\Delta A_k)^{1/4}$ , with  $k$  calibrated from a short walk. The step displacement used by all fusion back-ends is

$$\Delta \mathbf{p}_k = \begin{bmatrix} \Delta s_k \sin(\psi_k) \\ \Delta s_k \cos(\psi_k) \end{bmatrix}, \quad (1)$$

with the heading  $\psi_k$  obtained from the IMU/magnetometer-based orientation estimate. This shared PDR design ensures that performance differences mainly originate from the probabilistic back-ends rather than the motion front-end.

GNSS provides low-rate global fixes outdoors, while indoor absolute positions are produced by a UWB trilateration node and delivered to the fusion back-ends as `NavSatFix`. Each absolute measurement is transformed into ENU for fusion. In the subsequent evaluation, GNSS and UWB updates are selectively used according to their availability per scenario.

To provide a fair and compact comparison under identical sensing conditions, we implement three probabilistic back-ends that share the same chest-mounted PDR front-end and consume the same absolute position fixes from GNSS (outdoor) and UWB trilateration (indoor). The ESKF operates as a lightweight recursive filter: instead of integrating raw inertial acceleration, it treats the PDR displacement as the primary motion input, and uses sporadic GNSS/UWB fixes as position corrections to bound drift. This design yields high computational efficiency and stable real-time behavior when the error distribution is approximately unimodal. In contrast, the sliding-window FGO formulates the problem as a local MAP smoothing task over a compact pose chain. Consecutive nodes are connected by PDR between-factors, while GNSS/UWB updates are injected as absolute position factors; a weak anchor and window pruning maintain numerical stability and limit computational growth. This allows later absolute information to retroactively refine recent steps, which is beneficial during intermittent updates or transitions. Finally, the PF represents the posterior by a set of weighted particles and propagates them only upon detected steps using the same PDR increments with stochastic perturbations. Particle weights are updated by the GNSS/UWB likelihood, and resampling is triggered when degeneracy is detected via the effective sample size. This sampling-based formulation is more

**Algorithm 1** PDR-driven ESKF

---

```

1: State:  $(\mathbf{p}_{\text{est}}, \mathbf{v}_{\text{est}}, q_{\text{est}}, \mathbf{P})$ 
2: Error:  $[\delta\mathbf{p}, \delta\mathbf{v}, \delta\theta], \mathbf{P} \in \mathbb{R}^{9 \times 9}$ 
3: Init from first accepted absolute fix
4: for each IMU/PDR sample do
5:    $\Delta\mathbf{p} \leftarrow \mathbf{p}^{\text{PDR}} - \mathbf{p}_{\text{prev}}^{\text{PDR}}$ 
6:    $\mathbf{p}_{\text{est}} \leftarrow \mathbf{p}_{\text{est}} + \Delta\mathbf{p}; q_{\text{est}} \leftarrow q^{\text{PDR}}$ 
7:    $\mathbf{P} \leftarrow \mathbf{FPF}^T + \mathbf{LQL}^T$ 
8: end for
9: for each absolute fix (GNSS/UWB) do
10:   $\mathbf{z} \leftarrow \text{geodetic2enu}(\mathbf{z}_{\text{LLA}})$ 
11:   $\mathbf{K} \leftarrow \mathbf{PH}^T (\mathbf{HPH}^T + \sigma_{\text{sensor}}^2 \mathbf{I})^{-1}$ 
12:   $\delta\mathbf{x} \leftarrow \mathbf{K}(\mathbf{z} - \mathbf{p}_{\text{est}})$ 
13:   $\mathbf{p}_{\text{est}} \leftarrow \mathbf{p}_{\text{est}} + \delta\mathbf{p}; \mathbf{v}_{\text{est}} \leftarrow \mathbf{v}_{\text{est}} + \delta\mathbf{v}$ 
14:   $q_{\text{est}} \leftarrow \text{Exp}(\delta\theta) \otimes q_{\text{est}}$ 
15:   $\mathbf{P} \leftarrow (\mathbf{I} - \mathbf{KH})\mathbf{P}$ 
16: end for

