PhySynth: A Physics-Based Synthetic UWB/IMU Data Generator for Training of ML-based Tracking

Tehmina Bibi and Falko Dressler
School of Electrical Engineering and Computer Science, TU Berlin, Germany
{bibi,dressler}@tkn.tu-berlin.de

***Abstract*—Deep learning-based localization and tracking systems require large amounts of labeled training data, which are often difficult to acquire in challenging outdoor environments such as avalanche monitoring. To address this limitation, we present PhySynth, a physics-based synthetic data generator for synchronized ultra-wideband (UWB) ranging and inertial measurement unit (IMU) measurements. PhySynth is calibrated using real-world data collected during a cable-car experiment and models trajectory dynamics, sensor noise, measurement availability, and age of information effects within a unified framework. The generated traces preserve the physical characteristics of the underlying localization system while enabling systematic variation of UWB anchor deployments and channel conditions. Validation against independent descending and ascending trajectories demonstrates close agreement between real and synthetic measurements, achieving Kolmogorov–Smirnov statistics of 0.147 and 0.057, respectively. A comprehensive study further reveals that channel conditions have a greater impact on data fidelity than anchor density, with diminishing returns observed beyond three anchors. Finally, we evaluate the usefulness of the generated data through a case study using AoI-FusionNet, a deep learning-based tracking framework. The results show that PhySynth provides a practical and reproducible approach for generating realistic multimodal datasets to support the development, training, and benchmarking of machine learning-based localization and tracking systems.**

***Index Terms*—Synthetic data generation, deep learning, localization, tracking, ultra-wideband, inertial measurement unit**

I. Introduction

Sensor fusion using ultra-wideband (UWB) and inertial measurement unit (IMU) sensors has emerged as an effective approach for localization and tracking, particularly in environments where global navigation satellite system (GNSS) signals are unavailable or unreliable [1]. UWB technology provides accurate range measurements with low power consumption and strong robustness against multipath propagation, while IMUs deliver high-frequency motion information independent of external infrastructure. By combining the complementary strengths of both sensing modalities, UWB/IMU fusion systems can achieve accurate and robust position estimation, especially in dynamic outdoor settings [2].

Recent advances in machine learning and deep learning techniques, particularly recurrent neural network (RNN) and long short-term memory (LSTM) architectures [3], have demonstrated potential in sensor fusion methods for combining UWB and IMU data. By learning complex temporal relationships from multimodal data, these approaches improve the robustness of localization in the presence of noise, missing measurements, and dynamic environmental conditions. Nevertheless, their performance remains highly dependent on the availability of extensive and representative training datasets that accurately reflect real-world motion patterns, sensor imperfections, and environmental variability [4].

Data collection for outdoor tracking experiments is quite challenging [5], [6]. This, of course, also holds for a scenario we are working on: tracking of avalanches in Alpine environments [7], [8]. Our aim is studying particle tracking inside snow avalanches. Field experiments conducted in the Austrian Alps, using a cable-car counterweight to emulate avalanche-like motion, provide valuable real-world UWB and IMU measurements under controlled conditions. However, the amount of data available for the development, training, and evaluation of deep learning (DL)-based tracking algorithms is naturally quite limited.

To address this challenge, this paper presents `PhySynth`, a physics-based synthetic data generator for synchronized UWB and IMU measurements to support learning-based sensor fusion approaches. The proposed framework models trajectory dynamics, IMU noise characteristics, UWB ranging errors, anchor visibility, and age of information (AoI) effects within a unified simulation environment. Unlike purely data-driven augmentation approaches, the generated traces preserve the physical characteristics of the underlying localization system while allowing systematic variation of anchor deployments and channel conditions.

The proposed generator is validated using real AvaRange cable-car measurements collected during descending and ascending trajectories. In addition, the system is used to investigate the influence of anchor count, anchor geometry, and ranging conditions on the statistical properties of generated data. Finally, the usefulness of the generated traces is demonstrated through a case study based on our AoI-FusionNet algorithm [9], a deep learning framework for UWB/IMU localization.

Our main contributions can be summarized as follows:

- We propose a physics-based generator of synchronized UWB and IMU measurements for localization and tracking applications;
- we validated the generated data against real descending and ascending cable-car trajectories from the AvaRange project;
- we provide a systematic analysis of anchor deployment and channel-condition effects on synthetic data characteristics;
- and, we demonstrate the applicability of the generated data for training and evaluating AI-based localization methods.

