

Multimodal Indoor Localization Using Crowdsourced Radio Maps

Zhaoguang Yi^{§†}, Xiangyu Wen^{§†}, Qiyue Xia^{§†}, Peize Li^{§†}, Francisco Zampella[¶],
 Firas Alsehly[¶], Chris Xiaoxuan Lu^{*‡}

Abstract—Indoor Positioning Systems (IPS) traditionally rely on odometry and building infrastructures like WiFi, often supplemented by building floor plans for increased accuracy. However, the limitation of floor plans in terms of availability and timeliness of updates challenges their wide applicability. In contrast, the proliferation of smartphones and WiFi-enabled robots has made crowdsourced radio maps – databases pairing locations with their corresponding Received Signal Strengths (RSS) – increasingly accessible. These radio maps not only provide WiFi fingerprint-location pairs but encode movement regularities akin to the constraints imposed by floor plans. This work investigates the possibility of leveraging these radio maps as a substitute for floor plans in multimodal IPS. We introduce a new framework to address the challenges of radio map inaccuracies and sparse coverage. Our proposed system integrates an uncertainty-aware neural network model for WiFi localization and a bespoke Bayesian fusion technique for optimal fusion. Extensive evaluations on multiple real-world sites indicate a significant performance enhancement, with results showing $\sim 25\%$ improvement over the best baseline.

I. INTRODUCTION

Indoor location awareness is a key capability for smart devices and a fundamental need in embodied intelligence. There is a surge in applications of Indoor Positioning Systems (IPS) ranging from smart retail navigation to space utilization analysis.

Due to the non-line-of-sight issue of GPS in indoor environments, existing IPS leverage inertial measurement units (IMU) and/or building-dependent infrastructures such as WiFi, Bluetooth or magnetic fingerprints to localize smart devices. More recently, building floor plans emerge to be used as an effective input modality for IPS and can be fused with IMU and other building-dependent infrastructures to improve the localization accuracy [1] [2]. Intuitively, floor plans constrain the movement path of a moving agent - people cannot walk through walls and can only enter a room through a door, and can serve as strong cues in correcting localization offsets. However, the floor plans of many environments can be absent or lack timely updates after map creation. As a result, the usefulness of such indoor localization methods is largely limited by the availability of accurate maps.

On the other side, with the proliferation of smartphones and WiFi-enabled mobile robots, the radio maps of WiFi become largely available from crowdsourcing. A radio map

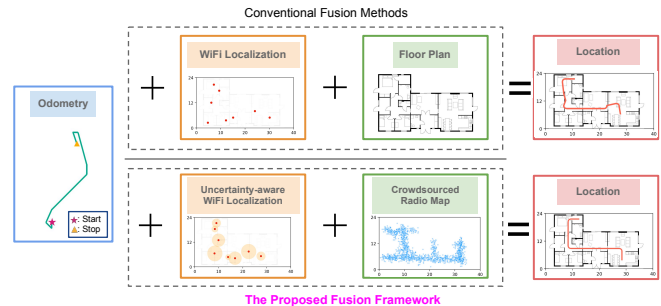


Fig. 1: Conventional IPS uses accurate radio maps for WiFi localization and floor plans in fusion. To fully utilize the noisy crowdsourced radio map, we propose an uncertainty-aware WiFi localization module and a bespoke Bayesian fusion method that replaces floor plans with crowdsourced radio maps in fusion.

broadly refers to a database consisting of a number of paired locations and their respective Received Signal Strengths (RSS) from different wireless access points. A research question of our interest is - *can we replace floor plans with radio maps in multimodal IPS?* Indeed, crowdsourced radio maps are often derived from tons of walking/moving trajectories by various users. Therefore, they not only contain wireless and location pairs but also carry side information about traversing regularity by inspecting the locations of fingerprints. With sufficient crowdsourced trajectories, we hypothesize the regularity embedded in a radio map implicitly reveals the allowable paths and can be used as an effective cue for localization, sharing a similar spirit of floor plans in constraining localization results. However, due to the nature of crowdsourcing, radio maps created this way inevitably suffer from location inaccuracy and sparse position coverage. Prior arts in fingerprint-based WiFi IPS [3] found that the accuracy of the WiFi localization systems can be far from satisfying and impose challenges in fusion if built upon noisy crowdsourced radio maps. Meanwhile, the sparse position coverage of radio maps cannot provide location constraints as explicit as that of floor plans, limiting the adoption of map-matching-like fusion methods [1], [2].

In this work, we propose a new framework to deal with the issues of radio map issues and unravel their full potential in IPS. Our framework, illustrated in Fig. 1, includes an uncertainty-aware neural network model for WiFi fingerprinting with radio maps and a new Bayesian filtering method admissible to the fusion of radio maps. Qualitative and quantitative evaluations on three real-world sites demonstrate that the proposed system is able to produce $\sim 25\%$ performance gain compared to the state-of-the-art IPS.

*Corresponding author. Email: xiaoxuan.lu@ucl.ac.uk

§Denotes equal contribution.

