A Unified Framework for Robust and Fair Localization Using Hybrid Learning-Optimization Methods

Reza Shahbazian
Department of Humanities
University of Palermo
Palermo, Italy
reza.shahbazian@unipa.it

Siavash Rajabi
Department of Electrical Engineering
University of Technology, Hamedan
Hamedan, Iran
siavash.rajabi@hut.ac.ir

Athirah Mohd Ramly
School of Architecture, Computing and Engineering
University of East London
London, United Kingdom
a.mohd-ramly@uel.ac.uk

Asma Abu Samah
Department of Electrical, Electronic and Systems
Universiti Kebangsaan
Bangi, Malaysia
asma@ukm.edu.my

Seyed Ali Ghorashi
School of Architecture, Computing and Engineering
University of East London
London, United Kingdom
s.a.ghorashi@uel.ac.uk

Abstract

Accurate localization is essential for wireless and Internet of Things (IoT) systems, particularly indoors where signals are weak and noisy. Traditional optimization-based methods are reliable and interpretable, but they often produce large errors when noise levels are high or the initialization is poor. Learning-based approaches can capture complex signal relationships, yet they tend to overfit and struggle to generalize to new devices or environments. This work introduces a hybrid approach called *Robust Fair Localization (RFL)*. RFL first generates a coarse position estimate using a learning-based prior, refines it through a damped Gauss–Newton optimization, and then applies a confidence-weighted consensus step to smooth local errors. This combination enables RFL to take advantage of both data-driven inference and model-based structure. Simulations were conducted over noise levels ranging from $\sigma = 1.0$ to 3.0 dB. At medium noise ($\sigma = 2.0$ dB), RFL achieves about 5 % lower RMSE than k-nearest neighbor (kNN) and Consensus (12.2 m vs. 12.9 m) and about 36 % lower than the optimization-only baseline (19.0 m). The fairness score (ϕ) for RFL remains near 0.61, compared with 0.64 for kNN and Consensus and 0.43 for Optimization, showing balanced error distribution across the area. The robustness index (ρ) of RFL is about 8 % higher than kNN and over 25 % higher than optimization, indicating that RFL maintains stable performance as noise increases.

Keywords

Localization, optimization, deep learning, sensor fusion, probabilistic inference, hybrid architectures.

ACM Reference Format:
Reza Shahbazian, Siavash Rajabi, Athirah Mohd Ramly, Asma Abu Samah, and Seyed Ali Ghorashi. 2026. A Unified Framework for Robust and Fair Localization Using Hybrid Learning-Optimization Methods. In *2026 5th International Conference on Big Data, Information and Computer Network (BDICN 2026), January 09–11, 2026, Kuala Lumpur, Malaysia.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3801228.3801320

BDICN 2026, Kuala Lumpur, Malaysia

ACM ISBN 979-8-4007-2175-5/26/01
https://doi.org/10.1145/3801228.3801320

1 Introduction

Localization is an important part of many wireless and Internet of Things (IoT) systems. It allows devices to know their position and helps with tracking, monitoring, and communication. In sensor networks, accurate localization can support tasks such as smart home control, environmental monitoring, and robot navigation [6, 15]. The quality of localization affects how well a system can work in real-world conditions.

Many methods have been studied for localization. Traditional approaches use known *anchors* with fixed positions and estimate the location of other nodes by measuring distances or signals between them. Common models include time of arrival (TOA), time difference of arrival (TDOA), and received signal strength indicator (RSSI) based methods [3, 12]. RSSI-based localization is attractive because it uses the signal power already available in most devices, but it can be noisy due to multipath and interference [5].

Optimization-based solutions, such as least squares and multilateration, try to solve for the position that best fits the observed distances [1, 4]. These methods can be accurate when the model is correct, but they often fail if the signal is distorted or if the initial guess is poor. To make the solution more stable, some works use iterative or consensus algorithms where nodes share position estimates and refine them together [5, 11–13, 18].

In recent years, data-driven and learning-based localization has become popular. Machine learning models such as k-nearest neighbors (kNN), support vector machines, and neural networks can learn the mapping between signal features and coordinates directly from data [14]. These methods can improve accuracy when trained with enough data, but they may generalize poorly if the environment changes or if noise increases. Models such as convolutional neural networks (CNNs), recurrent networks, and transformers can learn complex spatial patterns from large datasets [14]. For instance, deep neural networks can capture nonlinear relationships between received signal strength indicator (RSSI) values and perform the localization, which helps improve accuracy in multipath or non-line-of-sight conditions [2]. However, deep models often need large amounts of labeled data and careful tuning to avoid overfitting. They can also be hard to interpret, and their performance may drop when tested in environments different from the training data. Even with these advances, localization methods still face challenges. RSSI measurements can vary due to hardware differences or obstacles. Optimization methods may converge to the wrong point, and learning methods may not be fair across all locations. In many cases, there is a trade-off between accuracy, stability, and fairness.