```

---

**Algorithm 2** Sliding-window FGO

---

```

1: Node:  $\mathbf{T}_k = [x_k, y_k, z_k, \psi_k]^T$ , window  $W$ 
2: Wait first absolute fix; add prior  $\mathbf{r}_{\text{prior}} = \mathbf{T}_0 - \hat{\mathbf{T}}_0$ 
3: for each confirmed step  $k$  do
4:    $\Delta\mathbf{u}_k = [\Delta x, \Delta y, \Delta z, \Delta\psi]^T$ 
5:   Add node guess  $\mathbf{T}_k \leftarrow \mathbf{T}_{k-1} \oplus \Delta\mathbf{u}_k$ 
6:   Add Between:  $\mathbf{r}_{\text{PDR}} = \text{between}(\mathbf{T}_{k-1}, \mathbf{T}_k) - \Delta\mathbf{u}_k$ 
7:   if absolute fix available then
8:     Add Position:  $\mathbf{r}_{\text{abs}} = [x_k, y_k, z_k]^T - \hat{\mathbf{p}}_k$ 
9:   end if
10:  if nodes  $> W$  then
11:    Prune stale factors; weak-anchor oldest node
12:  end if
13:  Optimize:  $\min \sum \|\mathbf{r}_{\text{prior}}\|^2 + \sum \|\mathbf{r}_{\text{PDR}}\|^2 + \sum \|\mathbf{r}_{\text{abs}}\|^2$ 
14: end for

```

---

**Algorithm 3** Step-driven Particle Filter

---

```

1: Particle:  $\mathbf{x}^{(i)} = [x, y, z, \psi]^T$ , weights  $w^{(i)}$ ,  $N$ 
2: Init at first absolute ENU fix;  $w^{(i)} \leftarrow 1/N$ 
3: for each confirmed step  $k$  do
4:    $\Delta\mathbf{p} \leftarrow \mathbf{p}^{\text{PDR}} - \mathbf{p}_{\text{prev}}^{\text{PDR}}, \Delta\psi \leftarrow \text{wrap}(\psi^{\text{PDR}} - \psi_{\text{prev}}^{\text{PDR}})$ 
5:   for  $i \leftarrow 1 \rightarrow N$  do
6:      $\mathbf{p}^{(i)} \leftarrow \mathbf{p}^{(i)} + \Delta\mathbf{p} + \epsilon_p$ 
7:      $\psi^{(i)} \leftarrow \text{wrap}(\psi^{(i)} + \Delta\psi + \epsilon_\psi)$ 
8:   end for
9:   if absolute fix  $\hat{\mathbf{p}}$  available then
10:     $d_M^{2(i)} \leftarrow (\mathbf{p}^{(i)} - \hat{\mathbf{p}})^T \mathbf{R}^{-1} (\mathbf{p}^{(i)} - \hat{\mathbf{p}})$ 
11:     $w^{(i)} \leftarrow w^{(i)} \exp(-\frac{1}{2}d_M^{2(i)})$ ; normalize
12:     $\text{ESS} \leftarrow 1 / \sum_i (w^{(i)})^2$ ; resample if  $\text{ESS} < \tau N$ 
13:   end if
14:    $\hat{\mathbf{x}} \leftarrow \sum_i w^{(i)} \mathbf{x}^{(i)}$ 
15: end for

```

---

tolerant to non-Gaussian noise and occasional outliers, at the cost of higher computational demand. Together, these three back-ends cover filtering, smoothing, and sampling perspectives while keeping the motion and measurement interfaces consistent for systematic evaluation across outdoor, indoor, and transition scenarios.

### 3.2. Experimental Setting and Visualization

The wearable platform integrates an MTi-300 IMU, a GNSS receiver, and a DWM1001 UWB tag on a rigid chest-mounted plate driven by a Raspberry Pi 5, as shown in Fig. 2c. All data are timestamped and logged via ROS 2, and the three fusion back-ends publish their outputs using consistent message interfaces for fair comparison.

Experiments are conducted on a campus site that includes outdoor walkways and an adjacent instrumented laboratory building. We evaluate three scenarios: (i) outdoor walks with GNSS+PDR, (ii) indoor walks with UWB+PDR, and (iii) repeated outdoor–indoor transitions with both absolute sources available intermittently. Outdoor reference is provided by an RTK-capable solution, while the indoor reference is provided by a motion-capture system aligned to the same local ENU frame.