†School of Informatics, University of Edinburgh, United Kingdom.

¶ Huawei Technologies (UK).

‡ Department of Computer Science, University College London.

II. RELATED WORK

Indoor Localization via Wireless Networks Indoor localization based on existing wireless communication networks, such as WiFi [4] and Bluetooth [5], based on their RSS, has gained popularity because it does not require additional hardware [4]. For such systems, two primary methodologies are employed: 1) Utilizing an access-point database that comprises access-point IDs and their respective geo-coordinates. This setup facilitates the determination of user location through trilateration or weighted averaging of access-point coordinates [6]. 2) Employing a fingerprint database, a radiomap, which records the received signal strength indicators and the corresponding geo-coordinates at their observation points [7], [8]. Such an approach stands out for its inherent scalability, amplified by the prospects of crowd-sourced data collection [9].

Uncertainty Estimation Estimating the uncertainty, or more specifically, the suspicious evidence cues derived from data is crucial for the integration in real-world applications and as input for data fusion processes. To quantify uncertainty within deep learning, two primary distinctions emerge: epistemic uncertainty, associated with the model, and aleatoric uncertainty, related to the data [10]. Estimating aleatoric uncertainty can enhance the robustness of the network when training on noisy datasets. By introducing an additional aleatoric uncertainty prediction head, the reliability of neural networks has been significantly strengthened across various applications [11] [12] [13] [14]. Our research primarily concentrates on enhancing localization performance on noisy radio maps by predicting uncertainty, which in turn aids in data fusion processes.

Multi-modal Sensor Fusion Indoor positioning and navigation techniques using smartphones have seen significant advances with the incorporation of various data sources and modeling methods. WiFi-SLAM [15], [16] relies on WiFi-based loop-closing constraints, which effectively minimize the accumulation errors commonly found in inertial navigation. Bayesian filter-based algorithms can also be used to perform the multi-modal sensor fusion to estimate the online position, such as particle filter [17]–[19], Kalman filter and its variant [20]–[22]. Several studies have incorporated floor plan information into their fusion systems via offline image processing techniques [1], [23], while other methods use the floor plan as a constraint for particle filtering [18], [19] and conditional random fields [2]. However, there is no method to properly fuse IMU, uncertainty-aware WiFi localization results, and the crowdsourced WiFi radio map data.

III. PROBLEM FORMULATION

Presume we have an offline crowdsourced radio map containing N fingerprints and their associated locations $\mathcal{M} = \{(\mathbf{G}_i, \mathbf{L}_i) | i = 1, 2, \dots, N\}$, where \mathbf{L}_i denotes the location coordinates and $\mathbf{G}_i = \{(MAC_j, RSS_j) | j = 1, 2, \dots, |\mathbf{G}_i|\}$ is called a WiFi fingerprint represented as a list of MAC addresses of the WiFi Access Points (APs) and their corresponding RSS values at location \mathbf{L}_i . Note that the number of WiFi APs $|\mathbf{G}_i|$ varies across locations and the locations

\mathbf{L}_i considered in this work are two-dimensional, namely $\mathbf{L}_i = (x_i, y_i)$. As the radio map is crowdsourced rather than manually surveyed, the association between WiFi fingerprints and location coordinates can be inaccurate, rendering a **noisy radio map**. During the online stage and timestamp t , a mobile agent at an unknown location will measure a fresh WiFi fingerprint measurement \mathbf{G}_t and have a location/pose change $\Delta\tilde{\mathbf{L}}_{t-1:t}$ predicted from an odometry module. Our objective here is to obtain an updated location coordinate $\hat{\mathbf{L}}_t$ by utilizing the noisy crowdsourced map \mathcal{M} , the fresh fingerprint measurement \mathbf{G}_t and the location change $\Delta\tilde{\mathbf{L}}_{t-1:t}$. Without losing the generality, the mobile agent considered in our task can be either a smartphone or a WiFi-enabled mobile robot. In the case of smartphones, the predicted location or pose change can come from the smartphone inertial sensor and off-the-shelf inertial odometry e.g., [24]; In the case of mobile robots, such prediction change can be accessed from their wheel odometry module.

We propose a new method to address the problem above. At a high level, our method consists of an uncertainty-aware WiFi localization neural network model (c.f. Sec. IV) and a bespoke Bayesian fusion approach (c.f. Sec. V) that combines the WiFi localization results, odometry estimates and the noisy radio map to derive a more robust location.

IV. UNCERTAINTY-AWARE WiFi LOCALIZATION

A. Design Motivation

We start with the motivation behind our WiFi localization model. As found in [25], the lengths of WiFi fingerprints \mathbf{G}_i vary across locations and time even at the same location. This requires the WiFi localization model to take WiFi fingerprints in different lengths but still be able to output location coordinates. This requirement is similar to the one of natural language processing, where language models often take input as the sentences in different token lengths, but are still able to perform inference. Inspired by recent successes of transformers in natural language processing [26], [27], we propose utilizing the transformer’s encoder in our network design to handle the varying-length WiFi fingerprints. Moreover, to deal with the inherent noisy match in the crowdsourced radio map, we further introduce a network block to model the prediction uncertainty.