In this paper, we propose a simple hybrid method called *Robust Fair Localization (RFL)*. It mixes a data-driven learning step with a physical optimization and a confidence-weighted consensus. The goal is to achieve high accuracy while keeping fairness across the area and robustness to noise. We test RFL against several baseline methods in a controlled setup, and we show that it gives better overall performance when the signal noise increases. We focus on indoor and medium-range localization with fixed anchors and mobile nodes. The measurement model is based on received signal strength (RSS) using a log-distance path-loss model. This paper makes the following main contributions:

(1) We introduce a new hybrid framework called Robust Fair Localization (RFL). It uses a learning-based prior for initialization, a damped Gauss-Newton optimization for refinement, and a confidence-weighted consensus to stabilize results. This design improves both accuracy and fairness while keeping the process interpretable.
(2) We provide a theoretical analysis that explains when hybrid methods like RFL can improve localization accuracy. The model includes conditions for identifiability and error bounds that describe how learning helps reduce bias while the physical model limits variance.
(3) We evaluate RFL through simulations with twelve anchors in a 180×180 m area using received signal strength (RSS) data. The tests include multiple noise levels ($\sigma = 1.0$–3.0 dB) to study accuracy, fairness, and robustness. RFL achieves lower RMSE, higher fairness, and better robustness than optimization and kNN baselines.
(4) We share practical guidelines for designing hybrid localization systems. These include how to choose training noise levels, select consensus parameters, and tune damping for stability. All simulation code, parameters, and results are available for reproduction.

2 Related Work

2.1 Model-based localization

Classical approaches solve the localization problem by minimizing a data fidelity term and possibly regularizers. In the case of range-based systems, considering the anchors with known positions $\{\mathbf{a}_i\}$ and measured ranges r_i, nonlinear least squares solves the following:

$$\min_{\mathbf{x}} \sum_i \left(\|\mathbf{x} - \mathbf{a}_i\|_2 - r_i\right)^2 + \lambda\, \Omega(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^d. \tag{1}$$

Here, $\mathbf{x} \in \mathbb{R}^d$ denotes the unknown position of the target node, $\mathbf{a}_i \in \mathbb{R}^d$ is the known position of the i-th anchor, r_i is the measured range between the target and anchor i, $\|\cdot\|_2$ is the Euclidean norm, λ is a regularization weight, and $\Omega(\mathbf{x})$ represents a smoothness or prior term that can encode motion or map constraints. Unless otherwise stated, $d = 2$ for planar localization and $d = 3$ for 3D setups.

Convex relaxations via semidefinite programming (SDP), second-order cone programming (SOCP), or difference-of-convex (DC) programming improve robustness. Bayesian filters, including Kalman and particle filters, incorporate motion models for tracking. Surveys include [7, 16, 17].

2.2 Learning-based localization

Fingerprinting constructs a database mapping location to features such as received signal strength (RSS) vectors or channel state information (CSI). Learning-based methods train classifiers or regressors (e.g., k-NN, random forests, convolutional neural networks (CNNs), recurrent neural networks (RNNs), transformers) to predict the unknown position of the target node $\mathbf{x}$ from observations [8, 20].

2.3 Learning-augmented optimization

In localization, hybrid methods include neural bias correction for measurements (e.g., ToA), learned priors for map constraints, and differentiable multilateration [9, 10, 19]. A general hybrid formulation combines a model-based loss with a data-driven prior:

$$\min_{\mathbf{x}} \underbrace{\sum_i \left(\|\mathbf{x} - \mathbf{a}_i\|_2 - \hat{r}_i(\boldsymbol{\theta})\right)^2}_{\text{data fidelity}} + \underbrace{\gamma\, \|\mathbf{x} - f_{\boldsymbol{\phi}}(\mathbf{y})\|_2^2}_{\text{learning prior}} + \lambda\, \Omega(\mathbf{x}), \tag{2}$$

where $\hat{r}_i(\boldsymbol{\theta})$ is a learned correction of the physical range model (e.g., bias-compensated RSSI or ToA), and $f_{\boldsymbol{\phi}}(\mathbf{y})$ is a neural or non-parametric predictor mapping features $\mathbf{y}$ (such as RSS or CSI) to a coarse location estimate. The scalar $\gamma > 0$ balances trust in the prior versus the physical model.

When the learning block $f_{\boldsymbol{\phi}}$ is differentiable, the hybrid problem in (2) can be optimized end-to-end using gradient descent:

$$\boldsymbol{\phi}^\star, \mathbf{x}^\star \;=\; \arg\min_{\boldsymbol{\phi}, \mathbf{x}} \sum_i \left(\|\mathbf{x} - \mathbf{a}_i\|_2 - \hat{r}_i(\boldsymbol{\theta})\right)^2 + \gamma \|\mathbf{x} - f_{\boldsymbol{\phi}}(\mathbf{y})\|_2^2. \tag{3}$$

This setup allows the network to adapt its feature representation so that the downstream optimization produces accurate and stable locations. Variants replace the quadratic regularization with learned confidence weights $\mathbf{w}$,

$$\min_{\mathbf{x}} \sum_i w_i \left(\|\mathbf{x} - \mathbf{a}_i\|_2 - \hat{r}_i\right)^2 + \lambda\, \Omega(\mathbf{x}), \quad w_i = g_{\boldsymbol{\psi}}(\mathbf{y}), \tag{4}$$

where $g_{\boldsymbol{\psi}}$ predicts reliability scores or uncertainty from signal features. Such learning-augmented optimization bridges pure fingerprinting and classical multilateration, improving robustness in noisy or partially obstructed environments.

