Less is More: Improving WiFi Localization Accuracy with Fewer Scans

Yamini Shankar, Aravindh Sriram Kumar, Ayon Chakraborty
SeNSE Lab, Department of Computer Science and Engineering
IIT Madras

***Abstract*—WiFi-based indoor localization provides a practical alternative to vision and sensor-based tracking, but its adoption is limited by high energy costs from repeated signal scans. To address this, we propose LowFI, a layout-aware, low-energy WiFi localization framework that jointly optimizes the number of scans and access point (AP) selection. LowFI leverages fading in RF channels to adaptively budget scans and rank APs, reducing redundancy while maintaining accuracy. Experiments with various WiFi-enabled devices show that localization error in nLoS conditions decreases from ≈8 m with a single scan to ≈2 m with much limited number of scans, after which accuracy gains saturate. By selecting minimal scans and a modest set of APs, LowFI achieves sub-meter accuracy while cutting energy cost by up to 65%. These results demonstrate that careful scan and AP orchestration can deliver reliable localization without compromising device battery.**

I. Introduction

WiFi-based indoor localization has matured significantly over the past two decades, with both data-driven approaches and ranging-based solutions (e.g., 802.11mc Fine Time Measurement or FTM) being commercialized [1], [2]. However, WiFi FTM at typical operating bandwidths of 20 or 40 MHz remains too coarse-grained and becomes increasingly unreliable in heavy multipath environments. As a result, fine-grained location tracking often relies on dense measurements and learning-based methods.

Building on this trend, most recent advances in WiFi localization have aimed to push the limits of positioning accuracy. State-of-the-art techniques achieve this either by designing increasingly complex models, or by leveraging higher-dimensional features such as channel state information (CSI), or in some cases by customizing hardware. However, a common thread across these approaches is the assumption that localization models are consistently supplied with *high-quality RF measurements* (e.g., refined signal strength or channel state estimates), typically obtained through repeated channel scans. While such additional measurements improve accuracy, they also increase energy usage and inference latency. Beyond aggressive scanning, location inference itself is energy-intensive. As models grow in complexity, running them continuously on resource-constrained devices becomes increasingly impractical [3], [4]. Offloading computations to the edge is an alternative, but streaming measurement introduces its own costs, consuming additional power and adding to network contention. Thus, energy efficiency emerges as a critical challenge for mobile and IoT devices that require continuous localization support while preserving battery life.

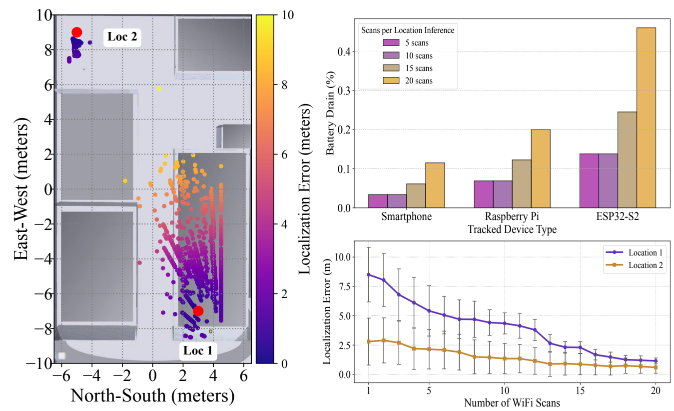

Fig. 1. Motivating experiments on scan efficiency for WiFi localization. **Left:** Lab floorplan with Loc1 (cluttered) and Loc2 (near APs), showing predicted positions as scatter plots. **Right (bottom):** Accuracy vs. number of scans, highlighting non-linear and location-dependent convergence. **Right (top):** Battery usage across three device classes (smartphone, ESP32-S2, Raspberry Pi), showing higher scan counts lead to greater energy drain.

In this paper, we address the critical challenge of energy-efficient WiFi localization, where the number of signal scans plays a central role. Triggering a high number of scans without regard to environmental context is wasteful: while additional scans may improve data fidelity, they also incur significant battery drain and add to overall delay. Importantly, what qualifies as a *statistically significant* set of measurements is highly context-dependent. Multipath-rich environments often require more scans, whereas line-of-sight regions or locations close to access points may need far fewer. Applying a uniform scanning strategy across such diverse conditions leads to inefficiency. This raises a key question: *Can the quality of RF data be assessed in real time, enabling context-aware channel scans that minimize energy consumption while sustaining reliable localization performance?*