B. Input Embedding of WiFi Fingerprint

A fingerprint in the list format of (MAC_j, RSS_j) cannot be directly employed as the input to a transformer-like network. Akin to the idea of distributed representations of words in natural language processing [28], a fingerprint needs to be represented as a continuous vector before being fed into the model. Pre-trained word embedding layers are often available for language models by which the derived word embeddings are close to each other if two words have similar semantic meanings. However, such a semantic-preserving embedding layer is absent in the case of WiFi fingerprints esp. for the MAC address of APs. The challenge lies in that it is difficult to relate two nearby APs by their MAC address manifestation

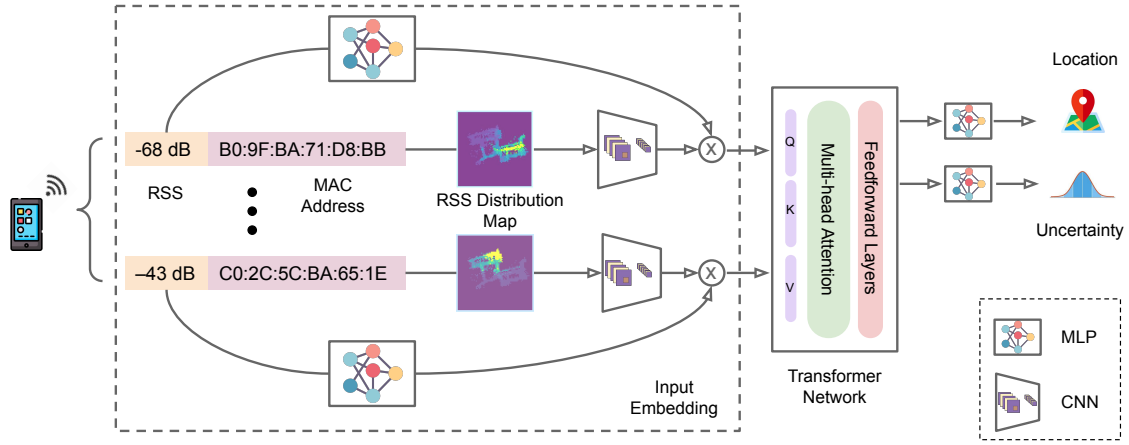


Fig. 2: The proposed *end-to-end trainable* WiFi localization model. Given a WiFi fingerprint list measured by a WiFi-enabled device (the smartphone in the figure), the MAC address part is input-embedded by applying a CNN on its AP-associated RSS distribution map; The RSS part of the fingerprint is embedded via an MLP, and multiplied with the MAC address embedding (c.f. Sec. IV-B). The constituted fingerprint embeddings are then fed into the transformer featuring a multi-head attention block in it (c.f. IV-C). After obtaining the encoded fingerprint features by the transformer, we append two MLPs to regress the location coordinates and quantify the prediction uncertainty (c.f. Sec. IV-D) for subsequent fusion.

- which is simply a 12-character alphanumeric attribute that is used to identify individual electronic devices.

To inject the location vicinity semantics in MAC address embedding, we propose to utilize the RSS distribution map. For each MAC address, we retrieve the radio map and construct a heatmap using the RSS signal strength received from that MAC address at all available locations. Intuitively, the highlighted area approximately indicates the location of the MAC address, and nearby WiFi APs tend to have similar RSS distribution maps. CNN blocks are used to process the distribution maps and extract the dense vectors as MAC address embeddings. As we will see in Sec. VII-C.1 the physical vicinity can be conveyed in the similarity of the embeddings. Even for the same AP, its RSS values measured by a mobile agent can vary across places. To make use of the RSS value in the WiFi fingerprint, we introduce a learnable MLP module that maps the RSS value and multiplies with the MAC address embedding to derive the final fingerprint embedding, \mathbf{E}_i . Fig. 2 illustrates the input embedding.

C. Multi-Head Attention Block

We aim to accurately predict location coordinates based on WiFi fingerprint embeddings. Given the dynamic and complex nature of WiFi fingerprints, it's crucial to capture dependencies between different features of the fingerprint for a precise location prediction. We thus adopt the self-attention transformer. Essentially, it allows each feature in the WiFi fingerprint embedding to weigh the importance of other features based on their content. This is particularly advantageous for our task because certain WiFi fingerprints might carry more relevance in determining specific location coordinates than others. Diving deeper into the mechanics, let's first revisit the single-head attention mechanism:

$$\begin{aligned} head_h &= Attention(\mathbf{E}_i \mathbf{W}_{Q_h}, \mathbf{E}_i \mathbf{W}_{K_h}, \mathbf{E}_i \mathbf{W}_{V_h}) \\ &= softmax\left(\frac{\mathbf{E}_i \mathbf{W}_{Q_h} (\mathbf{E}_i \mathbf{W}_{K_h})^T}{\sqrt{d_k}}\right) \mathbf{E}_i \mathbf{W}_{V_h} \end{aligned} \quad (1)$$