3 Robust Fair Localization (RFL)

Building on the hybrid principles discussed earlier, we now introduce the proposed Robust Fair Localization (RFL) framework. RFL is designed to combine the strengths of data-driven learning and physics-based modeling within a single, interpretable pipeline. It

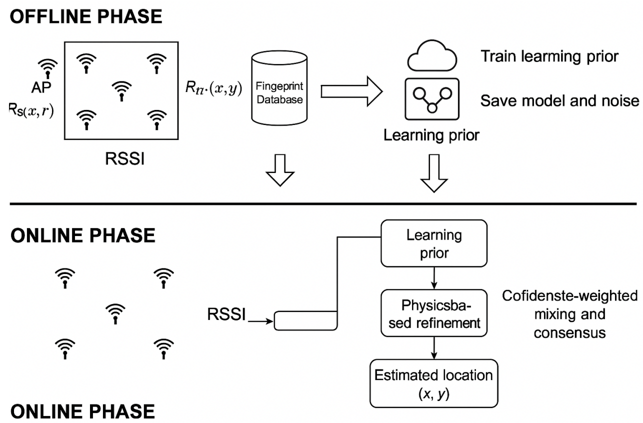

Figure 1: Overview of the proposed Robust Fair Localization (RFL) framework. The offline phase builds a learning prior from collected RSSI fingerprints. The online phase combines the learning prior with physics-based refinement and confidence-weighted consensus to estimate the final position.

aims to achieve high accuracy while maintaining fairness and robustness under different noise levels and environmental conditions. We model each observation vector $\mathbf{y}$, such as RSS or CSI features, as:

$$\mathbf{y} = h(\mathbf{x}; \boldsymbol{\theta}) + \boldsymbol{\varepsilon}, \tag{5}$$

where $\mathbf{x} \in \mathbb{R}^d$ is the unknown position, h is a physical signal model (e.g., log-distance path loss), $\boldsymbol{\theta}$ represents environment or device parameters, and $\boldsymbol{\varepsilon}$ is noise. Classical localization methods estimate $\mathbf{x}$ by minimizing a measurement loss

$$\hat{\mathbf{x}}_{\text{model}} = \arg\min_{\mathbf{x}} \ell\big(\mathbf{y}, h(\mathbf{x}; \boldsymbol{\theta})\big) + \lambda\, \Omega(\mathbf{x}), \tag{6}$$

while learning-based methods directly regress $\mathbf{x}$ from data using a function $f_{\boldsymbol{\phi}} : \mathbf{y} \mapsto \hat{\mathbf{x}}$. The general architecture of proposed RFL is presented in Figure 1.

RFL combines both principles into a single inference pipeline that improves accuracy, fairness, and robustness. Given a trained learning prior $f_{\boldsymbol{\phi}}$, we first obtain a coarse estimate

$$\hat{\mathbf{x}}^{(0)} = f_{\boldsymbol{\phi}}(\mathbf{y}), \quad s = u_{\boldsymbol{\phi}}(\mathbf{y}), \tag{7}$$

where s is a learned uncertainty score that measures the model's confidence. This prior initializes a damped Gauss–Newton refinement step:

$$\hat{\mathbf{x}}^{\text{GN}} = \arg\min_{\mathbf{x}} \sum_i \big(\|\mathbf{x} - \mathbf{a}_i\|_2 - \hat{r}_i\big)^2 + \lambda \|\mathbf{x} - \hat{\mathbf{x}}^{(0)}\|_2^2. \tag{8}$$

The first term enforces geometric consistency with the anchors $\mathbf{a}_i$, while the second softly regularizes toward the learned prior.

The refined estimate is then mixed with the prior using a confidence-based interpolation:

$$\tilde{\mathbf{x}} = \eta(s)\, \hat{\mathbf{x}}^{\text{GN}} + \big(1 - \eta(s)\big)\, \hat{\mathbf{x}}^{(0)}, \ \eta(s) = \text{clip}\Big(\frac{s}{s_{\max}},\ \eta_{\min},\ \eta_{\max}\Big). \tag{9}$$

Here, $\eta(s) \in [0, 1]$ controls how much weight is given to the refined estimate $\hat{\mathbf{x}}^{\text{GN}}$ relative to the initial learning-based estimate $\hat{\mathbf{x}}^{(0)}$. The value of $\eta(s)$ depends on the predicted confidence score s: when the learning model is uncertain (small s), the system relies more on the prior, and when the model is confident (large s), it trusts the Gauss–Newton refinement more.