To motivate this problem, we conducted experiments in a 20 m×12 m area in and around our lab. We first consider the effect of using only a single WiFi scan as input to a pre-trained localization model. Fig. 1 (left) illustrates the spatial variability in localization performance: at one test location (Loc_1), which is shielded by walls and surrounded by metal shelving, the median error is around 8 m. In contrast, at another location (Loc_2) situated close to two access points and with minimal surrounding clutter, the median error was reduced to about 2 m. This demonstration highlights how the same number of scans can lead to very different levels of accuracy depending on the underlying channel's attenuation and multipath characteristics. We further analyze how localization accuracy improves as additional scans are accumulated (Fig. 1 *(right bottom)*). Our results show a non-linear improvement curve, with the rate of convergence varying across locations: environments with heavy multipath require more samples to yield a representative RF signature, while cleaner line-of-sight zones achieve stable accuracy with far fewer scans. We observe that averaging scans up to 7 for $\approx 6-7$ APs reduces the localization error by nearly 75% in nLoS (from $\approx 8\,\mathrm{m}$ to $\approx 2\,\mathrm{m}$). Beyond this point, the accuracy gain is marginal—less than 10% additional improvement - while the power consumption continues to rise (Fig. 4). Finally, we study the energy impact of scanning across three classes of devices - a smartphone, an embedded microcontroller (ESP32-S2), and an embedded computer (Raspberry Pi). In all cases, we observed a clear trend: the greater the number of scans used per location update, the higher the battery drain. Given that localization services typically run as background processes on battery-constrained mobile and IoT devices, it is crucial that scan scheduling be performed strategically to extend device lifetime. In fact, increasing the scans from 7 to 20 raises the energy cost by over 190% (from $\approx 0.08\,\mathrm{mAh}$ to $\approx 0.232\,\mathrm{mAh}$ (Fig. 7)). This demonstrates that indiscriminate scanning wastes substantial energy without yielding commensurate accuracy benefits.

In this paper, we introduce LowFI, an end-to-end system that addresses the problem of energy-efficient localization by dynamically orchestrating WiFi scans and access point (AP) selection. It operates in three stages, building progressively from a coarse estimate to a refined prediction while minimizing redundant measurements.

Coarse-grained localization. The process begins with a single WiFi scan that records RSS values from neighboring APs, producing a coarse-grained location estimate via a pre-trained localization model. The precision of this estimate naturally varies depending on local multipath conditions. Given that the layout of the environment is assumed to be known, we use this coarse prediction to approximate the expected level of fading and uncertainty in RF measurements at that region.

Fading estimation via lightweight ray tracing. Since the initial prediction is coarse, we expand our search to the neighborhood around the predicted location. For each candidate point in this zone, we perform lightweight ray tracing to approximate multipath effects between APs and the receiver. This yields a channel impulse response (CIR), which we convert to the frequency domain to obtain CSI-like representations. We then compute the *spatial variance* of CSI across the neighborhood, which captures the heterogeneity of the zone in terms of clutter and obstacles. This variance is input to a pre-trained regression model that estimates the *temporal fading coefficient*, reflecting how unstable the channel is likely to be over time.

Scan budgeting and AP selection. The fading coefficient quantifies the uncertainty in RF measurements for each AP, enabling us to determine the minimum number of scans needed using a z-score–based analysis. Crucially, not all APs are equally useful: including all of them may introduce redundancy or noise. To address this, we define a joint utility that simultaneously (i) minimizes the number of scans, and (ii) selects the subset of APs that contribute the most independent information. Independence among APs is enforced using the Fisher score, ensuring that selected APs are not highly correlated. Once the optimal number of scans and AP set are chosen, the system collects the required measurements, which are then fed into the localization model to yield a fine-grained position estimate.

In this paper, we make the following contributions:

- We cast WiFi localization as a joint problem of determining the minimum number of scans and selecting informative access points, balancing positioning accuracy against energy cost.
- We develop an end-to-end pipeline that (i) derives coarse location from a single scan, (ii) characterizes environmental heterogeneity via lightweight ray tracing and CSI variance, and (iii) uses fading-aware analysis to determine scan budgets and AP subsets in real time.
- We implement the system on multiple device platforms (smartphone, ESP32-S2, Raspberry Pi) and evaluate it across diverse indoor environments. Our results demonstrate significant energy savings without compromising localization accuracy.