where \mathbf{E}_i is the input MAC embedding; \mathbf{W}_{Q_i} , \mathbf{W}_{K_i} , and \mathbf{W}_{V_i} are the projection matrices for query, key, and value respectively; $\sqrt{d_k}$ is the dimension of the key k and query q . However, relying on a single attention head might restrict the model from capturing multiple types of relationships or dependencies between the features. This is where the multi-head attention mechanism, as proposed in [26] is employed:

$$\begin{aligned} MultiHead(\mathbf{E}_i \mathbf{W}_{Q_h}, \mathbf{E}_i \mathbf{W}_{K_h}, \mathbf{E}_i \mathbf{W}_{V_h}) \\ = Concat(head_1, \dots, head_H) \cdot \mathbf{W}_O \end{aligned} \quad (2)$$

where $head_1, \dots, head_H$ represent the H single-head attention results. By using multiple attention heads, the model can simultaneously attend to various subspace representations of the input. Each head potentially captures different types of dependencies within the fingerprint embeddings. These concatenated results are then combined using a linear transformation matrix \mathbf{W}_O , to produce an attended embedding. Each attended embedding is applied a non-linear transformation before output. After processing through several of such multi-head attention layers, the output from the last layer is further refined by a feed-forward network to regress the location coordinates.

D. Uncertainty Estimation

As mentioned earlier, crowdsourced radio maps often come with inherent noise. Using them directly to train a transformer localization model can lead to high susceptibility to errors in predictions, given the model might learn from the noise rather than the actual signal. As a natural response to this challenge, one might consider ways to make the model more resilient to noise and simultaneously offer a more informative prediction. Thus, instead of merely giving a point estimate that provides a singular location prediction, it would be more beneficial if our model could quantify the level of confidence in its predictions.

To this end, we reformulate the network's output of location coordinates as a *probability distribution*. We choose

the Gaussian distribution for this purpose, given its well-understood properties and ease of integration in subsequent fusion steps:

$$\mathbf{L}_i^G \sim \mathcal{N}(f_\mu(\mathbf{G}_i), f_\sigma(\mathbf{G}_i)) \quad (3)$$

where $f_\sigma(\cdot)$ is an extra neural network head capturing the variance which is pivotal as it provides an estimate of the uncertainty associated with the prediction. By transitioning to this probabilistic framework, we ensure the transformer model does not blindly commit to a potentially noisy sample. Instead, it offers a range of possibilities, weighted by their likelihood. This not only makes the model less prone to overfitting to the noise but also enriches the output, as the variance gives an insight into the reliability of the prediction.

Altogether, given crowdsourced a radio map $\mathcal{M} = \{(\mathbf{G}_i, \mathbf{L}_i) | i = 1, 2, \dots, N\}$, we train our WiFi localization neural network with an uncertainty-aware loss function:

$$\mathcal{L} = \sum_i^N \frac{\|f_\mu(\mathbf{G}_i) - \mathbf{L}_i\|_2}{f_\sigma(\mathbf{G}_i)} + \ln f_\sigma(\mathbf{G}_i) \quad (4)$$

During training, the first term drives predictions closer to the radio map location coordinates. The numerator of the first term is to ensure that a larger variance is given for samples with significant prediction errors, mitigating large gradient impacts. The second term acts as a regularizer, discouraging the network from consistently yielding high variance.

V. MULTIMODAL FUSION

A. Design Motivation

We begin by presenting the motivation behind our multimodal fusion method which combines the strengths of the extended Kalman filter and particle filters. Prior arts of particle filters have shown the effectiveness of using particle filters to combine the WiFi localization results, inertial odometry and floor plans. In our context, however, the floor plan is unavailable and we can only access the radio map via crowdsourcing. A straightforward idea is to use the kernel density estimation algorithm [29] to generate a continuous likelihood distribution over the map area, indicating the frequency and accessibility of specific locations. This information can act as a constant similar to floor plans by adjusting the weight of importance in particle filters. However, since radio maps provide only a sparse estimate of location frequency, the resulting distribution is uneven. Consequently, without particles in true but unknown areas, the *kidnapped robot problem* [30], [31] is likely to happen if we are confined to changing the importance weight only (see the supplementary video). To address this, we propose incorporating the Kalman correction step into particle filters, creating an Extended Kalman Particle Filter (EKPF) that includes model prediction, importance weight adjustment, and WiFi measurement correction. This approach aims to ensure particles are more aligned with WiFi localization findings.

