This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

3D UWB Anchor Autocalibration via MDS and Multilateration with Paired-Anchor Height Decoupling

Marios Raspopoulos, *Senior Member, IEEE* and Iacovos Ioannou, *Senior Member, IEEE*

***Abstract*— Accurate anchor positions are required for UWB indoor localization, yet manual surveying is impractical in opportunistic deployments. This paper proposes a lightweight, purely anchor-based 3D autocalibration framework for sparse multihop UWB networks requiring no mobile agents, inertial sensors, or learning components. Two algorithms are developed: an MDS-based approach using Floyd-Warshall gap-fill, Procrustes alignment and iterative pivot least-squares refinement; and a multilateration-based greedy cascade. Both are evaluated in full 3D and Semi-3D modes via a physics-based Monte Carlo simulator modelling realistic non-line-of-sight (NLOS) ranging errors. Standalone MDS and multilateration achieve approximately 1 m 3D RMSE, with height estimation identified as the dominant bottleneck. To address this, a paired-anchor deployment is introduced in which two anchors share each XY position with a fixed vertical separation. Through cross-pair ranging, a closed-form height estimator is obtained by which horizontal position error is cancelled exactly. Results indicate that 3D RMSE is reduced by up to 39% for MDS and 41% for multilateration relative to standalone baselines, reaching 0.54 m and 0.53 m at $n_k = 7$ with negligible computational overhead. The cross-pair differential estimator further provides inherent NLOS/multipath resilience typical of UWB deployments. The reported results constitute a simulation-based proof-of-concept of the paired-anchor autocalibration approach and are intended to provide a quantitative foundation for subsequent experimental validation.**

***Index Terms*— Anchor Autocalibration, Ultra-wideband (UWB), 3D Indoor Localization, Multidimensional Scaling (MDS), Multilateration**

I. Introduction

Indoor localization has become a key enabling technology for a wide range of applications, including industrial automation, mobile robotics, asset tracking, emergency response and smart environments [1]. Among the available radio-frequency technologies, ultra-wideband (UWB) has emerged as a leading solution due to its fine temporal resolution, robustness to multipath fading and ability to provide centimetre-level ranging precision [2], [3]. These properties render UWB particularly attractive for infrastructure-based indoor localization systems, where fixed reference nodes, known as anchors, are used to localize mobile tags.

A critical requirement for accurate UWB localization is the precise knowledge of anchor coordinates. Any error in anchor placement propagates directly into the estimated tag positions. In controlled deployments, anchors are typically surveyed manually using laser rangefinders or map-matching tools. However, in many real-world scenarios — such as emergency response, temporary industrial installations, or ad-hoc deployments in unknown environments — manual surveying is impractical or impossible. Anchors may be placed opportunistically, under time pressure and without access to accurate surveying equipment. In such cases, the localization infrastructure itself must be capable of autonomously determining anchor positions from inter-node UWB measurements alone. This process is commonly referred to as *anchor autocalibration* or *self-calibration*.

Over the past two decades, a wide range of autocalibration techniques has been proposed, spanning optimization-based formulations, distributed algorithms, probabilistic inference and, more recently, multi-sensor fusion within simultaneous localisation and mapping (SLAM) and factor-graph frameworks. While many of these methods achieve impressive accuracy under controlled conditions, their underlying assumptions limit applicability in large-scale or opportunistic deployments. Single-hop methods assume dense pairwise connectivity and typically operate in 2D [4], [5]. Multihop methods relax the connectivity assumption but often rely on iterative global optimisation [6] or probabilistic inference with substantial computational overhead [7]. Multi-sensor approaches achieve centimetre-level accuracy but require mobile robots, LiDAR, or inertial sensors that may not be available at the deployment site [8]–[10].

A further limitation shared by most existing methods is the assumption of two-dimensional deployment. In practice, UWB anchors are mounted at varying heights throughout a building to provide adequate three-dimensional coverage. The vertical component of anchor position is consistently the hardest to estimate from inter-anchor ranging alone, because anchors in typical indoor deployments span a narrow height range relative to their horizontal separation, creating a poorly conditioned vertical geometry. This limitation is largely overlooked in the

This paper has been submitted for review on XX/XX/XX. This work was co-funded by the European Union under the programme of social cohesion "THALIA 2021-2027", through the Research and Innovation Foundation (Project: EXCELLENCE/0524/0218).

M. Raspopoulos is with the School of Sciences at UCLan Cyprus and the INterdisciplinary Science Promotion & Innovative Research Exploration (INSPIRE) Centre, Cyprus (e-mail: mraspopoulos@uclan.ac.uk).

I. Ioannou is with the CYENS Centre of Excellence, Nicosia Cyprus (e-mail: iacovos@cyens.org.cy).

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

literature, where 3D autocalibration results are rare and height estimation accuracy is seldom reported independently.

This paper addresses these limitations through two complementary contributions. The *first* is a pair of lightweight 3D autocalibration algorithms — one based on multidimensional scaling (MDS) and one on multilateration — that operate exclusively on inter-anchor ranging measurements, tolerate missing links due to limited radio range or non-line-of-sight (NLOS) blockage and require no optimisation solver, inertial sensors, or machine learning components. The MDS algorithm recovers the global anchor geometry via Floyd-Warshall gap-fill, classical MDS embedding, Procrustes alignment to the known-anchor frame and iterative pivot least-squares refinement with a residual gate. The multilateration algorithm propagates the known-anchor set outward through the network via a confidence-ordered greedy cascade of pivot least-squares solutions. Both operate in a full 3D mode and a Semi-3D mode that exploits inexpensively measured anchor heights as side information. The *second* contribution is a structured *paired-anchor* deployment strategy that places two anchors at each XY position with a fixed known vertical separation Δh. Cross-pair ranging measurements between a pair of unknown anchors and a pair of known references yield a *closed-form height estimator* that exactly cancels the horizontal inter-pair distance, making height estimation decoupled from horizontal position error for the first time in anchor autocalibration. This estimator, whose noise variance decreases as $1/\Delta h^2$, is applied as a post-processing step to either backbone algorithm with negligible computational overhead.

The framework is evaluated via a physics-based Monte Carlo simulation platform that models UWB link budgets, building wall-crossing NLOS effects and distance-dependent ranging noise in a realistic multi-room indoor environment. Three independent parameter sweeps characterise the effect of known-anchor count, ranging noise level and network size on autocalibration accuracy. The paired-anchor algorithms are additionally evaluated across various vertical separations to quantify the accuracy–deployment-constraint trade-off.

The main contributions of this work are as follows: (1) *MDS-based 3D autocalibration* via Floyd-Warshall gap-fill, classical MDS embedding, Procrustes alignment and iterative pivot least-squares refinement with a residual gate, (2) *Multilateration-based 3D autocalibration* via a greedy cascade of pivot least-squares solutions and (3) *Paired-anchor deployment with a closed-form height estimator* derived from cross-pair differential ranging, by which the 3D RMSE is significantly reduced relative to the standalone 3D baseline.

The remainder of the paper is organised as follows. Section II reviews existing autocalibration methods and positions the proposed framework within this landscape. Section III formalises the system model and the autocalibration problem. Section IV presents the MDS-based, multilateration-based and paired-anchor algorithms together with the simulation environment. Section V reports Monte Carlo evaluation results for both algorithm families across various parameter sweeps and discusses them in the context of related work. Section VI concludes the paper and outlines future research directions.

II. Related Work

A. General Landscape

Anchor self-calibration methods can be broadly classified into *single-hop* and *multihop* approaches. Single-hop methods assume that all anchors can directly measure pairwise distances, forming a fully connected ranging graph. Under this assumption, calibration reduces to a constrained geometric reconstruction problem, typically solved using MDS, trilateration, optimization, or robust estimation.

Multihop methods operate on sparse and irregular connectivity graphs, where anchors may only measure distances to a subset of neighbours. Such conditions are common in large-scale deployments and in environments affected by NLOS propagation. Multihop calibration is closely related to cooperative localization in sensor networks, where nodes must infer their positions from incomplete and noisy inter-node measurements. These methods must address ambiguity, error accumulation and the absence of a global reference frame.

B. Single-Hop Self-Calibration

Early and representative single-hop approaches focus on robust and efficient reconstruction under dense connectivity. Random Sample Consensus (RANSAC)-based minimal solvers for Time-of-Arrival (TOA) calibration are presented in [11], [12], achieving decimeter-level accuracy in small volumes. Time Difference of Arrival (TDOA)-based systems introduce clock asynchrony, addressed in [13] via joint synchronization and calibration. Optimization-based schemes refine anchor positions from rough initial estimates using a mobile tag [14], while iterative centroid and trilateration methods [15] demonstrate performance comparable to manual calibration.

A class of methods closely related to the approach proposed in this work is based on MDS. MDS reconstructs node coordinates from pairwise distances by solving a global distance-geometry problem [16]. While MDS-based approaches provide an elegant and computationally efficient way to recover a globally consistent geometry, they are primarily designed for relative positioning and assume sufficiently dense inter-node connectivity. Moreover, they do not explicitly address the challenges posed by indoor 3D deployments, where vertical observability is weak and measurement graphs are often highly fragmented. These limitations become critical in opportunistic UWB deployments, motivating the multihop and paired-anchor extensions proposed in this work.

A probabilistic alternative to MDS and trilateration-based single-hop calibration is presented by Schwarzbach and Jung [7], who evaluate a grid-based Bayesian uncertainty propagation framework for collaborative UWB self-calibration in a fully meshed network of 12 nodes deployed in a $44 \times 30\,\mathrm{m}^2$ industrial hall. Rather than collapsing anchor position uncertainty to a covariance matrix — as parametric approaches such as EKF do — the proposed method maintains and propagates discrete probability mass distributions across all nodes. During the measurement update, the likelihood of each grid cell is computed as a weighted superposition of contributions from the top-$\tau\%$ probability hypotheses of every refer-

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

ence node, thereby preserving multimodal and non-Gaussian uncertainty structure. The key practical advantage emerges under NLOS-contaminated initialisation: while a closed-form baseline degrades catastrophically (median error 1.48–$2.43\,\mathrm{m}$, with extreme outliers exceeding $30\,\mathrm{m}$), the grid-based method degrades gracefully (median error 0.62–$0.99\,\mathrm{m}$), with the mean-error ratio rising from $1.6\times$ under LOS to 3.7–$7.4\times$ under NLOS initialisation. The method also tolerates partial connectivity by incorporating only available reference nodes in the measurement update. These advantages come at the cost of significant computational overhead and a centralised architecture that requires aggregating all pairwise measurements at a single coordinator. Furthermore, the approach operates in 2D and targets near-fully meshed networks, limiting applicability to sparse or large-scale 3D deployments.

Distributed formulations further reduce centralization. In [4], anchors establish a coordinate system by fixing four reference nodes and minimizing a global cost over inter-anchor distances using gradient descent, achieving sub-decimeter RMSE. Hybrid schemes combine optimization with learning-based error correction [5], or fuse UWB with inertial sensing in Extended Kalman Filtering (EKF) frameworks [17]. While effective, these methods rely on dense connectivity and are therefore unsuitable for sparse or large-scale deployments.

C. Multihop and Network-Scale Methods

Multihop self-calibration relaxes the full connectivity assumption and targets large-scale networks. Classical solutions include iterative MDS [16], semidefinite programming (SDP) relaxations [18], distributed self-positioning [19], nonparametric belief propagation [20] and simulated annealing [21]. These methods establish relative coordinate systems under incomplete distance information and scale to hundreds of nodes in simulation. A UWB multihop solution is presented in [6], combining linear least-squares initialization with global optimization using Limited-memory Broyden–Fletcher–Goldfarb–Shanno with Box constraints (L-BFGS-B) and optional refinement via a mobile node. In an indoor environment with 18 anchors, a mean absolute error of $21.6\,\mathrm{cm}$ is reported under severe NLOS. Although effective, this framework relies on heavy nonlinear optimization, requires mobile assistance in its complete form and is also limited to 2D.

D. Recent Multi-Sensor and SLAM-Based Approaches

Recent research increasingly integrates UWB with additional sensing modalities and SLAM-style optimization. Factor Graph Optimization (FGO) frameworks jointly estimate anchor positions and robot trajectories using UWB and LiDAR [8]. Two-stage UWB–inertial measurement unit (IMU)–odometer pipelines employ MDS for initialization followed by graph-based refinement and learning-based NLOS mitigation [9]. SLAM-inspired inverse positioning [10] and Kalman-filter-based UWB–IMU fusion [22] achieve anchor RMSEs below $10\,\mathrm{cm}$. System-level calibration using anchor-to-anchor and anchor-to-tag ranging with nonlinear least squares is reported in [23], substantially reducing ranging bias and deployment effort. These approaches demonstrate high accuracy and robustness, but at the cost of increased system complexity. They typically require mobile robots, inertial sensors, LiDAR, or pre-defined trajectories and rely on computationally intensive optimization.

E. Positioning of This Work

In contrast to the literature, this work proposes a *purely anchor-based*, lightweight framework that evaluates two complementary autocalibration strategies — MDS-based and multilateration-based — each operating independently in both full 3D and Semi-3D modes and extends both with a structured *paired-anchor* deployment that enables a closed-form, range-decoupled height estimator. By exploiting the vertical pair geometry to decouple Z estimation from horizontal ranging error — without any optimisation solver, inertial sensor, or learning component — the proposed approach bridges the gap between graph-theoretic multihop localization and optimization-heavy, multi-sensor pipelines, making it well-suited for rapid, opportunistic deployments in sparse 3D environments where only UWB infrastructure is available, such as emergency response and industrial installations.

III. System Model and Problem Description

An overview of the proposed 3D UWB anchor autocalibration framework is shown in Fig. 1. The indoor environment and physical anchor deployment constitute the input. A physics-based channel model generates the sparse measured distance matrix $\widetilde{\mathbf{D}}$ from the link budget and ranging error models. The core algorithms recover an initial XY warm start together with a closed-form, height-decoupled estimate and a final refinement stage produces the 3D anchor positions. This section formalises the system model and the autocalibration problem, while the algorithms themselves are presented in Section IV.

