AoI-FusionNet: Age-Aware Tightly Coupled Fusion of UWB-IMU under Sparse Ranging Conditions

Tehmina Bibi, *Student Member, IEEE*, Anselm Köhler, Jan-Thomas Fischer, Falko Dressler, *Fellow, IEEE*

***Abstract*—Accurate motion tracking of snow particles in avalanche events requires robust localization in global navigation satellite system (GNSS)-denied outdoor environments. This paper introduces `AoI-FusionNet`, a tightly coupled deep-learning-based fusion framework that directly combines raw ultra-wideband (UWB) time of flight (ToF) measurements with inertial measurement unit (IMU) data for 3D trajectory estimation. Unlike loose-coupled pipelines based on intermediate trilateration, the proposed approach operates directly on heterogeneous sensor inputs, thereby enabling localization even under insufficient ranging availability. This framework integrates an age of information (AoI)-aware decay module to reduce the influence of stale UWB range measurements and a learned attention gate mechanism that adaptively balances trust between UWB and IMU modalities as a function of measurement availability and temporal freshness. To evaluate performance under limited training data and measurement variability, a diffusion-based residual augmentation strategy is applied during training, producing an augmented variant termed `AoI-FusionNet-DGAN`. Performance of the proposed model is assessed using offline post-processing of real-world measurement data from an alpine environment and benchmarked against geometric, model-based, learning-based, and tightly coupled UWB-IMU fusion baselines.The results demonstrate that `AoI-FusionNet` substantially reduces mean and tail localization errors under intermittent and degraded sensing conditions.**

***Index Terms*—Ultra-wideband (UWB), inertial measurement unit (IMU), tightly coupled fusion, deep learning**

I. Introduction

Motion tracking on a particle level inside extreme natural flow phenomena such as snow avalanches [1] is particularly challenging, as dynamic conditions, limited visibility of satellites, and poor geometric configurations can lead to reduced localization accuracy. Harsh terrain, signal attenuation, chaotic motion, and significant snow depth severely limit the applicability of global navigation satellite system (GNSS), making real-time localization with a single sensor unreliable [2]. Also, size and energy constraints of high-precision GNSS systems are often prohibitive for use in mobile devices.

Ultra-wideband (UWB) has emerged as a very promising alternative in GNSS-denied environments since it uses short pulses, wider bandwidth, and achieves centimeter-level ranging accuracy [3]. However, non-line-of-sight (NLOS) conditions and multipath effects affect UWB, ultimately leading to limited connectivity and intermittent rangings. Consumer-grade inertial measurement units (IMUs), in contrast, deliver high-frequency motion data enabling accurate short-term inertial dead-reckoning but suffer from accumulated drift over time [4] and therefore are unable to maintain long-term precision. To achieve accurate positioning under these conditions, fusion of complementary sensing modalities is therefore required. While many existing fusion techniques adopt a loosely coupled architecture that first processes UWB-based time of flight (ToF) measurements using trilateration to generate position estimates to later fuse with inertial data [5] via Bayesian filtering or learning-based fusion models [6], [7], tightly coupled approaches [8], [9] that directly incorporate UWB range measurements into the estimation process have also been investigated. However, these approaches typically rely on explicit state-transition and measurement models, together with predefined assumptions about sensor uncertainty. Such modeling can become challenging in environments with sparse anchor availability and complex, non-Gaussian UWB ranging errors, particularly under NLOS and multipath conditions. These challenges motivate the investigation of data-driven tightly coupled fusion under sparse UWB ranging conditions.

In the scope of the AvaRange project, we developed an in-flow sensor system in which wireless sensor nodes are embedded in snow for data acquisition during avalanches [10]. To track individual particles within the flow, a sensor network is created by placing static anchors (called AvaAnchors) across a large alpine area. Snow particle-shaped, 3D-printed mobile sensor nodes (called AvaNode) are equipped with a multi-sensor setup where ToF ranging is performed by UWB sensor, and an IMU sensor produces motion data. The position estimation of the mobile node is performed by measuring distances to the static anchors. The tracking results from this experiment confirmed that UWB-based ranging using ToF measurements has better accuracy than conventional localization methods such as GNSS [11]. Unfortunately, line-of-sight (LOS) conditions were not always achievable, and localization accuracy is significantly affected by multipath effects, NLOS conditions, or LOS signal blockage. This leads to intermittent and unreliable UWB ranging, motivating the need for sensor fusion that can account for varying measurement availability and reliability.

Existing mitigation and correction strategies either depend on dense anchor deployments, computationally expensive post-processing algorithms [12], [13], [14], or identification of NLOS measurements and their subsequent classification as outliers by excluding them using localization filters [15]. Recent intelligent positioning approaches have also explored multi-frequency smartphone positioning and A-GNSS services, including 3D mapping-aided techniques for challenging outdoor and NLOS-prone environments [16]. Classical tightly coupled

T. Bibi and F. Dressler are with School of Electrical Engineering and Computer Science, TU Berlin, Germany (e-mail: {bibi,dressler}@ccs-labs.org).

A. Köhler and J.-T. Fischer are with the Austrian Research Centre for Forests, Austria (e-mail: {anselm.koehler,jt.fischer}@bfw.gv.at).

This work was supported by the project AvaRange funded by German Research Foundation (DFG), grant DR 639/22-1 and Austrian Science Fund (FWF), grant I 4274-N29.

Kalman filter (KF)-based fusion methods [17], [18], [19] rely on recursive state estimates, making them sensitive to model inaccuracies and difficult to apply in nonlinear or non-Gaussian scenarios [20]. More recent learning-based ranging compensation methods mitigate NLOS bias without discarding measurements; however, they still assume a sufficient number of simultaneous ranges for trilateration and therefore remain vulnerable under sparse anchor visibility or intermittent ranging conditions [21]. Together, these observations highlight the need for a fusion framework that directly processes raw measurements and models their temporal reliability.

Motivated by these challenges, this work develops a data-driven, tightly coupled UWB-IMU fusion framework that performs joint temporal reasoning over heterogeneous sensor measurements and remains operational even when an insufficient number of UWB measurements are available to produce standalone positioning. In particular, we propose `AoI-FusionNet`, a deep-learning-based fusion architecture that directly integrates raw UWB range measurements with IMU acceleration measurements without relying on intermediate trilateration. Under sparse and intermittent ranging conditions, the temporal freshness of UWB measurements plays a role comparable to their availability. To incorporate this, `AoI-FusionNet` uses an age of information (AoI)-aware decay mechanism that down-weights stale UWB rangings, combined with a learned attention gate that adaptively balances the relative trust between UWB and IMU measurements. The proposed model learns to predict incremental 3D displacements over a sliding temporal window and reconstructs the full trajectory through accumulation, thereby improving trajectory stability and reducing long-horizon drift. To mitigate overfitting due to data scarcity, we introduced a data augmentation strategy that injects physically plausible UWB perturbations during training to improve stress-testing of robustness without changing the underlying motion trajectory.

Our main contributions can be summarized as follows:

- We propose `AoI-FusionNet`, a tightly coupled deep learning-based sensor fusion framework that fuses raw UWB range measurements with time-synchronized IMU acceleration data without relying on trilateration.
- We introduce an AoI-aware decay mechanism and a learned attention gate that dynamically modulate the influence of UWB and IMU measurements based on anchor availability, decayed UWB quality, and motion context, under sparse, intermittent, and stale measurements.
- We integrate a diffusion-based UWB data augmentation strategy at the residual level to improve robustness and generalization under limited training data.
- We validate the proposed approach using real-world experimental data and benchmark its performance against geometric, model-based, learning-based, and tightly-coupled UWB-IMU fusion baselines.

The remainder of this paper is organized as follows: Section II reviews related work, including localization based on UWB and IMU and fusion algorithms comprising filter-based and learning-based fusion strategies. Section III presents the system model and reviews state-of-the-art fusion approaches. Section IV introduces our new `AoI-FusionNet` algorithm and provides details of the proposed model. Section V presents the evaluation results and discusses the lessons learned from our study. Finally, Section VI summarizes the article and outlines future research directions.

II. Related Work

A. UWB Range-based Localization

Recent research has demonstrated that UWB has great potential for wireless ranging and localization, offering centimeter-level spatial resolution. UWB localization systems typically rely on ToF, angle of arrival (AoA), and time difference of arrival (TDoA) measurements with algorithms such as least squares, the Taylor algorithm, CHAN algorithm for position calculation. High positioning accuracy can be achieved under LOS situations between the mobile node and the static anchors. However, UWB faces limitations in complex outdoor environments where NLOS conditions deteriorate ToF-based ranging. Additionally, multipath effects and poor anchor placement (high geometric dilution of precision (GDOP)) further degrade positioning accuracy and hinder seamless localization performance [22].

Poulose et al. [23] proposed a UWB-based positioning framework using extended Kalman filter (EKF) to enhance indoor localization in a multi-anchor configuration. The algorithm processes time of arrival (ToA) measurements from multiple anchors to address the inherent system nonlinearities, and its performance was tested using MATLAB simulations in a 2D environment under both LOS and NLOS conditions. While the proposed EKF-based approach achieved positioning errors of less than ±50 cm, its use in dynamic scenarios is constrained by its reliance on predefined noise models, thus limiting its adaptability in real-world conditions.

Van Herbruggen et al. [24] introduced two heuristic anchor selection algorithms for two-way ranging (TWR)-based UWB indoor positioning systems to improve accuracy and update rates. One method selects anchors based on link quality and GDOP, while the other method is designed as a lightweight, real-time algorithm for deployment on low-power UWB tags. The performance has been evaluated using a Kalman filter in industrial environments, achieving positioning accuracy in the order of 15–20 cm. However, the non-constrained algorithm requires continuous communication with all anchors, and these methods rely on static link quality metrics and do not account for motion dynamics or temporal variations, which is essential for reliable localization in dynamic environments.

B. IMU-based Localization

Data-driven approaches based on IMU measurements have shown promising results for localization based on pedestrian dead-reckoning (PDR) and have been investigated to improve individual components of PDR. Tiwari and Jain [25] introduced a step detection method using the dynamic weight integrated fuzzy C-means (DWIFCM) algorithm for pedestrian dead-reckoning navigation, where it exploits features extracted from tri-axial accelerometer data and dynamically adjusts feature weights corresponding to changes in walking patterns and

terrain. Liu et al. [26] proposed an F-LSTM-based heading estimation technique that fuses IMU and magnetometer data to correct PDR heading angle estimation. This fully connected long short-term memory (LSTM) model learns to regress the non-linear sensor error when it is trained on data collected in different heading directions.

While these algorithms demonstrate improved inertial navigation accuracy, the dependence on supervised training with restricted motion patterns and limited environmental variations, such as dynamic outdoor environments, limits their adaptability and generalization to unseen users and magnetic disturbances. To address gyrometer saturation in orientation estimation, Neurauter et al. [27] proposed two fusion algorithms using low-cost IMU gyrometer and magnetometer readings, later fused with the state-of-the-art Madgwick's algorithm. Despite testing on real snow avalanche datasets, the dependence on magnetometer data limits its robustness in environments sensitive to magnetic disturbances.

C. Model-based Sensor Fusion

IMUs provide good localization performance over short step sequences, and several studies have explored its potential in multisensor fusion combining IMU with other sensor modalities, including vision [28], LiDAR [29], or UWB [30] to achieve better localization accuracy. Among wireless communication technologies in outdoor scenarios, UWB stands out as a cost-effective and compact solution capable of maintaining sub-meter accuracy. Classical filtering methods such as KF and its nonlinear variants have been extensively used so far [31]. Nonetheless, the reliance of such models on Gaussian noise models reduces their usefulness in nonlinear and non-Gaussian settings. Hybrid solutions such as [32], [33] that integrate vision with UWB further improve robustness but are sensitive to avalanche-like environments where other factors, such as visibility and occlusions due to the surrounding powder cloud, adversely affect localization accuracy.