II. Related Work

A. Kalman filter (KF)-based Sensor Fusion

Classical UWB/IMU fusion techniques are typically based on Bayesian filtering methods that combine short-term motion information from IMUs with absolute ranging measurements provided by UWB sensors. KF-based approaches, such as adaptive unscented Kalman filter (AUKF) [10], variational Bayesian UKF (VBUKF), and error-state Kalman filter (ESKF) [11], have demonstrated improved localization accuracy and robustness in GNSS-denied environments. Although these model-based methods are effective, they rely on predefined motion and noise models and often struggle in highly dynamic environments characterized by nonlinear motion, intermittent measurements, and non-Gaussian measurement errors.

B. AI-based Sensor Fusion

Data-driven models, in contrast, learn complex relationships directly from multimodal sensor measurements. Muthineni et al. [12] proposed a two-stage neural network for UWB/IMU fusion in industrial positioning scenarios to achieve sub-decimeter accuracy, followed by graph neural network (GNN) to handle varying anchor visibility under NLOS conditions. Similarly, Bibi et al. [13] proposed a bidirectional LSTM (Bi-LSTM) fusion algorithm to combine UWB and IMU sensor datasets in avalanche-like environments. These studies highlight the potential of deep learning methods in terms of handling the non-linearity of the sensor measurements while improving localization accuracy. These models also need extensive synchronized UWB and IMU measurements together with accurate ground-truth trajectories for training and validation. Therefore, a key limitation of modern deep learning models is their dependence on a large volume of labeled training datasets [14]. Meanwhile, acquiring large-scale datasets is often impractical in complex outdoor environments, such as snow-avalanche experiments and alpine terrains, where particle-level measurements require specialized instrumentation and field campaigns are constrained by challenging weather conditions, accessibility, and safety considerations [8].

C. Synthetic Data Generation

To overcome the challenge of data scarcity, several studies have explored data augmentation to generate synthetic data for different sensor applications. Diffusion-based models have been proposed for high-fidelity radar measurements for recognition and classification tasks. Lang et al. [15] augmented radar spectrograms for human motion recognition, improving diversity and classification performance in few-shot settings, and Yun et al. [16] combined synthetic minority over-sampling technique (SMOTE) with diffusion models to generate realistic IR-UWB radar signals for object classification. Variational autoencoder (VAE) [17] and generative adversarial network (GAN) [18] have been introduced to improve the classification performance under data limitation.

Despite their advantages for target classification problems, these augmentation strategies are data-driven approaches and designed to replicate the original data distribution and are less suitable for exploring scenarios beyond the observed data, since they learn statistical patterns rather than the underlying physical processes. This limitation motivates the need for a physics-based data generation framework. `PhySynth` combines motion dynamics, UWB ranging behavior, and IMU sensor characteristics within a unified simulation framework, allowing reproducible experimentation across diverse operating conditions while remaining statistically consistent with real field measurements.

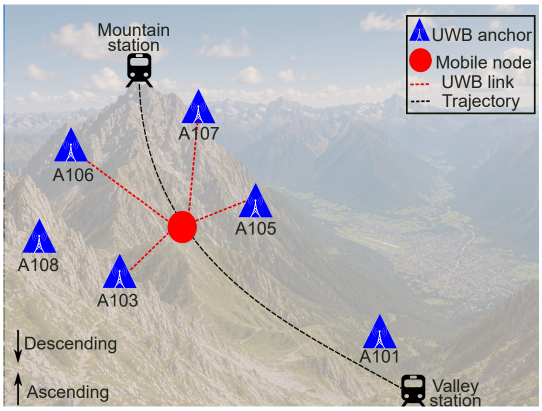

Figure 1. Depiction of Avarange experimental scenario containing a mobile node equipped with the sensor package (UWB, IMU) alongside EMLID GNSS mounted onto a cable car counterweight that travels along the Nordkettenbahn cable car test-site. Six static anchors (A101, A103, A105, A106, A107, A108) are deployed along the route.