B. Model Prediction

In EKPF, every particle maintains states of the particle $\mathbf{s}_t = [x_t, y_t, \theta_t, v_t]^T$, an importance weight w and a corresponding error covariance matrix \mathbf{P}_t , where x_t, y_t are the map frame 2D position, θ represents the odometry frame to the map frame rotation. v is the scale ratio of the estimated velocity from IMU to the filter velocity. Whenever relative position updates $\mathbf{O}_t = [O_{tx}, O_{ty}]^T$ in odometry frame is received, the states are propagated as follows:

$$\mathbf{s}_t = \mathbf{s}_{t-1} + \mathbf{M} \cdot (\mathbf{O}_t \cdot (1 + \phi(n_p))) + [0, 0, \phi(n_r), 0]^T \quad (5)$$

$$\mathbf{M} = \begin{bmatrix} v_t \cos \theta_{t-1} & -v_t \sin \theta_{t-1} \\ v_t \sin \theta_{t-1} & v_t \cos \theta_{t-1} \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \quad (6)$$

where $\phi(n)$ represents random Gaussian noise with zero mean and variance n .

$$\check{\mathbf{P}}_t = \mathbf{F}_{t-1} \check{\mathbf{P}}_{t-1} \mathbf{F}_{t-1}^T + \mathbf{N}_{t-1} \mathbf{Q}_{t-1} \mathbf{N}_{t-1}^T \quad (7)$$

where $\check{\mathbf{P}}_t$ is the state covariance at time t , \mathbf{Q}_{t-1} is the noise covariance at time t . \mathbf{F}_{t-1} is the motion model Jacobian with respect to the last state while \mathbf{N}_{t-1} is the motion model Jacobian with respect to the noise of the last state.

C. Importance Weight Update

The weight of every particle is updated using its current location and the likelihood map during this stage. To obtain a smooth prior map f_{map} , the crowdsourced radio map is fed into the Gaussian kernel density estimation algorithm [29] and then normalized:

$$f_{map}(\mathbf{L}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(L-L_i)^2} + \beta \quad (8)$$

where β is a small constant to control the smoothness of the generated prior map. The importance weight of every particle is updated using its current location states. Like standard particle filter algorithms, the ratio of effective particles is monitored during the process and the standard systematic method [32] is used for the particle resampling. Note that during the particle resampling phase, the error covariance matrix \mathbf{P} remains unchanged.

D. Measurement Model

Similar to the usage in the original Kalman filter, in EKPF, the measurement model relates the current state $\check{\mathbf{s}}_t$ to a fresh WiFi localization result at timestamp t , $\mathbf{L}_t^G \sim \mathcal{N}(f_\mu(\mathbf{G}_t), f_\sigma(\mathbf{G}_t))$. In the map frame, we have:

$$f_\mu(\mathbf{G}_t) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \mathbf{s}_t + \mathbf{m}_t \quad (9)$$

\mathbf{m}_t is the measurement noise. To compute the Kalman Gain, we have:

$$\mathbf{K}_t = \check{\mathbf{P}}_t \mathbf{H}_t^T (\mathbf{H}_t \check{\mathbf{P}}_t \mathbf{H}_t^T + \mathbf{R}_t)^{-1} \quad (10)$$

where R_t is the measurement noise covariance where we used the estimated uncertainty $f_\sigma(\mathbf{G}_t)$ from the aforementioned WiFi localization model; \mathbf{H}_t is the measurement model Jacobian with respect to the last state. The Kalman Gain K_t is used to correct the predicted state:

$$\hat{\mathbf{s}}_t = \check{\mathbf{s}}_t + \mathbf{K}_t (\mathbf{L}_t - f_\mu(\mathbf{G}_t)) \quad (11)$$

Finally, the covariance correction is updated as follows:

$$\hat{\mathbf{P}}_t = (\mathbf{I} - \mathbf{K}_t \mathbf{H}_t) \check{\mathbf{P}}_t \quad (12)$$

where \mathbf{I} represents the identity matrix.

VI. EXPERIMENTAL SETUP

A. Data Collection

In diverse campus environments, we gathered WiFi fingerprints and IMU data using multiple smartphones over varying durations. This data was refined using a graph optimization algorithm [33] to produce radio maps for different structures. We focused on Building A (2,800 m^2 commercial space), Building B (3,300 m^2 educational facility), and Building C (4,500 m^2 library). Ground truth was acquired using Velodyne VLP-16 LiDAR integrated with an Xsens IMU, offering precise location data at 200Hz and we ran with Lio-SAM [34] to obtain the true pose. The smartphones collected IMU data at 100Hz and WiFi fingerprints every 2-3 seconds. We gathered around 2 hours of test data across the buildings, with trajectories lasting 5-10 minutes and covering a total of 4.2 km.

B. Implementation Details

1) *WiFi Localization Network*: The model was end-to-end trained on the radio map dataset, split 9:1 for training and validation. The best model is selected based on the performance achieved on the validation set. Note that the test set for performance evaluation is separate and independent of the radio map set. Coordinates were min-max normalized for faster convergence. Training used the AdamW optimizer [35] at a $1e^{-4}$ learning rate, betas of (0.9, 0.999), and a weight decay of $1e^{-3}$ over 200 epochs. Unless otherwise specified, the ReLU [36] activation function was utilized throughout the model. The CNN used for MAC address embedding consisted of two layers with kernel sizes 20 and 5 and strides 4 and 2, raising channel counts from 1 to 16. An MLP then reduced dimensions to 256×1 . The RSS embedding MLP has neurons arranged as (1, 128, 64, 1), using Sigmoid activation at the output. The multi-head attention transformer has 4 layers and 4 heads, a feed-forward size of 256, and a 0.2 dropout rate. Both output layers's hidden dimensions are set as (256, 128, 64). The uncertainty-focused MLP's output adopts the Sigmoid function mapping the probability from [0, 1]. The localization MLP's last layer only has the linear transformation.