Xu et al. [34] introduced a loosely coupled integrated localization system based on a least-squares ac SVM- assisted unbiased finite impulse response (UFIR); however, the NLOS situation is not considered. Fan et al. [35] proposed an ac EKF-based classical tight coupling algorithm for indoor navigation that integrates raw UWB with IMU to mitigate NLOS effects. While the proposed EKF-based approach combines carrier motion characteristics to improve robustness, the manually designed probabilistic models and threshold-dependent NLOS handling limit its performance for outdoor scenarios. Similarly, Li et al. [36] introduced a tightly coupled indoor localization UWB-PDR solution using a factor graph. The proposed method improves localization accuracy by jointly optimizing raw UWB ranges and pedestrian motion constraints; however, it lacks adaptive noise modeling to account for environmental dynamics, and its application is limited to 2-D scenarios. The explicit assumptions about motion dynamics and measurement noise statistics underlying classical model-based frameworks become difficult to satisfy under sparse anchor availability and multipath-induced non-Gaussian UWB errors, thereby limiting their robustness in our scenario and motivating the data-based approach.

D. Learning-based Sensor Fusion

Recent studies have investigated data-driven learning approaches to improve sensor fusion and localization performance compared to traditional model-based filtering, particularly in complex, non-linear, and non-Gaussian environments.

Machine learning techniques have also been investigated for multi-sensor fusion to obtain precise localization. Using an extreme gradient boosting (XGBoost)-enhanced fusion framework, Tommingas et al. [37] proposed fusing UWB and GNSS to handle indoor-outdoor transitions by predicting sensor confidence levels. However, the approach relies on GNSS availability and hand-crafted features, which may limit its scalability. Deep neural network architectures such as convolutional neural network (CNN)-LSTM have been proposed by Zhi et al. [38], taking advantage of CNN and LSTM architectures simultaneously to improve the performance of the integrated navigation system during GPS outages, but the reliance on IMU measurements makes the system susceptible to sensor drift and synchronization challenges.

Muthineni et al. [39] proposed a cascaded deep learning (DL) architecture for UWB-IMU fusion in industrial environments. While the approach outperforms the EKF baseline, the model is trained for a fixed number of inputs and does not explicitly address multipath-rich outdoor environments with intermittent and varying LOS/NLOS conditions. Fontaine et al. [40] presented a transfer learning (TL)-based automatic optimization solution to develop a high precision UWB localization system. Similarly, Yang et al. [41] proposed a tightly coupled-bidirectional encoder representations from transformers (BERT) algorithm for accurate indoor positioning. Ren et al. [42] combine an attention-based LSTM for UWB pseudo-range generation with a tightly coupled navigation filter to explore attention-based temporal models for NLOS-affected ranging. Tran et al. [43] proposed a generative adversarial network (GAN)-based data augmentation technique for enhancing the training dataset so that the model extracts multiple features of channel impulse response (CIR) to reduce the localization error caused by NLOS conditions.

These approaches demonstrate the potential of learning-based methods for UWB localization and sensor fusion, addressing UWB error correction, NLOS identification, and data augmentation. However, the temporal freshness of intermittent UWB measurements is not explicitly considered. In contrast, `AoI-FusionNet` directly fuses raw UWB ranges and IMU measurements while incorporating AoI-aware decay and learned attention for adaptive sensor weighting.

E. Our Solution: `AoI-FusionNet`

In this work, we propose `AoI-FusionNet`, an AoI-aware tightly coupled LSTM-based UWB-IMU fusion framework designed for 3D trajectory estimation under intermittent range measurements, sparse anchor visibility, high GDOP and NLOS conditions. Unlike existing fusion architectures that rely solely on LOS UWB ranges by explicitly identifying and excluding NLOS measurements, our proposed architecture uses both LOS and NLOS measurements within a unified fusion process. To the best of our knowledge, the explicit modeling of temporal

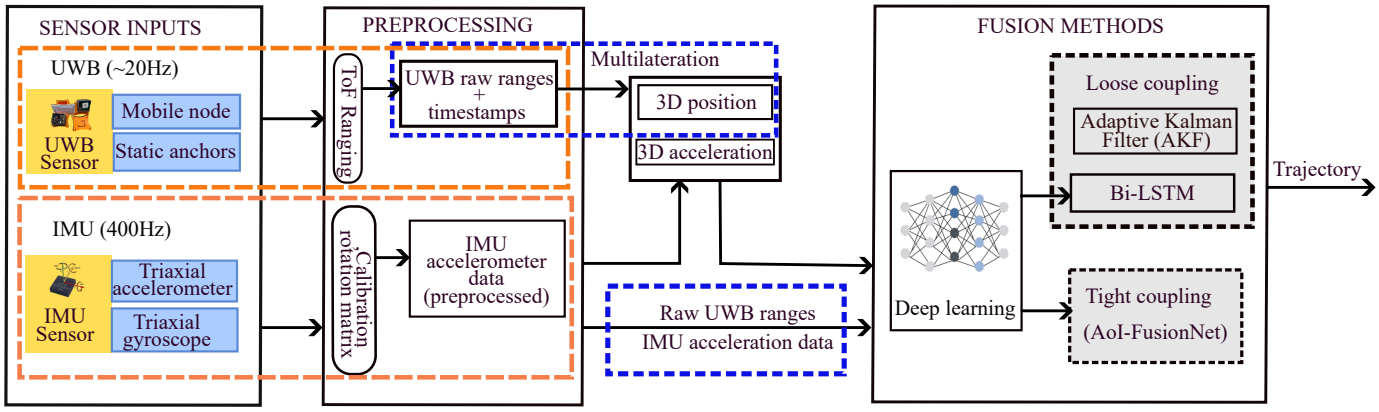


Figure 1. System model and data flow for UWB-IMU fusion. The mobile node is equipped with a UWB transceiver and a triaxial IMU and communicates with static anchors in a GNSS-denied outdoor environment. Raw UWB ranges with timestamps and preprocessed inertial measurements are provided as inputs to multiple localization pipelines. Classical baselines include loose-coupled fusion using AKF and a Bi-LSTM model operating on trilaterated UWB positions. In contrast, the proposed `AoI-FusionNet` performs tightly coupled fusion directly on raw UWB and IMU measurements, incorporating AoI-aware decay, a learned fusion gate, and LSTM-based temporal modeling. The resulting 3D trajectory is evaluated against ground truth.

measurement reliability within a tightly coupled learning-based fusion framework has not been investigated in this specific formulation. The model integrates an AoI-aware decay mechanism in the fusion algorithm that adaptively down-weighs stale UWB measurements, and a learned attention gate dynamically balances trust between UWB and IMU features based on the UWB measurement freshness and anchor reliability. By modeling temporal dependencies through the recurrent architecture, our model learns temporal motion patterns directly from sequential sensor data and maintains stable localization performance under sparse anchor measurements and NLOS conditions.

III. System Model

In the following, we describe our system model that uses UWB and IMU systems as its key building modules. The system integrates UWB ranging and inertial measurements to construct a multi-sensor dataset for both classical filtering and DL-based fusion methods. As shown in Fig. 1, the system adopts a modular architecture in which synchronized data collected from the mobile node and anchor network undergo preprocessing before being processed by the sensor fusion and trajectory evaluation modules. The system model is structured into four primary stages: (1) experimental setup using the mobile node (AvaNode) and static anchors (AvaAnchors) for data acquisition, (2) sensor-level preprocessing that includes UWB-based ranging and localization, (3) IMU calibration followed by frame transformation, (4) multi-sensor fusion methods involving loosely coupled UWB-IMU fusion, and tightly coupled UWB-IMU fusion. Each stage of the system is described in detail in the following.

A. AvaRange Scenario

The sensor network deployed for this experiment comprises two principal components: the AvaNode [44], which functions as the mobile sensing unit, and multiple AvaAnchors that are statically distributed across the test area to support position estimation [11]. The sensor payload integrated into the AvaNode included a UWB transceiver and an IMU, and



Figure 2. Tracking experiment at Innsbruck Nordkettenbahn cable car to derive the trajectory along the avalanche path. An IMU and UWB sensor node, together with a reference GNSS, is traversing high alpine terrain that is equipped with static UWB anchors.

was further supported by a high-precision real-time kinematic (RTK) GNSS receiver to provide reference data for ground-truth information. In total, six anchors were deployed where one anchor was the gateway anchor responsible for managing the backhaul WiFi mesh network connecting the remaining anchors. This configuration enables synchronized acquisition of multi-modal sensor data from the UWB module, an IMU, and a high-precision RTK GNSS sensor, for subsequent localization and sensor fusion analysis.

In this work, the measurement data used for post-processing were collected during a field experiment conducted in the Austrian Alps [11]. The experimental arrangement consisted of an AvaNode mounted onto a cable-car counterweight of the Hafelekar section of the Innsbruck Nordkettenbahn, which travels along the Nordkette avalanche test-site, and was used herein to emulate realistic avalanche terrain geometry (see Fig. 2). The moving node was used to investigate UWB-based tracking under realistic alpine geometry. The setup provides a controlled and reproducible trajectory through alpine terrain and enables the investigation of UWB-based tracking under dynamic motion, intermittent ranging, and GNSS-denied conditions. The cable car experiment therefore serves as a controlled proxy for localization conditions relevant to avalanche-particle tracking and terrain geometries, rather than as a direct avalanche-field experiment.

1) AvaNode: The AvaNode represents the tracked device and is designed to simulate the motion of a snow granule within

an avalanche. It consists of multiple key sensing elements. UWB ToF-based ranging is performed using a Decawave DW1000 transceiver interfaced with an Adafruit Feather M0 microcontroller. The DW1000 chip supports timestamping precision down to 15 ps, thus allowing highly accurate distance measurements. Inertial sensing has been performed by the MPU9250 IMU that integrates a triaxial accelerometer and gyroscope. The IMU is housed within a 3D-printed cube with an edge length of 160 mm, ensuring alignment between sensor coordinate frame and the geometric center of the enclosure. IMU records motion data at a sampling frequency of 400 Hz and the measurements are logged to an onboard SD card via the Feather M0 microcontroller, while UWB ranging data is sampled at a frequency of 20 Hz. Despite the incorporation of additional components in the AvaNode platform that includes low-cost GPS module and a recovery mechanism, only the UWB and IMU measurements are considered in this article for tracking analysis.

2) AvaAnchor: Six AvaAnchors equipped with DW1000 UWB chips serve as static anchors with known fixed locations installed along the cable car route. In addition, the anchors establish an IEEE 802.11 WiFi mesh network, while LoRa is used for sending commands to the node. The anchor positions were precisely surveyed with an RTK GNSS system to allow accurate multilateration. In a nutshell, this configuration enables reliable acquisition of high-resolution motion data under controlled conditions to evaluate sensor fusion techniques for 3D trajectory reconstruction in avalanche-like environments.

3) Ground truth: To establish high-accuracy ground truth, an RTK GNSS (Emlid Reach RS2+) unit was mounted. The reference positions were recorded at 10 Hz. Over the synchronized segment, the start-to-end displacement is approximately 700–720 m (horizontal: $\approx 620\,\mathrm{m}$; vertical: $\approx 343\,\mathrm{m}$), corresponding to an average speed of $\approx 22\,\mathrm{km/h}$.

To extend the evaluation beyond one recorded trajectory, synthetic UWB-IMU trajectories generated with the physics-based *PhySynth* framework [45] are additionally considered. The framework allows controlled variation of trajectory geometry, motion profile, and UWB ranging conditions. A total of 180 independent traces are generated using the same six-anchor deployment as the recorded campaign, covering three ranging regimes: Dense, Sparse, and Degraded with at least two anchors visible for 12.7%, 2.4%, and 0.5% of samples, respectively.

B. UWB Ranging and Localization

The alternative double-sided two-way ranging (AltDS-TWR) method [46] was employed for UWB-based distance estimation. Even though AltDS-TWR reduces timing uncertainty and mitigates clock drift without depending upon tight clock synchronization, it does not eliminate the multipath and NLOS bias. The time-of-flight (T_f) is calculated as:

$$T_f = \frac{R_a R_b - D_a D_b}{2(R_a + D_a)}, \tag{1}$$

where R_a and R_b denote round-trip delays measured at each device, and D_a, D_b are the respective known reply delays

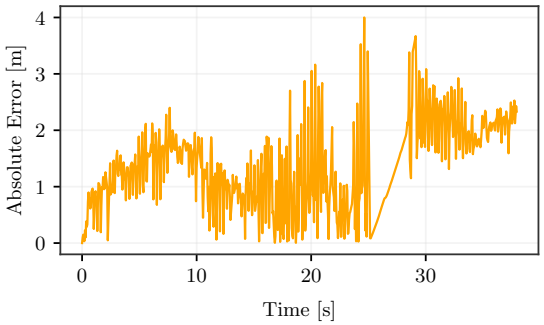


Figure 3. Relative error computed between UWB-only trilaterated positions and RTK GNSS ground truth along Z-axis.