III. `PhySynth`: Synthetic UWB/IMU Data Generator

A. Experimental Dataset

Experimental data from multimodal sensors have been collected during a cable-car experiment conducted on the Nordkettenbahn in Austria. The scenario as illustrated in Fig. 1 consists of a cable-car counterweight traveling 710.2 m between the mountain top and valley stations, with an elevation difference of 342 m. The mobile platform equipped with a UWB tag of burst rate $\sim$23 Hz, a 400 Hz IMU, and an EMLID Reach RS+ GNSS receiver at 10 Hz used as ground truth is mounted to this cable-car. Six static UWB anchors, referred to as A101, A103, A105, A106, A107, and A108, are deployed at known positions, surveyed prior to the experiment using GNSS [7]. On the 100 Hz evaluation grid, UWB measurements from at least one anchor are available in only 17.8% of timesteps, making the dataset both challenging and representative for synthetic data generation.

B. `PhySynth`

Our proposed framework `PhySynth` generates realistic synchronized UWB, IMU, and AoI traces by extracting key parameters from real cable-car measurements, including trajectory geometry, speed profile, IMU noise characteristics, UWB ranging statistics, and measurement availability patterns, and uses them as inputs to a physics-based simulation framework. Each synthetic trace is generated independently, where real measurements are not copied or interpolated. As illustrated

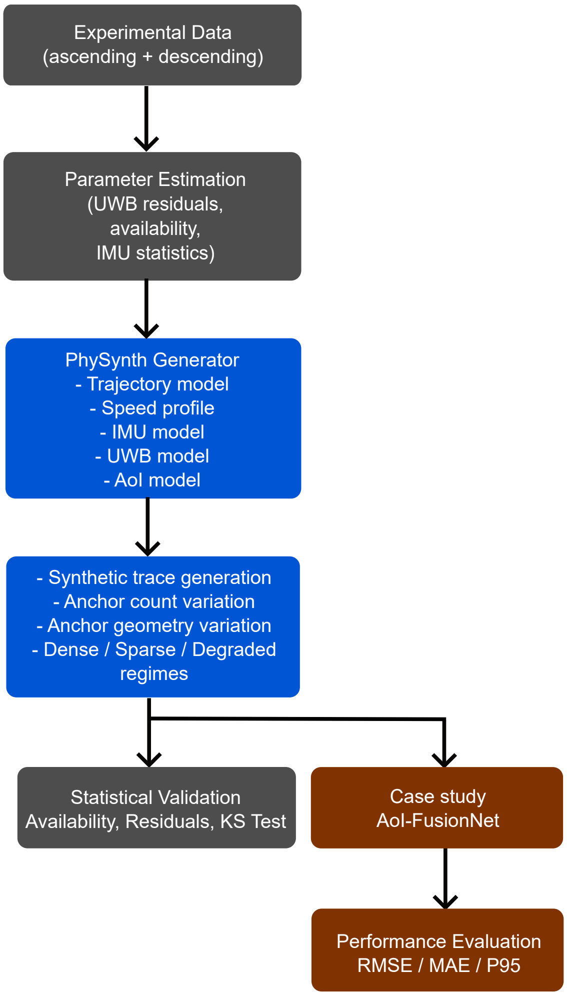

Figure 2. Overview of `PhySynth` generator using parameters calibrated from real cable-car measurements to generate synthetic UWB, IMU, availability, and AoI traces under different deployment scenarios. The generated data is subsequently used for statistical validation and localization benchmarking.

in Fig. 2, the framework reproduces the statistical characteristics of the real data while enabling controlled variation of anchor count, placement, and UWB propagation conditions. `PhySynth` consists of five main components: a trajectory model, a motion model, an IMU model, a UWB ranging model, and an availability/AoI model described in the following subsections.

C. Trajectory Model

The geometry of the cable-car is modeled by a degree-10 polynomial fitted to the measured EMLID ground-truth positions, resulting in an RMSE of 0.087 m over the trajectory length 710.2 m. The trajectory is parameterized by the arc length $s \in [0, L]$ and mapped to ENU coordinates as

$$\mathbf{p}(s) = \big(E(s), N(s), U(s)\big). \tag{1}$$

To capture realistic run-to-run variability, each synthetic trajectory is generated by perturbing the nominal catenary with random lateral offsets ($\sigma = 2\,\mathrm{m}$) and vertical sag variations ($\sigma = 1\,\mathrm{m}$), reflecting natural cable deformation between runs. The resulting ensemble of trajectories is physically plausible, while spanning a realistic range of cable positions.