2) *EKPF Fusion*: 400 particles were used in all experiments. Empirically, the location noise covariance, $\phi(n_p)$, was set to 0.1, and the rotation covariance, $\phi(n_r)$, was set to 0.05. Two critical EKPF parameters are: the softening constant β for the prior map likelihood and the measurement noise covariance constant γ adjusted by $f_\sigma(\mathbf{G}_t)$. Given the inconsistent reliability of the prior map and WiFi model across buildings, optimal parameters differ. For each site, a validation trajectory was used for a grid search to determine these settings, with β ranging from $1e^{-4}$ to $1e^{-5}$ and γ from 100 to 200. The kernel density function processed the radio map using a Gaussian kernel with a bandwidth of 1.0.

VII. RESULTS

A. Multimodal IPS Performance

We now compare our multimodal localization approach with the following competing methods: (1) **FusionDHL** [1]: We adapt the FusionDHL to a version without floorplan on our dataset. FusionDHL is an offline optimization algorithm based on least squares and the closest baseline to ours. (2) **EKF** [37]: The normal extended Kalman filter combines the odometry and the uncertainty-aware WiFi localization results. (3) **PF** [32]: The particle filter fusion method uses an implementation similar to [17], fusing the odometry and the uncertainty-aware WiFi localization results.

Figure 3 illustrates the cumulative distribution function for localization errors in meters, excluding data from the first minute to allow filter methods time to stabilize. Our Extended Kalman Particle Filter (EKPF) fusion approach outperforms all baseline methods across three sites. Compared to FusionDHL, EKPF better integrates radio map data with WiFi/IMU fusion, enhancing localization accuracy by approximately 10.6%, 21.5%, and 43.3% in three different buildings. Figure 4 further highlights the qualitative improvements our method achieves in the largest building.

B. WiFi Localization Performance

We evaluate our model against both traditional and state-of-the-art machine-learning localization techniques. For the conventional method, we used the Weighted K-nearest neighbours (WKNN) [38], which employs Euclidean distance with $k = 50$ neighbours. Additionally, we compared against a method utilizing a stacked denoising autoencoder for regression (DSAE) [39], which first trains an autoencoder and then fine-tunes with an MLP for improved localization.

In Tab. I, our model's performance outshines both WKNN and DSAE across three buildings, excelling in mean, median, and 95th percentile errors. Although we observed a slightly higher maximum error in Building B compared to WKNN, our model's uncertainty prediction effectively reduces outlier impact, ensuring dependable outcomes overall.

C. Ablation Study

1) *Input MAC Embedding*: In our study conducted on Building A, we compared the use of the RSS Distribution Map against directly using MAC address indices, akin to neural word embeddings. [28]. The results revealed a notable




Fig. 3: Overall performance of our multimodal IPS.

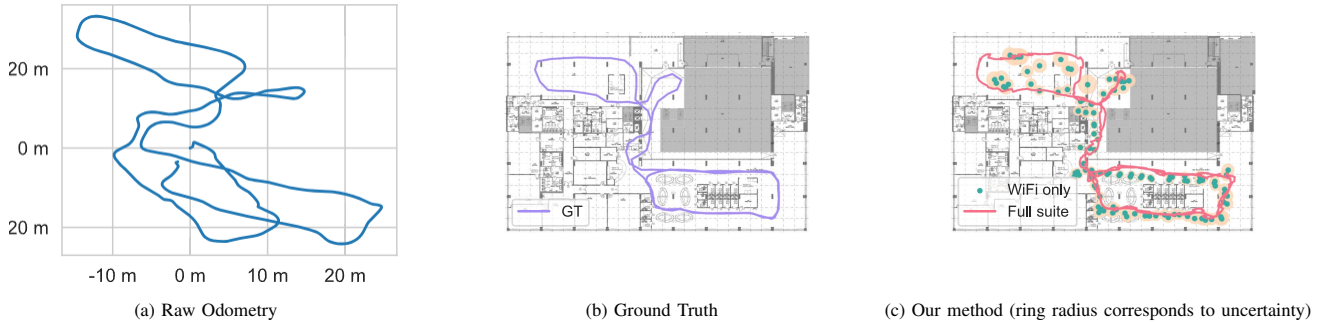


Fig. 4: Qualitative results of our method in the library (Building C).