A. Network Topology and Problem Formulation

Consider a set of N UWB anchors deployed in a three-dimensional indoor environment. Let $\mathbf{p}_i = [x_i,\, y_i,\, z_i]^\top \in \mathbb{R}^3$ denote the true position of anchor i. The anchors are partitioned into two disjoint sets: the *known* set $\mathcal{K} \subset \{1,\dots,N\}$, whose positions are assumed to be known with bounded uncertainty and the *unknown* set $\mathcal{U} = \{1,\dots,N\} \setminus \mathcal{K}$, whose positions are to be recovered. The measured distance matrix $\widetilde{\mathbf{D}} \in \mathbb{R}^{N\times N}$ is sparse: entry $\widetilde{d}_{ij}$ is available only when anchors i and j share a radio link. Unavailable entries are represented as `NaN`. The *connectivity graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ has vertex set $\mathcal{V} = \{1,\dots,N\}$ and edge set $\mathcal{E} = \{(i,j) : \widetilde{d}_{ij} \text{ is available}\}$. The autocalibration problem can be stated as:

$$\hat{\mathbf{p}}_i = \arg\min_{\mathbf{p}_i} \sum_{(i,j)\in\mathcal{E}} w_{ij} \left(\|\mathbf{p}_i - \mathbf{p}_j\| - \widetilde{d}_{ij} \right)^2, \quad \forall\, i \in \mathcal{U}, \quad (1)$$

subject to the known-anchor constraints $\hat{\mathbf{p}}_i = \mathbf{p}_i + \boldsymbol{\eta}_i$ for all $i \in \mathcal{K}$, where $\boldsymbol{\eta}_i \sim \mathcal{N}(\mathbf{0}, \sigma_{\mathrm{KA}}^2 \mathbf{I})$ models the positional measurement uncertainty of the known anchors with standard deviation σ_{KA}. The scalars $w_{ij} > 0$ are measurement reliability weights derived from the link budget, as detailed below.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

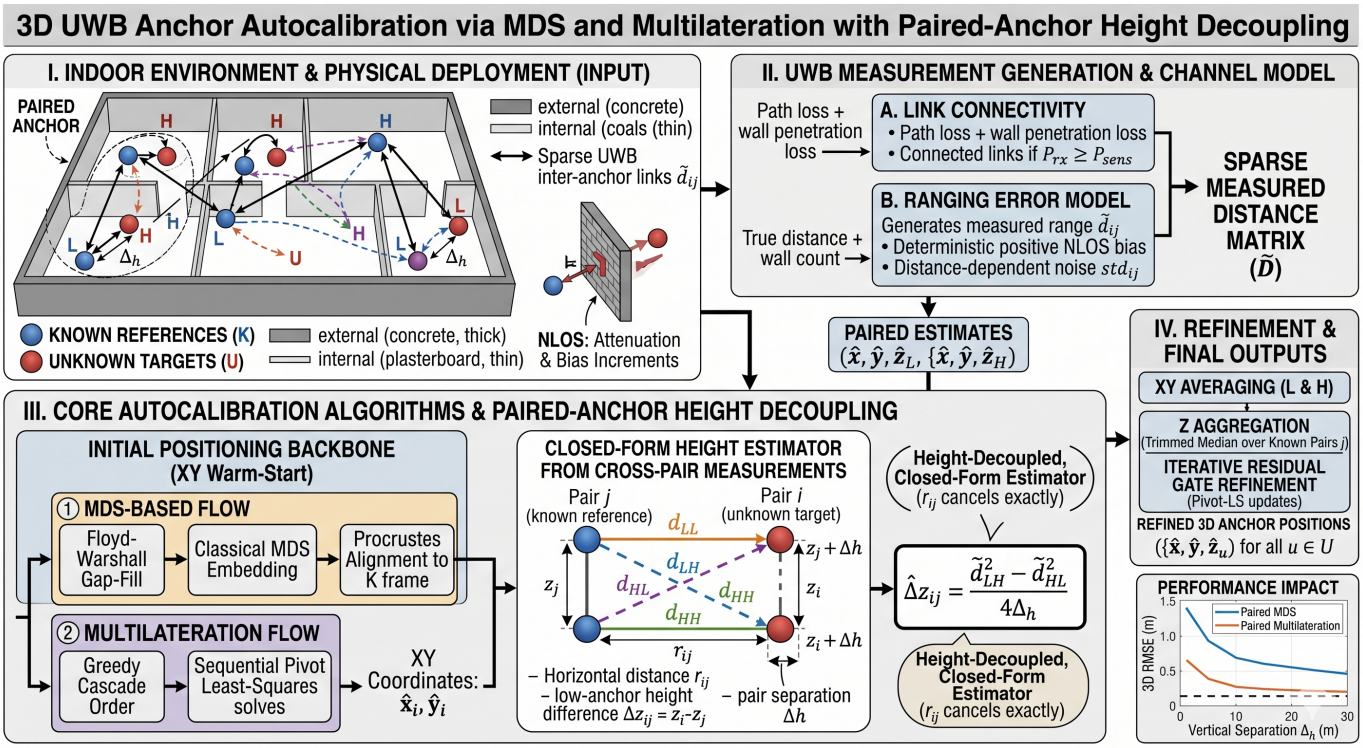


Fig. 1. Overview of the proposed 3D UWB anchor autocalibration framework. (I) The indoor environment and physical anchor deployment, with paired low and high anchors separated by a fixed vertical offset $\boldsymbol{\Delta h}$. (II) UWB measurement generation and channel model, comprising link connectivity and the ranging error model, yielding the sparse measured distance matrix $\tilde{\mathbf{D}}$. (III) The MDS-based and multilateration-based autocalibration algorithms providing the XY warm start, together with the closed-form paired-anchor height estimator derived from cross-pair measurements. (IV) The refinement stage producing the final 3D anchor positions.

B. UWB Link Budget and Connectivity Model

The received power at anchor j from a transmission by anchor i is computed using the standard log-distance path loss model augmented with wall penetration losses:

$$P_{\text{rx}}^{(ij)} = P_{\text{tx}} - \underbrace{PL_0 + 10\,n\,\log_{10}\left(\frac{d_{ij}}{d_0}\right)}_{\text{free-space path loss}} - \underbrace{n_{\text{int}}^{(ij)} L_{\text{int}} + n_{\text{ext}}^{(ij)} L_{\text{ext}}}_{\text{wall penetration loss}} \quad (2)$$

where P_{tx} is the transmit power in dBm, $PL_0 = 20\log_{10}(4\pi d_0 f_c/c)$ is the free-space path loss at the reference distance $d_0 = 1\,\text{m}$ for carrier frequency f_c and speed of light c, n is the path loss exponent, $n_{\text{int}}^{(ij)}$ and $n_{\text{ext}}^{(ij)}$ are the numbers of internal and external wall intersections along the ray from i to j and L_{int}, L_{ext} are the corresponding per-wall attenuation coefficients in dB. A link between anchors i and j exists only if the received power exceeds the receiver sensitivity threshold P_{sens}: $(i,j) \in \mathcal{E} \iff P_{\text{rx}}^{(ij)} \geq P_{\text{sens}}$. The wall intersection count $n_{\text{int}}^{(ij)} + n_{\text{ext}}^{(ij)}$ is computed geometrically by testing the segment $\overline{\mathbf{p}_i\mathbf{p}_j}$ against each facet of the building model.

C. Ranging Error Model

For each available link $(i,j) \in \mathcal{E}$, the measured range $\widetilde{d}_{ij}$ deviates from the true Euclidean distance $d_{ij} = \|\mathbf{p}_i - \mathbf{p}_j\|$ due to multipath, NLOS propagation and hardware timing errors according to:

$$\widetilde{d}_{ij} = d_{ij} + b_{ij} + \varepsilon_{ij}, \quad (3)$$

where b_{ij} is a deterministic positive NLOS bias and ε_{ij} is a zero-mean Gaussian noise term. Under LOS conditions ($n_{\text{int}}^{(ij)} + n_{\text{ext}}^{(ij)} = 0$): $b_{ij} = 0$ and $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_{\text{LOS}}^2)$. Under NLOS conditions, each wall crossing introduces additional bias and increases the noise variance:

$$b_{ij} = n_{\text{int}}^{(ij)} b_{\text{int}} + n_{\text{ext}}^{(ij)} b_{\text{ext}}, \quad (4)$$

$$\sigma_{ij} = \sigma_{\text{NLOS}} + \left(n_{\text{int}}^{(ij)} + n_{\text{ext}}^{(ij)}\right)\sigma_{\text{wall}}, \quad (5)$$

where b_{int}, b_{ext} are the per-wall bias increments and σ_{wall} is the per-wall noise increment. The weight w_{ij} in (1) is defined as the inverse of the link variance: $w_{ij} = \sigma_{ij}^{-2}$.

D. Semi-3D Variant: Known Heights

In many practical indoor deployments, anchor heights can be measured inexpensively (e.g. with a tape measure or laser rangefinder or a barometer) even when horizontal positions are unknown. We therefore consider a *Semi-3D* operating mode in which the height z_i of every anchor is available as a noisy measurement:

$$\tilde{z}_i = z_i + \delta_i, \quad \delta_i \sim \mathcal{N}(0, \sigma_Z^2), \quad (6)$$

with σ_Z characterising the height measurement accuracy. In the Semi-3D mode, the estimation problem reduces to recovering only the horizontal coordinates (x_i, y_i) for each $i \in \mathcal{U}$, with the vertical coordinate substituted by $\tilde{z}_i$ after

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

every estimation step. This decoupling substantially improves accuracy in quasi-planar deployments where UWB has weak vertical observability, as anchors typically span a narrow height range relative to the room height.

IV. Methodology

This section presents the complete methodological framework of the proposed anchor autocalibration system. To support a rigorous and reproducible evaluation, we have developed a custom MATLAB-based simulation platform that synthesises realistic UWB measurement data under configurable deployment conditions. The simulator allows the user to define the three-dimensional geometry of the indoor environment as well as the positions and heights of all UWB anchors. The radio front-end is characterised by user-defined parameters including transmit power, carrier frequency, path loss exponent, receiver sensitivity and per-wall attenuation coefficients for both internal partitions and external structural walls. The channel model generates per-link received signal strength (RSS) values from a log-distance path loss formula augmented with wall-crossing losses, determines link connectivity based on a received power threshold and produces noisy range measurements according to a distance-dependent LOS/NLOS error model with positive range bias and Gaussian noise whose variance increases with the number of wall intersections along each path. This simulation environment is used to drive a Monte Carlo evaluation framework in which anchor positions and known-anchor subsets are randomised across hundreds of independent trials, enabling statistically robust performance assessment under diverse deployment geometries and connectivity conditions.

Within this framework, two autocalibration algorithms are developed and evaluated, each in a standard full 3D mode and a Semi-3D variant that exploits known anchor heights. The first, *MDS-based autocalibration*, reconstructs the global anchor geometry from the incomplete distance matrix using multidimensional scaling followed by iterative pivot least-squares refinement. The second, *multilateration-based autocalibration*, estimates anchor positions sequentially via a greedy cascade of pivot least-squares solutions, propagating the known set outward through the network. Both algorithms are then extended to a *paired-anchor* deployment in which anchors are installed in vertically separated pairs at each XY position. This structured deployment enables a closed-form, range-decoupled height estimator that substantially reduces vertical localisation error while preserving the horizontal accuracy of the underlying algorithm.

A. MDS-Based Autocalibration

Classical MDS recovers a Euclidean embedding from a complete matrix of pairwise squared distances. Applying MDS directly to $\widetilde{\mathbf{D}}$ is problematic because: (i) many entries are missing due to limited radio range; (ii) measurement noise violates the Euclidean distance constraint; and (iii) the solution is defined only up to a rigid-body transformation. We address each of these challenges in turn.

1) Floyd-Warshall Gap-Fill: Missing entries in $\widetilde{\mathbf{D}}$ are imputed using all-pairs shortest paths computed by the Floyd-Warshall algorithm [24], [25]. Defining the initialisation $\mathbf{F}^{(0)}$ with $F_{ij}^{(0)} = \widetilde{d}_{ij}$ for $(i,j) \in \mathcal{E}$, $F_{ii}^{(0)} = 0$ and $F_{ij}^{(0)} = \infty$ otherwise, the algorithm iterates:

$$F_{ij}^{(k)} = \min\Big(F_{ij}^{(k-1)},\ F_{ik}^{(k-1)} + F_{kj}^{(k-1)}\Big), \quad k = 1, \ldots, N \tag{7}$$

for all i, j. Entries that were originally missing and have been filled by a shortest path of length $\ell \geq 2$ hops are scaled by a shrinkage factor $\alpha \in (0,1)$:

$$\hat{F}_{ij} = \begin{cases} \widetilde{d}_{ij}, & (i,j) \in \mathcal{E}, \\ \alpha\, F_{ij}^{(N)}, & (i,j) \notin \mathcal{E}, \end{cases} \tag{8}$$

yielding the imputed matrix $\hat{\mathbf{F}}$. The shrinkage corrects the systematic overestimation introduced by multi-hop path routing, which accumulates intermediate ranging noise along the path. The resulting symmetric matrix $\hat{\mathbf{F}}$ is made symmetric by averaging: $\hat{F}_{ij} \leftarrow (\hat{F}_{ij} + \hat{F}_{ji})/2$.

2) Classical MDS Embedding: Given the imputed squared-distance matrix $\hat{\mathbf{F}}^{\circ 2}$, the double-centring operator produces the Gram matrix.

$$\mathbf{B} = -\tfrac{1}{2}\,\mathbf{J}\,\hat{\mathbf{F}}^{\circ 2}\,\mathbf{J} \tag{9}$$

,where $\mathbf{J} = \mathbf{I}_N - \frac{1}{N}\mathbf{1}\mathbf{1}^\top$ is the centring matrix. If the imputed distances were Euclidean, $\mathbf{B}$ would equal $\mathbf{X}\mathbf{X}^\top$ for some coordinate matrix $\mathbf{X}$. In practice, $\mathbf{B}$ is symmetrised and its eigenvalue decomposition is computed as $\mathbf{B} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$, where the columns of $\mathbf{V}$ are unit eigenvectors and $\mathbf{\Lambda} = \operatorname{diag}(\lambda_1, \ldots, \lambda_N)$ with $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$. Retaining the top three eigenpairs $(\lambda_k, \mathbf{v}_k)$, $k = 1, 2, 3$, the 3D embedding is:

$$\hat{\mathbf{X}} = \mathbf{V}_{1:3}\,\operatorname{diag}\Big(\sqrt{\max(\lambda_1, 0)},\ \sqrt{\max(\lambda_2, 0)},\ \sqrt{\max(\lambda_3, 0)}\Big), \tag{10}$$

where negative eigenvalues (arising from non-Euclidean imputed distances) are clamped to zero. The fraction of total variance explained by the 3D embedding is $\sum_{k=1}^{3} \max(\lambda_k, 0) / \sum_{k=1}^{N} \max(\lambda_k, 0)$, which serves as an indicator of embedding quality.