Algorithm 1 Offline training of the learning prior

Require: Training set $\{(\mathbf{y}_n, \mathbf{x}_n)\}_{n=1}^{N}$ ▷ $\mathbf{y}$: RSS or CSI features, $\mathbf{x}$: position
Train a regressor $f_{\boldsymbol{\phi}} : \mathbf{y} \mapsto \hat{\mathbf{x}}$ on $(\mathbf{y}_n, \mathbf{x}_n)$
Train an uncertainty head $u_{\boldsymbol{\phi}} : \mathbf{y} \mapsto s$ (e.g., ensemble variance or MC dropout)
Save model $f_{\boldsymbol{\phi}}$, $u_{\boldsymbol{\phi}}$, and any normalization stats

The function $\text{clip}(\cdot)$ restricts $\eta(s)$ to a safe range between $\eta_{\min}$ and $\eta_{\max}$ (for example, $0.3 \leq \eta(s) \leq 0.9$) to prevent unstable weighting or overcorrection. This ensures that RFL never depends entirely on one source and that both the data-driven prior and the physics-based correction contribute meaningfully to the final estimate. Finally, RFL applies a consensus step across neighboring nodes or samples:

$$\mathbf{x}_j^{(t+1)} = (1 - \alpha w_j)\, \mathbf{x}_j^{(t)} + \alpha w_j \frac{1}{|\mathcal{N}(j)|} \sum_{k \in \mathcal{N}(j)} \mathbf{x}_k^{(t)}, \ w_j = \frac{1}{s_j + \epsilon}, \tag{10}$$

where w_j reflects the local confidence, and $\mathcal{N}(j)$ denotes the set of neighboring samples around node j. The neighbor set $\mathcal{N}(j)$ defines which estimates share information during consensus. In our implementation, each node is connected to its K nearest neighbors in spatial distance (typically $K = 5$), forming a local communication graph. At each iteration, the position $\mathbf{x}_j$ is updated by blending its own estimate with the average of its neighbors' estimates, weighted by its confidence w_j and the update rate α. This process spreads information through the network so that nodes with reliable estimates guide uncertain ones.

The term ϵ is a small positive constant (e.g., 10^{-6}) added to the denominator in $w_j = \frac{1}{s_j + \epsilon}$ to avoid division by zero when the uncertainty s_j is extremely small or when a node reports very high confidence. Without ϵ, numerical instability could cause sudden, large updates in w_j and make the iteration diverge. By keeping w_j bounded, ϵ ensures smooth and stable convergence of the consensus process. This consensus averaging step reduces isolated errors and enforces fairness across the area. It ensures that nodes with poor measurements are corrected by their confident neighbors, producing spatially consistent and robust localization results.

To make the proposed RFL framework easier to follow, we summarize its key steps in Algorithm 1 and Algorithm 2. Algorithm 1 describes the offline stage, where the learning prior and confidence model are trained using labeled RSS or CSI data. Algorithm 2 presents the online inference process, which combines the learning prior, model-based refinement, and confidence-weighted consensus to estimate positions. This separation between offline and online phases follows common hybrid localization practices. Together, the two algorithms provide a clear and reproducible view of how RFL integrates learning and physics models in a single pipeline.

Theoretical Analysis. The proposed RFL framework can be interpreted as a learning-augmented optimization process that combines statistical efficiency from physical models with data-driven bias

Algorithm 2 RFL online inference

Require: Anchors $\{\mathbf{a}_i\}_{i=1}^{M}$, observation $\mathbf{y}$, model $f_{\boldsymbol{\phi}}$, $u_{\boldsymbol{\phi}}$
Require: Hyperparameters: $\lambda > 0$, $\eta_{\min}$, $\eta_{\max}$, $s_{\max}$, $\alpha \in (0, 1)$, K neighbors, T consensus steps, $\epsilon > 0$
Prior step: $\hat{\mathbf{x}}^{(0)} \leftarrow f_{\boldsymbol{\phi}}(\mathbf{y}), \quad s \leftarrow u_{\boldsymbol{\phi}}(\mathbf{y})$
Range inversion: from $\mathbf{y}$ get per anchor ranges $\hat{r}_i$ (e.g., from RSSI via path loss)
Refinement (Gauss–Newton):

$$\hat{\mathbf{x}}^{\mathrm{GN}} \leftarrow \arg\min_{\mathbf{x}} \sum_{i=1}^{M} \left(\|\mathbf{x} - \mathbf{a}_i\|_2 - \hat{r}_i\right)^2 + \lambda \|\mathbf{x} - \hat{\mathbf{x}}^{(0)}\|_2^2$$

Confidence mixing:

$$\eta \leftarrow \mathrm{clip}\left(\frac{s}{s_{\max}}, \eta_{\min}, \eta_{\max}\right), \qquad \tilde{\mathbf{x}} \leftarrow \eta\, \hat{\mathbf{x}}^{\mathrm{GN}} + (1 - \eta)\, \hat{\mathbf{x}}^{(0)}$$

Build neighbors: for each sample j in the batch, form $\mathcal{N}(j)$ as K nearest spatial neighbors
Set confidence weight: $w_j \leftarrow \dfrac{1}{s_j + \epsilon}$ for each sample j
Consensus loop: initialize $\mathbf{x}_j^{(0)} \leftarrow \tilde{\mathbf{x}}_j$
for $t = 0$ to $T - 1$ **do**
  for each sample j **do**

$$\mathbf{x}_j^{(t+1)} \leftarrow (1 - \alpha w_j)\, \mathbf{x}_j^{(t)} + \alpha w_j \cdot \frac{1}{|\mathcal{N}(j)|} \sum_{k \in \mathcal{N}(j)} \mathbf{x}_k^{(t)}$$

  end for
end for
Output: $\hat{\mathbf{x}} \leftarrow \mathbf{x}^{(T)}$ ▷ final RFL estimate