[10]. Distances d_i can be calculated from time-of-flight (T_f) measurements as:

$$d_i = c \cdot T_f, \tag{2}$$

Here, c denotes the speed of light, which generally depends on temperature, humidity, and the snow-air density inside a powder cloud, but is kept constant here. This formulation results in accurate range measurements even in the presence of timing asymmetries.

The AvaNode position was estimated using multilateration formulated as a non-linear optimization problem. Specifically, the objective of this optimization is to minimize the mean squared error (MSE) between the measured UWB distances and the corresponding Euclidean distances to a set of fixed AvaAnchors. Let $\mathbf{a}_i = (x_i, y_i, z_i)$ denote the known coordinates of anchor i, and let d_i represent the corresponding measured distance. The estimated position of the mobile node $\mathbf{p} = (x, y, z)$ is obtained by solving the following problem:

$$\mathrm{MSE}(\mathbf{p}) = \frac{1}{N} \sum_{i=1}^{N} \left(\|\mathbf{p} - \mathbf{a}_i\| - d_i \right)^2 \tag{3}$$

where N represents the number of anchors providing distance measurements. A minimum of four anchors is required to obtain a unique solution in three-dimensional space.

Fig. 3 shows the absolute localization error along the vertical axis for UWB-only multilateration against the RTK GNSS reference. Over time, an increasing deviation of up to 4 m is observed, particularly along the vertical axis, where the impact of unfavorable anchor geometry reflected by high GDOP is more significant.

C. IMU Calibration and Preprocessing

The calibration of IMU has been adopted from the procedure proposed by Neurauter and Gerstmayr [47] and Winkler et al. [48]. These methods explicitly model sensor errors, estimate orientation from angular velocity measurements, and compute gravity-compensated integration of accelerometer measurements. Precise calibration of IMU sensors is imperative for mitigation of deterministic measurement errors. After proper calibration, the accelerometer data reflect only inertial and gravitational components, while the gyroscope accurately tracks angular motion independent of any systematic drift.

The resulting calibrated inertial measurements are illustrated in Fig. 4. During the experiment, the IMU sensor experienced

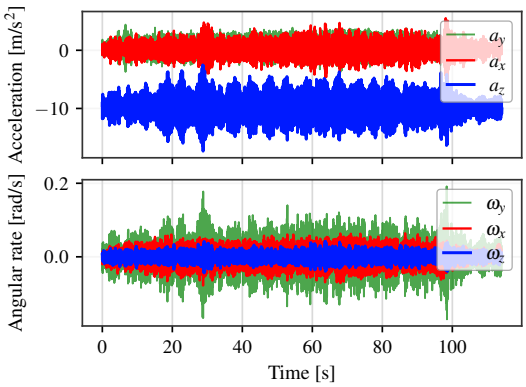


Figure 4. Calibrated accelerometer and gyroscope measurements along X-, Y-, and Z-axes.

predominantly translational motion with limited rotational activity. Therefore, the accelerometer data exhibit consistent behavior along the X- and Y-axes; however, the measurements along the Z-axis are largely influenced by gravity and centered around −9.81 m/s^2, consistent with theoretical expectations. In contrast, the gyroscope measurements remain close to zero across all axes, with only minor fluctuations, highlighting the absence of significant rotational motion.

Following calibration, the next stage is motion reconstruction, where inertial data, specifically translational acceleration and angular velocity, are used to reconstruct the 3D trajectory of the mobile node. The complete kinematic state of the mobile node is estimated by integrating calibrated accelerometer and gyroscope measurements over time. To transform measurements from the local sensor (body) frame to the global frame of reference, rotation matrices are computed to estimate the orientation from gyroscope data. Based on the estimated orientation and corresponding rotation matrices, the linear motion is reconstructed by the transformation of calibrated translational acceleration to the global reference frame. Subsequently, gravitational acceleration is subtracted to obtain translational acceleration, followed by the numerical integration of global acceleration using the explicit Euler method to obtain velocity and position. The resulting gravity-compensated translational acceleration is used as an input to fusion methods.

D. Loose-coupled UWB-IMU Fusion Using Kalman Filter

We first adopted an AKF, an extension of standard KF, as our baseline method for loose coupling-based fusion of data from UWB and IMU sensors. By dynamically adjusting the measurement noise covariance under time-varying sensor noise and uncertainties, the filter estimates 3D position, velocity, and acceleration. The state vector can be defined as:

$$\mathbf{x}_k = \begin{bmatrix} x_k & y_k & z_k & \dot{x}_k & \dot{y}_k & \dot{z}_k & \ddot{x}_k & \ddot{y}_k & \ddot{z}_k \end{bmatrix}^\top \tag{4}$$

State prediction follows the standard KF prediction model as:

$$\mathbf{x}_{k|k-1} = \mathbf{F}\mathbf{x}_{k-1|k-1}, \tag{5}$$

$$\mathbf{P}_{k|k-1} = \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^\top + \mathbf{Q}_k, \tag{6}$$

where $\mathbf{F}$ is the constant acceleration 9×9 state transition matrix based on sampling interval Δt, and $\mathbf{Q}_k = \mathbf{G}\mathbf{Q}_0\mathbf{G}^\top$ is the process noise covariance. Measurement updates incorporating actual sensor data are performed using:

$$\mathbf{z}_k = \mathbf{H}_k\mathbf{x}_{k|k-1} + \mathbf{v}_k, \quad \mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R}_k). \tag{7}$$

The Kalman gain and state update equations are given by:

$$\mathbf{y}_k = \mathbf{z}_k - \mathbf{H}_k\mathbf{x}_{k|k-1}, \tag{8}$$

$$\mathbf{S}_k = \mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^\top + \mathbf{R}_k, \tag{9}$$

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}_k^\top\mathbf{S}_k^{-1}, \tag{10}$$

$$\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \mathbf{K}_k\mathbf{y}_k, \tag{11}$$

$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1}. \tag{12}$$

To account for time-varying uncertainties, AKF adaptively updates measurement noise covariance $\mathbf{R}_k$ using the empirical innovation covariance over a sliding window of size N. Let $\mathbf{y}_i$ be the innovation vector at time step i over a fixed window of size N. The sample innovation covariance is:

$$\hat{\mathbf{C}} = \frac{1}{N-1}\sum_{i=k-N+1}^{k}\mathbf{y}_i\mathbf{y}_i^\top \tag{13}$$

The measurement noise covariance is updated as:

$$\mathbf{R}_k = \hat{\mathbf{C}} - \mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^\top \tag{14}$$

In addition, GDOP-based dynamic scaling to $\mathbf{R}_k$ is applied to UWB measurement noise to account for geometric uncertainty.

$$\mathbf{R}_{\text{UWB}}^{\text{scaled}} = \alpha(g)^2 \cdot \mathbf{R}_{\text{UWB}} \tag{15}$$

where the scaling factor $\alpha(g)$ is a function of the GDOP value g. This combined residual-based adaptation and geometry-aware scaling improves reliability under degraded anchor configurations, especially along the vertical axis, where UWB measurements are often less accurate, due to the limited height variation in anchor positions.

E. Loose-coupled UWB-IMU Fusion using Bi-LSTM

In our preliminary work [49], we proposed a learning-based loosely coupled fusion technique using a Bi-LSTM-based deep recurrent network on the same experimental dataset. In this method, UWB position estimates are fused with IMU-derived inertial measurements, consistent with loose-coupling pipeline. The designed model learns the complex temporal relationships between IMU acceleration measurements and position estimates obtained from UWB multilateration. The model thereby learns to predict the frame-wise displacement vectors referred to as the delta position $\Delta\mathbf{p}_t$ between consecutive timestamps, where each frame contains a single-time instance incorporating input features that contain both IMU and UWB data. Instead of predicting absolute positions directly, the model focuses on learning incremental motion patterns to improve generalization and reduces long-term drifts.

Input sequences are taken in a sliding window of 200 consecutive time steps, where each time step contains six features that are three-axis accelerometer readings and three-axis UWB positions, creating a $[200 \times 6]$ input tensor. The architecture consists of three stacked Bi-LSTM layers; each layer has a hidden dimension of 512 units per direction (forward

and backward), followed by a fully connected layer producing 3D displacement vectors. During inference, these predicted displacement deltas are cumulatively summed to reconstruct a full 3D trajectory. The training of the model is performed using adaptive moment estimation (Adam) optimizer with a learning rate of 0.001 and mini-batch size of 64. Early stopping with a patience of 30 epochs prevents model from overfitting. To account for axis-dependent error characteristics, a weighted mean squared-error (WMSE) loss function is applied, placing greater emphasis on the vertical dimension. It can be expressed as follows:

$$\mathcal{L}_{\text{WMSE}}(t) = \left\|\mathbf{W}^{1/2}\big(\widehat{\Delta\mathbf{p}}_t - \Delta\mathbf{p}_t\big)\right\|_2^2, \tag{16}$$

where $\widehat{\Delta\mathbf{p}}_t = [\widehat{\Delta x}_t, \widehat{\Delta y}_t, \widehat{\Delta z}_t]^T$ and $\Delta\mathbf{p}_t = [\Delta x_t, \Delta y_t, \Delta z_t]^T$ denotes the predicted and ground-truth incremental displacement vectors, respectively, and $\mathbf{W} = \text{diag}(w_x, w_y, w_z)$ shows a diagonal weighting matrix that emphasizes axis-dependent errors. For a mini batch of size N, the averaged WMSE $\mathcal{L}$ can be calculated as:

$$\mathcal{L} = \frac{1}{N}\sum_{n=1}^{N} \mathcal{L}_{\text{WMSE}}(\widehat{\Delta\mathbf{p}}^{(n)}, \Delta\mathbf{p}^{(n)}). \tag{17}$$

The model configuration was empirically optimized to find the sweet spot and achieve optimal performance. More details can be found in our earlier work [49].

IV. AoI-FusionNet

Contrary to the loosely coupled (LC-Bi-LSTM) approach, a tightly coupled learning-based localization framework (`AoI-FusionNet`) is proposed that combines inertial and UWB ranging data while explicitly accounting for ranging sparsity, temporal staleness, and availability of anchors. This method is specifically designed to operate under UWB degradation, including scenarios where only one or two anchors are intermittently visible. `AoI-FusionNet` operates on the sliding window of multi-sensor temporal measurements and performs incremental motion prediction rather than absolute position estimation. Such an approach mitigates long-horizon drift and improves localization performance under missing or unreliable UWB measurements. The predicted motion deltas are then accumulated to recover the full trajectory.

Before feeding inertial data into the tightly coupled fusion pipeline, the acceleration bias optimization method proposed by Neurauter and Gerstmayr [47] is applied to mitigate the integration drift in the calibrated translational acceleration data. The measured acceleration is corrected by a first-order polynomial function of time:

$$\tilde{\mathbf{a}}(t) = \mathbf{a}(t) + \mathbf{a}_0 + \mathbf{a}_1 t,$$

where $\mathbf{a}_0, \mathbf{a}_1 \in \mathbb{R}^3$ are the optimized correction parameters. Rotation matrices are then applied to the corrected acceleration so that it is rotated into the global frame of reference, after which global velocity and position are obtained via numerical integration. The correction parameters $\mathbf{a}_0$ and $\mathbf{a}_1$ are selected by minimizing a cost function:

$$F = \|\mathbf{v}_{\text{final}}\|^2 + \|\mathbf{x}_{\text{final}} - \mathbf{x}_{\text{ref}}\|^2,$$

Table I
KEY HYPERPARAMETERS OF AoI-FusionNet: TIGHTLY-COUPLED IMU-UWB FUSION MODEL

Parameter	Value	Description
IMU input dim.	3	3-axis accel.
UWB anchors (A)	Variable	Available anchors
Output dim.	3	$\Delta x, \Delta y, \Delta z$
Hidden size (H)	128	Shared latent dim.
Temporal model	LSTM	Tight temporal coupling
Number of Layers	1	LSTM layer
Anchor embed. (E)	3	Anchor identity encoding
Learned attention	Yes	IMU-UWB gating
Loss function	Composite Huber	Robust multi-term loss
Optimizer	AdamW	Weight-decayed Adam
Learning rate	3×10^{-4}	Initial LR
Early stopping	20 epochs	Val-based patience
Data split	70/15/15	Train/Val/Test

where $\mathbf{v}_{\text{final}}$ and $\mathbf{x}_{\text{final}}$ denote the terminal velocity and position, while $\mathbf{x}_{\text{ref}}$ is a known reference endpoint. This optimization encourages the reconstructed motion to end at rest and match the known final position, thereby removing the sensor drift.