3) Procrustes Alignment: The embedding $\hat{\mathbf{X}}$ is defined in an arbitrary coordinate frame. It must be aligned to the global frame defined by the known anchors $\{\mathbf{p}_i : i \in \mathcal{K}\}$. Let $\mathbf{P} = \hat{\mathbf{X}}_{\mathcal{K}} \in \mathbb{R}^{K \times 3}$ be the embedded positions of the known anchors and $\mathbf{Q} \in \mathbb{R}^{K \times 3}$ be their true (noisy) positions. The orthogonal Procrustes problem [26], [27] seeks the rotation matrix $\mathbf{R} \in \mathrm{SO}(3)$ and translation vector $\mathbf{t} \in \mathbb{R}^3$ that minimise:

$$\min_{\mathbf{R},\,\mathbf{t}} \ \big\|\mathbf{P}\mathbf{R}^\top + \mathbf{1}\mathbf{t}^\top - \mathbf{Q}\big\|_F^2 . \tag{11}$$

Denoting the centred matrices $\bar{\mathbf{P}} = \mathbf{P} - \bar{\mathbf{p}}\mathbf{1}^\top$ and $\bar{\mathbf{Q}} = \mathbf{Q} - \bar{\mathbf{q}}\mathbf{1}^\top$, the solution is obtained via SVD of the cross-covariance matrix $\mathbf{M} = \bar{\mathbf{P}}^\top\bar{\mathbf{Q}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$:

$$\mathbf{R} = \mathbf{V}\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \det(\mathbf{V}\mathbf{U}^\top) \end{bmatrix}\mathbf{U}^\top, \qquad \mathbf{t} = \bar{\mathbf{q}} - \bar{\mathbf{p}}\,\mathbf{R}^\top, \tag{12}$$

where the determinant term corrects for improper rotations (reflections), ensuring $\det(\mathbf{R}) = +1$. The aligned embedding is $\hat{\mathbf{X}}^* = \hat{\mathbf{X}}\mathbf{R}^\top + \mathbf{1}\mathbf{t}^\top$. The Procrustes residual

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

$\varepsilon_{\text{proc}} = \|\hat{\mathbf{X}}^*_{\mathcal{K}} - \mathbf{Q}\|_F / \sqrt{K}$ quantifies the alignment quality and is used as a proxy for embedding accuracy in the residual-gated refinement stage.

4) Iterative Pivot Least-Squares Refinement with Residual Gate: The MDS embedding provides a warm start for each unknown anchor. Each anchor position is then refined iteratively using a weighted pivot least-squares (pivot-LS) procedure. For anchor $u \in \mathcal{U}$, let $\mathcal{R}_u = \{j : (u,j) \in \mathcal{E},\ \hat{\mathbf{p}}_j \text{ available}\}$ denote its reference set (known or already estimated anchors with measured links). Starting from (3) and squaring:

$$\|\mathbf{p}_u - \mathbf{p}_j\|^2 = \widetilde{d}^2_{uj} + \text{error terms}. \tag{13}$$

Selecting anchor $r_1 \in \mathcal{R}_u$ as the pivot and subtracting its equation from that of each $r_j \in \mathcal{R}_u \setminus \{r_1\}$ eliminates the unknown $\|\mathbf{p}_u\|^2$, yielding the linear system:

$$\mathbf{A}\begin{bmatrix} x_u \\ y_u \end{bmatrix} = \mathbf{b}, \tag{14}$$

where for $j = 2, \ldots, |\mathcal{R}_u|$:

$$\mathbf{A} = 2\left[\mathbf{p}_{r_j,\text{xy}} - \mathbf{p}_{r_1,\text{xy}}\right], \tag{15}$$

$$b_j = d^2_{\text{xy},r_j} - d^2_{\text{xy},r_1} + \|\mathbf{p}_{r_1,\text{xy}}\|^2 - \|\mathbf{p}_{r_j,\text{xy}}\|^2. \tag{16}$$

Here $d_{\text{xy},r_j} = \sqrt{\max(0,\ \widetilde{d}^2_{ur_j} - (z_u - z_{r_j})^2)}$ is the projected horizontal range after accounting for the known or estimated height difference. The jth row of $\mathbf{A}$ and the scalar b_j are weighted by $w_j = 1/\max(d_{\text{xy},r_j}, \epsilon)$ to downweight references that are nearly co-planar with u in the horizontal plane. The system is solved in the minimum-norm least-squares sense:

$$\begin{bmatrix} \hat{x}_u \\ \hat{y}_u \end{bmatrix} = (\mathbf{W}^{1/2}\mathbf{A})^\dagger\, \mathbf{W}^{1/2}\mathbf{b}, \tag{17}$$

where $\mathbf{W} = \text{diag}(w_2, \ldots, w_{|\mathcal{R}_u|})$ and $(\cdot)^\dagger$ denotes the Moore-Penrose pseudoinverse.

In **3D mode**, the height z_u is estimated by inverse-distance weighting (IDW) over the known anchors $\mathcal{K}_u = \mathcal{K} \cap \mathcal{R}_u$:

$$\hat{z}_u = \begin{cases} \dfrac{\sum_{k\in\mathcal{K}_u} \widetilde{d}^{-1}_{uk} z_k}{\sum_{k\in\mathcal{K}_u} \widetilde{d}^{-1}_{uk}}, & \mathcal{K}_u \neq \emptyset, \\ \bar{z}_{\mathcal{K}}, & \mathcal{K}_u = \emptyset, \end{cases} \tag{18}$$

where $\bar{z}_{\mathcal{K}}$ is the mean height of all known anchors. In the **Semi-3D mode**, $\hat{z}_u = \tilde{z}_u$ from (6) is substituted directly, bypassing the IDW estimation entirely. A *residual gate* controls the acceptance of each refined estimate. Let $\rho(\mathbf{p}, \mathcal{R}_u)$ denote the root mean squared ranging residual of position $\mathbf{p}$ against the reference set $\mathcal{R}_u$:

$$\rho(\mathbf{p}, \mathcal{R}_u) = \sqrt{\frac{1}{|\mathcal{R}_u|}\sum_{j\in\mathcal{R}_u}\left(\|\mathbf{p} - \hat{\mathbf{p}}_j\| - \widetilde{d}_{uj}\right)^2}. \tag{19}$$

The refined position $\hat{\mathbf{p}}^{\text{new}}_u$ is accepted only if it reduces the residual relative to the current estimate $\hat{\mathbf{p}}^{\text{old}}_u$:

$$\hat{\mathbf{p}}_u \leftarrow \begin{cases} \hat{\mathbf{p}}^{\text{new}}_u, & \rho(\hat{\mathbf{p}}^{\text{new}}_u, \mathcal{R}_u) < \rho(\hat{\mathbf{p}}^{\text{old}}_u, \mathcal{R}_u), \\ \hat{\mathbf{p}}^{\text{old}}_u, & \text{otherwise}. \end{cases} \tag{20}$$

Algorithm 1 MDS Autocalibration

Require: $\widetilde{\mathbf{D}}$, $\mathcal{K}$, $\{\mathbf{p}_i\}_{i\in\mathcal{K}}$, $\{\tilde{z}_i\}_{i=1}^N$ (Semi-3D only), α, P
Ensure: $\{\hat{\mathbf{p}}_i\}_{i\in\mathcal{U}}$

```
Compute F̂ via Floyd-Warshall with shrinkage α (7)–(8)
Compute Gram matrix B (9)
Eigendecompose B = VΛV^⊤; form X̂ (10)
Compute R, t via Procrustes; align X̂* ← X̂R^⊤ + 1t^⊤ (12)
p̂_u ← clip(X̂*_u, 0, r_max) for all u ∈ U
if Semi-3D then ẑ_u ← z̃_u for all u ∈ U
end if
for p = 1, ..., P do
    Sort U by |{(u, j) ∈ E}| descending
    for each u ∈ U (sorted) do
        Compute ẑ_u via IDW (18) (3D) or ẑ_u ← z̃_u (Semi-3D)
        Solve (17) to get p̂^new_u
        Apply residual gate (20)
    end for
end for
return {p̂_i}_{i∈U}
```

This gate prevents poorly estimated neighbours from degrading a well-placed anchor through successive refinement passes. Anchors are processed in descending order of link count to maximise the number of reliable references available at each step. The process is repeated for P passes (typically $P = 3$). Algorithm 1 summarises the MDS autocalibration procedure.

B. Multilateration-Based Autocalibration

While MDS recovers the global geometric structure simultaneously, it requires a sufficiently dense and connected measurement graph to produce a reliable embedding. Multilateration takes the complementary approach: it estimates one anchor at a time from locally available measurements, propagating the known set outwards through the network via a greedy cascade. This makes multilateration resilient to fragmented connectivity, as each anchor only requires a small number of reliable neighbours rather than a globally well-connected graph.

1) Greedy Cascade Warm Start: The multilateration autocalibration begins with a greedy cascade that builds the reference set incrementally. At each step, the unresolved anchor with the most available references (known or already estimated anchors connected by a measured link) is selected for estimation. Ties are broken by preferring the anchor with the smallest mean distance to its references, as closer references typically yield better conditioned geometry. Formally, the selection criterion at iteration t is:

$$u^* = \arg\max_{u\in\mathcal{U}^{(t)}} \left|\mathcal{R}^{(t)}_u\right|, \tag{21}$$

with tie-breaking on $\bar{d}_u = |\mathcal{R}^{(t)}_u|^{-1}\sum_{j\in\mathcal{R}^{(t)}_u}\widetilde{d}_{uj}$. The selected anchor u^* is immediately estimated using (17) and promoted to the reference set: $\mathcal{K}^{(t+1)} = \mathcal{K}^{(t)} \cup \{u^*\}$. The cascade continues until either all anchors are estimated or no unresolved

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

anchor has $|\mathcal{R}_u^{(t)}| \geq M_{\min}$, where $M_{\min} \geq 3$ is the minimum number of references required.

The greedy cascade differs from MDS-based initialisation in a key respect: each anchor is estimated using the *most recently updated* reference positions rather than a globally computed embedding. This means that errors in early estimates can propagate to later ones — the cascade order critically affects accuracy. The strategy of processing the most-connected anchor first mitigates this by building the reference set from anchors with the strongest geometric constraints before attempting anchors with fewer or more distant links.

2) Pivot-LS Formulation and Z Estimation: The pivot-LS formulation follows exactly as described in Section IV-A through (17). The key difference in the multilateration context is that the reference set $\mathcal{R}_u$ includes both originally known anchors and anchors estimated by the cascade, rather than only the MDS-aligned positions. This allows the cascade to reach anchors that have no direct link to any originally known node, provided a path exists through the network.

Height estimation in 3D mode follows (18), restricted to known anchors in $\mathcal{K}_u$: only originally known anchors (not cascade estimates) contribute to the IDW average for z_u, as their heights are known with bounded uncertainty σ_{KA}, whereas cascade-estimated heights may carry accumulated error. In Semi-3D mode, $\hat{z}_u = \tilde{z}_u$ regardless of the reference set composition.

3) Iterative Residual Gate Refinement: Following the greedy cascade, all estimated positions are refined using the same iterative residual gate procedure described in Section IV-A. Crucially, at this stage the reference set for each anchor u includes *all* anchors with finite estimated positions and a measured link to u — not only originally known ones — since the cascade estimates have been stabilised by the greedy ordering. This enables mutual refinement between adjacent estimated anchors, progressively reducing the accumulated cascade error. Algorithm 2 summarises the full procedure.

Algorithm 2 Multilateration Autocalibration

Require: $\widetilde{\mathbf{D}}$, $\mathcal{K}$, $\{\mathbf{p}_i\}_{i\in\mathcal{K}}$, $\{\tilde{z}_i\}_{i=1}^{N}$ (Semi-3D only), $M_{\min}$, P
Ensure: $\{\hat{\mathbf{p}}_i\}_{i\in\mathcal{U}}$
$\mathcal{K}^{(1)} \leftarrow \mathcal{K}$; $\mathcal{U}^{(1)} \leftarrow \mathcal{U}$; $t \leftarrow 1$
while $\mathcal{U}^{(t)} \neq \emptyset$ **do**
  Select u^* via (21)
  if no $u \in \mathcal{U}^{(t)}$ has $|\mathcal{R}_u^{(t)}| \geq M_{\min}$ **then break**
  end if
  Estimate $\hat{z}_{u^*}$ via IDW (18) (3D) or $\leftarrow \tilde{z}_{u^*}$ (Semi-3D)
  Solve (17) to get $\hat{\mathbf{p}}_{u^*}$
  $\mathcal{K}^{(t+1)} \leftarrow \mathcal{K}^{(t)} \cup \{u^*\}$; $\mathcal{U}^{(t+1)} \leftarrow \mathcal{U}^{(t)} \setminus \{u^*\}$; $t \leftarrow t+1$
end while
for $p = 1, \ldots, P$ **do** ▷ Gate refinement
  Sort $\mathcal{U}$ by $|\{(u,j) \in \mathcal{E}\}|$ descending
  for each $u \in \mathcal{U}$ **do**
    Compute $\hat{z}_u$ (3D) or use $\tilde{z}_u$ (Semi-3D)
    Solve (17); apply gate (20)
  end for
end for
return $\{\hat{\mathbf{p}}_i\}_{i\in\mathcal{U}}$

C. Paired-Anchor Autocalibration

The standalone MDS and multilateration methods described in Sections IV-A and IV-B treat each anchor independently and rely on the IDW average (18) to recover height, which is sensitive to ranging noise in the vertical component. We propose a structured *paired-anchor* deployment that breaks the height–range ambiguity exactly, yielding a closed-form height estimate that is decoupled from horizontal position error.

1) Deployment Model and Notation: Instead of placing N anchors at arbitrary 3D locations, anchors are deployed in pairs: each of $M = N/2$ distinct XY positions hosts one *low* anchor L_i at height z_i and one *high* anchor H_i at height $z_i + \Delta h$, where $\Delta h > 0$ is a fixed, known vertical separation shared by all pairs. Let $\mathcal{M} = \{1, \ldots, M\}$ index the XY positions and let $\mathrm{pm} \in \mathbb{Z}^{M\times 2}$ be the pair-membership matrix whose row i contains the node indices $[L_i, H_i]$.

A subset $\mathcal{K}_p \subset \mathcal{M}$, $|\mathcal{K}_p| = n_k$, denotes the *known pairs*: positions for which both L_i and H_i are calibrated references with positions known to within uncertainty σ_{KA}. The remaining $M - n_k$ positions form the unknown set $\mathcal{U}_p = \mathcal{M} \setminus \mathcal{K}_p$. Intra-pair links (between L_i and H_i at the same XY position) are excluded from the range matrix $\widetilde{\mathbf{D}}$, as their short separation amplifies noise in the differential formulas derived below.