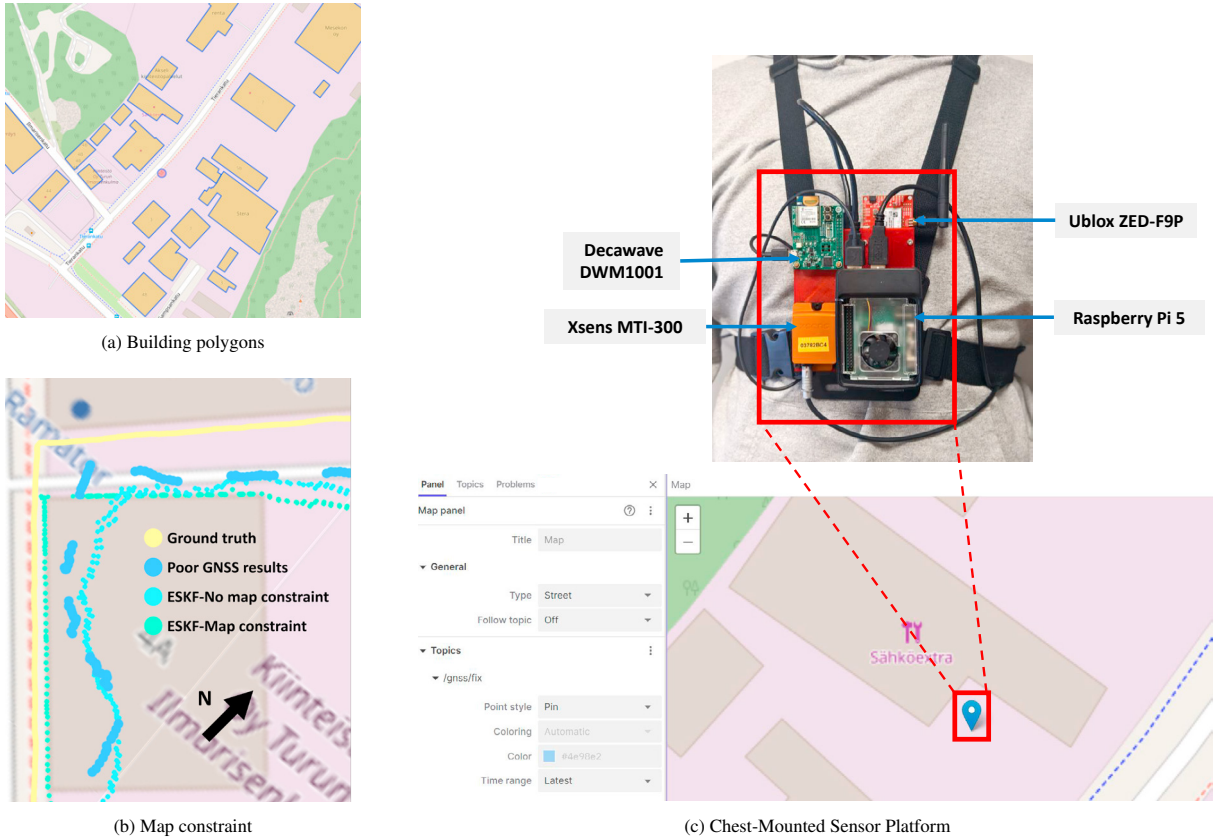


Fig. 2: Overview of map constraints and the chest-mounted sensor platform.

Table 1: Key back-end parameters used in the released ROS 2 implementations.

Item	ESKF	FGO	PF
<b>Model params</b>	$var_f = 1, var_\omega = 1$	$\sigma_{odom,xy} = 0.25m, \sigma_{odom,z} = 0.4m,$ $\sigma_{odom,\psi} = 10^\circ$	$\sigma_{prop,xy} = 0.18m, \sigma_{prop,z} = 0.05m,$ $\sigma_{prop,\psi} = 6^\circ$
<b>Measurement noise</b>	$var_{GNSS} = 10$ (std $\approx$ 3.16m), $var_{UWB} = 0.25$ (std=0.5m)	GNSS: $\sigma_{xy} = 3m, \sigma_z = 5m$ ; UWB: $\sigma_{xy} = 1m, \sigma_z = 1m$	GNSS: $\sigma_{xy} = 3m, \sigma_z = 5m$ ; UWB: $\sigma_{xy} = 1m, \sigma_z = 1m$
<b>Backend hyper-params</b>	$N_{skip} = 5$ (UWB); skip GNSS once UWB available	$W = 80$ ; iters=5; anchor: $\sigma_{xy} =$ $2m, \sigma_z = 3m, \sigma_\psi = 45^\circ$ ; Huber $\delta = 1$	$N = 250$ ; resample if $ESS < 0.5N$