TABLE I: WiFi Localization Performance Comparison

Site	Method	Mean	Median	Max	95th
Building A	WKNN [38]	6.0200	5.0613	40.5234	14.1351
	DSAE [39]	5.2531	4.4546	19.8080	11.8786
	Ours	2.5706	2.1238	14.4922	6.3661
Building B	WKNN [38]	4.3299	3.4623	19.5425	10.3376
	DSAE [39]	5.5253	5.1184	22.1047	11.5780
	Ours	3.3942	2.4951	25.3909	9.1316
Building C	WKNN [38]	7.8455	6.4902	24.4383	18.9681
	DSAE [39]	2.7263	2.3757	21.7769	5.7564
	Ours	2.1596	1.7336	9.4093	5.5427

increase in testing mean error by approximately 1.5 m, suggesting the model was overfitting to the radio map. The introduction of the RSS Distribution Map as an input appears to serve both as an initializer and as a regularizer. As shown in Fig.5, we present a 2D projection of the input embeddings using UMAP [40] for ease of visualization. we demonstrated that MACs detected in close temporal proximity were similarly colored. The RSS Distribution Map enabled the model to capture both global and local topological structures and correlations between proximate MACs, unlike the baseline approach, which isolated MACs into separate spaces.

2) *Uncertainty Estimation*: We conducted another experiment on one building to examine the impact of uncertainty on our models' performances. Table II highlights the enhancement of over 10% for both WiFi and Fusion positioning results. By using the metric of Pearson product-moment correlation, we found that the correlation between errors and uncertainty in all environments exceeded 0.3, indicating a relatively strong relationship. Incorporating uncertainty offers dual advantages: it prevents the WiFi model from overfitting to the noisy radio map data and also refines the precision of Fusion's location predictions. The fusion model trusts the WiFi-based location estimates more while

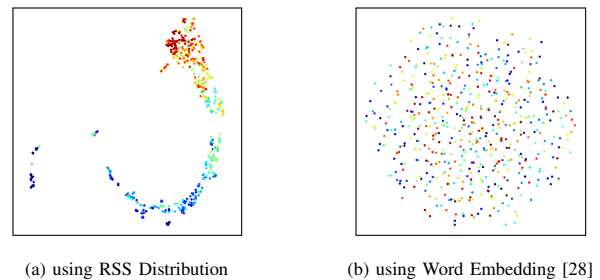


Fig. 5: By using the RSS distribution for input embedding, the resultant MAC embeddings are more discriminative where we find physically closer MAC addresses (denoted in similar colors) are closer in the projected 2D feature space.

TABLE II: Uncertainty Integration Performance Comparison

Methods	Mean	Median	Max	95th
WiFi (with Unc)	2.5706	2.1238	14.4922	6.3661
WiFi (without Unc)	3.1406	2.6301	21.9450	8.0191
Fusion (with Unc)	1.8558	1.4207	7.2425	4.1152
Fusion (without Unc)	2.1558	1.5880	7.1923	4.8716

remaining cautious of those with high uncertainty.

VIII. CONCLUSION

This research explored the potential of replacing traditional floor plans in Indoor Positioning Systems (IPS) with crowdsourced radio maps. Our findings highlighted radio maps as a viable alternative, given their inherent movement regularities. Despite challenges like inaccuracies and sparse coverage, our proposed framework, combining an uncertainty-aware neural network with Bayesian filtering, showcased a promising improvement of $\sim 25\%$ over current IPS methods. As the digital landscape evolves, leveraging radio maps could redefine indoor localization. Future endeavors will focus on refining uncertainty models and exploring advanced fusion techniques to further enhance IPS efficacy.