2) Closed-Form Height Estimation from Cross-Pair Measurements: Consider unknown pair $i \in \mathcal{U}_p$ and any reference pair $j \in \mathcal{K}_p$. Let r_{ij} denote the horizontal distance between the two XY positions and $\Delta z_{ij} = z_i - z_j$ the height difference between their low anchors. The four inter-pair ranges satisfy:

$$d_{LL,ij}^2 = r_{ij}^2 + \Delta z_{ij}^2, \tag{22}$$
$$d_{LH,ij}^2 = r_{ij}^2 + \big(\Delta z_{ij} + \Delta h\big)^2, \tag{23}$$
$$d_{HL,ij}^2 = r_{ij}^2 + \big(\Delta z_{ij} - \Delta h\big)^2, \tag{24}$$
$$d_{HH,ij}^2 = r_{ij}^2 + \Delta z_{ij}^2. \tag{25}$$

Subtracting (24) from (23) and expanding:

$$d_{LH,ij}^2 - d_{HL,ij}^2 = \big(\Delta z_{ij} + \Delta h\big)^2 - \big(\Delta z_{ij} - \Delta h\big)^2 = 4\,\Delta z_{ij}\,\Delta h. \tag{26}$$

Solving for the height difference gives the *closed-form estimator*:

$$\widehat{\Delta z}_{ij} = \frac{\tilde{d}_{LH,ij}^2 - \tilde{d}_{HL,ij}^2}{4\,\Delta h}, \tag{27}$$

where $\tilde{d}_{LH,ij}$ and $\tilde{d}_{HL,ij}$ are the measured ranges. Note that r_{ij} cancels exactly in (26), so $\widehat{\Delta z}_{ij}$ is independent of horizontal position error. The estimated height of unknown pair i is then:

$$\hat{z}_i = z_j + \widehat{\Delta z}_{ij}, \tag{28}$$

for each reference pair j that provides a valid estimate. Multiple reference pairs yield multiple estimates; a 20% trimmed median is used to aggregate them robustly:

$$\hat{z}_i = \mathrm{TrimMed}_{0.20}\big(\{z_j + \widehat{\Delta z}_{ij}\}_{j\in\mathcal{K}_p}\big). \tag{29}$$

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

Algorithm 3 CROSSPAIRDIST: Cross-pair distance matrices

Require: Range matrix $\widetilde{\mathbf{D}} \in \mathbb{R}^{N\times N}$, pair-membership $\mathrm{pm} \in \mathbb{Z}^{M\times 2}$, vertical separation Δh
Ensure: $\mathbf{R}_{xy} \in \mathbb{R}^{M\times M}$, $\boldsymbol{\Delta}\mathbf{Z} \in \mathbb{R}^{M\times M}$

Initialise $\mathbf{R}_{xy} \leftarrow \mathbf{NaN}$, $\boldsymbol{\Delta}\mathbf{Z} \leftarrow \mathbf{NaN}$
for each pair (i, j) with $i \neq j$ **do**
  $L_i, H_i \leftarrow \mathrm{pm}(i,:)$; $L_j, H_j \leftarrow \mathrm{pm}(j,:)$
  Read $\tilde{d}_{LH} \leftarrow \widetilde{D}_{L_i,H_j}$, $\tilde{d}_{HL} \leftarrow \widetilde{D}_{H_i,L_j}$, $\tilde{d}_{LL} \leftarrow \widetilde{D}_{L_i,L_j}$
  if both $\tilde{d}_{LH}$ and $\tilde{d}_{HL}$ are valid **then**
    $\widehat{\Delta Z}_{ij} \leftarrow (\tilde{d}^2_{LH} - \tilde{d}^2_{HL}) \,/\, (4\,\Delta h)$ ▷ (27)
    $R_{xy,ij} \leftarrow \sqrt{\max(0,\ \tilde{d}^2_{LL} - \widehat{\Delta Z}^2_{ij})}$ ▷ (31)
  else if $\tilde{d}_{LL}$ is valid **then**
    $\widehat{\Delta Z}_{ij} \leftarrow \mathrm{NaN}$; $R_{xy,ij} \leftarrow \tilde{d}_{LL}$ ▷ single-link fallback
  end if
end for
return $\mathbf{R}_{xy}$, $\boldsymbol{\Delta}\mathbf{Z}$

The noise sensitivity of (27) is governed by the ranging noise variance σ^2 and the separation Δh. Linearising around the true ranges:

$$\mathrm{Var}\Big(\widehat{\Delta z}_{ij}\Big) \approx \frac{d^2_{LH,ij} + d^2_{HL,ij}}{8\,\Delta h^2}\,\sigma^2, \tag{30}$$

which decreases quadratically with Δh, providing the principal motivation for maximising the vertical separation within the available ceiling height.

Once $\widehat{\Delta z}_{ij}$ is available, the horizontal inter-pair distance is recovered from (22):

$$\hat{r}_{ij} = \sqrt{\max\big(0,\ \tilde{d}^2_{LL,ij} - \widehat{\Delta z}^2_{ij}\big)}. \tag{31}$$

These two quantities, $\widehat{\Delta z}_{ij}$ and $\hat{r}_{ij}$, are assembled into pair-level matrices $\boldsymbol{\Delta}\mathbf{Z}$ and $\mathbf{R}_{xy}$ by the CROSSPAIRDIST subroutine (Algorithm 3).

3) Paired-Anchor MDS Autocalibration: The paired MDS algorithm (Algorithm 4) applies the original 3D MDS pipeline of Section IV-A to all N individual anchors, treating the low and high anchor of each known pair as two independent references at their respective 3D positions. Intra-pair links are set to NaN so that they do not contaminate the Floyd-Warshall embedding. The pair structure enters only in the post-processing step, where it provides two independent position estimates for each unknown XY location.

Because the same horizontal position is estimated by both $\hat{\mathbf{p}}_{L_i}$ and $\hat{\mathbf{p}}_{H_i}$, their XY components are averaged:

$$\hat{\mathbf{x}}_i = \tfrac{1}{2}\big(\hat{\mathbf{p}}^{(xy)}_{L_i} + \hat{\mathbf{p}}^{(xy)}_{H_i}\big), \tag{32}$$

which reduces the XY variance by a factor of two relative to either anchor alone, provided the two MDS estimates are approximately independent.

The height of pair i is estimated by applying (27)–(29) with reference pairs $j \in \mathcal{K}_p$, using the range measurements directly from $\widetilde{\mathbf{D}}$. This replaces the IDW estimate (18) with the closed-form differential formula, which is immune to the horizontal position error that would otherwise propagate into the IDW weights.

Algorithm 4 Paired-Anchor MDS Autocalibration

Require: $\widetilde{\mathbf{D}}$, pm, $\mathcal{K}_p$, known positions $\{\mathbf{p}_j\}_{j\in\mathcal{K}_p}$, Δh, mode, parameters Θ
Ensure: $\{\hat{\mathbf{x}}_i, \hat{z}_i\}_{i\in\mathcal{U}_p}$

Build individual-anchor reference set: $\mathcal{K}_{\mathrm{ind}} \leftarrow \{L_j, H_j : j \in \mathcal{K}_p\}$
Set $\widetilde{D}_{L_i,H_i} \leftarrow \mathrm{NaN}$ and $\widetilde{D}_{H_i,L_i} \leftarrow \mathrm{NaN}$ for all $i \in \mathcal{M}$ ▷ exclude intra-pair links
Run the MDS pipeline (Algorithm 1) on $\widetilde{\mathbf{D}}$ with references $\mathcal{K}_{\mathrm{ind}} \to \hat{\mathbf{P}} \in \mathbb{R}^{N\times 3}$
for each unknown pair $i \in \mathcal{U}_p$ **do**
  $\hat{\mathbf{x}}_i \leftarrow \frac{1}{2}\big(\hat{\mathbf{P}}^{(xy)}_{L_i} + \hat{\mathbf{P}}^{(xy)}_{H_i}\big)$ ▷ (32)
  $\hat{\mathbf{x}}_i \leftarrow \mathrm{clip}(\hat{\mathbf{x}}_i, [0, \ell_x] \times [0, \ell_y])$
  Compute $\{\boldsymbol{\Delta}\mathbf{Z}, \mathbf{R}_{xy}\} \leftarrow$ CROSSPAIRDIST$(\widetilde{\mathbf{D}}, \mathrm{pm}, \Delta h)$
  $\hat{z}_i \leftarrow \mathrm{TrimMed}_{0.20}\big(\{z_j + \Delta Z_{ij}\}_{j\in\mathcal{K}_p,\ \Delta Z_{ij}\neq\mathrm{NaN}}\big)$ ▷ (29)
  if no valid ΔZ_{ij} exists **then**
    $\hat{z}_i \leftarrow \frac{1}{2}\big(\hat{P}_{L_i,z} + \hat{P}_{H_i,z} - \Delta h\big)$ ▷ MDS fallback
  end if
  $\hat{z}_i \leftarrow \mathrm{clip}(\hat{z}_i, [0,\ \ell_z - \Delta h])$
end for
return $\{\hat{\mathbf{x}}_i, \hat{z}_i\}$

4) Paired-Anchor Multilateration Autocalibration: The paired multilateration algorithm (Algorithm 5) follows an identical post-processing strategy but replaces the MDS backbone with the greedy cascade of Section IV-B. The original 3D multilateration runs on all N individual anchors; the pair structure is again used only in post-processing.

Multilateration is particularly suited to the paired setting because the cascade can exploit the doubled reference count: both anchors of each known pair enter $\mathcal{K}_{\mathrm{ind}}$, providing more well-placed references and improving the conditioning of the pivot-LS system, especially for unknown anchors that have limited links to known nodes.

Table I lists all configurable parameters. The paired approach introduces two new parameters: the vertical separation Δh and the number of known pairs n_k.

V. RESULTS

A. Environment and UWB Simulation Model

All evaluations in this section are conducted using the custom physics-based MATLAB simulator introduced in Section III. This subsection describes the building geometry, UWB radio parameters and measurement generation procedure used throughout the Monte Carlo studies.

1) Building Geometry: The simulated environment represents a multi-room indoor floor plan with overall dimensions $22.3\,\mathrm{m} \times 26.3\,\mathrm{m} \times 3.0\,\mathrm{m}$, modelled as a set of 21 vertical planar facets. Seven facets represent load-bearing external walls (reinforced concrete) and fourteen represent internal partitions (plasterboard). The layout includes a cluster of small rooms and a series of larger open areas, separated by corridors and

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

Algorithm 5 Paired-Anchor Multilateration Autocalibration

Require: $\widetilde{\mathbf{D}}$, pm, $\mathcal{K}_p$, known positions $\{\mathbf{p}_j\}_{j\in\mathcal{K}_p}$, Δh, mode, parameters Θ
Ensure: $\{\hat{\mathbf{x}}_i, \hat{z}_i\}_{i\in\mathcal{U}_p}$
Build $\mathcal{K}_{\text{ind}} \leftarrow \{L_j, H_j : j \in \mathcal{K}_p\}$
Set $\widetilde{D}_{L_i,H_i} \leftarrow \text{NaN}$ for all $i \in \mathcal{M}$ ▷ exclude intra-pair links
Run multilateration cascade (Alg. 2) on $\widetilde{\mathbf{D}}$ with references $\mathcal{K}_{\text{ind}} \to \hat{\mathbf{P}} \in \mathbb{R}^{N\times 3}$
for each unknown pair $i \in \mathcal{U}_p$ **do**
  if both $\hat{\mathbf{P}}_{L_i}$ and $\hat{\mathbf{P}}_{H_i}$ valid **then**
    $\hat{\mathbf{x}}_i \leftarrow \frac{1}{2}\big(\hat{\mathbf{P}}^{(xy)}_{L_i} + \hat{\mathbf{P}}^{(xy)}_{H_i}\big)$ ▷ (32)
  else if only one of $\hat{\mathbf{P}}_{L_i}$, $\hat{\mathbf{P}}_{H_i}$ valid **then**
    $\hat{\mathbf{x}}_i \leftarrow$ XY of the valid estimate
  else
    $\hat{\mathbf{x}}_i \leftarrow [\ell_x/2,\, \ell_y/2]$ ▷ room-centre fallback
  end if
  $\hat{\mathbf{x}}_i \leftarrow \text{clip}(\hat{\mathbf{x}}_i,\, [0,\ell_x]\times[0,\ell_y])$
  Compute $\{\mathbf{\Delta Z}, \mathbf{R}_{xy}\} \leftarrow$ CrossPairDist$(\widetilde{\mathbf{D}}, \text{pm}, \Delta h)$
  $\hat{z}_i \leftarrow \text{TrimMed}_{0.20}\big(\{z_j + \Delta Z_{ij}\}_{j\in\mathcal{K}_p,\, \Delta Z_{ij}\neq\text{NaN}}\big)$ ▷ (29)
  if no valid ΔZ_{ij} exists **then**
    $\hat{z}_i \leftarrow$ IDW average of $\hat{P}_{L_i,z}$ and $\hat{P}_{H_i,z} - \Delta h$ ▷ multilat fallback
  end if
  $\hat{z}_i \leftarrow \text{clip}(\hat{z}_i,\, [0,\, \ell_z - \Delta h])$
end for
return $\{\hat{\mathbf{x}}_i, \hat{z}_i\}$

TABLE I
ALGORITHM PARAMETERS AND RECOMMENDED VALUES

Symbol	Description	Value	Eq.
α	FW shrinkage factor	0.90	(8)
P	Gate refinement passes	3	(20)
$M_{\min}$	Min references (multilat)	3	(21)
M_{target}	Reference saturation count	3	—
Δh	Vertical pair separation	0.5–2.0 m	(27)
n_k	Known pairs	3–7	(28)
σ_{res}	Residual scale (multilat)	0.55 m	—
σ_{res}	Residual scale (MDS)	1.20 m	—
γ	Geometry condition scale	10	—
P_{weak}	RSSI lower bound	−100 dBm	—
P_{strong}	RSSI upper bound	−60 dBm	—
δ_{room}	Room bound margin	1.0 m	—

TABLE II
UWB SIMULATION PARAMETERS

Parameter	Symbol	Value
Carrier frequency	f_c	6.5 GHz
Transmit power	P_{tx}	−10 dBm
Path loss at $d_0 = 1$ m	PL_0	43.2 dB
Path loss exponent	n	1.8
Rx sensitivity	P_{sens}	−93 dBm
Internal wall loss	L_{int}	5 dB
External wall loss	L_{ext}	12 dB
LOS range noise std	σ_{LOS}	0.20 m
NLOS range noise std	σ_{NLOS}	0.30 m
Per-wall noise increment	σ_{wall}	0.01 m
NLOS bias (internal wall)	b_{int}	0.01 m
NLOS bias (external wall)	b_{ext}	0.02 m
Known-anchor noise std	σ_{KA}	0.10 m
Height measurement std	σ_Z	0.10 m

partial walls with door openings. This heterogeneous topology creates a mix of LOS and NLOS conditions between anchor pairs, with wall intersection counts ranging from zero (clear LOS) to 3 or more (multi-wall NLOS), challenging both autocalibration algorithms. This floor plan is representative of a medium-sized commercial office building, selected to provide a realistic range of LOS and NLOS propagation conditions without resorting to adversarial geometry. A representative deployment within this environment is illustrated in Fig. 2.

2) UWB Radio and Ranging Parameters: The UWB link budget follows (2) with the parameters listed in Table II. The log-distance path loss model in (2) is parameterised using the wall penetration losses listed in Table II.

Range measurements are generated according to (3): each connected link receives a measured range equal to the true Euclidean distance plus a positive NLOS bias (proportional to the number of wall crossings) and an additive zero-mean Gaussian noise term whose standard deviation increases by σ_{wall} per wall crossing beyond LOS. A link is included in the measurement graph if and only if the received power computed by (2) exceeds P_{sens}; links below threshold are recorded as missing entries (NaN) and are unavailable to all algorithms.