D. Motion Model

The motion of the cable-car is represented by an acceleration-cruise-deceleration profile using a degree-4 polynomial fitted to the real ground-truth speed trace. Each synthetic trace independently perturbs the cruise speed (±12 %, uniform) and total path length (±5 %, uniform) to reproduce run-to-run variability. The speed profile is smoothed prior to numerical differentiation to ensure smooth acceleration signals suitable for subsequent IMU simulation.

E. IMU Measurement Model

Synthetic IMU acceleration data is achieved by double differentiation of the generated trajectory. Three calibrated noise components are added: Velocity random walk (VRW) modeled as a per-axis Gaussian process with power spectral density $\boldsymbol{\sigma}_w = [0.161,\ 0.086,\ 0.106]\,\mathrm{m\,s^{-2}\,Hz^{-1/2}}$ (East, North, Up), reflecting the dominant East-Up cable motion; bias drift follows a stochastic random-walk process with $\sigma_b = 0.002\,\mathrm{m\,s^{-2}}$; and mechanical vibrations are modeled as band-limited noise at 2.8 Hz with harmonics, amplitude $\sigma_v = 0.12\,\mathrm{m\,s^{-2}}$, fitted from the real IMU power spectrum. Synthetic standard deviations by axis match real measurements to within 2.6 %.

F. UWB Ranging Model

For each anchor i located at position $\mathbf{a}_i$, the geometric range is computed as

$$r_{\mathrm{geom}}(t) = |\mathbf{p}(t) - \mathbf{a}_i|\,. \tag{2}$$

A two-component noise model generates the measured range as

$$\begin{aligned} r_{\mathrm{meas}} &= r_{\mathrm{geom}} + (1-z)\,\mathcal{N}(\mu_i, \sigma_i^2) \\ &\quad + z\,\delta_i, \\ z &\sim \mathrm{Bernoulli}(p_{\mathrm{NLOS}}). \end{aligned} \tag{3}$$

Here, $\mathcal{N}(\mu_i, \sigma_i^2)$ refers to per-anchor LOS Gaussian noise and δ_i is a bimodal NLOS offset fitted to real residual clusters per anchor. Temporary NLOS bursts are introduced near support pylons while ranging follows a burst-based acquisition process at approximately 17–20 Hz, matching the update rate of the real deployment. Per-anchor parameters are estimated directly from measured residual distributions.

G. Availability and AoI Model

UWB measurements are sparse due to a limited update rate, visibility constraints, and temporary outages. Availability is therefore modeled using per-anchor dropout and outage processes, with per-burst dropout probability $p_d = 0.04$ and outage rate $p_o = 0.003$ per burst, both calibrated from real per-anchor availability fractions. Each anchor alternates between visible and unavailable states, while extended outages are sampled from exponentially distributed holding times ($\bar{\tau} = 16\,\mathrm{s}$), reproducing the prolonged measurement gaps observed in the real deployment. The AoI for anchor i is defined as

$$\mathrm{AoI}_i(t) = t - t_i^{\mathrm{last}}, \tag{4}$$

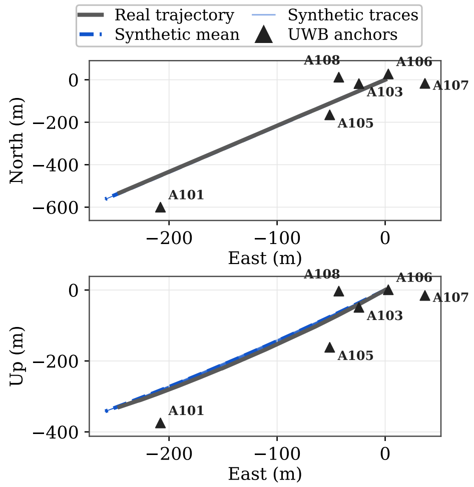

Figure 3. Horizontal (East-North) and vertical (East-Up) projections of the measured cable-car trajectory and synthetic trajectories generated by `PhySynth`.