REFERENCES

- [1] S. Herath, S. Irandoust, B. Chen, Y. Qian, P. Kim, and Y. Furukawa, "Fusion-dhl: Wifi, imu, and floorplan fusion for dense history of locations in indoor environments," 2021.
- [2] Z. Xiao, H. Wen, A. Markham, and N. Trigoni, "Lightweight map matching for indoor localisation using conditional random fields," in *IPSN-14 proceedings of the 13th international symposium on information processing in sensor networks*. IEEE, 2014, pp. 131–142.
- [3] B. Huang, Z. Xu, B. Jia, and G. Mao, "An online radio map update scheme for wifi fingerprint-based localization," *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 6909–6918, 2019.
- [4] C. Yang and H.-r. Shao, "Wifi-based indoor positioning," *IEEE Communications Magazine*, vol. 53, no. 3, pp. 150–157, 2015.
- [5] Z. Chen, Q. Zhu, and Y. C. Soh, "Smartphone inertial sensor-based indoor localization and tracking with ibeacon corrections," *IEEE Transactions on Industrial Informatics*, vol. 12, no. 4, pp. 1540–1549, 2016.
- [6] Wigle - wireless geographic logging engine. [Online]. Available: <https://wigle.net/>
- [7] P. Bahl and V. Padmanabhan, "Radar: an in-building rf-based user location and tracking system," in *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064)*, vol. 2, 2000, pp. 775–784 vol.2.
- [8] C. Wu, Z. Yang, Y. Liu, and W. Xi, "Will: Wireless indoor localization without site survey," *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 4, pp. 839–848, 2013.
- [9] M. Lee, H. Yang, D. Han, and C. Yu, "Crowdsourced radiomap for room-level place recognition in urban environment," *2010 8th IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pp. 648–653, 2010.
- [10] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" 2017.
- [11] K. Cai, C. X. Lu, and X. Huang, "Stun: Self-teaching uncertainty estimation for place recognition," 2022.
- [12] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia, "On the uncertainty of self-supervised monocular depth estimation," 2020.
- [13] A. Taha, Y.-T. Chen, T. Misu, A. Shrivastava, and L. Davis, "Unsupervised data uncertainty learning in visual retrieval systems," 2019.
- [14] D. Feng, L. Rosenbaum, and K. Dietmayer, "Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection," 2018.
- [15] B. Ferris, D. Fox, and N. D. Lawrence, "Wifi-slam using gaussian process latent variable models." in *IJCAI*, vol. 7, no. 1, 2007, pp. 2480–2485.
- [16] P. Mirowski, T. K. Ho, S. Yi, and M. MacDonald, "Signalslam: Simultaneous localization and mapping with mixed wifi, bluetooth, lte and magnetic signals," in *International Conference on Indoor Positioning and Indoor Navigation*. IEEE, 2013, pp. 1–10.
- [17] H. Zou, Z. Chen, H. Jiang, L. Xie, and C. Spanos, "Accurate indoor localization and tracking using mobile phone inertial sensors, wifi and ibeacon," in *2017 IEEE International Symposium on Inertial Sensors and Systems (INERTIAL)*, 2017, pp. 1–4.
- [18] F. Hong, Y. Zhang, Z. Zhang, M. Wei, Y. Feng, and Z. Guo, "Wap: Indoor localization and tracking using wifi-assisted particle filter," in *39th Annual IEEE Conference on Local Computer Networks*, 2014, pp. 210–217.
- [19] J. L. Carrera, Z. Zhao, T. Braun, and Z. Li, "A real-time indoor tracking system by fusing inertial sensor, radio signal and floor plan," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2016, pp. 1–8.
- [20] H. Hellmers, A. Norrdine, J. Blankenbach, and A. Eichhorn, "An imu/magnetometer-based indoor positioning system using kalman filtering," in *International Conference on Indoor Positioning and Indoor Navigation*, 2013, pp. 1–9.
- [21] A. Poullose and D. S. Han, "Indoor localization using pdr with wi-fi weighted path loss algorithm," in *2019 International Conference on Information and Communication Technology Convergence (ICTC)*, 2019, pp. 689–693.
- [22] 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, 2020.
- [23] L. Ma, Y. Fan, Y. Xu, and Y. Cui, "Pedestrian dead reckoning trajectory matching method for radio map crowdsourcing building in wifi indoor positioning system," in *2017 IEEE International Conference on Communications (ICC)*. IEEE, 2017, pp. 1–6.
- [24] M. Li and A. Mourikis, "High-precision, consistent ekf-based visual-inertial odometry," *The International Journal of Robotics Research*, vol. 32, pp. 690–711, 05 2013.
- [25] S. Xia, Y. Liu, G. Yuan, M. Zhu, and Z. Wang, "Indoor fingerprint positioning based on wi-fi: An overview," *ISPRS International Journal of Geo-Information*, vol. 6, no. 5, 2017. [Online]. Available: <https://www.mdpi.com/2220-9964/6/5/135>
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [27] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [28] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," 2013.
- [29] B. W. Silverman, *Density estimation for statistics and data analysis*. Routledge, 2018.
- [30] S. Thrun, W. Burgard, and D. Fox, "Probabilistic robotics," 2005.
- [31] T. Hester and P. Stone, "Negative information and line observations for monte carlo localization," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 2764–2769.
- [32] A. Doucet, N. De Freitas, N. J. Gordon, et al., *Sequential Monte Carlo methods in practice*. Springer, 2001, vol. 1, no. 2.
- [33] I. V. Rory Hughes, Lei Tao and F. Alshely, "Calibration-free radiomap construction based on graph map matching," 2023.
- [34] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and R. Daniela, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5135–5142.
- [35] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," 2019.
- [36] A. F. Agarap, "Deep learning using rectified linear units (relu)," 2019.
- [37] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [38] K. Hechenbichler and K. Schliep, "Weighted k-nearest-neighbor techniques and ordinal classification," *discussion paper*, vol. 399, 01 2004.
- [39] R. Wang, Z. Li, H. Luo, F. Zhao, W. Shao, and Q. Wang, "A robust wi-fi fingerprint positioning algorithm using stacked denoising autoencoder and multi-layer perceptron," *Remote Sensing*, vol. 11, no. 11, 2019. [Online]. Available: <https://www.mdpi.com/2072-4292/11/11/1293>
- [40] L. McInnes, J. Healy, and J. Melville, "Umap: Uniform manifold approximation and projection for dimension reduction," 2020.