Fig. 2 illustrates a representative anchor deployment with $N = 15$ nodes positioned throughout the building at heights between 0.2 m and 2.4 m. Connected links are drawn as coloured dotted lines, colour-coded by the received power at the receiver. Unconnected links below the sensitivity threshold are shown as black solid lines. It is thereby shown that links blocked by multiple external or internal walls fall below −93 dBm and are absent from the ranging graph.

3) Monte Carlo Framework: For the standalone evaluations in Sections V-B and V-C, three parameter sweeps are defined as detailed at the opening of Section V. In each trial the simulator is re-invoked with fresh random anchor positions and a new random noise realisation, but all radio parameters remain fixed at the values in Table II. Identical random seeds are used across both standalone algorithms so that performance differences reflect solely the algorithm rather than the deployment geometry. The resulting measurement matrices ($\widetilde{\mathbf{D}}$, **RSSI**, $\mathbf{W}$) are passed directly to each algorithm without any preprocessing or outlier removal, reflecting realistic deployment conditions where no a priori knowledge of the channel is available.

The standalone and paired-anchor autocalibration algorithms are evaluated via three independent Monte Carlo sweeps, each comprising 100 trials with identical random seeds across methods, enabling direct numerical comparison. In every trial, N anchors are placed at uniformly random positions (minimum separation $d_{\min} = 3.5$ m, heights $\in$

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

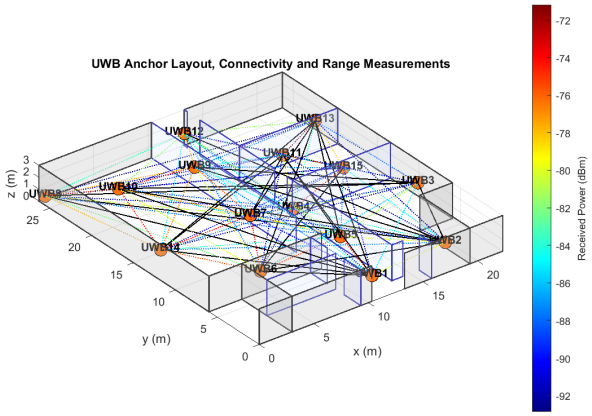


Fig. 2. Illustrative UWB anchor deployment ($N = 15$).

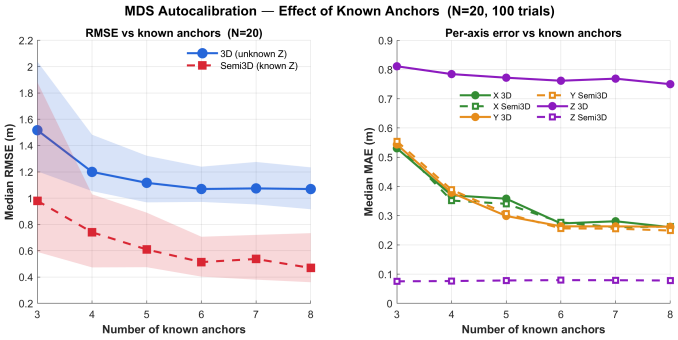


Fig. 3. MDS: RMSE (left) and per-axis MAE (right) vs n_{known}, $N = 20$, $\sigma_{\text{LOS}} = 0.20$ m. Shaded bands: IQR over 100 trials.

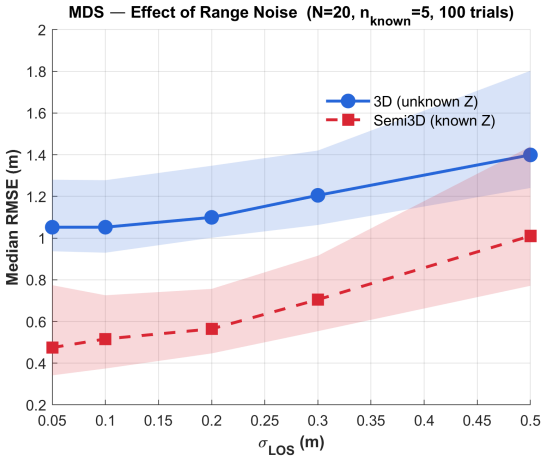


Fig. 4. MDS: RMSE vs σ_{LOS}, $N = 20$, $n_{\text{known}} = 5$. Shaded bands: IQR over 100 trials.

$[0.2, 2.8]$ m) and n_{known} of them are selected at random as reference nodes; a full UWB measurement set is synthesised by the simulator described in Section III. Performance is reported as the median RMSE and per-axis MAE over unknown anchors. Both the *3D mode* (heights estimated by IDW of known neighbours) and the *Semi-3D mode* (heights known to within $\sigma_Z = 0.10$ m) are evaluated throughout. Three sweeps are defined:

- **Sweep A**: $n_{\text{known}} \in \{3 - 8\}$, $N = 20$, $\sigma_{\text{LOS}} = 0.20$ m.
- **Sweep B**: $\sigma_{\text{LOS}} \in \{0.05 - 0.50\}$ m, $N = 20$, $n_{\text{known}} = 5$.
- **Sweep C**: $N \in \{8 - 20\}$, $n_{\text{known}} = 4$, $\sigma_{\text{LOS}} = 0.20$ m.

B. MDS Autocalibration

1) Effect of Known-Anchor Count (Sweep A): Fig. 3 shows RMSE and per-axis MAE. In 3D mode, RMSE falls from 1.52 m at $n_{\text{known}} = 3$ to 1.07 m at $n_{\text{known}} = 6$ (30%), then plateaus — additional reference anchors provide negligible benefit once the Procrustes alignment is sufficiently constrained. Semi-3D mode improves more substantially, from 0.98 m to 0.47 m (52%) and continues declining through $n_{\text{known}} = 8$. The persistent gap between modes is explained by the per-axis breakdown: Z MAE in 3D mode is constant at 0.75–0.81 m regardless of n_{known}, confirming that IDW height estimation is the primary accuracy bottleneck. Semi-3D Z MAE is ≈ 0.08 m (the measurement noise floor σ_Z). Horizontal errors (X,Y) decrease monotonically in both modes. A practical operating point is $n_{\text{known}} = 5$: it captures most of the attainable accuracy (1.12 m / 0.61 m in 3D / Semi-3D) with diminishing returns beyond it.

2) Sensitivity to Range Noise (Sweep B): Both modes exhibit monotonically increasing RMSE with noise (Fig. 4), but at different rates. The 3D mode rises from 1.05 m to 1.40 m (+33%): a modest increase because the height estimation floor (≈ 0.75 m Z MAE) already dominates and masks the effect of ranging noise. Semi-3D starts considerably lower (0.48 m) but climbs more steeply to 1.01 m (+113%), as it is more directly exposed to horizontal ranging errors once height uncertainty is removed. IQR bands widen substantially at $\sigma_{\text{LOS}} \geq 0.30$ m for Semi-3D, while 3D maintains a relatively stable spread.

3) Scalability (Sweep C): Counter-intuitively, RMSE *improves* as N grows (Fig. 5). The 3D RMSE decreases from 1.68 m at $N = 8$ to 1.23 m at $N = 20$ and Semi-3D from 0.98 m to 0.74 m. Larger networks produce denser measurement graphs, reducing Floyd-Warshall path inflation globally. However, variance is strongly N-dependent: IQR half-widths of 1.73 m (3D) and 1.94 m (Semi-3D) at $N = 8$ collapse to 0.24 m at $N = 20$, indicating that MDS is fragile under poorly conditioned small graphs. The anomalous peak at $N = 12$ (IQR ± 1.03 m) reflects ill-conditioned topologies where FW gap-fill introduces significant distortion. Networks of $N \leq 12$ should be treated with caution.

C. Multilateration Autocalibration

1) Effect of Known-Anchor Count (Sweep A): The cascade achieves 100% recovery in every trial across all n_{known} values: the dense connectivity of a 20-anchor network always provides each anchor with $\geq M_{\min} = 3$ reachable references. RMSE decreases from 1.17 m to 0.93 m in 3D and from 0.67 m to 0.39 m in Semi-3D, as shown in Fig. 6. Multilateration outperforms MDS at every operating point — at $n_{\text{known}} = 3$ by 23% in 3D and 32% in Semi-3D — because it uses *measured*

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

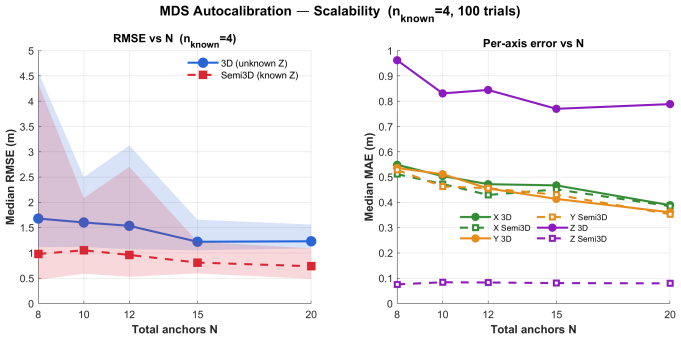


Fig. 5. MDS: RMSE (left, IQR bands) and per-axis MAE (right) vs N, $n_{\text{known}} = 4$, $\sigma_{\text{LOS}} = 0.20$ m.

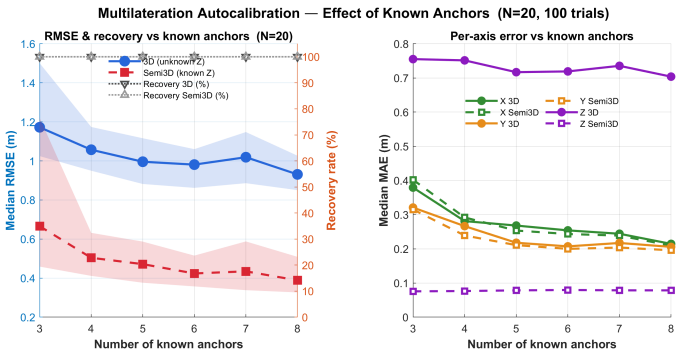


Fig. 6. Multilateration: RMSE and recovery (left, dual y-axis) and per-axis MAE (right) vs n_{known}, $N = 20$, $\sigma_{\text{LOS}} = 0.20$ m.

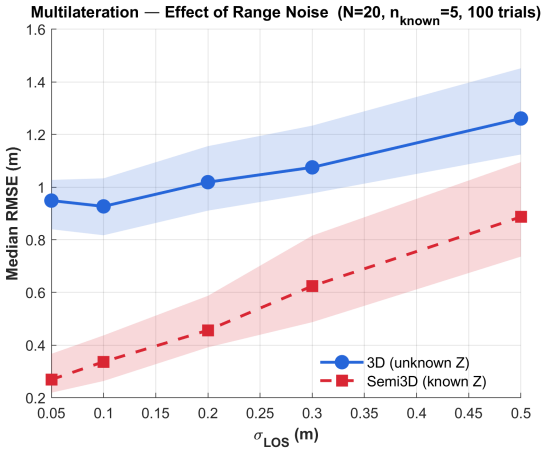


Fig. 7. Multilateration: RMSE vs σ_{LOS}, $N = 20$, $n_{\text{known}} = 5$. Shaded bands: IQR over 100 trials.

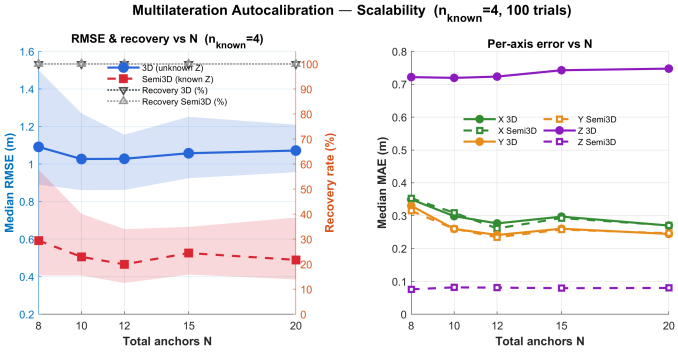


Fig. 8. Multilateration: RMSE and recovery (left, dual y-axis) and per-axis MAE (right) vs N, $n_{\text{known}} = 4$, $\sigma_{\text{LOS}} = 0.20$ m.

distances rather than multi-hop Floyd-Warshall paths. The per-axis pattern is qualitatively identical to MDS: Z MAE constant at 0.70–0.76 m in 3D (IDW floor), $\approx$ 0.08 m in Semi-3D; horizontal errors decrease monotonically with n_{known}.

2) Sensitivity to Range Noise (Sweep B): Recovery remains 100% at all noise levels. The 3D RMSE rises from 0.95 m to 1.26 m and Semi-3D from 0.27 m to 0.89 m, as shown in Fig. 7. At $\sigma_{\text{LOS}} = 0.05$ m, Semi-3D multilateration (0.27 m) is 43% better than Semi-3D MDS (0.475 m), reflecting the complete elimination of FW distortion. This advantage narrows to 12% at $\sigma_{\text{LOS}} = 0.50$ m as both methods become noise-dominated. The 3D mode's flat response is again explained by the IDW height floor masking horizontal ranging noise.

3) Scalability (Sweep C): Unlike MDS, multilateration RMSE is nearly flat with N: 3D varies between 0.93 m and 1.09 m with no consistent trend; Semi-3D improves modestly from 0.59 m to 0.49 m (Fig. 8). This is a structural property: each anchor is estimated locally, so additional network anchors do not improve the local geometric sub-problem.

D. Paired-Anchor Autocalibration Evaluation

The paired-anchor autocalibration algorithms described in Section IV-C are evaluated via the same three Monte Carlo sweeps used for the standalone methods, over 100 independent trials. In each trial, N anchors are placed in $M = N/2$ vertically separated pairs at randomly chosen XY positions. Then, n_k pairs are selected as calibrated references and a full UWB measurement set is generated by the physics-based simulator. Four vertical separations $\Delta h \in \{0.5, 1.0, 1.5, 2.0\}$ m are evaluated in parallel for each trial configuration. Performance is reported as the median 3D RMSE over the unknown pairs. The standalone 3D and Semi-3D results from Sections V-B–V-C are overlaid as reference baselines.

1) Effect of Known-Pair Count (Sweep A): Sweep A fixes $N = 20$ and $\sigma_{\text{LOS}} = 0.20$ m while varying $n_k \in \{3, 4, 5, 6, 7\}$. Results for MDS-Paired and ML-Paired are shown in Figs. 9a and 10a respectively.