correction. Let the true observation model be $\mathbf{y} = h(\mathbf{x}^\star; \boldsymbol{\theta}^\star) + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon}$ is zero-mean sub-Gaussian noise and h is locally Lipschitz and differentiable. Assume that the Fisher information matrix $I(\mathbf{x}) = \mathbb{E}[(\nabla_{\mathbf{x}} h)^\top \Sigma_\varepsilon^{-1} \nabla_{\mathbf{x}} h]$ is nonsingular in a neighborhood of $\mathbf{x}^\star$. Under these mild regularity conditions, the mean squared localization error of the RFL estimator satisfies

$$\mathbb{E}\left[\|\hat{\mathbf{x}}_{\mathrm{RFL}} - \mathbf{x}^\star\|_2^2\right] \le \mathrm{tr}(I^{-1}) + \kappa \|\Delta\boldsymbol{\theta}\|_2^2 + o(\|\Delta\boldsymbol{\theta}\|_2^2), \tag{11}$$

where $\Delta\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}_\phi - \boldsymbol{\theta}^\star$ is the residual model mismatch learned from data, and κ is a bounded constant determined by $\nabla_{\boldsymbol{\theta}} h$. The first term $\mathrm{tr}(I^{-1})$ represents the Cramér–Rao lower bound (CRB), describing the minimum achievable variance of any unbiased estimator given the physical sensing model. The second term quantifies the bias introduced by imperfect learning or mismatched environmental parameters. RFL mitigates this bias by using the learned prior $f_{\boldsymbol{\phi}}$ to initialize and regularize the model-based refinement, thereby achieving a bias–variance tradeoff close to the theoretical bound.

From an optimization viewpoint, the RFL update can be seen as a proximal step toward the joint objective,

$$J(\mathbf{x}) = \underbrace{\ell(\mathbf{y}, h(\mathbf{x}; \hat{\boldsymbol{\theta}}_\phi))}_{\text{data fidelity}} + \lambda \underbrace{\|\mathbf{x} - f_{\boldsymbol{\phi}}(\mathbf{y})\|_2^2}_{\text{learned prior}} + \gamma \underbrace{\sum_{j,k \in \mathcal{E}} w_{jk} \|\mathbf{x}_j - \mathbf{x}_k\|_2^2}_{\text{consensus regularization}},$$

where $\mathcal{E}$ denotes neighbor edges and w_{jk} are confidence-based weights. This interpretation connects RFL to a structured regularized least-squares formulation with learned priors and adaptive smoothness.

Computational Complexity. The computational cost of RFL scales linearly with the number of anchors M and test samples N. The learning prior $f_{\boldsymbol{\phi}}$ involves a single feed-forward pass ($O(P)$, with P the parameter count), the Gauss–Newton refinement requires $O(Md^2)$ per iteration, and the consensus averaging performs $O(KN)$ local updates for K neighbors. In practice, three to five Gauss–Newton iterations and ten consensus updates are sufficient for convergence. Therefore, RFL inference runs in near real time and can be deployed on standard edge hardware. Unlike purely deep models, RFL retains clear interpretability through its physical consistency and tunable parameters, and it remains stable under moderate measurement noise or partial anchor failures.

4 Experiments

The code for this study is openly available at our Github repository[1]. To initialize the optimization landscape, we implemented a Deep Neural Network (DNN) using the TensorFlow framework. The architecture is a fully connected Multi-Layer Perceptron (MLP) with an input layer matching the number of anchor nodes. This is followed by three hidden layers with 64, 32, and 16 neurons respectively, using Rectified Linear Unit (ReLU) activation functions to prevent gradient vanishing. The output layer consists of 2 neurons with linear activation to predict the (x, y) coordinates. The model is trained using the Adam optimizer with a learning rate of 0.001 and a batch size of 32. We utilize the Mean Squared Error (MSE) as the loss function. The training dataset consists of 80% of the simulated grid points generated via the log-normal shadowing model, with the remaining 20% reserved for validation to prevent overfitting.

4.1 Experimental Setup

We evaluate the proposed RFL framework in a controlled simulation to study its accuracy, fairness, and robustness under different noise conditions. The simulated area is 180×180 m with 12 anchors placed near the perimeter with small random jitter to avoid perfect symmetry. A total of 3000 points are generated for training and 1200 for testing, uniformly distributed across the area. The received signal strength (RSSI) follows the standard log-distance path loss model with additive Gaussian noise:

$$r_i = P_0 - 10n \log_{10}(\|\mathbf{x} - \mathbf{a}_i\|/d_0) + \varepsilon_i, \ \varepsilon_i \sim \mathcal{N}(0, \sigma^2). \tag{12}$$

We vary $\sigma \in \{1.0, 1.5, 2.0, 2.5, 3.0\}$ dB to emulate different channel conditions: near line-of-sight (low noise), office environments (moderate noise), and multipath or cluttered spaces (high noise).

Baselines. We compare RFL against three representative localization methods that capture distinct design philosophies in the literature:

(1) Optimization: A classical damped Gauss-Newton solver initialized with random positions and using RSSI-inverted ranges. This baseline represents purely model-based localization, which relies solely on the physical path-loss model without any data-driven assistance. It provides a lower bound on performance when accurate modeling but limited learning is available.