A. AoI-FusionNet Model Construction

`AoI-FusionNet` follows a four-stage processing pipeline designed to (1) encode modality-specific features, (2) model the time-varying reliability of UWB measurements using AoI, (3) perform adaptive cross-modal fusion in a reliability-aware manner over time, and (4) predict incremental motion updates. This fusion framework includes an explicit reliability model for UWB measurements using learned AoI-decay, reducing the impact of stale ranging updates. The age-weighted UWB features are integrated with IMU data through a learned attention gate that modulates the relative contribution of each sensor modality based on measurement availability, temporal freshness, and motion context. The fused sequence is then processed by a recurrent model to enforce temporal dependencies, while learned anchor embeddings provide an anchor-specific representation without explicitly encoding anchor coordinates.

Algorithm 1 presents the pseudo code for the proposed model and Table I summarizes its hyperparameters.

1) IMU Feature Encoding: The proposed model operates on synchronized input data from UWB and IMU sensors over a temporal window of length T. At each discrete time step t, the model processes optimized tri-axial acceleration measurements $\mathbf{u}_t \in \mathbb{R}^3$. The complete representation of the IMU sequence over the window can be denoted by

$$\mathbf{U} = \{\mathbf{u}_t\}_{t=1}^{T}. \tag{18}$$

A shared IMU encoder $f_{\text{imu}}(\cdot)$ is applied independently at each time step to obtain latent features $\mathbf{h}_t^{\text{imu}} \in \mathbb{R}^H$ as

$$\mathbf{h}_t^{\text{imu}} = f_{\text{imu}}(\mathbf{u}_t), \tag{19}$$

where $f_{\text{imu}}(\cdot)$ is implemented as a multilayer perceptron with parameters shared across time and $\mathbf{h}_t^{\text{imu}}$ is the learned feature embedding of IMU acceleration data used for fusion. The encoded sequence is therefore

$$\mathbf{H}^{\text{imu}} = \{\mathbf{h}_t^{\text{imu}}\}_{t=1}^{T}. \tag{20}$$

Algorithm 1 Pseudo code for tightly-coupled UWB-IMU Fusion: `AoI-FusionNet`

Require: IMU $\{\mathbf{u}_t\}_{t=1}^{T}$, UWB ranges $\{d_{t,a}\}_{t=1,a=1}^{T,A}$, mask $\{m_{t,a}\}_{t=1,a=1}^{T,A}$, AoI $\{\tau_{t,a}^{\text{age}}\}_{t=1,a=1}^{T,A}$, window length T
Require: (Training) ground-truth trajectory $\mathbf{P} = \{\mathbf{p}_t\}_{t=1}^{T+1}$ from RTK-GNSS, optimizer Opt, scheduler Sch
Ensure: (Training) best parameters $\theta^\star$; (Testing) reconstructed positions $\{\hat{\mathbf{p}}_t\}_{t=1}^{T+1}$

Training Phase
Split data into train/val/test; compute train-only normalization (μ, σ)
Initialize fusion network f_θ (LSTM) with parameters θ
for epoch $e = 1$ to E **do**
  Set model to train mode
  for each mini-batch of windows **do**
    Extract $(\mathbf{U}, \mathbf{D}, \mathbf{M}, \boldsymbol{\tau}, \mathbf{P})$
    Causal-fill and normalize $\mathbf{D}$
    $\Delta\hat{\mathbf{P}} \leftarrow f_\theta(\mathbf{U}, \mathbf{D}, \mathbf{M}, \boldsymbol{\tau})$
    $\hat{\mathbf{p}}_1 \leftarrow \mathbf{p}_1$
    for $t = 1$ to T **do**
      $\hat{\mathbf{p}}_{t+1} \leftarrow \hat{\mathbf{p}}_t + \Delta\hat{\mathbf{p}}_t$
    end for
    Compute total loss $\mathcal{L}$ and update θ via Opt
  end for
  Validate endpoint error; save best $\theta^\star$
  Update learning rate via Sch
end for

Testing Phase
Load $\theta^\star$ and set eval mode
for each test window **do**
  Extract $(\mathbf{U}, \mathbf{D}, \mathbf{M}, \boldsymbol{\tau})$
  Causal-fill and normalize $\mathbf{D}$
  $\Delta\hat{\mathbf{P}} \leftarrow f_{\theta^\star}(\mathbf{U}, \mathbf{D}, \mathbf{M}, \boldsymbol{\tau})$
  Set $\hat{\mathbf{p}}_1 \leftarrow \mathbf{0}$ ▷ relative trajectory within the window
  for $t = 1$ to T **do**
    $\hat{\mathbf{p}}_{t+1} \leftarrow \hat{\mathbf{p}}_t + \Delta\hat{\mathbf{p}}_t$
  end for
end for
return reconstructed trajectory $\{\hat{\mathbf{p}}_t\}_{t=1}^{T+1}$

2) UWB Representation with AoI-Based Reliability: In parallel, the network receives UWB ranging measurements to a set of fixed anchors indexed by $\mathcal{A} = \{1, \ldots, N_A\}$ where $N_A = |\mathcal{A}|$ is the number of deployed anchors. These measurements are arranged into a matrix

$$\mathbf{D} = \{d_{t,a}\} \in \mathbb{R}^{T \times N_A}, \quad a \in \mathcal{A}. \tag{21}$$

where $d_{t,a}$ is the measured distance between the mobile node and anchor a at time step t. However, UWB measurements may be missing or unreliable due to intermittent communication and NLOS conditions. To explicitly model UWB availability, a binary measurement mask

$$\mathbf{M} = \{m_{t,a}\} \in \{0,1\}^{T \times N_A}, \quad a \in \mathcal{A}. \tag{22}$$

is introduced, where $m_{t,a} = 1$ indicates the presence of valid UWB measurement and $m_{t,a} = 0$ denotes a missing or invalid measurement. This mask will later be combined with AoI decay to estimate the time-varying reliability of UWB input.

The proposed model also accounts for the temporal freshness of UWB rangings using AoI representation. In this context, AoI measures the elapsed time since the last available UWB range update from each anchor. For each anchor $a \in \mathcal{A}$ and time step t, the AoI $\tau_{t,a}^{\text{age}}$ can be defined as

$$\boldsymbol{\tau} = \{\tau_{t,a}^{\text{age}}\} \in \mathbb{R}_{+}^{T \times N_A}, \tag{23}$$

where $\mathbb{R}_+$ represents the set of non-negative real numbers. The AoI variable is defined recursively for each anchor $a \in \mathcal{A}$ where $m_{t,a} \in \{0,1\}$ denotes the UWB measurement availability indicator. The AoI evolves as

$$\tau_{t,a}^{\text{age}} = \begin{cases} 0, & \text{if } m_{t,a} = 1, \\ \tau_{t-1,a}^{\text{age}} + 1, & \text{if } m_{t,a} = 0, \end{cases} \tag{24}$$

with initialization $\tau_{0,a}^{\text{age}} = 0$. The resulting AoI tensor $\boldsymbol{\tau} = \{\tau_{t,a}^{\text{age}}\}$ provides an explicit, physically interpretable measure of temporal freshness, enabling the network to attenuate stale UWB observations during fusion.

Each anchor $a \in \mathcal{A}$ is assigned a learnable embedding vector $\mathbf{e}_a \in \mathbb{R}^E$, where E denotes the dimensionality of the anchor embedding space and it is set to $E = 3$, as this provided stable performance without gains from larger embeddings. These anchor embeddings allow the network to capture systematic anchor-dependent biases rather than explicitly encoding anchor coordinates or identities. To handle intermittent measurements, missing values in $d_{t,a}$ are causally filled by propagating the last valid observation, producing $\bar{d}_{t,a}$. The resulting signal is standardized using per-anchor statistics computed from the training data:

$$\tilde{d}_{t,a} = \frac{\bar{d}_{t,a} - \mu_a}{\sigma_a + \varepsilon}, \tag{25}$$

where μ_a and σ_a are the training-set mean and standard deviation for anchor a, and ε is a small constant for numerical stability.

The composite UWB feature is then formed as

$$\mathbf{x}_{t,a}^{\text{uwb}} = \left[\tilde{d}_{t,a},\ \exp\!\left(-\tau_{t,a}^{\text{age}}/\lambda_a\right),\ \mathbf{e}_a\right], \tag{26}$$

where $\tau_{t,a}^{\text{age}}$ is the observed age of the most recent measurement and λ_a is a learned anchor-specific decay constant controlling the influence of stale observations.

Then, a shared anchor-wise encoder $f_{\text{uwb}}(\cdot)$ is applied to each feature vector, producing a latent representation $\mathbf{h}_{t,a}^{\text{uwb}} \in \mathbb{R}^H$ as

$$\mathbf{h}_{t,a}^{\text{uwb}} = f_{\text{uwb}}(\mathbf{x}_{t,a}^{\text{uwb}}), \tag{27}$$

where H denotes the size of the learned latent representation used by the network internally. To account for measurement availability and temporal freshness, an AoI-weighted reliability coefficient is defined as

$$\tilde{m}_{t,a} = m_{t,a} \exp\!\left(-\frac{\tau_{t,a}^{\text{age}}}{\lambda_a}\right), \tag{28}$$

which models reliability decay as the measurement age increases. This exponential form is consistent with the forgetting

behavior used in recursive state estimation, providing continuous attenuation of stale observations while preserving recent measurements. These anchor-level features are then aggregated across anchors in a mask-weighted manner to form a per-step UWB feature representation.

$$\mathbf{h}_t^{\text{uwb}} = \frac{\sum_{a=1}^{N_A} \tilde{m}_{t,a} \mathbf{h}_{t,a}^{\text{uwb}}}{\sum_{a=1}^{N_A} \tilde{m}_{t,a} + \varepsilon} \tag{29}$$

where $\varepsilon > 0$ is a small constant added for numerical stability.

To construct a scalar proxy reflecting UWB measurements support and temporal freshness, `AoI-FusionNet` computes two complementary scalar quality indicators at each time step t. The first indicator reflects the instantaneous measurement availability q_t^{raw} as

$$q_t^{\text{raw}} = \frac{1}{N_A} \sum_{a=1}^{N_A} m_{t,a}, \tag{30}$$

where $m_{t,a} \in \{0, 1\}$ is the availability of the UWB measurement from anchor a at time t. The second indicator incorporates temporal freshness by considering age-dependent reliability q_t^{decay} as

$$q_t^{\text{decay}} = \frac{1}{N_A} \sum_{a=1}^{N_A} \tilde{m}_{t,a}, \tag{31}$$

The final scalar UWB quality score is obtained by combining both indicators

$$q_t = \frac{1}{2} q_t^{\text{raw}} + \frac{1}{2} q_t^{\text{decay}}. \tag{32}$$

The first term reflects how many anchors are currently visible, while the second reflects how recent their measurements are. Equal weighting is adopted to balance instantaneous availability and temporal freshness without introducing additional hyperparameters. The network subsequently learns how to exploit this reliability signal during training.