MDS-Paired. The vertical separation Δh is the dominant performance parameter. At $\Delta h = 0.5$ m the paired approach is uniformly worse than the standalone 3D baseline across all values of n_k, with RMSE ranging from 1.84 m at $n_k = 3$ to 1.23 m at $n_k = 7$. This confirms the theoretical prediction of (30): with small Δh the differential estimator amplifies ranging noise and the Z improvement is insufficient to offset the loss of XY reference diversity due to the halved number of anchor positions. Increasing the separation to $\Delta h = 1.0$ m produces a crossover at $n_k = 5$, above which the paired algorithm outperforms the standalone 3D mode (1.12 m vs. 1.17 m at $n_k = 5$, improving to 0.89 m vs. 1.05 m at $n_k = 7$). At $\Delta h = 1.5$ m the crossover moves to $n_k = 4$ and the gap widens substantially at higher known-pair counts, reaching 0.63 m at $n_k = 7$ — a 40% reduction over the 3D baseline.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

The largest separation $\Delta h = 2.0\,\text{m}$ is more variable at low n_k ($1.57\,\text{m}$ at $n_k = 3$ vs. $1.42\,\text{m}$ for Orig-3D), but converges to the lowest RMSE of any configuration at $n_k \geq 5$, achieving $0.54\,\text{m}$ at $n_k = 7$ — within 8% of the Semi-3D reference ($0.50\,\text{m}$) and 48% below the standalone 3D result.

ML-Paired. The ML-Paired algorithm follows the same qualitative trends with Δh playing the same dominant role. The standalone ML-3D baseline is lower than its MDS counterpart throughout (1.19–$0.96\,\text{m}$ vs. 1.42–$1.05\,\text{m}$), reflecting the cascade's superior XY accuracy. At $\Delta h = 2.0\,\text{m}$ and small n_k the paired approach is again unstable ($1.50\,\text{m}$ at $n_k = 3$ vs. $1.19\,\text{m}$ Orig-ML-3D), indicating that the cascade requires at least four to five reference pairs to initialise reliably. From $n_k = 4$ onwards, $\Delta h = 1.5\,\text{m}$ and $2.0\,\text{m}$ both beat the standalone baseline, with $\Delta h = 2.0\,\text{m}$ achieving $0.53\,\text{m}$ at $n_k = 7$ — 45% below the $0.96\,\text{m}$ baseline and approaching the Semi-3D floor of $0.41\,\text{m}$.

2) Noise Robustness (Sweep B): Sweep B fixes $N = 20$ and $n_k = 5$ while varying $\sigma_{\text{LOS}} \in \{0.05, 0.10, 0.20, 0.30, 0.50\}\,\text{m}$. Results are shown in Figs. 9b and 10b.

MDS-Paired. The $\Delta h = 2.0\,\text{m}$ curve is the most noise-robust configuration, ranging from $0.63\,\text{m}$ at $\sigma = 0.05\,\text{m}$ to $0.91\,\text{m}$ at $\sigma = 0.50\,\text{m}$ — a monotonically increasing but consistently low-error response that remains well below the standalone 3D baseline (1.09–$1.38\,\text{m}$ over the same range). The relative improvement over the standalone 3D mode is 34%–42% and grows modestly with noise level, indicating that the height-decoupling property of the differential estimator (27) is effective and increasingly advantageous as ranging noise increases. At $\Delta h = 0.5\,\text{m}$ the paired algorithm is again worse than standalone 3D.

ML-Paired. A very similar pattern emerges for ML-Paired. At $\Delta h = 2.0\,\text{m}$ the RMSE spans 0.48–$0.93\,\text{m}$ compared to 0.93–$1.25\,\text{m}$ for the standalone ML-3D baseline, an improvement of 26%–48%. The $\Delta h = 1.5\,\text{m}$ curve achieves RMSE below $1.10\,\text{m}$ at all noise levels, consistently outperforming the standalone 3D mode. The $\Delta h = 0.5\,\text{m}$ curve reaches $1.39\,\text{m}$ at $\sigma_{\text{LOS}} = 0.20\,\text{m}$ and rises to $1.65\,\text{m}$ at $\sigma_{\text{LOS}} = 0.50\,\text{m}$, confirming that the small-separation regime is detrimental at all noise levels.

3) Network Size Sensitivity (Sweep C): Sweep C fixes $n_k = \min(5, M-1)$ and $\sigma_{\text{LOS}} = 0.20\,\text{m}$ while varying $N \in \{10, 12, 16, 20\}$. Results are shown in Figs. 9c and 10c.

MDS-Paired. At all network sizes, $\Delta h = 2.0\,\text{m}$ yields the lowest RMSE. The largest relative gain occurs at $N = 12$, where $\Delta h = 2.0\,\text{m}$ achieves $0.63\,\text{m}$ against the standalone 3D result of $1.30\,\text{m}$ — a 51% reduction. The trend is non-monotonic at larger N: RMSE increases slightly from $N = 12$ to $N = 16$ before stabilising at $N = 20$, partly because n_k saturates at 5 while the total number of unknown pairs grows, diluting the relative reference density. The paired method nonetheless closes the gap between standalone 3D and Semi-3D substantially for all $N \geq 12$.

ML-Paired. The ML-Paired results mirror the MDS-Paired trends. At $N = 12$, $\Delta h = 2.0\,\text{m}$ achieves $0.64\,\text{m}$ vs. $1.00\,\text{m}$ for standalone ML-3D (36% improvement). At $N = 20$ the gap is $0.60\,\text{m}$ vs. $1.00\,\text{m}$ (40% improvement). Unlike MDS-Paired, the ML-Paired RMSE at $\Delta h = 2.0\,\text{m}$ decreases monotonically with N, reflecting the cascade's natural benefit from a denser connectivity graph as more anchors are added. Both methods with $\Delta h = 2.0\,\text{m}$ remain below $0.65\,\text{m}$ for $N \geq 12$, substantially outperforming their respective 3D baselines.

4) NLOS Severity Sensitivity (Sweep D): In practice, UWB anchor-to-anchor ranges are affected not only by zero-mean Gaussian noise but also by a systematic positive bias caused by multipath propagation and NLOS wall crossings. The parameter b_{int}, which appears in the physics-based channel model described in Section III (equations 4 and 5), controls the per-internal-wall range bias added to each obstructed link; the per-external-wall counterpart is set to $b_{\text{ext}} = 2\,b_{\text{int}}$ throughout, consistent with the higher attenuation of external masonry walls. Because b_{int} is the single parameter that governs the severity of systematic NLOS distortion in the simulated building, sweeping it provides a direct measure of how the paired-anchor framework degrades as NLOS conditions worsen.

The closed-form cross-pair estimators in equations 4 and 5 exploit the differential geometry between the low and high anchors of each pair. When both cross-links between two pairs traverse the same obstructing walls, the NLOS bias terms cancel in the differencing operation, leaving the estimator unaffected. Residual sensitivity arises only when the two cross-links experience *asymmetric* wall blockage, i.e. when only one of them is NLOS-affected, so that the cancellation is incomplete. Sweep D is designed to make this mechanism visible experimentally. The sweep fixes $N = 20$ anchors, $n_k = 5$ known pairs, and $\Delta h = 2.0\,\text{m}$, and varies b_{int} over $\{0.01, 0.05, 0.10, 0.30, 0.50, 1.00\}\,\text{m}$. The LOS range noise is held at $\sigma_{\text{LOS}} = 0.20\,\text{m}$. The lower end of the sweep ($b_{\text{int}} \leq 0.10\,\text{m}$) corresponds to light NLOS typical of modern IEEE 802.15.4z UWB transceivers in furnished office environments; the upper end ($b_{\text{int}} = 1.00\,\text{m}$) represents heavily obstructed NLOS well beyond what compliant devices are expected to encounter.

Results are presented in Fig. 11 for the MDS-Paired and ML-Paired variants. At practical bias levels ($b_{\text{int}} \leq 0.10\,\text{m}$), both paired variants retain their full accuracy advantage over the Original 3D baseline, achieving a reduction of approximately 35–40% in median 3D RMSE. This confirms that the differential estimator successfully suppresses the dominant common-mode bias component under realistic NLOS conditions. As bias grows, the advantage narrows due to increasing asymmetric NLOS, with the three methods converging at extreme values ($b_{\text{int}} \geq 0.50\,\text{m}$); nonetheless, even at $b_{\text{int}} = 1.00\,\text{m}$ the MDS-Paired variant ($2.00\,\text{m}$ RMSE) remains below both Semi-3D ($2.21\,\text{m}$) and Original 3D ($2.37\,\text{m}$). These results confirm that the paired-anchor architecture provides meaningful NLOS resilience across the full range of bias levels likely to be encountered in practice, without requiring explicit NLOS identification or mitigation.

5) Computational Overhead: Fig. 12 compares the median per-trial execution times of the standalone 3D, standalone Semi-3D and paired algorithms for both modes. The paired curve reports the mean over all four Δh values, as execution time is insensitive to Δh.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

TABLE III
3D RMSE (M) AT $N = 20$, $n_k = 5$, $\sigma_{\text{LOS}} = 0.20$ M; LOS CONDITIONS ($b_{\text{int}} = 0$). NLOS SENSITIVITY IS CHARACTERISED SEPARATELY IN SWEEP D (FIG. 11).

Method	Orig. 3D	Orig. Semi-3D	$\Delta h = 1.5$ m	$\Delta h = 2.0$ m
MDS-Paired	1.167	0.602	0.839	0.710
ML-Paired	0.991	0.460	0.750	0.588

For MDS-Paired (Fig. 12a), execution times range from 1.04 ms to 1.75 ms and closely track the standalone 3D curve, confirming that the post-processing step (CROSSPAIRDIST, pair averaging, trimmed-median Z) contributes negligible overhead relative to the MDS backbone. The slight decrease with n_k reflects the reduced number of unknown anchors that the pivot-LS refinement must process. For ML-Paired (Fig. 12b), execution times are substantially higher at $n_k = 3$ (≈16–18 ms) due to cascade instability under sparse references, but converge to 1–2 ms at $n_k \geq 6$ — on a par with the standalone ML-3D. The steep decline for all three ML curves reflects the cascade's ability to resolve unknown anchors in fewer iterations when more references are available. In both cases the paired approach introduces no material computational penalty over the corresponding standalone algorithm.

6) Summary: Table III summarises the best paired configurations at the common operating point $N = 20$, $n_k = 5$, $\sigma_{\text{LOS}} = 0.20$ m alongside the corresponding standalone and Semi-3D baselines.

The paired approach with $\Delta h = 2.0$ m achieves the best results for both backbone algorithms, reducing the 3D RMSE by 39% (MDS) and 41% (ML) relative to the standalone 3D baseline at the $n_k = 5$ operating point and by approximately 49% (MDS) and 45% (ML) at $n_k = 7$. The improvement increases monotonically with Δh, confirming that the noise variance of the differential height estimator (30) — which decreases as $1/\Delta h^2$ — is the principal driver of performance. The minimum effective separation is approximately $\Delta h = 1.0$ m for MDS-Paired and $\Delta h = 1.5$ m for ML-Paired with $n_k \geq 5$ pairs. With a practical ceiling height of 3.0 m and a low-anchor height range of $[0.2, 1.0]$ m, a separation of $\Delta h = 2.0$ m is readily achievable and is recommended as the deployment target. Sweep D further demonstrates that the paired-anchor framework is resilient to realistic NLOS conditions: at practical per-internal-wall bias levels ($b_{\text{int}} \leq 0.10$ m), the 35–40% RMSE advantage over the standalone 3D baseline is fully preserved, owing to the common-mode bias cancellation inherent in the cross-pair differential estimators (27). Even under extreme multipath distortion ($b_{\text{int}} = 1.00$ m), MDS-Paired achieves 2.00 m RMSE, remaining below both the Semi-3D (2.21 m) and standalone 3D (2.37 m) baselines, without requiring explicit NLOS identification or mitigation.

7) Single-Trial 3D Visualisation: To complement the statistical Monte Carlo results, Fig. 13 presents a single representative trial of the paired-anchor deployment for both MDS-Paired (left) and ML-Paired (right). The trial uses $N = 40$ individual anchors ($M = 20$ pairs) with $\Delta h = 1.5$ m, $\sigma_{\text{KA}} = 0.10$ m, $\sigma_{\text{LOS}} = 0.20$ m, $\sigma_{\text{NLOS}} = 0.30$ m and $n_k = 5$ randomly selected known pairs (green diamonds). Unknown pairs are shown as black circles at their true low-anchor positions, with the vertical bar to the corresponding high anchor rendered in grey. Error vectors (dashed grey lines) connect each true position to its estimate and each label reports the pair index and displacement in metres.

Both algorithms recover the spatial layout of the network with good visual fidelity. The majority of estimates cluster tightly around their true positions, with error vectors short enough that estimates and true positions often overlap, confirming the sub-metre accuracy reported in Sections V-D. ML-Paired achieves a lower per-trial RMSE (0.588 m vs. 0.650 m), consistent with the statistical advantage of the cascade's greedy nearest-neighbour resolution over the global MDS embedding for this network size and noise level.

E. Discussion

1) Summary of Key Findings: This work has developed and evaluated two lightweight anchor autocalibration algorithms — MDS-based and multilateration-based — in both full 3D and Semi-3D operating modes and has proposed a structured paired-anchor extension that substantially reduces height estimation error through a closed-form differential formula. The results across four parameter sweeps yield several clear conclusions.

First, *height estimation is the dominant accuracy bottleneck in 3D mode.* In every configuration evaluated, the Z MAE remains between 0.70 m and 0.82 m regardless of the number of known anchors, ranging noise level, or network size. The horizontal components (X, Y) respond readily to more known anchors and better connectivity, but the 3D RMSE is dominated throughout by the vertical term.

Second, *multilateration consistently outperforms MDS in XY accuracy*. At $n_{\text{known}} = 5$, $N = 20$, $\sigma_{\text{LOS}} = 0.20$ m, multilateration achieves 0.99 m 3D RMSE compared to 1.17 m for MDS — a 15% advantage that persists across all sweep configurations. This arises because multilateration uses direct measured distances whereas MDS relies on Floyd-Warshall multi-hop paths whose accumulated overestimation, despite shrinkage correction, introduces global distortion. However, MDS has a structural robustness advantage: it provides a simultaneous global embedding that is less sensitive to cascade ordering effects and it is substantially more stable at small N where multilateration recovery rates degrade.

Third, *the paired-anchor deployment effectively decouples height estimation from ranging noise.* At $\Delta h = 2.0$ m and $n_k = 5$ known pairs, the paired approach reduces 3D RMSE by 39% (MDS) and 41% (multilateration) relative to the standalone 3D baselines. At $n_k = 7$, the reduction reaches approximately 49% (MDS) and 45% (multilateration), with the best results approaching within 8%–10% of the Semi-3D ceiling. The $1/\Delta h^2$ dependence of the differential estimator variance (30) explains why larger vertical separations are so effective: doubling Δh from 1.0 to 2.0 m reduces the height estimation variance by a factor of four, yielding improvements that outweigh the XY reference density reduction caused by halving the number of distinct anchor positions.