II. Background and Problem Formulation

A. Related Works and Existing Gaps

Wi-Fi based localization systems are power-hungry, primarily due to two factors: the overhead of active scanning and the energy demands of running inference models, either locally on the device or through offloading to edge/cloud servers [5]–[7]. Even for RSSI-based systems [8], active scanning incurs significant power consumption, and this cost escalates dramatically in CSI-based approaches, where fine-grained channel measurements are collected at high frequency [9]. Moreover, CSI-based models are considerably larger and more complex, which further increases the inference energy footprint [10].

Several studies have quantified these inefficiencies: [11] report that aggressive probing in dense deployments accelerates device energy drain by up to 44%. On the other hand, reducing the number of scans can indeed conserve energy, but doing so risks outdated channel information and degraded inference accuracy, ultimately leading to larger localization errors [12], [13]. This illustrates the unavoidable trade-off between lowering energy consumption and sustaining reliable accuracy.

To mitigate this, researchers have proposed adaptive scan scheduling techniques [6], [14] and sensor-fusion approaches that trigger scans conditionally using inertial cues [15]. While effective to some extent, these strategies primarily focus on reducing scan counts alone, without explicitly addressing redundancy in access points (APs) or anchors. A parallel research direction has therefore explored AP/anchor selection, where only a subset of stable or strategically placed anchors are used for inference. Works [16]–[18] show that discarding unstable APs, prioritizing long-term stable signals, or optimizing anchor placement based on geometric diversity can significantly improve efficiency. More recent works [18]–[20] further demonstrate that anchor optimization based on criteria such as minimum standard deviation or minimum error propagation, often combined with particle filtering, enhances robustness and reduces error spread in distributed systems.

Despite these advances, existing methods remain fragmented: scan reduction and anchor selection are usually considered in isolation. What is missing is a unified algorithmic framework that jointly addresses both dimensions to achieve deployment-time efficiency. In this work, we address precisely this gap. Given a pre-trained localization model and layout information, our framework adaptively determines how many APs to use and how many RSSI samples to collect to meet a target confidence.

B. Motivation and Problem Statement

The crux of our work is to understand how the number of signal measurements influences localization accuracy, and to delve deeper into the principles that govern such tradeoff. In the arena shown in Fig. 2, we first measure how many samples are required at each location to achieve 90% average confidence (across all APs) in estimating the true RSS mean. The reference mean was computed by collecting 50 measurements at each of 150 grid points (1 m × 1 m). The resulting heatmap reveals that some regions require far more scans than others: areas with strong multipath demand many samples due to high signal variance, while line-of-sight regions converge quickly with only a few scans.

A second observation concerns the role of access points (APs). Not all APs are equally informative for a given location. For instance, at the three marked locations in Fig. 2, AP 1 fails to discriminate among them, whereas AP 4 provides clearly separable distributions. Thus, localization accuracy depends not only on the number of scans but also on which APs are considered. Including all APs is not always beneficial: with a limited scan budget, noisy or redundant APs may increase confusion rather than improve accuracy.

Problem statement. Let $\mathcal{S} = \{1, \ldots, N\}$ be the set of AP indices and $L \in \mathcal{L}$ the (unknown) true location. For a subset $\mathcal{S}' \subseteq \mathcal{S}$ and $k \in \mathbb{N}$ scans, the localizer outputs

$$\hat{L}^k\mathcal{S}' \;=\; \mathcal{M}_{\text{loc}}\big(\overline{\mathbf{RSS}}^k_{\mathcal{S}'} \text{ or } \overline{\mathbf{CSI}}^k_{\mathcal{S}'}\big),$$

where $\mathcal{M}_{\text{loc}}$ is the localization model described in Sec. III.

Our goal is to identify: (i) the smallest number of scans $k^\star$, and (ii) a subset of APs $\mathcal{S}^\star \subseteq \mathcal{S}$, such that the resulting estimate $\hat{L}^{k^\star}\mathcal{S}^\star$ achieves the desired confidence in predicting L. In addition, $\mathcal{S}^\star$ must be informative and non-redundant, i.e., satisfy an independence criterion $\mathcal{J}\mathcal{S}^\star \geq \beta$, where β is a user-defined threshold on acceptable diversity. This ensures that $\mathcal{S}^\star$ includes only APs that contribute complementary information rather than redundant or highly correlated signals.

The outputs $\mathcal{S}^\star, k^\star$ thus jointly specify *which APs to consider* and *how many scans to perform*. This sets the stage for algorithmic pipeline behind LowFI.