3) Reliability-Aware Cross-Modal Fusion: At each time step t, `AoI-FusionNet` constructs a fused feature representation by concatenating modality-specific embeddings and reliability cues $\mathbf{z}_t$ as

$$\mathbf{z}_t = \left[\mathbf{h}_t^{\text{imu}},\ \mathbf{h}_t^{\text{uwb}},\ q_t,\ q^{\text{prior}}\right], \tag{33}$$

where q^{prior} represents a time-invariant anchor-quality prior computed from training statistics and concatenated to $\mathbf{z}_t$ at each time step. The sequence $\mathbf{z}_t$ is processed by an LSTM model

$$\mathbf{h}_t^{\text{rnn}} = \text{LSTM}(\mathbf{z}_t). \tag{34}$$

To dynamically balance inertial and ranging information, `AoI-FusionNet` employs a learned attention gate that adapts relative sensor contribution under varying UWB availability, while maintaining stable motion estimation during sparse ranging conditions. The gating input $\mathbf{g}_t$ is defined as

$$\mathbf{g}_t = \left[\mathbf{h}_t^{\text{imu}},\ \mathbf{h}_t^{\text{uwb}},\ q_t\right], \tag{35}$$

The attention gate predicts a scalar IMU weight $\alpha_t \in [0, 1]$ using a multilayer perceptron followed by a sigmoid activation as

$$\alpha_t = \sigma(f_{\text{att}}(\mathbf{g}_t)), \tag{36}$$

where $\sigma(\cdot)$ denotes the sigmoid function. The attended fusion feature is then computed as

$$\mathbf{h}_t^{\text{att}} = \alpha_t \mathbf{h}_t^{\text{imu}} + (1 - \alpha_t) \mathbf{h}_t^{\text{uwb}}. \tag{37}$$

This creates a convex combination of inertial and ranging features, allowing the network to emphasize IMU information during UWB degradation and rely more on ranging when reliable measurements are available. A simple dominance constraint is applied to prevent degenerate behavior during complete UWB outages and does not otherwise restrict the learned fusion behavior

$$\alpha_t \geq \alpha_{\min} \quad \text{if } q_t^{\text{raw}} = 0, \tag{38}$$

where $\alpha_{\min}$ is a predefined lower bound.

4) Incremental Motion Prediction: The final prediction head estimates incremental motion updates rather than absolute positions, improving stability under sparse measurements. The predicted displacement $\Delta\hat{\mathbf{p}}_t$ at time t is given as

$$\Delta\hat{\mathbf{p}}_t = s \cdot \tanh\left(\frac{f_\Delta\left([\mathbf{h}_t^{\text{rnn}}, \mathbf{h}_t^{\text{att}}]\right)}{s}\right), \tag{39}$$

where s is a step-scale parameter derived from training statistics. The complete trajectory is constructed by cumulative integration of the predicted motion increments denoted by $\Delta\hat{\mathbf{P}} = \{\Delta\hat{\mathbf{p}}_t\}_{t=1}^{T}$.

B. Computational Complexity

`AoI-FusionNet` comprises 580 124 trainable parameters, requiring approximately 2.21 MByte of FP32 storage. On a single central processing unit (CPU) thread without GPU acceleration, one inference pass over a window of $T = 100$ time steps requires a mean latency of 2.85 ms (median 2.84 ms, $P95$ 2.91 ms), while the window represents 1 s of sensor data. Thus, the inference latency is approximately $351\times$ shorter than the input-window duration. This indicates that the model computation is sufficiently lightweight for practical window-based processing on a general-purpose CPU. Finally, the windowed architecture is intended for window-based rather than sample-by-sample streaming inference. The current system is designed for post-processing at an anchor node or central server rather than on the mobile node. However, this does not necessarily preclude online operation: the measured inference latency is substantially shorter than the one-second input window, indicating that windowed inference could be performed as measurements become available. Practical real-time deployment would require validation on target hardware, including communication, buffering, and energy constraints.

C. `AoI-FusionNet-DGAN`: Overfitting Avoidance

Given that our field campaign is limited to one trajectory, a key concern is that `AoI-FusionNet` may overfit to trajectory-specific UWB error characteristics such as anchor-dependent biases, multipath-induced heavy tails, and bursty outage patterns rather than learning transferable sensor fusion behavior. To probe this risk, a diffusion-based UWB residual augmentation strategy, referred to as `AoI-FusionNet-DGAN`, is

introduced that generates alternative, physically plausible measurement corruptions for the same underlying motion trajectory. This enables a controlled assessment of the model's sensitivity to measurement variability without altering the ground-truth motion. Specifically, we train a conditional diffusion model on UWB residuals, defined as the deviation between measured ranges and their geometric ground-truth values. During `AoI-FusionNet` training, we occasionally perturb real residuals as follows:

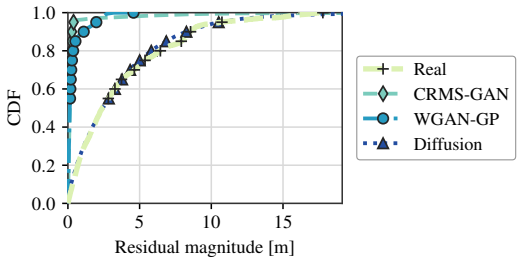
1) Compute real residuals $\epsilon_{t,a}^{\text{real}} = d_{t,a} - r_{t,a}^{\text{geom}}$ for valid measurements, where $r_{t,a}^{\text{geom}}$ refers to the ground-truth geometric range computed using ground-truth positions.
2) Sample synthetic residuals $\epsilon_{t,a}^{\text{fake}}$ from the diffusion model, conditioned on motion and anchor visibility patterns.
3) For a random subset (about 10%) of valid points, form a soft mixture

$$\epsilon_{t,a}^{\text{mix}} = (1 - \alpha_{\text{gan}})\, \epsilon_{t,a}^{\text{real}} + \alpha_{\text{gan}}\, \epsilon_{t,a}^{\text{fake}},$$

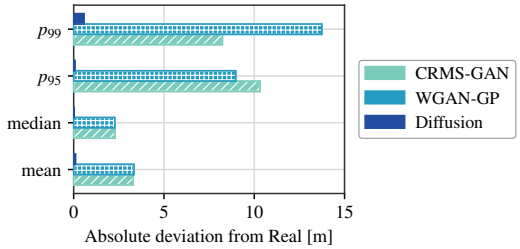
where $\alpha_{\text{gan}} \in [0, 1]$ controls the strength of the synthetic perturbation, and reconstruct augmented ranges $d_{t,a}^{\text{aug}} = r_{t,a}^{\text{geom}} + \epsilon_{t,a}^{\text{mix}}$.

This augmentation strategy preserves the original sparsity pattern of the real UWB data while injecting diversity into the noise structure. Consequently, `AoI-FusionNet` observes a broader family of realistic UWB degradations, reducing its sensitivity to any single noise fingerprint present in the original dataset. This augmentation does not replace evaluation using additional trajectories; rather, it provides evidence that `AoI-FusionNet` does not rely on a specific measurement realization to achieve good performance.

The proposed diffusion-based generator (DGAN) produces residual patterns that more accurately reflect the heavy-tailed, anchor-dependent nature of real UWB noise compared to the two baseline generative models we evaluated for data augmentation (CRMS-GAN and WGAN-GP). As shown in Fig. 5a, the DGAN-generated residuals exhibit a CDF that more closely matches the real data, with a residual distribution that concentrates more probability mass near zero, hence a steeper rise of the CDF curve. The deviation plot in Fig. 5b further quantifies the absolute difference between model-generated and real residual statistics, where the diffusion model shows the lowest deviation across evaluated residual statistics, particularly at high percentiles compared to the GAN baselines. These results indicate that DGAN better reproduces the statistical characteristics of real UWB residuals and therefore provides more realistic perturbations for data augmentation. The resulting augmentation improves performance of `AoI-FusionNet` under the evaluated sparsity patterns and UWB degradation regimes compared to training without data augmentation.



(a) Cumulative distribution function (CDFs) of UWB measurement residuals obtained from real data and different GAN models.



(b) Absolute deviation of key residual statistics (mean, median, $P95$, $P99$) from real-data values; smaller deviation indicates closer agreement.

Figure 5. Comparison of UWB residual distribution for different GAN models.

D. Loss Function and Training Strategy

The training objective of `AoI-FusionNet` is defined as a composite robust regression. It promotes accurate short-term motion estimation, long-term motion consistency, and robust sensor fusion under sparse or degraded UWB conditions. The objective also encourages temporally stable attention behavior without requiring explicit attention regularization. Training relies on Huber (Smooth L1) loss, which provides a balance between sensitivity to small errors and robustness to large deviations. For a scalar error e, the Huber loss is

$$\mathcal{L}_{\text{Huber}}(e) = \begin{cases} \frac{1}{2}e^2, & |e| \le \delta, \\ \delta\left(|e| - \frac{1}{2}\delta\right), & |e| > \delta, \end{cases} \tag{40}$$

where δ denotes a transition threshold. In our setting, $\mathcal{L}_{\text{Huber}}$ is applied element-wise to the 3D residuals between predicted and ground-truth incremental displacements and their cumulative positions, while a separate tail term operates on the Euclidean endpoint error. The loss behaves quadratically for small residuals and linearly for large errors, providing robustness to large outliers and intermittent measurement dropouts.

To avoid temporal leakage and preserve causality, data are split at the sequence level rather than at individual time steps. Sliding windows of length T are extracted from each trajectory, and exclusively assigned to training, validation, and testing set. Normalization statistics (μ, σ) are computed using only the training set and reused unchanged for validation and test data. Missing UWB measurements within a window are causally filled using the hold of the last-observation-carried-forward strategy. So, the missing value of UWB at t is replaced by the most recent valid measurement at $t' < t$ to ensure causal consistency. No information from future timesteps $t' > t$ is used during preprocessing.

The network is trained end-to-end using mini-batch stochastic optimization with an adaptive learning rate. Fixed-duration temporal segments sampled from continuous trajectories are employed for training. To ensure stable convergence, the recurrent backbone and age-aware decay parameters are first optimized during a brief warm-up stage before auxiliary regularization terms, such as reliability constraints and attention smoothness,

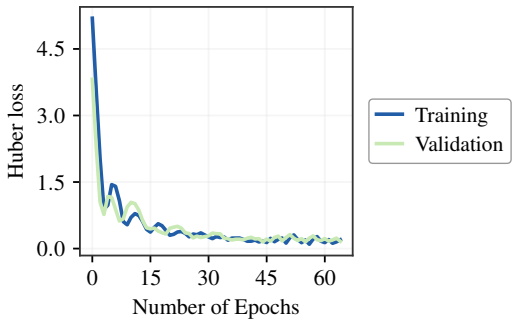


Figure 6. Training and validation endpoint loss as a function of training epochs. The close alignment of curves indicates stable optimization and minimal overfitting.

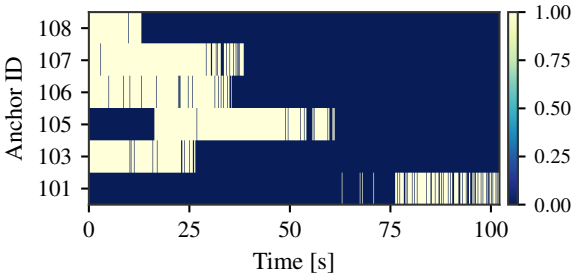


Figure 7. Temporal variation of UWB anchors illustrating intermittent anchor visibility along the experimental trajectory. Although six anchors are deployed, only a subset is visible at any time, with no more than four anchors simultaneously available.

are activated. An optional diffusion-based generative augmentation module is used during training to assess robustness under measurement variability. This module produces more realistic UWB range residuals conditioned on motion platform and anchor visibility and is applied only to the training set. Rather than replacing real measurements, the generative model injects small, physically plausible perturbations into valid UWB observations, encouraging the fusion network to generalize beyond the limited set of observed UWB patterns. To avoid destabilizing early training, a curriculum strategy gradually increases the probability of augmented windows across training epochs, ensuring that the network first learns from clean data before being exposed to synthetic perturbations. Model selection and early stopping are based on validation endpoint error, and all normalization parameters are derived solely from the training data to prevent information leakage.