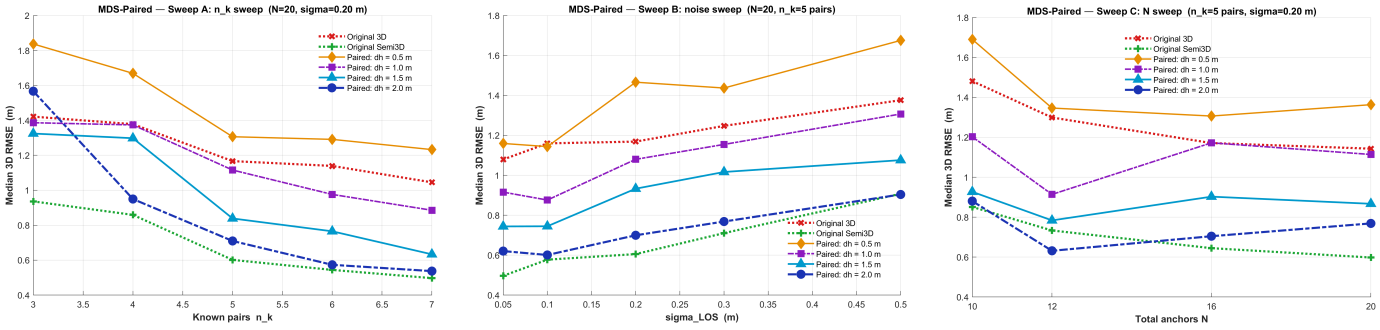
This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

Fourth, *the paired-anchor differential estimator provides inherent resilience to realistic NLOS conditions.* Sweep D varies the per-internal-wall ranging bias b_{int} over $[0.01, 1.00]$ m and demonstrates that at bias levels typical of IEEE 802.15.4z UWB transceivers in furnished office environments ($b_{\text{int}} \leq 0.10$ m), the 35–40% RMSE advantage of the paired approach over the standalone 3D baseline is fully preserved. This robustness arises naturally from the common-mode cancellation property of the cross-pair differencing in (27): when both cross-links between two pairs traverse the same obstructing walls, the NLOS bias terms cancel exactly, leaving the estimator unaffected. Residual sensitivity is limited to cases of asymmetric wall blockage and, even at the extreme value $b_{\text{int}} = 1.00$ m, the MDS-Paired variant (2.00 m RMSE) remains below both the Semi-3D (2.21 m) and standalone 3D (2.37 m) baselines without requiring any explicit NLOS identification or mitigation.

2) Comparison with the Literature: Placing these results in the context of closely related published work requires careful accounting for differences in dimensionality, environment and measurement methodology.

a) Comparison with Van Herbruggen et al. [6]: The most directly comparable multihop UWB self-calibration system is that of Van Herbruggen et al. [6], who report a mean absolute error (MAE) of 21.6 cm in a real-world industrial environment ($25 \times 11\,\text{m}^2$, 18 nodes, metallic NLOS obstacles) using a three-stage algorithm combining LLS initialisation, L-BFGS-B global optimisation and an optional Kalman-filter refinement with a mobile node. This figure is substantially lower than the XY errors reported in the present work. However, several fundamental differences temper direct comparison. The most consequential is *dimensionality*: Van Herbruggen et al. solve a 2D problem, with all anchors mounted at a fixed height of 1.60 m on tripods. In the present work, anchor heights span $[0.2, 2.8]$ m and must be estimated jointly with horizontal coordinates, creating an intrinsically harder problem. A second consideration is *the optional mobile node in Stage 3* of [6], which collects additional measurements from a handheld tag walked through the deployment area. Without this stage, the algorithm in [6] achieves 28.2 cm MAE (Stage 2 only), which is a fairer comparator to the present work's anchor-only approach. The proposed paired algorithm, which adds no mobile node and no additional sensing, achieves 2D XY RMSE values below 0.50 m at $n_k \geq 5$ in the horizontal plane, suggesting comparable XY accuracy relative to Stage 2 of [6] when evaluated on a 2D basis — though a definitive quantitative comparison would require evaluation on real hardware in a shared environment. A further structural distinction is *computational complexity*. Stage 2 of [6] uses L-BFGS-B global optimisation, an iterative nonlinear solver whose convergence depends on the cost function landscape and can require many gradient evaluations in large networks. The proposed MDS

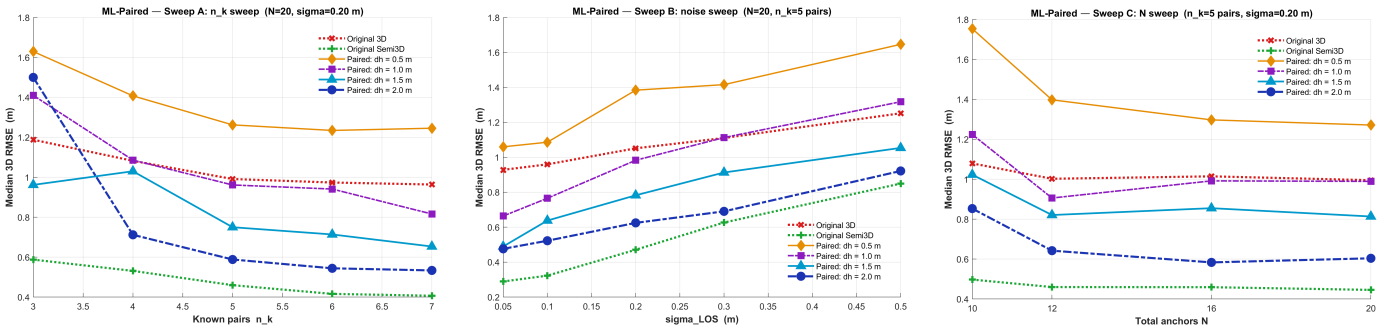


(a) Sweep A: n_k sweep ($N = 20$, $\sigma_{\text{LOS}} = 0.20$ m)

(b) Sweep B: noise sweep ($N = 20$, $n_k = 5$ pairs)

(c) Sweep C: N sweep ($n_k = 5$ pairs, $\sigma_{\text{LOS}} = 0.20$ m)

Fig. 9. MDS-Paired autocalibration: median 3D RMSE for $\Delta h \in \{0.5, 1.0, 1.5, 2.0\}$ m (coloured lines) compared against the standalone MDS 3D baseline (red dotted) and Semi-3D upper bound (green dotted), over 100 Monte Carlo trials.



(a) Sweep A: n_k sweep ($N = 20$, $\sigma_{\text{LOS}} = 0.20$ m)

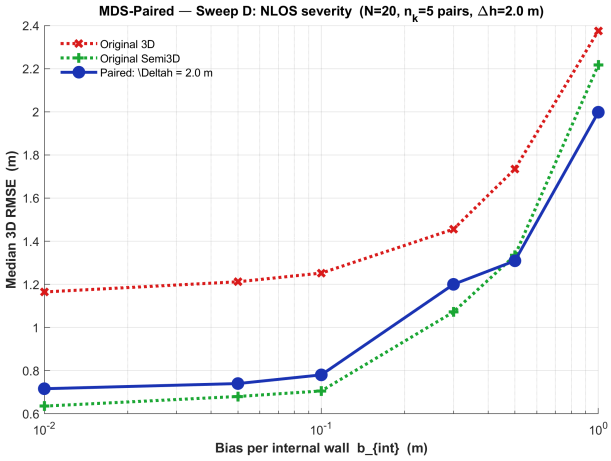
(b) Sweep B: noise sweep ($N = 20$, $n_k = 5$ pairs)

(c) Sweep C: N sweep ($n_k = 5$ pairs, $\sigma_{\text{LOS}} = 0.20$ m)

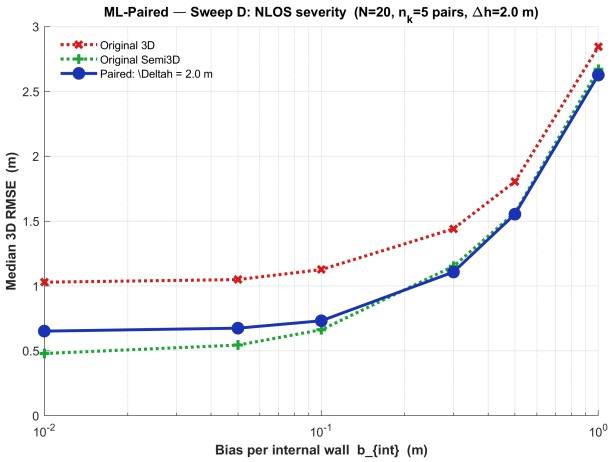
Fig. 10. ML-Paired autocalibration: median 3D RMSE for $\Delta h \in \{0.5, 1.0, 1.5, 2.0\}$ m (coloured lines) compared against the standalone ML 3D baseline (red dotted) and Semi-3D upper bound (green dotted), over 100 Monte Carlo trials.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

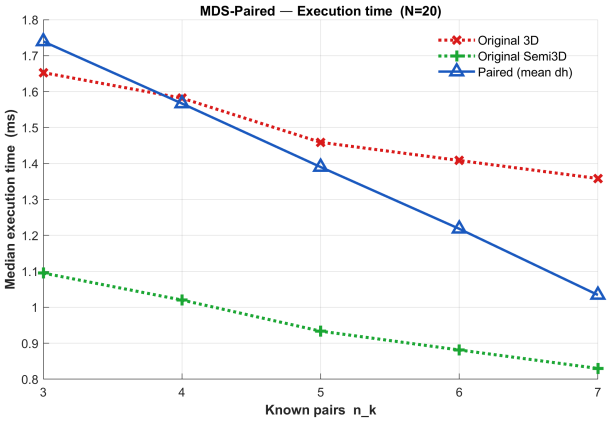


(a) MDS-Paired

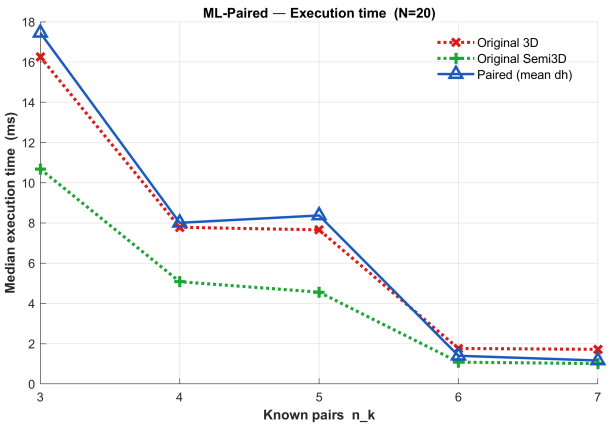


(b) ML-Paired

Fig. 11. 3D RMSE versus per-internal-wall NLOS bias $\boldsymbol{b_{\mathrm{int}}}$ for all $\boldsymbol{\Delta h}$ configurations. Top: MDS-Paired; Bottom: ML-Paired. Fixed parameters: $\boldsymbol{N = 20}$, $\boldsymbol{n_k = 5}$, $\boldsymbol{\sigma_{\mathrm{LOS}} = 0.20}$ m.



(a) MDS-Paired



(b) ML-Paired

Fig. 12. Median per-trial execution time vs. known-pair count $\boldsymbol{n_k}$. The paired curve is the mean over $\boldsymbol{\Delta h \in \{0.5,\ 1.0,\ 1.5,\ 2.0\}}$ m.

and multilateration algorithms are non-iterative in their core stages: MDS costs $O(N^3)$ for a single eigendecomposition, and multilateration costs $O(M_{\max}^2 N)$ for N pivot-LS solves. Both are fully deterministic given the measurement inputs, with no optimisation hyperparameters to tune.

b) Comparison with Ridolfi et al. [5]: Ridolfi et al. [5] report a best-case accuracy of 2.4 cm average error (95th percentile: 5.7 cm) for UWB anchor self-calibration in realistic NLOS conditions, achieved through a combination of iterative gradient descent, machine learning classification of channel impulse response (CIR) signals for NLOS detection and adaptive physical layer (PHY) configuration that selects per-link physical settings to maximise ranging quality. This accuracy level is well beyond what the present work achieves. However, the comparison is not straightforward. The accuracy of [5] is conditioned on access to the full CIR signal from each UWB link and on the availability of a convolutional neural network (CNN) trained on measurements from the specific environment and hardware platform. The present work intentionally targets *infrastructure-only autocalibration with no learning components*: the algorithms operate on scalar ranging measurements only, require no channel impulse response data and carry no trained models that would need to be re-trained when the hardware platform or environment changes. This makes the proposed approach significantly more portable and easier to deploy in new environments without a prior data collection phase. The 2.4 cm accuracy of [5] should therefore be understood as an upper bound on what is achievable when full CIR processing and adaptive PHY are available, rather than as a baseline for lightweight anchor-only autocalibration. A second consideration is again *dimensionality*: the evaluation in [5] is 2D, with anchors in a fixed-height deployment, whereas the present work targets 3D environments with variable anchor heights. Furthermore, the environment in [5] contains up to 8 anchors in a relatively small volume, whereas the present work targets networks of 15–20 anchors in a $22.3 \times 26.3 \times 3.0$ m building with multi-room NLOS structure.

c) Comparison with Schwarzbach and Jung [7]: The probabilistic grid-based approach of Schwarzbach and Jung [7] is the most recent comparable work, operating in 2D on a fully meshed 12-node network in a $44 \times 30\,\mathrm{m}^2$ industrial hall. Under LOS initialisation, their probabilistic grid-based positioning (PGP) achieves median errors of 0.34–0.68 m, comparable to the horizontal RMSE of the present work's algorithms in their Semi-3D modes (0.39–0.61 m). The PGP method's key advantage over deterministic approaches is its robustness under NLOS-contaminated initialisation, where it

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

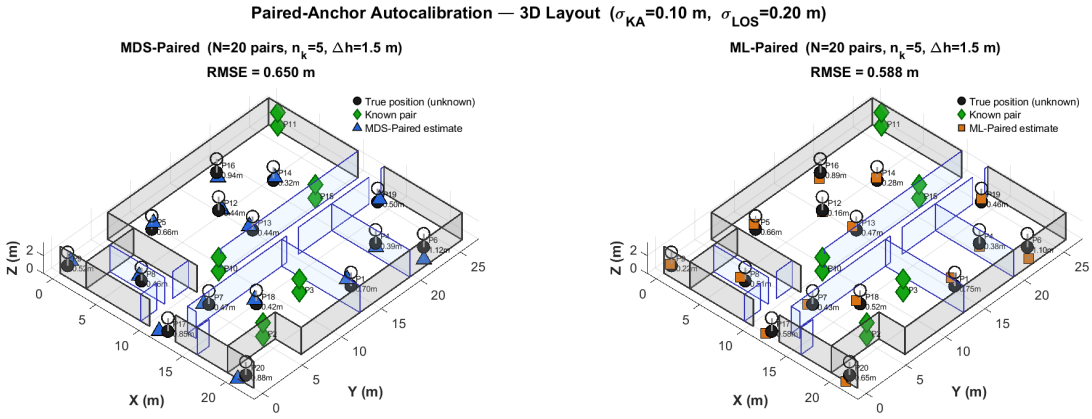


Fig. 13. Single-trial 3D visualisation of paired-anchor autocalibration ($N = 40$ anchors, $M = 20$ pairs, $n_k = 5$, $\Delta h = 1.5$ m, $\sigma_{\text{KA}} = 0.10$ m, $\sigma_{\text{LOS}} = 0.20$ m). Left: MDS-Paired (blue triangles), RMSE = 0.650 m. Right: ML-Paired (orange squares), RMSE = 0.588 m. Green diamonds: known pairs. Black circles: true unknown positions. Dashed lines: error vectors from true to estimated position.

maintains 0.62–0.99 m median accuracy while closed-form baselines collapse to 1.48–2.43 m. The present work does not evaluate NLOS-contaminated reference anchor initialisation — the known anchor positions are assumed to be available with bounded Gaussian uncertainty $\sigma_{\text{KA}} = 0.10$ m — so this robustness dimension cannot be directly compared. However, the NLOS robustness of PGP comes at a computational cost of $\approx 2.3 \times 10^7$ likelihood evaluations per epoch ($O(\tau^2 M^2 N^2)$), whereas the present algorithms complete in 1–2 ms per trial on a standard workstation, a difference of several orders of magnitude that becomes significant for large-scale or battery-constrained deployments.