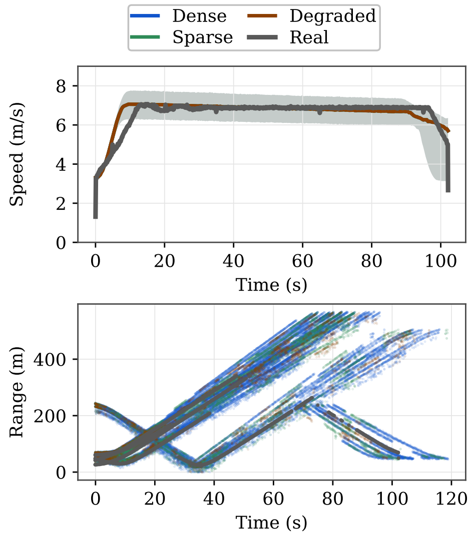

Figure 4. Generated speed profiles and corresponding UWB range realizations under different deployment regimes. Variations in speed, trajectory duration, and channel conditions produce diverse yet physically plausible synthetic traces.

where t_i^{last} is the timestamp of the most recent valid measurement. The resulting AoI matrix captures the freshness of available observations and is directly compatible with learning-based fusion approaches such as AoI-FusionNet [9], where measurement age modulates anchor attention weights.

IV. Evaluation

In the following, we evaluate the ability of the `PhySynth` generator to match the statistical properties of the real deployment of the Avarange scenario (see Fig. 2). The evaluation is conducted in three parts: a (mountain-to-valley) descending trajectory of the cable-car experiment used for calibration; a (valley-to-mountain) ascending trajectory used to assess generalizability without retuning; and a controlled anchor-deployment study quantifying the effects of anchor count, anchor geometry, and channel conditions.

A. Descending Trajectory

Fig. 3 demonstrates the generated trajectory ensemble closely matching the measured cable-car path and reproducing realistic variability observed across repeated experimental runs. Fig. 4 illustrates the resulting speed profiles and corresponding UWB range realizations. The synthetic trajectories captures the acceleration, cruise, and deceleration phases observed in the real deployment while yielding diverse measurement patterns across dense, sparse, and degraded channel regimes. The IMU model is validated by comparing acceleration statistics along all three axes as provided in Table I, where the synthetic standard deviations match the measured values within 2.6%, which confirms correct noise calibration.

We also compared the empirical cumulative distribution functions (CDFs) of the absolute ranging residuals that show close agreement across the low-error and tail regions. This confirms that the proposed model effectively reproduces the measured error characteristics.

Fig. 5 compares the residual distributions for each anchor. The generated traces capture the anchor-specific variability, including the larger residual spreads associated with anchors A103, A105, and A107 consistent with their higher non-line-of-sight (NLOS) probability and heavier bias components. This confirms that the generator preserves anchor-dependent sensing behavior rather than relying on a global noise model.

Table II summarizes the validation results of the descending and ascending trajectories. For the descending (calibration) trajectory, the generated traces reproduce the residual statistics measured within 12 % and mean AoI to within 8.9 %, while the Kolmogorov-Smirnov (KS) value of 0.147 indicates good distributional agreement. The larger deviation (45–50 %) reflects conservative outage calibration (cf. Section III-G) to intentionally produce lower synthetic availability in order to train AoI-aware deep networks under conditions worse than real deployment for better generalization. `PhySynth` with the same parameter set generalizes to ascending trajectory without retuning. The residual median, P95, and P99 values match to within 10 % while the mean AoI agree to within 12 % and KS statistic improves to 0.057 suggesting that the

Table I
IMU Acceleration Noise Validation: Real vs. Synthetic

Axis	σ_{real} (m/s^2)	σ_{synth} (m/s^2)	$\lvert\Delta\sigma\rvert$	Rel. (%)
East	1.616	1.616	0.000	0.0
North	0.880	0.903	0.023	2.6
Up	1.066	1.070	0.004	0.4
Mean	1.187	1.196	0.009	1.0

Rel. (%) $= \lvert\sigma_{\text{synth}} - \sigma_{\text{real}}\rvert \,/\, \sigma_{\text{real}} \times 100$.

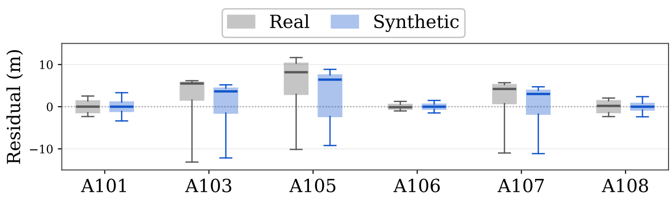

Figure 5. Comparison of per-anchor UWB residual distributions for real and synthetic data.

Table II
GENERATOR VALIDATION ON DESCENDING AND ASCENDING TRAJECTORIES. PARAMETERS CALIBRATED ON DESCENDING ONLY; ASCENDING DEMONSTRATES TRANSFER WITHOUT RETUNING.