Fig. 6 shows the evolution of training and validation loss over epochs during model training, computed using the Huber loss defined in Equation (40) The initial rapid decrease in loss during early epochs corresponds to efficient learning of motion dynamics, and the subsequent plateau indicates convergence before early stopping. The close agreement of the two curves suggests stable convergence with no signs of overfitting. This stability is further supported by stochastic residual augmentation, which empirically reduces sensitivity to trajectory-specific noise patterns.

V. Performance Evaluation

This section evaluates the performance of the proposed `AoI-FusionNet` under sparse and intermittent UWB anchor observability.

Fig. 7 displays the temporal availability of UWB anchors along the experimental trajectory. Multipath reflections, terrain occlusions, and varying LOS conditions cause fluctuations in anchor visibility over time, leading to extended intervals when fewer than four anchors are available. Under such intermittent-observability conditions, standalone 3D trilateration is frequently infeasible, and the position estimates are often affected by high GDOP. This highlights the limitation of trilateration-based and loosely coupled fusion methods, motivating tightly coupled ranging-inertial fusion.

The evaluation of the proposed model is performed on the test segment, and localization accuracy is examined using root mean square error (RMSE), mean absolute error (MAE), and percentile errors ($P50$, $P95$, $P99$), capturing both average performance and tail robustness. We evaluate `AoI-FusionNet` under two training configurations: (i) the standard variant trained on real measurements, and (ii) a version trained with diffusion-based augmentation referred to as `AoI-FusionNet-DGAN`. Both variants use the same training and inference pipeline; the DGAN variant differs only during training, where augmentation is applied and is not used during inference. Performance is compared against the three baselines used in the original evaluation, together with two additional tightly coupled methods selected from recent literature to strengthen the comparison with model-based and learning-based tightly coupled approaches.

- UWB-only Multilateration: Direct position estimation using available UWB ranges.
- LC-AKF: A loose-coupling-based adaptive Kalman filtering method where UWB trilaterated positions are fused with IMU acceleration.
- LC-Bi-LSTM: A data-driven loose-coupling approach that fuses IMU data with UWB-derived positions using a bidirectional LSTM.
- TC-MC-RUKF: A tightly coupled maximum-correntropy unscented Kalman filter (UKF) that directly incorporates raw UWB range measurements into the state estimation [31].
- TC-Att-LSTM: A tightly coupled learning-based approach that uses an attention-based LSTM to generate UWB pseudo-range measurements, which are subsequently incorporated into the navigation filter [42].

TC-MC-RUKF is implemented by extending the published 2-D formulation to 3-D, using the same bias-calibrated ranges and gravity-compensated acceleration as `AoI-FusionNet`. Its noise parameters are tuned on non-test data. TC-Att-LSTM combines an attention-based LSTM for UWB pseudo-range generation with a tightly coupled navigation filter. The method is extended to 3-D, using the available gravity-compensated acceleration, per-anchor ranges, measurement ages, and availability indicators as inputs. Both baselines and `AoI-FusionNet` are evaluated on the same test windows, with position initialized from the reference at the beginning of each window and velocity left uninitialized. The same evaluation metrics are used for all methods.

Table II
LOCALIZATION ERROR STATISTICS (METERS)

Method	RMSE	MAE	$P50$	$P95$	$P99$
UWB-only	1.954	1.742	1.645	3.186	3.949
LC-AKF	1.701	1.485	1.369	3.372	3.638
LC-Bi-LSTM	0.312	0.279	0.389	0.416	0.439
TC-MC-RUKF	1.485	1.076	0.631	3.284	3.748
TC-Att-LSTM	2.684	2.316	2.040	4.403	5.099
`AoI-FusionNet`	0.136	0.110	0.086	0.276	0.383
`AoI-FusionNet-DGAN`	**0.129**	**0.106**	**0.082**	**0.259**	**0.355**

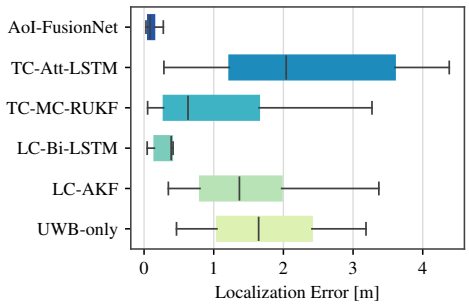


Figure 8. Distribution of absolute localization error for UWB-only, LC-AKF, LC-Bi-LSTM, TC-MC-RUKF, TC-Att-LSTM, and `AoI-FusionNet`. The boxes indicate the interquartile range (25th-75th percentiles), the central line denotes the median error, and the whiskers indicate the lower and upper whisker limits.

A. Comparison of Evaluation Results

Fig. 8 and Table II summarize the localization error statistics across all methods. Classical UWB-only multilateration and AKF-based loose coupling exhibit large errors and heavy tails, primarily due to unreliable geometry and intermittent UWB coverage. The LC-Bi-LSTM baseline considerably improves performance by learning temporal motion patterns; however, it remains sensitive to UWB outages because it relies on pre-computed UWB positions requiring at least four anchors for reliable 3D localization. The tightly coupled baselines provide additional context: TC-MC-RUKF and TC-Att-LSTM achieve RMSE values of 1.49 m and 2.68 m, respectively. `AoI-FusionNet` achieves the lowest RMSE and MAE, while also reducing tail errors ($P95$ and $P99$). RMSE drops from 0.31 m for LC-Bi-LSTM to 0.14 m for `AoI-FusionNet`. These results indicate that tightly coupling raw UWB ranges with IMU data through temporal modeling and explicit measurement-reliability weighting can provide more stable trajectory estimates under sparse and noisy ranging conditions. Incorporating diffusion-based augmentation further improves robustness, particularly in the upper error percentiles. For example, $P99$ error is reduced from 0.43 m LC-Bi-LSTM to 0.38 m and 0.35 m for `AoI-FusionNet` and `AoI-FusionNet-DGAN` respectively.

B. Error Distribution and Tail Robustness

We compare the absolute localization error of `AoI-FusionNet` with the multilateration and fusion baselines. Fig. 8 shows that our proposed method exhibits the lowest median error and the narrowest interquartile range, indicating higher precision and greater consistency compared to the baselines. UWB-only and LC-AKF show

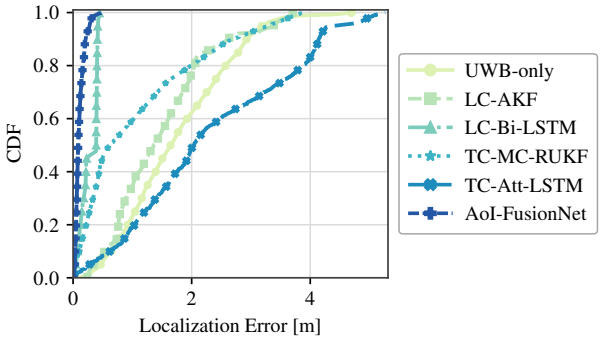


Figure 9. CDFs of absolute localization error for UWB-only, LC-AKF, LC-Bi-LSTM, TC-MC-RUKF, TC-Att-LSTM, and `AoI-FusionNet` indicating that `AoI-FusionNet` achieves higher probability mass at lower error thresholds.

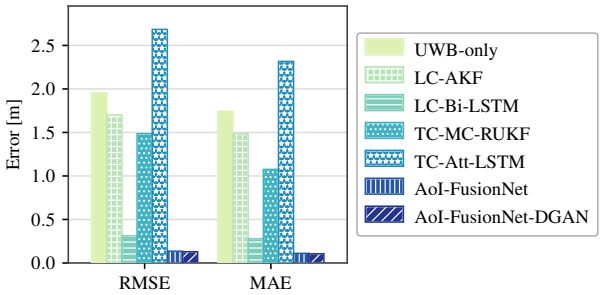


Figure 10. RMSE and MAE of 3D localization across all methods. `AoI-FusionNet` reduces RMSE from approximately 0.31 m (LC-Bi-LSTM) to about 0.14 m, while `AoI-FusionNet-DGAN` further improves performance to about 0.13 m. Similar trends are observed for MAE, confirming consistent gains in both average and squared error measures.

substantially larger interquartile ranges and long upper tails with wider spreads, reflecting increased error and variability under intermittent anchor visibility. The tightly coupled baselines show wider error distributions, with TC-MC-RUKF and TC-Att-LSTM exhibiting substantially larger tails than `AoI-FusionNet`, indicating that tight coupling alone does not ensure low-error or consistent localization under the considered conditions. Although LC-Bi-LSTM improves over traditional methods, its error distribution remains notably broader than `AoI-FusionNet`. Overall, the proposed method achieves reduced error dispersion across the evaluated conditions. Fig. 9 illustrates the distribution and cumulative behavior of localization errors. While multilateration and baseline fusion methods exhibit heavy tails, the CDF of `AoI-FusionNet` shows a steeper rise and reduced tail, highlighting improved robustness.

Fig. 10 and Fig. 11 summarize quantitative comparisons of the average and tail localization errors across all methods. The DGAN variant of the proposed model is illustrated here through these metrics since it acts as a training-time augmentation mechanism rather than a separate inference model. Fig. 10 reports RMSE and MAE, showing that the proposed model and its diffusion-augmented version achieve lower mean errors than baselines. In particular, `AoI-FusionNet-DGAN` achieves the lowest RMSE, indicating that diffusion-based augmentation has provided improvement without degrading the proposed model. Fig. 11 further examines robustness through percentile-errors at $P50$, $P95$, and $P99$ reflecting tail behavior. Baseline methods exhibit substantial degradation in high-percentile

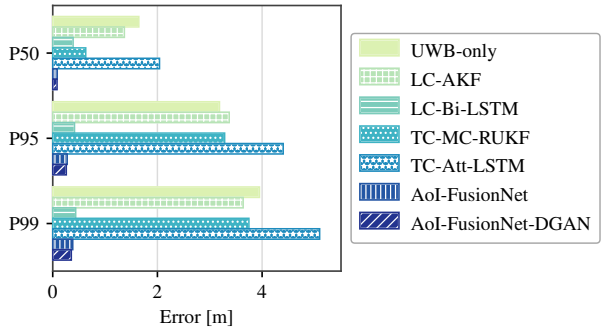


Figure 11. Median ($P50$), 95th percentile ($P95$), and 99th percentile ($P99$) 3D localization errors (Euclidean norm). `AoI-FusionNet` and `AoI-FusionNet-DGAN` exhibit reduced tail errors compared to all baselines, indicating improved robustness under adverse sensing conditions.

regimes, indicating sensitivity to UWB intermittent availability and measurement unreliability. However, `AoI-FusionNet` and `AoI-FusionNet-DGAN` maintain lower high-percentile errors, demonstrating resilience to degrading ranging conditions.

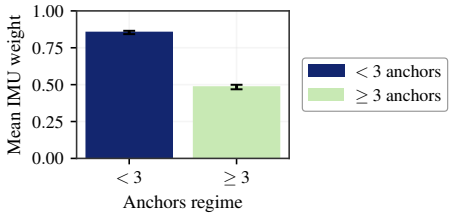
C. Ablation of Attention (ATT) and AoI Modules

To evaluate the individual contribution of each architectural module independently, we have performed an ablation study summarized in Table III. The performance is reported in mean error metrics RMSE, MAE, and tail metrics ($P95$, $P99$). The table shows that removing either the attention gate module or the AoI-aware decay mechanism results in decreased mean localization and higher percentile errors. Overall, the primary impact of the AoI-aware decay and attention gate is observed in the tail metrics rather than only in the mean errors, suggesting that these modules contribute most to preventing severe localization failures under intermittent measurement availability.