3) Limitations and Scope: The present work is best understood as a *proof-of-principle study* for the paired-anchor autocalibration concept: its primary contribution is to demonstrate, through systematic Monte Carlo simulation, that a structured vertical deployment combined with a closed-form differential height estimator can substantially close the gap between lightweight 3D autocalibration and the Semi-3D performance ceiling. Several limitations of this simulation-based validation must be acknowledged before the approach can be recommended for deployment.

The most significant limitation is the *simulation-only evaluation*. No real UWB hardware was used to validate the algorithms. The physics-based simulator captures the dominant propagation effects — Gaussian ranging noise, NLOS wall-crossing bias, and the building's multi-room geometry — but real UWB transceivers exhibit additional systematic errors, including antenna directional gain patterns, clock drift, hardware-specific ranging biases that vary with received power, and intermittent packet loss due to fading, which are not modelled. It is therefore expected that absolute RMSE values observed on real hardware will be higher than those reported here. Experimental validation on real hardware in a realistic indoor environment is the most pressing next step, and one that would also enable a direct, fair comparison with the experimentally obtained results of [6].

A second limitation is the *single building geometry used throughout*. All Monte Carlo evaluations use the same 22.3×26.3 m multi-room floor plan. While anchor positions and known-anchor subsets are randomised across trials, the wall layout and NLOS pattern are fixed. Buildings with different aspect ratios, more irregular room layouts, or higher wall densities may produce different connectivity graphs and different asymmetric-NLOS incidence rates for the cross-pair links, affecting the robustness results reported in Sweep D. Evaluation across a broader range of building geometries remains an open direction.

For the paired-anchor approach specifically, the practical requirement of installing *two* anchors at each XY position doubles the hardware count relative to a conventional single-layer deployment, which may constrain applicability where cost or physical mounting is limited. Additionally, the benefit of large Δh requires sufficient ceiling height: in low-ceiling environments (below 2.5 m) the achievable separation is correspondingly limited, reducing the height-estimation gain. The minimum effective separations identified in the simulation ($\Delta h \approx 1.0$ m for MDS-Paired and $\Delta h \approx 1.5$ m for ML-Paired with $n_k \geq 5$) should be treated as deployment guidelines to be re-verified on real hardware rather than as universal thresholds.

Notwithstanding these limitations, the simulation results provide a clear and consistent demonstration of the principle: pairing anchors vertically and exploiting the differential geometry of the cross-pair measurements is an effective strategy for decoupling height estimation from horizontal geometry error in 3D UWB autocalibration.

VI. Conclusion

A lightweight, purely anchor-based 3D autocalibration framework for sparse multihop UWB networks has been presented, with the primary aim of establishing the proof of principle for a structured paired-anchor deployment as a practical path to closing the accuracy gap between 3D and Semi-3D autocalibration without requiring additional infrastructure or learning components. Two complementary algorithms were developed and evaluated: an MDS-based approach

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

that recovers the global anchor geometry via Floyd-Warshall gap-fill, Procrustes alignment and iterative pivot least-squares refinement; and a multilateration-based greedy cascade that propagates the known-anchor set outward through the network. Both were evaluated in full 3D and Semi-3D modes across four Monte Carlo parameter sweeps, achieving 3D RMSE of $1.17\,\mathrm{m}$ and $0.99\,\mathrm{m}$ respectively at the reference configuration ($N = 20$, $n_{\mathrm{known}} = 5$, $\sigma_{\mathrm{LOS}} = 0.20\,\mathrm{m}$). Height estimation was identified as the dominant accuracy bottleneck in 3D mode, with the IDW-based Z estimator contributing a near-constant $\approx 0.75\,\mathrm{m}$ MAE irrespective of network parameters.

To address this bottleneck, a structured *paired-anchor* deployment was introduced in which two anchors share each XY position with a fixed vertical separation Δh. Cross-pair ranging yields a closed-form height estimator that exactly cancels the horizontal inter-pair distance, making Z estimation independent of horizontal position error. The noise variance of this estimator scales as $1/\Delta h^2$, providing a clear deployment guideline: maximise the vertical separation within the available ceiling height. At $\Delta h = 2.0\,\mathrm{m}$ and $n_k = 5$ known pairs the paired approach reduces 3D RMSE by 39% (MDS) and 41% (multilateration) over the standalone 3D baselines, reaching $0.54\,\mathrm{m}$ and $0.53\,\mathrm{m}$ at $n_k = 7$ — within 10% of the Semi-3D performance ceiling — with negligible computational overhead. Sweep D further demonstrates that this advantage is maintained under realistic NLOS conditions ($b_{\mathrm{int}} \leq 0.10\,\mathrm{m}$ typical of IEEE 802.15.4z transceivers in furnished environments), owing to the inherent common-mode bias cancellation of the cross-pair differential estimator.

These simulation results validate the paired-anchor concept and provide a quantitative basis for a physical deployment study. Future work will focus on three directions: (i) experimental validation on real UWB hardware in a representative indoor environment, which is essential to quantify the simulation-to-reality gap, account for hardware-specific ranging biases and packet loss, and enable direct comparison with the experimentally validated results of [6]; (ii) incorporation of lightweight NLOS detection via residual-based outlier gating to improve robustness in heavily obstructed deployments; and (iii) extension of the differential height estimator to anchor-to-tag measurements, providing additional vertical constraints without requiring a dedicated mobile survey node.

REFERENCES

[1] A. Sesyuk, S. Ioannou, and M. Raspopoulos, "A survey of 3d indoor localization systems and technologies," *Sensors*, vol. 22, no. 23, 2022. [Online]. Available: https://www.mdpi.com/1424-8220/22/23/9380

[2] M. Raspopoulos, I. Ioannou, and N. Paspallis, "Wi-fi ftm versus uwb for 3d indoor positioning," in *2026 IEEE 50th Annual Computers, Software, and Applications Conference (COMPSAC)*, 2026, pp. 1896–1901.

[3] A. Sesyuk, S. Ioannou, and M. Raspopoulos, "3d millimeter-wave sensing vs ultra-wideband positioning," in *2024 14th International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2024, pp. 1–6.

[4] M. Hamer and R. D'Andrea, "Self-calibrating ultra-wideband network supporting multi-robot localization," *IEEE Access*, vol. 6, pp. 22 292–22 304, 2018.

[5] Ridolfi, Matteo and Fontaine, Jaron and Van Herbruggen, Ben and Joseph, Wout and Hoebeke, Jeroen and De Poorter, Eli, "UWB anchor nodes self-calibration in NLOS conditions : a machine learning and adaptive PHY error correction approach," *WIRELESS NETWORKS*, vol. 27, no. 4, pp. 3007–3023, 2021.

[6] B. V. Herbruggen, S. Luchie, J. Fontaine, and E. De Poorter, "Multi-hop self-calibration algorithm for ultra-wideband (uwb) anchor node positioning," *IEEE Journal of Indoor and Seamless Positioning and Navigation*, vol. 1, pp. 1–11, 2023.

[7] P. Schwarzbach and A. Jung, "Evaluation of grid-based uncertainty propagation for collaborative self-calibration in indoor positioning systems," *IEEE Journal of Indoor and Seamless Positioning and Navigation*, vol. 4, pp. 161–170, 2026.

[8] X. Liu and M. Cao, "Robust simultaneous uwb-anchor calibration and robot localization for emergency situations," 2025. [Online]. Available: https://arxiv.org/abs/2503.22272

[9] J. Sun, W. Sun, J. Zheng, X. Fang, J. Liu, and A. Mian, "Uwb–imu–odometer fusion for simultaneous calibration and localization," *IEEE Internet of Things Journal*, vol. 12, no. 1, pp. 950–963, 2025.

[10] C. Hamesse, R. Vleugels, M. Vlaminck, H. Luong, and R. Haelterman, "Fast and cost-effective uwb anchor position calibration using a portable slam system," *IEEE Sensors Journal*, vol. 24, no. 16, pp. 26 496–26 505, 2024.

[11] K. Batstone, M. Oskarsson, and K. Åström, "Robust time-of-arrival self calibration and indoor localization using wi-fi round-trip time measurements," in *2016 IEEE International Conference on Communications Workshops (ICC)*, 2016, pp. 26–31.

[12] ——, "Towards real-time time-of-arrival self-calibration using ultra-wideband anchors," in *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, 2017, pp. 1–8.

[13] A. Vashistha, A. Gupta, and C. L. Law, "Self calibration of the anchor nodes for uwb-ir tdoa based indoor positioning system," in *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*, 2018, pp. 688–693.

[14] A. D. Preter, G. Goysens, J. Anthonis, J. Swevers, and G. Pipeleers, "Range bias modeling and autocalibration of an UWB positioning system," in *Proc. of the 2019 Intern. Conf. on Indoor Positioning and Indoor Navigation (IPIN)*. Pisa, Italy: IEEE, Sep. 2019, pp. 1–8.

[15] W. Yu, X. Zhao, and G. Sun, "Self-calibration of anchor positions for indoor localization," in *2017 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2017, pp. 581–586.

[16] X. Ji and H. Zha, "Sensor positioning in wireless ad-hoc sensor networks using multidimensional scaling," in *IEEE INFOCOM 2004*, vol. 4, 2004, pp. 2652–2661 vol.4.

[17] Q. Shi, S. Zhao, X. Cui, M. Lu, and M. Jia, "Anchor self-localization algorithm based on UWB ranging and inertial measurements," *Tsinghua Science and Technology*, vol. 24, no. 6, pp. 728–737, 2019.

[18] P. Biswas and Y. Ye, "Semidefinite programming for ad hoc wireless sensor network localization," in *Third Intern. Symposium on Information Processing in Sensor Networks, 2004. IPSN 2004*, 2004, pp. 46–54.

[19] S. Čapkun, M. Hamdi, and J.-P. Hubaux, "Gps-free positioning in mobile ad hoc networks," *Cluster Computing*, vol. 5, no. 2, pp. 157–167, Apr 2002. [Online]. Available: https://doi.org/10.1023/A:1013933626682

[20] A. Ihler, J. Fisher, R. Moses, and A. Willsky, "Nonparametric belief propagation for self-localization of sensor networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 809–819, 2005.

[21] A. Kannan, G. Mao, and B. Vucetic, "Simulated annealing based localization in wireless sensor network," in *The IEEE Conf. on Local Computer Networks 30th Anniversary (LCN'05)l*, 2005, pp. 2 pp.–514.

[22] M. Qi, B. Xue, and W. Wang, "Calibration and compensation of anchor positions for uwb indoor localization," *IEEE Sensors Journal*, vol. 24, no. 1, pp. 689–699, 2024.

[23] R. Yudanto, J. Cheng, E. Hostens, M. Van der Wilt, and M. Vande Cavey, "Ultra-wideband localization: Advancements in device and system calibration for enhanced accuracy and flexibility," *IEEE Journal of Indoor and Seamless Positioning and Navigation*, vol. 1, pp. 242–253, 2023.

[24] R. W. Floyd, "Algorithm 97: Shortest path," *Commun. ACM*, vol. 5, no. 6, p. 345, Jun. 1962. [Online]. Available: https://doi.org/10.1145/367766.368168

[25] J. A. Costa, N. Patwari, and A. O. Hero, "Distributed weighted-multidimensional scaling for node localization in sensor networks," *ACM Trans. Sen. Netw.*, vol. 2, no. 1, p. 39–64, Feb. 2006. [Online]. Available: https://doi.org/10.1145/1138127.1138129

[26] J. C. Gower and G. B. Dijksterhuis, "Introduction," in *Procrustes Problems*. Oxford University Press, 01 2004. [Online]. Available: https://doi.org/10.1093/acprof:oso/9780198510581.003.0001

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Journal of Indoor and Seamless Positioning and Navigation. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JISPIN.2026.3749195

[27] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-d point sets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 5, pp. 698–700, 1987.



Marios Raspopoulos received an MEng in Electronics and Mobile Technologies in 2003 and an MSc in Communication Networks and Software in 2004 and a PhD in Telecommunications in 2008 from the University of Surrey, UK. From 2004 to 2008 he has been a PhD researcher at the Centre for Communications and Systems Research (CCSR) of the University of Surrey receiving his PhD in Telecommunication in 2008. From 2008 until 2015 he was working in the industry while in 2015 he was appointed as an Assistant Professor at UCLan Cyprus where he is currently an Associate Professor (since 2021) being the coordinator for Engineering programmes. His research interests span a wide spectrum of ICT topics, with a special focus on telecommunications, wireless communications and wireless technologies. He has been heavily involved with Engineering and Telecommunication Projects in areas like Radio Propagation/Radio Planning, Electromagnetism, RF/Antenna Design and Wireless Channel Modelling but has also expanded in other areas of ICT Research, like localization, navigation and Tracking etc.



Iacovos I. Ioannou is a researcher at the Networks Research Laboratory of the University of Cyprus. He is also a Junior Researcher at the Smart Networked Systems Research Group of the CYENS Center. He holds a BSc in Computer Science from the University of Cyprus (2006) and an MSc in Computer and Network Security from the Open University of Cyprus (2017). His Ph.D. was received from the University of Cyprus in 2021, with a focus on telecommunications with AI/ML. His research interests include mobile and wireless communications, next-generation networks (5G) and Device-to-Device (D2D) communications using Artificial intelligence techniques. Iacovos is a highly skilled developer with 20 years of hands-on working experience in the IT industry, ranging in the spectrum of IT systems from analysis, development, installation and management. He has worked at the Philelefthero Publishing Group as an IT administrator and programmer for seven years, at the Cyprus Stock Exchange as an IT specialist and programmer for seven years and at Primetel for six years as an IT and software engineer in the Services Department. Iacovos has vast experience with Cellular network infrastructure and all modern development platforms and languages. He holds a certificate in ICONIX/SCRUM and is also CCNET certified by Cisco.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/