	Descending			Ascending		
Metric	Real	Synth	%	Real	Synth	%
Duration (s)	102.1	104.0	1.9	113.1	114.2	1.0
UWB $\geq$1 (frac.)	0.178	0.097	45.5†	0.183	0.157	14.2
UWB $\geq$2 (frac.)	0.112	0.056	50.0†	0.127	0.123	3.1
Residual std (m)	7.485	6.591	11.9	0.407	0.374	8.1
Residual med. (m)	0.947	0.960	1.4	0.235	0.211	10.2
$\lvert r\rvert$ P95 (m)	18.12	16.03	11.5	0.830	0.809	2.5
$\lvert r\rvert$ P99 (m)	21.29	21.62	1.5	1.208	1.114	7.8
KS statistic	–	0.147	–	–	0.057	–
Mean AoI (s)	14.6	13.3	8.9	14.6	16.35	12.0

% = |Synth − Real|/|Real| × 100.
†Conservative outage calibration; see Section III-G.

proposed generator captures general channel characteristics rather overfitting to one calibrated trajectory.

B. Anchor Deployment Study

One of the main strengths of `PhySynth` is its ability to explore deployment configurations that would be time-consuming, costly, and sometimes impractical to practice through field experiments. To study the effect of anchor density, $n = 50$ traces are generated for each configuration ($n = 1, \dots, 6$) using anchors selected in decreasing order of real availability, while all generator parameters remained unchanged. The resulting traces are then compared against full six-anchor real deployment using a composite fidelity score. This score is a weighted sum of deviations in availability, residual median and P95 from real data, quantifying how well the generator reproduces the real sensor behavior under hypothetical anchor layout.

Fig. 6 illustrates the impact of anchor count on fidelity score and measurement availability. The score drops sharply from $n = 1$ to $n = 3$ anchors and plateaus thereafter, indicating diminishing returns beyond three anchors (see the left panel of Fig. 6). Synthetic measurement availability follows a trend similar to real deployment; however, the gap remains due to the conservative outage calibration strategy adopted in the generator. Meanwhile, channel quality has an even

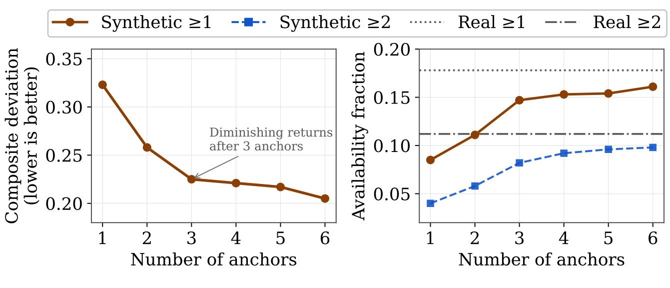

Figure 6. Impact of anchor count on synthetic trace generation. (left) Composite deviation from real data as a function of anchor count. Fidelity improves rapidly up to three anchors, after which diminishing returns are observed. (right) Fractions of measurements with at least one and at least two visible anchors compared with the real deployment. Availability increases monotonically with anchor count and gradually approaches the reference values

Table III
FIDELITY SCORE PER ANCHOR CONFIGURATION (LOWER IS BETTER). ROWS CORRESPOND TO ANCHOR COUNT AND COLUMNS REPRESENT UWB CHANNEL REGIME (50 SYNTHETIC TRACES EACH).

Anchors	Dense	Sparse	Degraded	Mean
6	**0.206**	0.260	0.311	0.259
4	**0.223**	0.281	0.332	0.279
3	**0.218**	0.288	0.324	0.277
2	**0.251**	0.307	0.346	0.301
Mean	**0.225**	0.284	0.328	-

Fidelity score $\mathcal{S} = \Delta_{\text{avail}} + \frac{1}{10}\Delta_{\text{med}} + \frac{1}{20}\Delta_{\text{P95}}$, where Δ denotes absolute deviation from real data. Anchors selected in decreasing availability order: A107, A106, A105, A103, A108, A101.

larger influence than anchor count. Table III reveals that a degraded six-anchor deployment (score 0.311) exhibits a high composite deviation and thus performs worse than a dense two-anchor deployment (score 0.251), confirming that measurement quality and availability dominate over infrastructure density alone. These results highlight the value of `PhySynth` in systematically studying deployment trade-offs under controlled and repeatable conditions.