For efficient online inspection and offline figure generation, a unified Foxglove dashboard visualizes raw GNSS, UWB positions, three fused trajectories (ESKF/FGO/PF), PDR step events, and essential diagnostic signals. This setup allows us to trace estimator behavior around challenging segments such as GNSS-degraded building proximities and UWB anchor-visibility changes.

To address reproducibility and fairness concerns, we summarize the key configuration choices shared across the three back-ends. All three methods share the same step-triggered PDR increments (approximately 50 Hz effective rate) and consume the same ENU-transformed absolute fixes from GNSS (outdoor) and UWB trilateration (indoor). Table 1 reports the main model parameters, the assumed measurement noise for GNSS/UWB updates, and the principal algorithm-specific hyper-parameters used in our released ROS 2 implementations. Our evaluation focuses on horizontal (2D) accuracy, consistent with the targeted pedestrian navigation use case.




Fig. 3: The comparison of CDF under different experimental settings.




Fig. 4: Trajectories of the compared methods under different environmental settings.

#### 4. Experimental Results

We evaluate three scenarios: indoor with IMU and UWB using motion capture (MoCap) as ground truth, outdoor with GNSS (u-blox ZED-F9P) and IMU using RTK as ground truth, and a seamless outdoor–indoor experiment with all three sensors where the MoCap system publishes latitude–longitude and the ground truth is the fusion of RTK when outside and MoCap when inside. For all scenarios, accuracy is assessed against ground truth using the horizontal error and reported as cumulative distribution functions (CDFs), shown in Fig. 3, together with summary statistics, shown in Table 2.

The pedestrian walked a square loop three times inside the allowed building. As shown in Fig. 3a and Table 2, ESKF achieves the best indoor accuracy (median 0.415 m, RMSE 0.499 m), closely followed by PF (median 0.447 m, RMSE 0.552 m), while FGO shows larger errors.

The outdoor trial covers sidewalks and building façades with GNSS/PDR fusion. ESKF again provides the most accurate distribution (median 1.491 m, RMSE 2.186 m), followed by FGO and PF (Table 2). The building-polygon constraint mainly improves robustness against multipath outliers, which is also reflected by cleaner tails in the CDF (Fig. 3b).

In the seamless trial, the estimator naturally transitions from GNSS to UWB updates at the doorway without visible discontinuities in the estimated trajectory (Fig. 4c). Across the full sequence, ESKF yields the best overall accuracy (median 0.997 m, RMSE 1.248 m), PF is second, and FGO is less stable in this mixed regime (Table 2). These results indicate that a tightly managed prediction–update cycle with consistent covariance handling (ESKF) is particularly effective for maintaining continuity across heterogeneous sensing conditions.

Table 2: Indoor/Outdoor experiment

	Indoor	Outdoor (Mean, Median, RMSE, STD, Max), Unit: m	Outdoor-Indoor
<b>FGO</b>	(0.836, 0.715, 0.992, 0.535, 3.303)	(1.995, 1.946, 2.389, 1.315, 7.619)	(2.070, 1.962, 2.302, 1.008, 5.351)
<b>Particle filter</b>	(0.491, 0.447, 0.552, 0.253, 1.374)	(2.284, 1.681, 2.816, 1.648, 7.357)	(1.292, 1.017, 1.653, 1.031, 5.108)
<b>ESKF</b>	(0.441, 0.415, 0.499, 0.235, 1.309)	(1.726, 1.491, 2.186, 1.341, 8.181)	(1.085, 0.997, 1.248, 0.617, 2.631)

## 5. Conclusion

This paper presented a seamless outdoor–indoor pedestrian positioning system that fuses GNSS/UWB/IMU under a unified chest-mounted PDR framework, and compared three estimators (ESKF, FGO, PF) under the same sensing streams. Lightweight OSM building-polygon constraints were integrated to improve feasibility and suppress building-proximity outliers.