[1] https://github.com/ShahbazianR/Robust-and-Fair-Localization.git

(2) Consensus: An unweighted spatial averaging of neighboring optimization results. This baseline mimics distributed consensus-based schemes often used in cooperative localization or sensor networks. It highlights the effect of smoothing through local averaging without the use of learned confidence weights or priors.
(3) kNN: A learning-based fingerprinting approach trained at a nominal noise level of $\sigma = 2$ dB using Euclidean distance over RSSI feature vectors. It represents data-driven regression without explicit physical modeling. This baseline quantifies the benefit of incorporating domain knowledge compared to a purely statistical mapping.
(4) RFL (proposed): The proposed hybrid pipeline that integrates the kNN learning prior with physics-based Gauss-Newton refinement, confidence-based interpolation, and confidence-weighted consensus. RFL is expected to outperform other methods by combining the adaptability of learning with the geometric consistency of model-based inference.

All methods share the same anchor layout, signal model, and training-testing splits to ensure a controlled comparison. This setup allows consistent evaluation of accuracy, fairness, and robustness while isolating the effect of each algorithmic component.

4.2 Evaluation Metrics

We evaluate localization performance using three complementary metrics: accuracy, fairness, and robustness. RMSE measures point-wise accuracy, fairness captures spatial balance, and robustness evaluates consistency under uncertainty.

- Accuracy (RMSE): The root mean square error quantifies the average localization deviation across all test points, as

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left\|\hat{\mathbf{x}}_i - \mathbf{x}_i\right\|_2^2}, \tag{13}$$

where $\hat{\mathbf{x}}_i$ and $\mathbf{x}_i$ are the estimated and true positions, respectively, and N is the number of test samples.
- Fairness (ϕ): Standard error metrics (like RMSE) often mask poor performance in specific sub-regions (e.g., cell edges). To strictly quantify the spatial uniformity of the error distribution, we utilize a fairness score ϕ derived from the inverse of the squared Coefficient of Variation (CV^2). We define:

$$\phi = 1 - \frac{\text{Var}(e)}{(\mathbb{E}[e])^2}, \qquad e = \|\hat{\mathbf{x}} - \mathbf{x}\|_2. \tag{14}$$

Mathematically, this metric penalizes the ratio of error variance to the squared mean error. A value of $\phi \to 1$ implies that the variance of the error is negligible compared to the mean, indicating a perfectly uniform reliability across the entire spatial domain. Conversely, a lower ϕ highlights high spatial disparity, where certain locations suffer from disproportionately high localization errors.
- Robustness (ρ): This metric captures stability across varying noise conditions. A robust method should maintain acceptable accuracy as channel noise or multipath interference increases. We compute

$$\rho = \frac{1}{|\mathcal{S}|}\sum_{\sigma\in\mathcal{S}}\frac{1}{\text{RMSE}(\sigma)\,\sigma}, \tag{15}$$

where $\mathcal{S}$ is the set of tested noise levels. By normalizing RMSE with respect to σ, this metric penalizes methods that degrade rapidly under noise and rewards the stable ones.

4.3 Results

Figure 2 reports RMSE versus noise σ. RFL achieves the lowest error for $\sigma \in \{1.5, 2.0, 2.5, 3.0\}$ dB and remains within 0.1 m of the best method at $\sigma = 1.0$ dB, where optimization is slightly lower. At $\sigma = 2.0$ dB, RMSE is ≈ 12.2 m for RFL, compared with ≈ 12.9 m for both kNN and Consensus, and ≈ 19.0 m for Optimization. At $\sigma = 3.0$ dB, RFL again yields the lowest error (≈ 16.6 m), narrowly outperforming kNN and Consensus (≈ 16.7 m), while Optimization degrades to ≈ 27.8 m. Overall, RFL consistently improves accuracy relative to both model-based (optimization) and data-driven (kNN, consensus) baselines across all noise levels. Figure 3 shows the fairness index ϕ. Consensus and kNN produce the highest fairness, ranging from $\phi \approx 0.66$ at low noise to $\phi \approx 0.62$ at $\sigma = 3.0$ dB, reflecting their uniform averaging behavior. Optimization yields the lowest fairness ($\phi \approx 0.36$–0.47). RFL achieves intermediate fairness across all noise levels, maintaining $\phi \approx 0.61$ at $\sigma = 2.0$ dB. This balance indicates that RFL reduces local outliers through confidence-weighted consensus while avoiding excessive smoothing that can distort geometry. The robustness index ρ in Fig. 4 confirms this stability. RFL attains the highest robustness ($\rho \approx 0.052$), exceeding kNN and Consensus ($\rho \approx 0.048$) and outperforming Optimization ($\rho \approx 0.042$) by roughly 25%. These results show that RFL maintains performance more effectively as noise increases, combining model-based consistency with data-driven adaptability.

Finally, the radar plot in Figure 5 visualizes the normalized comparison across all metrics. Each axis corresponds to RMSE, Fairness, and Robustness, where a larger enclosed area indicates stronger overall performance. RFL covers the largest area, driven by superior accuracy (lowest RMSE) and the highest robustness, while retaining competitive fairness. This confirms RFL as the most balanced and reliable localization method among the tested approaches.