C. Ascending Trajectory Generalization

To evaluate transferability to an unseen trajectory configuration, the generator is applied to the ascending direction of the same cable-car deployment with only the trajectory geometry and anchor configuration updated while all sensor and channel parameters remain the same. The ascending results included in Table II show that the residual statistics match within 10% and availability within 14% while achieving a lower statistic KS (0.057) than the descending case (0.147) due to more compact residual distribution of ascending trajectory (ascending real P95 $= 0.83$ m versus 18.12 m for descending), thus creating a simpler error profile for the generator to model. This indicates that the generator maintains strong agreement between the measured and synthetic residual distributions despite being calibrated on the descending run. In general, these results demonstrate that `PhySynth` can generalize across trajectory configurations, enabling realistic data generation for unseen deployment scenarios.

V. CASE STUDY

To demonstrate the practical utility of `PhySynth` for artificial intelligence (AI)-based tracking frameworks, we apply the generated synthetic traces as pre-training data to train using the tracking solution AoI-FusionNet [9], a deep-learning architecture for UWB-IMU sensor fusion under sparse UWB ranging conditions.

Three training strategies are evaluated; among them, two strategies are already mentioned in the referenced article:

- **NOAUG**: trains exclusively on real descending data for 200 epochs using adaptive moment estimation (Adam) optimizer and with learning rate 10^{-3}, therefore serving as the supervised baseline.
- **AUG_GAN**: augments GAN-generated synthetic UWB residuals with real data using the same schedule as

Table IV
LOCALIZATION PERFORMANCE ON THE DESCENDING TRAJECTORY. NOAUG IS THE REAL-ONLY SUPERVISED BASELINE. ΔRMSE IS RELATIVE TO NOAUG.

Method	**RMSE** (m)	**MAE** (m)	**P50** (m)	**P95** (m)	**ΔRMSE** (%)
NOAUG (real only)	0.136	0.110	0.086	0.276	-
AUG_GAN	0.129	0.106	0.082	0.259	-5.1
SYNTH_PHY (ours)	0.187	0.156	0.136	0.340	$+37.5$

[†]SYNTH_PHY uses **no real data** during pre-training (300 epochs on synthetic, then fine-tuned 200 epochs on real at lr=5×10^{-5}). P95 gap vs. NOAUG: $+23.2\%$.

NOAUG in order to stress test the AoI-FusionNet algorithm.
- **SYNTH_PHY**: follows a two-stage process. The first stage contains pre-training on synthetic traces generated by `PhySynth` for 300 epochs at learning rate 10^{-3}, followed by fine-tuning on real data for 200 epochs at a reduced learning rate of 5×10^{-5}. No real data is used during the first stage.

Table IV reports the localization performance evaluated against EMLID ground truth. Despite the RMSE gap, SYNTH_PHY achieves competitive P95 (0.340 m vs 0.276 m, $+23\%$) without any real pre-training data, demonstrating effective knowledge transfer from the physics-based generator to the real deployment scenario. The results show that models trained with the `PhySynth` data achieve competitive localization accuracy, without requiring additional real measurements. Although GAN-based augmentation provides the best overall performance, the proposed generator offers an important advantage: anchor count, geometry, channel conditions, availability, and AoI can be varied systematically, enabling controlled benchmarking of localization and tracking algorithms under deployment scenarios that are not represented in the original dataset.

VI. CONCLUSION

This paper presented `PhySynth`, a physics-based generator for synthesizing synchronized UWB, IMU, availability, and AoI measurements, addressing the challenge of data scarcity for training machine learning (ML)-based tracking algorithms. Using parameters extracted from a single descending AvaRange cable-car trajectory for calibration, the framework is able to produce realistic sensor traces. Statistical validation on both descending and ascending trajectories demonstrated close distributional agreement between real and synthetic measurements, achieving Kolmogorov-Smirnov statistics of 0.147 and 0.057, respectively. Furthermore, controlled anchor deployment experiments showed that high-fidelity synthetic data can be generated with as few as three anchors; however, channel quality was found to have a greater influence than anchor count. A case study of a learning-based fusion algorithm (AoI-FusionNet) further revealed effective transfer of synthetically generated data for training ML-based tracking. A limitation of the current work is its evaluation on a single deployment environment. Future work will extend this evaluation to multiple sites to establish broader parameter transferability and generalization across diverse scenarios.