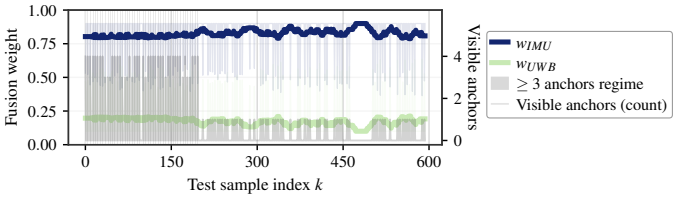
Finally, we examine the behavior of the learned fusion gate. As shown in Fig. 12a, the model assigns higher weight to IMU data under sparse conditions (< 3 visible UWB anchors), indicating greater reliance on inertial propagation when anchor geometry is unreliable and less weight when sufficient anchors are visible. Fig. 12b shows the temporal evolution of the fusion gate over the test sequence. The w_{IMU} and w_{UWB} refer to the learned fusion weights assigned to inertial and ranging measurements at each time step. The vertical spikes represent the instantaneous number of visible UWB anchors, providing a direct illustration of UWB measurement availability over time. The evolution of weights remains smooth with changes in anchor availability, confirming gradual adaptation rather than rapid temporal switching. This behavior demonstrates that

Table III
ABLATION RESULTS (LOCALIZATION ERRORS IN METERS) FOR ATTENTION (ATT) AND AOI MODULES USING RAW UWB RANGES AND IMU DATA.

ATT	AoI	RMSE	MAE	$P95$	$P99$
✓	✓	0.136	0.110	0.276	0.383
✓	✗	0.165	0.140	0.320	0.413
✗	✓	0.223	0.185	0.438	0.524
✗	✗	0.189	0.160	0.342	0.465



(a) Mean IMU attention weight conditioned by anchor regime



(b) Temporal evolution of fusion gate

Figure 12. Behavior of the learned fusion gate. (a) Under sparse UWB conditions (< 3 visible anchors), the estimator operates primarily in the inertial propagation regime, resulting in high w_{IMU}, while greater anchor availability leads to reduced IMU weighting. (b) Time-resolved attention weights w_{IMU}, w_{UWB} along with the number of visible anchors over the test sequence. The predominance of w_{IMU} reflects the intermittent and limited ranging geometry in this scenario, with UWB acting primarily as a corrective constraint.

the gating mechanism responds in a physically meaningful manner to measurement availability while preserving stable fusion dynamics.

D. Evaluation Across Trajectories and Ranging Regimes

To assess the reported behavior of `AoI-FusionNet` beyond the measurement conditions of the recorded campaign, `AoI-FusionNet` is additionally evaluated on synthetic trajectories that vary trajectory geometry, motion profile, and UWB ranging availability while preserving the physical characteristics of the deployment. The synthetic traces are partitioned at the trajectory level into training, validation, and test sets, ensuring that trajectories used for testing are not observed during training. The three ranging regimes are pooled during training, and performance is reported separately for each regime. These results are obtained from synthetic data and are therefore not directly comparable in absolute terms with the recorded-trajectory results of Table II; the baseline comparison in Table II is unaffected.

As shown in Table IV, the RMSE increases from 0.299 ± 0.017 m in the Dense regime to 0.351 ± 0.016 m in the Sparse regime and 0.385 ± 0.014 m in the Degraded regime. This corresponds to a 29% increase in RMSE across a $25\times$ reduction

Table IV
LOCALIZATION PERFORMANCE OF `AoI-FusionNet` ACROSS THREE *PhySynth* RANGING REGIMES.

Regime	**≥2 anchors**	**RMSE (m)**	**MAE (m)**	**P95 (m)**
Dense	12.7%	0.299 ± 0.017	0.216	0.623
Sparse	2.4%	0.351 ± 0.016	0.247	0.675
Degraded	0.5%	0.385 ± 0.014	0.270	0.733
Overall	–	$\mathbf{0.347 \pm 0.015}$	0.244	0.678

Table V
EFFECT OF THE NUMBER OF ACTIVE ANCHORS, GROUPED BY WHETHER FEWER THAN THREE OR AT LEAST THREE ANCHORS ARE ACTIVE.

Active-anchor group	$\geq$2 avail.	RMSE (m)	MAE (m)	P95 (m)
$n < 3$	2.1%	0.419 ± 0.017	0.291	0.794
$n \geq 3$	8.8%	0.391 ± 0.004	0.286	0.785

in the fraction of samples with at least two visible anchors. The results show a gradual increase in localization error as UWB ranging availability decreases, while the RMSE remains below $0.4\,\mathrm{m}$ across all three evaluated regimes. These results show that localization error increases as UWB ranging availability decreases, while the evaluation also exposes the model to variations in trajectory geometry and motion conditions. The evaluation therefore provides a controlled assessment of the behavior of `AoI-FusionNet` under a broader range of conditions than the recorded campaign. The synthetic evaluation therefore extends the range of evaluated trajectory and ranging conditions, while independent validation across deployment environments and anchor configurations remains outside the scope of this study.

In addition to ranging availability, the sensitivity to the number of active anchors is examined using the same synthetic evaluation framework. The model trained on the full deployment is evaluated without retraining on traces in which only a subset of the installed anchors is active. The configurations span one to four active anchors, corresponding to the maximum number observed within a single ranging burst in the recorded deployment. The active anchors are selected in descending order of their recorded ranging availability. For compactness, the results are grouped into configurations with fewer than three and at least three active anchors, as shown in Table V.

The configurations with at least three active anchors achieve an RMSE of $0.391 \pm 0.004\,\mathrm{m}$ compared with $0.419 \pm 0.017\,\mathrm{m}$ for configurations with fewer than three active anchors. This indicates a modest reduction in localization error when the number of active anchors increases beyond two under the considered synthetic conditions.

E. Discussion

`AoI-FusionNet` achieves lower localization errors than the evaluated baselines in the considered GNSS-denied outdoor scenario. The improvements are particularly evident at high percentiles, indicating reduced sensitivity to intermittent UWB observability in the evaluated dataset. The empirical analysis of anchor availability and AoI suggests that the informative UWB measurements primarily occur in short intermittent bursts, and these measurements are generally fresh when available. In this regime, the proposed AoI-aware decay mechanism provides a freshness prior that attenuates the contribution of stale measurements when fresh observations are unavailable. The learned fusion gate shows smooth and adaptive weighting of inertial and UWB cues, exploiting informative anchor geometry when available. When UWB measurements are insufficient, the model places greater reliance on learned inertial motion, while reliable UWB measurements receive greater weight when available, helping to correct accumulated inertial drift. An LSTM backbone is employed to model temporal dependencies while maintaining a relatively lightweight architecture. Due to the single-layer LSTM design and anchor-wise shared encoders, the model is computationally lightweight, making it suitable for offline post-processing. The AoI-aware decay and attention gate are formulated in terms of per-source availability and freshness, providing a basis for incorporating additional sensor modalities. Such an extension would, however, require modifications to the input representation and access to synchronized multi-modal training data.

To examine the model performance under variations in measurement noise, a diffusion-based residual augmentation technique has been used to produce physically plausible UWB noise patterns for the same underlying motion trajectory. While DGAN-based augmentation provides a controlled assessment under variations in UWB measurement noise, it does not substitute for evaluation using independent trajectories or different physical environments. The evaluation across synthetic trajectories further examines the behavior of `AoI-FusionNet` under variations in trajectory geometry, motion profile, and UWB ranging availability. The results show a gradual increase in localization error as ranging availability decreases, with RMSE remaining below $0.4\,\mathrm{m}$ across the evaluated regimes. The additional active-anchor analysis indicates a modest reduction in error when moving from fewer than three to at least three active anchors under the considered synthetic conditions.

Although performance improved in the evaluated scenario, several limitations remain. The real-world evaluation is based on one recorded trajectory and a fixed six-anchor deployment while synthetic measurements are generated by the *PhySynth* model calibrated from the recorded campaign. Consequently, systematic evaluation across different sites, including unseen motion dynamics, channel and multipath conditions, anchor geometries, and deployment environments, remains necessary to assess generalization. The cable-car experiment provides a controlled proxy for avalanche-particle tracking, but real avalanche conditions may involve substantially more chaotic motion, including abrupt accelerations and rotations that could potentially lead to IMU saturation. Dedicated field validation is therefore required to assess performance under such motion regimes.

Extending `AoI-FusionNet` to additional sensor modalities is a potential direction for future work, and factor-graph formulations for multi-modal fusion [50] illustrate how heterogeneous measurements can be combined within a unified framework. Relative ranges between mobile nodes could similarly be incorporated as additional ranging observations, together with their availability and AoI, allowing the same reliability-aware fusion mechanism to exploit cooperative constraints when static-anchor visibility is reduced. This direction will be investigated in future work as part of cooperative multi-node trajectory reconstruction.

Despite these limitations, the evaluated results indicate that the tightly coupled, AoI-aware fusion model can maintain localization performance under intermittent UWB sensing without relying on heuristic measurement filtering.

VI. Conclusion

This paper presented an adaptive tightly-coupled UWB-IMU fusion method, `AoI-FusionNet`, for 3D localization in GNSS-denied environments with sparse and intermittent UWB range measurements. The method integrates temporal motion dynamics, age-aware reliability modeling of UWB measurements using AoI, and adaptive sensor weighting through learned attention gate to achieve lower localization error, particularly in high-percentile error metrics, than multilateration, Kalman filtering, and deep learning-based loose coupling fusion methods, as well as the evaluated tightly coupled baselines. The evaluation conducted during offline post-processing of real-world data collected in a challenging alpine environment indicates that the proposed model maintains sub-meter localization error under the evaluated intermittent UWB conditions. The ablation results further demonstrate that AoI-aware decay and attention-based fusion provide complementary roles, with the largest gains observed in high-percentile error metrics. The evaluation is based on one recorded trajectory, and a synthetic evaluation using the same six-anchor deployment and measurement characteristics calibrated from that campaign. Cross-environment, cross-deployment, and embedded-device validation therefore remain important directions for future work.

Future work will focus on multi-node and cooperative localization, cross-environment validation, real avalanche-field evaluation under more complex motion dynamics, and explicit CIR-based LOS/NLOS modeling. Joint LOS/NLOS classification could further improve robustness under severe multipath conditions. Potential applications include avalanche-particle tracking, snow-glide monitoring, and other mobile trajectory reconstruction tasks in GNSS-denied environments. Together, these directions pave the way for a more general and transferable learning-based tracking model suitable for diverse GNSS-denied scenarios.

References

[1] M. Neuhauser, A. Köhler, R. Neurauter, M. S. Adams, and J.-T. Fischer, "Particle trajectories, velocities, accelerations and rotation rates in snow avalanches," *Annals of Glaciology*, pp. 1–18, Oct. 2023.

[2] S. Wang and N. S. Ahmad, "A Comprehensive Review on Sensor Fusion Techniques for Localization of a Dynamic Target in GPS-Denied Environments," *IEEE Access*, pp. 2252–2285, Dec. 2024.

[3] H. Arslan, Z. N. Chen, and M.-G. Di Benedetto, *Ultra Wideband Wireless Communication*. Jersey City, NJ: John Wiley & Sons, Inc., 2006, p. 528.

[4] X. Yang, J. Wang, D. Song, B. Feng, and H. Ye, "A Novel NLOS Error Compensation Method Based IMU for UWB Indoor Positioning System," *IEEE Sensors Journal*, vol. 21, no. 9, pp. 11 203–11 212, May 2021.

[5] L. Cheng, F. Zhao, P. Zhao, and J. Guan, "UWB/INS Fusion Positioning Algorithm Based on Generalized Probability Data Association for Indoor Vehicle," *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 1, pp. 446–458, Jan. 2024.

[6] Y. Zhang, "A Fusion Methodology to Bridge GPS Outages for INS/GPS Integrated Navigation System," *IEEE Access*, pp. 61 296–61 306, Apr. 2019.

[7] S. Chen, M. Xin, F. Yang, X. Zhang, J. Liu, G. Ren, and S. Kong, "Error Compensation Method of GNSS/INS Integrated Navigation System Based on AT-LSTM During GNSS Outages," *IEEE Sensors Journal*, vol. 24, no. 12, pp. 20 188–20 199, Jun. 2024.

[8] Z. Chen, S. Jin, G. Bian, X. Sui, and A. Xu, "A Robust, Real-Time, and Drift-Reduced LiDAR–Inertial–UWB Tightly Coupled Fusion," *IEEE Sensors Journal*, vol. 26, no. 1, pp. 477–486, Jan. 2026.

[9] Y. Yu, Y. Fan, G. Lai, C. Shi, F. Sun, and X. Meng, "Tight Fusion of Odometry, Kinematic Constraints, and UWB Ranging Systems for State Estimation of Integrated Aerial Platforms," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 18 297–18 313, 2025.