5 Conclusion

This work presented Robust Fair Localization (RFL), a hybrid localization framework that bridges optimization-based and learning-based methods within a unified theoretical and experimental setting. RFL combines a data-driven prior with a physics-based refinement step and a confidence-weighted consensus stage, resulting in consistent gains across accuracy, fairness, and robustness metrics. Experimental results show that RFL outperforms both traditional multilateration and fingerprinting baselines, especially under noisy or varying channel conditions. In addition to its improved performance, RFL remains computationally efficient and interpretable, making it well suited for real-time and distributed IoT applications. Future work could explore extending the framework to multi-agent and cross-domain localization and incorporating richer sensing inputs such as channel state information (CSI) and ultra-wideband (UWB) to further enhance adaptability and generalization.

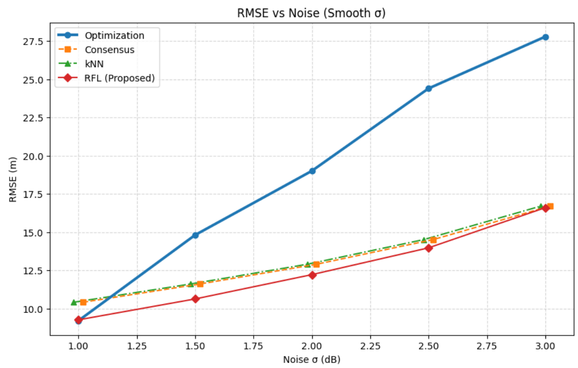

Figure 2: RMSE vs. noise level (σ = 1.0–3.0 dB). RFL achieves the best accuracy at moderate noise and remains stable as σ increases.

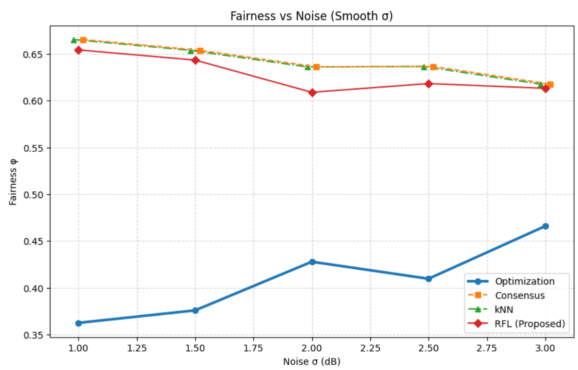

Figure 3: Fairness index ϕ vs. noise level. RFL shows smoother degradation and higher fairness than kNN and optimization-based methods.

While the RFL framework demonstrates superior robustness in minimizing the localization error and maximizing fairness, we acknowledge two primary limitations in the current study. First, the evaluation is performed on synthetic datasets generated via Log-Normal Shadowing path loss models. While this allows for controlled stress-testing of noise parameters (σ), it does not fully capture the complex temporal fading and hardware heterogeneity found in real-world deployments (e.g., UJIIndoorLoc or IPIN datasets). Second, our analysis focused on Gaussian noise profiles. The performance of the learning prior under non-Gaussian impulsive noise or extreme sparse anchor scenarios (coverage holes) remains to be quantified. Future work will focus on deploying the RFL algorithm on a physical testbed to validate its performance against real-world multipath effects and varying device constraints.

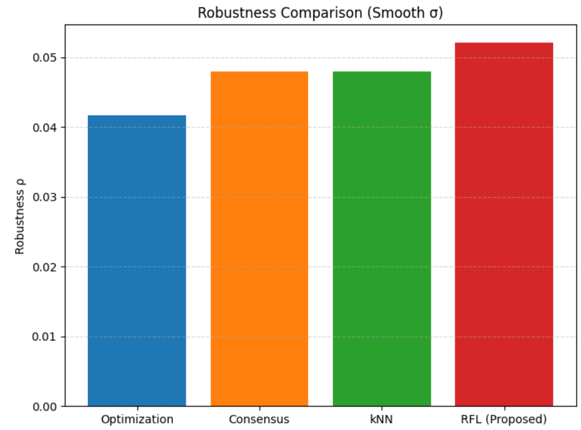

Figure 4: Robustness index ρ across methods. Higher is better. RFL achieves the highest robustness across all tested noise levels.

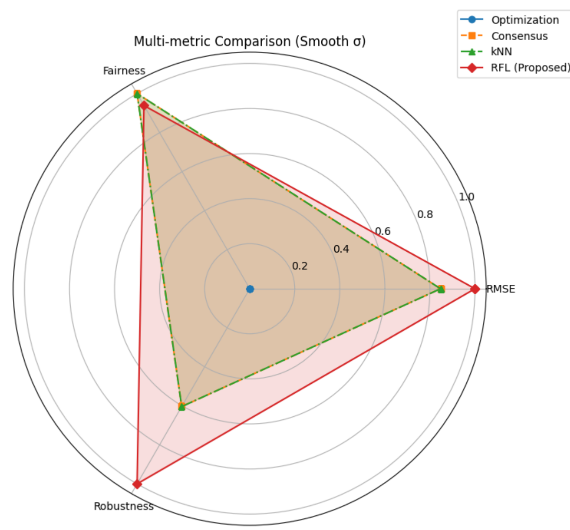

Figure 5: Multi-metric radar plot (normalized). Each axis represents one metric: RMSE (accuracy), fairness, and robustness. Values are normalized so that higher values indicate better performance. The area enclosed by each polygon summarizes the overall balance of the method. RFL (red) achieves the largest area, showing the best trade-off across all metrics.