ACKNOWLEDGMENTS

This work was supported by the project AvaRange funded by German Research Foundation (DFG), grant DR 639/22-1.

REFERENCES

[1] 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.

[2] Z. Zhaoxia, X. Zhongwei, and J. Xia, "Deep learning optimization positioning algorithm based on UWB/IMU fusion in complex indoor environments," *Elsevier Physical Communication*, vol. 71, Aug. 2025.

[3] K. Wei, X. Li, W. Zhu, W. Hu, Y. Hu, and H. Che, "Multistage Learning-Based Fusion Framework for Autonomous Vehicle Localization in GNSS-Denied Environments," *IEEE Sensors Journal*, vol. 25, no. 16, pp. 31 897–31 909, Aug. 2025.

[4] M. A. Bansal, D. R. Sharma, and D. M. Kathuria, "A Systematic Review on Data Scarcity Problem in Deep Learning: Solution and Applications," *ACM Computing Surveys*, vol. 54, no. 10s, pp. 1–29, Jan. 2022.

[5] A. Kerim, U. Celikcan, E. Erdem, and A. Erdem, "Using synthetic data for person tracking under adverse weather conditions," *Image and Vision Computing*, vol. 111, p. 104 187, 2021.

[6] G. Curnis, S. Fontana, and D. G. Sorrenti, "GTASynth: 3D synthetic data of outdoor non-urban environments.," *Data in Brief*, vol. 43, p. 108 412, Aug. 2022.

[7] 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.

[8] M. Neuhauser, A. Köhler, A. Wirbel, F. Oesterle, W. Fellin, J. Gerstmayr, F. Dressler, and J.-T. Fischer, "Particle and front tracking in experimental and computational avalanche dynamics," *Natural Hazards and Earth System Science*, vol. 25, pp. 4185–4202, Oct. 2025.

[9] T. Bibi, A. Köhler, J.-T. Fischer, and F. Dressler, "AoI-FusionNet: Age-Aware Tightly Coupled Fusion of UWB-IMU under Sparse Ranging Conditions," arXiv, eess.SP 2603.12849, Mar. 2026.

[10] Q. Zeng, Y. Song, R. Wang, and X. You, "A Sequential Fusion Localization Algorithm Based on AUKF for UWB/IMU/ODOM Systems," in *5th International Conference on Sensors and Information Technology*, Nanjing, China: IEEE, Mar. 2025, pp. 846–850.

[11] B. Cao, M. Jiang, M. Li, X. Ke, C. Zhang, H. Zhang, Q. Zeng, and B. Xu, "Improving Accuracy of the IMU/UWB Fusion Positioning Approach Utilizing ESEKF and VBUKF for Underground Coal Mining Working Face," *IEEE Internet of Things Journal*, vol. 12, no. 13, pp. 24 672–24 685, 2025.

[12] 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.

[13] 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.

[14] L. Alzubaidi, J. Bai, A. Al-Sabaawi, J. Santamaria, A. S. Albahri, B. S. N. Al-dabbagh, M. A. Fadhel, M. Manoufali, J. Zhang, A. H. Al-Timemy, Y. Duan, A. Abdullah, L. Farhan, Y. Lu, A. Gupta, F. Albu, A. Abbosh, and Y. Gu, "A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications," *Journal of Big Data*, vol. 10, no. 1, Apr. 2023.

[15] Y. Lang, W. Xia, Y. Fang, Y. Fu, S. Li, and Y. Yang, "SpecAug-DM: Few-Shot Radar Spectrogram Augmentation via Diffusion Model for Human Motion Recognition," *IEEE Sensors Journal*, pp. 1–1, May 2026.

[16] K. Yun, J. Park, J. Kim, and T. Kim, "SMOTE-Diffusion: A Combined Approach for Authentic Data Generation for Time-Domain Radar Signal in Intelligent Transportation System," *IEEE Sensors Journal*, vol. 25, no. 8, pp. 14 278–14 294, Apr. 2025.

[17] T.-H. Dang, J. Park, V.-T. Tran, and W. Y. Chung, "VAE-LSTM Data Augmentation for Cattle Behavior Classification Using a Wearable Inertial Sensor," *IEEE Sensors Letters*, vol. 8, no. 5, pp. 1–4, May 2024.

[18] 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.