Across indoor, outdoor, and outdoor–indoor experiments, ESKF delivered the most consistent and accurate performance overall. Indoors, where UWB provides absolute corrections to PDR, ESKF achieved the best error distribution, with PF as a close runner-up, while FGO exhibited larger residual errors in this setup. Outdoors, ESKF remained the most stable against RTK-referenced ground truth, and the map constraint effectively suppressed multipath-driven outliers near facades. In the seamless trial, the estimator naturally shifted from GNSS to UWB updates at the doorway without noticeable trajectory discontinuities. Overall, an ESKF-based pipeline with PDR-driven propagation, opportunistic GNSS/UWB updates, and modest feasibility checks offers a strong balance of accuracy, robustness, and implementation simplicity for practical pedestrian navigation.

While smoothing can outperform filtering in theory, FGO performance is sensitive to window/robustness choices and can be more affected by heavy-tailed absolute-fix errors when only step-wise PDR increments are used. Moreover, the back-ends employ different practical measurement-handling strategies (e.g., UWB-first switching and estimator-specific outlier responses), so our results should be interpreted as a comparison of practical implementations rather than a strict ablation of each design choice. This study focuses on horizontal (2D) positioning and a single instrumented building; extending to complex urban and multi-floor scenarios will require richer accessibility and indoor topology models and additional vertical/floor cues.

## Acknowledgements

This research work is supported by the Research Council of Finland's Digital Waters (DIWA) flagship (Grant No. 359247), DIWA Doctoral Training Pilot project funded by the Ministry of Education and Culture (Finland).

## References

- [1] S. Zhang, H. Tang, L. Chen, Y. Gao, A seamless pedestrian localization system based on gnss/imu/uwb/map integration, *IEEE Internet of Things Journal* (2025).
- [2] V. Di Pietra, P. Dabov, M. Piras, Loosely coupled gnss and uwb with ins integration for indoor/outdoor pedestrian navigation, *Sensors* 20 (21) (2020) 6292.
- [3] D. Yuan, J. Zhang, J. Wang, X. Cui, F. Liu, Y. Zhang, Robustly adaptive ekf pdr/uwb integrated navigation based on additional heading constraint, *Sensors* 21 (13) (2021) 4390.
- [4] Q. Liu, C. Gao, R. Shang, Z. Peng, R. Zhang, L. Gan, Environment perception based seamless indoor and outdoor positioning system of smartphone, *IEEE Sensors Journal* 22 (17) (2022) 17205–17215.
- [5] Z. Jin, Y. Li, Z. Yang, Y. Zhang, Z. Cheng, Real-time indoor positioning based on ble beacons and pedestrian dead reckoning for smartphones, *Applied Sciences* 13 (7) (2023) 4415.
- [6] J. Liu, L. Zhang, J. Xu, J. Shi, Dynamic feasible region-based imu/uwb fusion method for indoor positioning, *IEEE Sensors Journal* 24 (13) (2024) 21447–21457.
- [7] L. Wu, S. Guo, L. Han, C. A. Baris, Indoor positioning method for pedestrian dead reckoning based on multi-source sensors, *Measurement* 229 (2024) 114416.
- [8] J. Wang, J. Liu, X. Xu, Z. Yu, Z. Li, A yaw correction method for pedestrian positioning using two low-cost mimus, *Measurement* 217 (2023) 112992.
- [9] C. Jiang, Y. Chen, C. Chen, J. Jia, H. Sun, T. Wang, J. Hyppa, Implementation and performance analysis of the pdr/gnss integration on a smartphone, *GPS Solutions* 26 (3) (2022) 81.
- [10] I. Klein, Pedestrian inertial navigation: An overview of model and data-driven approaches, *Results in Engineering* (2025) 104077.
- [11] T. Tian, Y. Hu, X. Zhao, H. Zhao, Y. Wang, Z. Liang, An autonomous localization vest system based on advanced adaptive pdr with binocular vision assistance, *Micromachines* 16 (8) (2025) 890.
- [12] X. Hou, J. Bergmann, Rocop: robust continuous inertial position tracking for complex actions emerging from the interaction of human actors and environment, *Applied Intelligence* 55 (6) (2025) 1–10.
- [13] Z. Deng, H. Luo, X. Gao, P. Liu, Fusion-based localization system integrating uwb, imu, and vision, *Applied Sciences* 15 (12) (2025) 6501.