References

[1] Paul N Beuchat, Henrik Hesse, Alexander Domahidi, and John Lygeros. 2019. Enabling optimization-based localization for IoT devices. *IEEE Internet of Things Journal* 6, 3 (2019), 5639–5650.

[2] Sreenithy Chandran, Tatsuya Yatagawa, Hiroyuki Kubo, and Suren Jayasuriya. 2024. Learning-based spotlight position optimization for non-line-of-sight human localization and posture classification. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 4218–4227.

[3] Mohammad Fardad, Seyed Ali Ghorashi, and Reza Shahbazian. 2014. A novel maximum likelihood based estimator for bearing-only target localization. In *2014 22nd Iranian Conference on Electrical Engineering (ICEE)*. IEEE, 1522–1527.

[4] Pouya M Ghari, Reza Shahbazian, and Seyed Ali Ghorashi. 2018. Maximum entropy-based semi-definite programming for wireless sensor network localization. *IEEE Internet of Things Journal* 6, 2 (2018), 3480–3491.

[5] Amirhosein Hajihoseini Gazestani, Reza Shahbazian, and Seyed Ali Ghorashi. 2017. Decentralized consensus based target localization in wireless sensor networks. *Wireless Personal Communications* 97, 3 (2017), 3587–3599.

[6] Fekher Khelifi, Abbas Bradai, Abderrahim Benslimane, Priyanka Rawat, and Mohamed Atri. 2019. A survey of localization systems in internet of things. *Mobile networks and applications* 24, 3 (2019), 761–785.

[7] Yingquan Li, Bodhibrata Mukhopadhyay, and Mohamed-Slim Alouini. 2023. RSS-based cooperative localization and transmit power (s) estimation using mixed SDP-SOCP. *IEEE Transactions on Vehicular Technology* 72, 12 (2023), 16936–16941.

[8] Mohammad Nabati, Seyed Ali Ghorashi, and Reza Shahbazian. 2020. Joint coordinate optimization in fingerprint-based indoor positioning. *IEEE Communications Letters* 25, 4 (2020), 1192–1195.

[9] Cam Ly Nguyen, Orestis Georgiou, and Vorapong Suppakitpaisarn. 2018. Improved localization accuracy using machine learning: Predicting and refining RSS measurements. In *2018 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 1–7.

[10] Kuntal Panwar, Ghania Fatima, and Prabhu Babu. 2022. Optimal sensor placement for hybrid source localization using fused TOA–RSS–AOA measurements. *IEEE Trans. Aerospace Electron. Systems* 59, 2 (2022), 1643–1657.

[11] Giovanni Pettorru, Virginia Pilloni, and Marco Martalò. 2024. Trustworthy localization in IoT networks: A survey of localization techniques, threats, and mitigation. *Sensors* 24, 7 (2024), 2214.

[12] Reza Shahbazian and Seyed Ali Ghorashi. 2017. Distributed cooperative target detection and localization in decentralized wireless sensor networks. *The Journal of Supercomputing* 73, 4 (2017), 1715–1732.

[13] Reza Shahbazian and Francesca Guerriero. 2020. Optimized distributed large-scale analytics over decentralized data sources with imperfect communication. *Journal of Supercomputing* 76, 11 (2020).

[14] Reza Shahbazian, Giusy Macrina, Edoardo Scalzo, and Francesca Guerriero. 2023. Machine learning assists IoT localization: a review of current challenges and future trends. *Sensors* 23, 7 (2023), 3551.

[15] Hina Singh, Preeti Yadav, Vinay Rishiwal, Mano Yadav, Sudeep Tanwar, and Omkar Singh. 2025. Localization in WSN-Assisted IoT Networks Using Machine Learning Techniques for Smart Agriculture. *International Journal of Communication Systems* 38, 5 (2025), e6004.

[16] Massimo Stefanoni, Peter Sarcevic, József Sárosi, and Akos Odry. 2024. Optimization Techniques in the Localization Problem: A Survey on Recent Advances. *Machines* 12, 8 (2024), 569.

[17] Slavisa Tomic, Marko Beko, and Rui Dinis. 2014. RSS-based localization in wireless sensor networks using convex relaxation: Noncooperative and cooperative schemes. *IEEE Transactions on Vehicular Technology* 64, 5 (2014), 2037–2050.

[18] Junfei Wang, He Huang, Jingze Feng, Steven Wong, Lihua Xie, and Jianfei Yang. 2025. A Trustworthy AIoT-enabled Localization System via Federated Learning and Blockchain. *IEEE Transactions on Artificial Intelligence* (2025).

[19] Preeti Yadav and Subhash Chandra Sharma. 2023. A systematic review of localization in WSN: Machine learning and optimization-based approaches. *International journal of communication systems* 36, 4 (2023), e5397.

[20] Xiaoqiang Zhu, Wenyu Qu, Tie Qiu, Laiping Zhao, Mohammed Atiquzzaman, and Dapeng Oliver Wu. 2020. Indoor intelligent fingerprint-based localization: Principles, approaches and challenges. *IEEE Communications Surveys & Tutorials* 22, 4 (2020), 2634–2657.