[10] J. Kuß, A. Köhler, M. Neuhauser, J.-T. Fischer, R. Neurauter, J. Gerstmayr, and F. Dressler, "A Measurement System for Distributed UWB-based Ranging and Localization in Snow Avalanches," in *29th ACM International Conference on Mobile Computing and Networking (MobiCom 2023), Poster Session*, Madrid, Spain: ACM, Oct. 2023.

[11] J. Kuß, A. Köhler, M. Neuhauser, R. Neurauter, J. Gerstmayr, J.-T. Fischer, and F. Dressler, "Distributed UWB-based Ranging for Particle Tracking in Avalanches," in *19th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2024)*, Chamonix, France: IEEE, Jan. 2024, pp. 125–132.

[12] Z. Xiao, H. Wen, A. Markham, N. Trigoni, P. Blunsom, and J. Frolik, "Non-Line-of-Sight Identification and Mitigation Using Received Signal Strength," *IEEE Transactions on Wireless Communications*, vol. 14, no. 3, pp. 1689–1702, Mar. 2015.

[13] L. Zhang, C. Wang, M. Ma, and D. Zhang, "WiDIGR: Direction-Independent Gait Recognition System Using Commercial Wi-Fi Devices," *IEEE Internet of Things Journal*, vol. 7, no. 2, pp. 1178–1191, Feb. 2020.

[14] J. Zhang, J. Salmi, and E.-S. Lohan, "Analysis of Kurtosis-Based LOS/N-LOS Identification Using Indoor MIMO Channel Measurement," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 6, pp. 2871–2874, Jul. 2013.

[15] N. Dwek, M. Birem, K. Geebelen, E. Hostens, A. Mishra, J. Steckel, and R. Yudanto, "Improving the Accuracy and Robustness of Ultra-Wideband Localization Through Sensor Fusion and Outlier Detection," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 32–39, Jan. 2020.

[16] J. Wang, C. Shi, F. Zheng, C. Yang, X. Liu, S. Liu, M. Xia, G. Jing, T. Li, W. Cheng, Q. Li, Y. Hu, Y. Tian, and Y. Shan, "Multi-frequency smartphone positioning performance evaluation: insights into A-GNSS PPP-B2b services and beyond," *Satellite Navigation*, vol. 5, no. 1, Aug. 2024.

[17] J. Zhu and S. S. Kia, "Bias Compensation for UWB Ranging for Pedestrian Geolocation Applications," *IEEE Sensors Letters*, vol. 3, no. 9, pp. 1–4, Sep. 2019.

[18] D. Feng, C. Wang, C. He, Y. Zhuang, and X.-G. Xia, "Kalman-Filter-Based Integration of IMU and UWB for High-Accuracy Indoor Positioning and Navigation," *IEEE Internet of Things Journal*, vol. 7, no. 4, pp. 3133–3146, Apr. 2020.

[19] R. Ali, R. Liu, A. Nayyar, B. Qureshi, and Z. Cao, "Tightly Coupling Fusion of UWB Ranging and IMU Pedestrian Dead Reckoning for Indoor Localization," *IEEE Access*, vol. 9, pp. 164 206–164 222, Dec. 2021.

[20] H. A. Hashim, L. J. Brown, and K. McIsaac, "Nonlinear Stochastic Attitude Filters on the Special Orthogonal Group 3: Ito and Stratonovich," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 49, no. 9, pp. 1853–1865, Sep. 2019.

[21] M. Shalihan, Z. Cao, K. Pongsirijinda, B. K. Kiat Ng, B. P. L. Lau, R. Liu, C. Yuen, and U.-X. Tan, "Localization through mitigating and compensating UWB NLOS ranging error with neural network," *Elsevier Digital Signal Processing*, vol. 166, p. 105 397, Nov. 2025.

[22] G. Feng, C. Shen, C. Long, and F. Dong, "GDOP index in UWB indoor location system experiment," in *IEEE SENSORS 2015*, Busan, South Korea: IEEE, Nov. 2015, pp. 1–4.

[23] A. Poulose, Z. Emersic, O. Steven Eyobu, and D. S. Han, "An Accurate Indoor User Position Estimator For Multiple Anchor UWB Localization," in *11th International Conference on Information and Communication Technology Convergence (ICTC 2020)*, Jeju Island, South Korea: IEEE, Oct. 2020, pp. 478–482.

[24] B. Van Herbruggen, D. Van Leemput, J. Van Landschoot, and E. De Poorter, "Real-Time Anchor Node Selection for Two-Way-Ranging (TWR) Ultra-Wideband (UWB) Indoor Positioning Systems," *IEEE Sensors Letters*, vol. 8, no. 3, pp. 1–4, Feb. 2024.

[25] S. Tiwari and V. K. Jain, "A novel step detection technique for pedestrian dead reckoning based navigation," *ICT Express*, vol. 9, no. 1, pp. 16–21, Feb. 2023.

[26] F. Liu, K. Liu, X. Guo, G. Chen, P. Zhou, and J. Yang, "Pedestrian Heading Estimation Based on F-LSTM Neural Network," in *16th International Conference on Electronic Measurement & Instruments (ICEMI 2023)*, vol. 1, Harbin, China: IEEE, Aug. 2023, pp. 417–422.

[27] R. Neurauter, A. Holzinger, M. Neuhauser, J.-T. Fischer, and J. Gerstmayr, "Motion Reconstruction of Fast-rotating Rigid Bodies," *Journal of*

Computational and Nonlinear Dynamics, vol. 19, no. 1, p. 011 005, Jan. 2024.

[28] A. Beauvisage, K. Ahiska, and N. Aouf, "Robust Multispectral Visual-Inertial Navigation With Visual Odometry Failure Recovery," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 9089–9101, Jul. 2022.

[29] Q. Zou, Q. Sun, L. Chen, B. Nie, and Q. Li, "A Comparative Analysis of LiDAR SLAM-Based Indoor Navigation for Autonomous Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 6907–6921, Jul. 2022.

[30] S. Guan and X. Luo, "Fusing Ultra-wideband Range Measurements with IMU for Mobile Robot Localization," in *11th International Conference on Intelligent Control and Information Processing (ICICIP 2021)*, Dali, China: IEEE, Dec. 2021, pp. 107–111.

[31] X. Li, J. Ye, Z. Zhang, F. Liang, W. Xia, and B. Wang, "A Novel Low-Cost UWB/IMU Positioning Method With the Robust Unscented Kalman Filter Based on Maximum Correntropy," *IEEE Sensors Journal*, vol. 24, no. 18, pp. 29 219–29 231, Sep. 2024.

[32] Y. Xu, Z. Chen, M. Zhao, F. Tang, Y. Li, J. Liu, and N. Kato, "UVtrack: Multi-Modal Indoor Seamless Localization Using Ultra-Wideband Communication and Vision Sensors," *IEEE Open Journal of the Computer Society*, vol. 6, pp. 272–281, 2025.

[33] C. Hu, P. Huang, and W. Wang, "Tightly Coupled Visual-Inertial-UWB Indoor Localization System With Multiple Position-Unknown Anchors," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 351–358, Jan. 2024.

[34] Y. Xu, Y. Li, C. K. Ahn, and X. Chen, "Seamless indoor pedestrian tracking by fusing INS and UWB measurements via LS-SVM assisted UFIR filter," *Neurocomputing*, vol. 388, pp. 301–308, May 2020.

[35] M. Fan, J. Li, and W. Wang, "An IMU/UWB tightly coupled navigation algorithm to improve positioning accuracy under large-scale NLOS conditions," *Measurement Science and Technology*, vol. 36, no. 4, Mar. 2025.

[36] T. Li, Z. Deng, Y. Zhang, W. Dong, and H. Yu, "A Tightly Coupled UWB/PDR Fusion Positioning Algorithm for Indoor Environments," *IEEE Transactions on Instrumentation and Measurement*, vol. 73, pp. 1–12, Sep. 2024.

[37] M. Tommingas, T. Laadung, S. Varbla, I. Müürsepp, and M. Mahtab Alam, "UWB and GNSS Sensor Fusion Using ML-Based Positioning Uncertainty Estimation," *IEEE Open Journal of the Communications Society*, vol. 6, pp. 2177–2189, 2025.

[38] Z. Zhi, D. Liu, and L. Liu, "A performance compensation method for GPS/INS integrated navigation system based on CNN-LSTM during GPS outages," *Measurement*, vol. 188, Jan. 2022.

[39] K. Muthineni, A. Artemenko, J. Vidal, M. Nájar, M. Catalan, and J. Paradells, "Deep Learning-Based UWB-IMU Data Fusion for Indoor Positioning in Industrial Scenario," *IEEE Open Journal of Vehicular Technology*, vol. 6, pp. 1209–1221, May 2025.

[40] J. Fontaine, F. Che, A. Shahid, B. Van Herbruggen, Q. Z. Ahmed, W. Bin Abbas, and E. De Poorter, "Transfer Learning for UWB Error Correction and (N)LOS Classification in Multiple Environments," *IEEE Internet of Things Journal*, vol. 11, no. 3, pp. 4085–4101, Feb. 2024.

[41] H. Yang, Y. Wang, C. K. Seow, Z. Li, M. Sun, C. De Cock, J. Bi, W. Joseph, and D. Plets, "NLOS Identification and Ranging Trustworthiness for Indoor Positioning With LLM-Based UWB–IMU Fusion," *IEEE Transactions on Instrumentation and Measurement*, vol. 74, pp. 1–17, 2025.

[42] M. Ren, J. Wei, J. Qin, X. Guo, H. Wang, and S. Li, "Attention based LSTM framework for robust UWB and INS integration in NLOS environments," *Scientific Reports*, vol. 15, no. 1, 2025.

[43] D. H. Tran, B. Chung, and Y. M. Jang, "GAN-based Data Augmentation for UWB NLOS Identification Using Machine Learning," in *International Conference on Artificial Intelligence in Information and Communication (ICAIIC 2022)*, Jeju Island, South Korea: IEEE, Feb. 2022, pp. 417–420.

[44] R. Neurauter, P. Hergel, and J. Gerstmayr, "Evaluation of Inertial Measurement Units for Short Time Motion Tracking," in *17th International Conference on Multibody Systems, Nonlinear Dynamics, and Control (MSNDC 2021)*, Virtual Conference: ASME, Aug. 2021.

[45] T. Bibi and F. Dressler, "PhySynth: A Physics-Based Synthetic UWB/IMU Data Generator for Training of ML-based Tracking," in *IEEE International Conference on Future and Intelligent Networking (FINE 2026)*, Osaka, Japan: IEEE, Aug. 2026.

[46] D. Neirynck, E. Luk, and M. McLaughlin, "An alternative double-sided two-way ranging method," in *13th Workshop on Positioning, Navigation and Communications (WPNC 2016)*, Bremen, Germany: IEEE, Oct. 2016.

[47] R. Neurauter and J. Gerstmayr, "A novel motion-reconstruction method for inertial sensors with constraints," *Multibody System Dynamics*, vol. 57, no. 2, pp. 181–209, Dec. 2022.

[48] R. Winkler, M. Neuhauser, R. Neurauter, F. Erlacher, W. Steinkogler, and J.-T. Fischer, "Particle tracking in snow avalanches with in situ calibrated inertial measurement units," *Annals of Glaciology*, vol. 65, e14, Jan. 2024.

[49] T. Bibi, A. Köhler, J.-T. Fischer, and F. Dressler, "Deep Bidirectional LSTM-Based Sensor Fusion for 3D Localization in Dynamic Environments," in *22nd IEEE International Conference on Mobile Ad-Hoc and Smart Systems (MASS 2025)*, Chicago, IL: IEEE, Oct. 2025, pp. 1–9.

[50] J. Wang, K. Li, M. Xia, and C. Shi, "Enhancing Positioning Accuracy of Wrist-Worn Instrumentation: A Factor Graph-Based Fusion of GNSS and Inertial Measurements," *IEEE Transactions on Instrumentation and Measurement*, vol. 75, pp. 1–13, Apr. 2026.