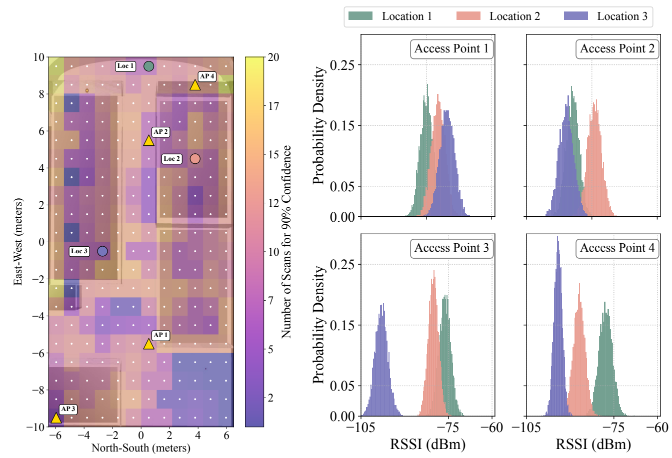

Fig. 2. Heatmap overlaid on the area layout, shows the number of scans required at each location to achieve 90% prediction confidence. Regions with higher multipath demand more scans due to increased signal variance, while low multipath regions require fewer scans for reliable localization

III. Methodology

As discussed earlier, conventional deployments rely on prolonged scanning to gather many measurements, improving inference confidence but at the expense of significant battery drain. We propose a power-aware framework: LowFI, that *adaptively* reduces energy in Wi-Fi localization by jointly selecting (i) which access points (APs) to use and (ii) how many RSSI scans to collect at run time, while meeting a target confidence for localization.

A. Measurement Database and Models

Our framework builds on pre-existing components established through an initial bootstrapping and crowd-sourcing phase. At its core is the *Measurement Database* ($\mathcal{D}_{\text{meas}}$), which stores wireless signatures and environmental layout information. This database, along with the pre-trained localization model ($\mathcal{M}_{\text{loc}}$) and fading predictor ($\mathcal{M}_{\text{fade}}$), forms the foundation for all subsequent operations.

Measurement Database ($\mathcal{D}_{\text{meas}}$). A repository of crowd-sourced signal signatures collected over time from multiple devices at known reference locations. For each location L, the database maintains an RSS vector $\mathbf{RSS}_L \in \mathbb{R}^N$, where N is the number of access points (APs), along with a CSI matrix $\mathbf{CSI}_L \in \mathbb{R}^{N\times 64}$, where each row corresponds to the CSI estimate from one AP across 64 subcarriers.

All measurements are processed (e.g., outlier removal, updating means) as new data arrives, ensuring that the repository reflects the latest channel conditions. In addition, we assume access to a floor-plan layout of the deployment area, represented as a fine-grained occupancy matrix MAP where each cell denotes either free space or obstruction. This assumption is reasonable in infrastructure-oriented setups (e.g., enterprise or campus Wi-Fi) and provides useful context for estimating multipath effects.

Localization model ($\mathcal{M}_{\text{loc}}$). The localization model $\mathcal{M}_{\text{loc}}$ maps incoming wireless measurements to a predicted location

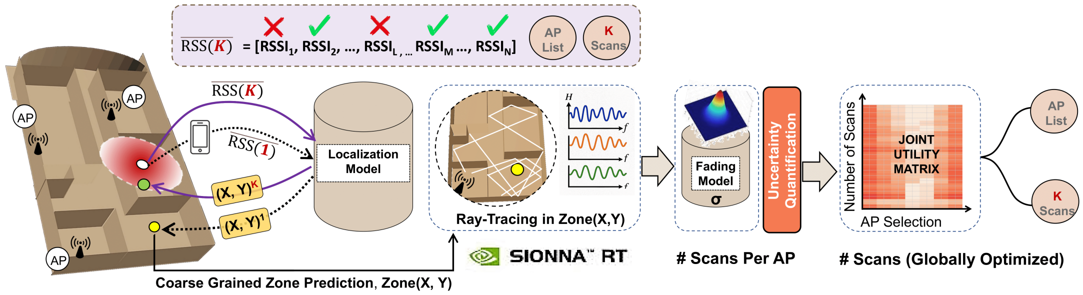

Fig. 3. Overall workflow of LowFI. A single Wi-Fi scan provides a coarse zone-level location estimate X, Y^1. This estimate guides a ray-tracing step to capture multipath effects within the zone, which determines scan variability. Using this information, LowFI quantifies uncertainty through the fading model and then jointly optimizes both the number of scans and the subset of access points (APs) to be used, yielding an energy-efficient yet reliable localization configuration.

$\hat{L}^k$. At runtime, the model operates on aggregated measurements: the aggregate of k RSS vectors, denoted by $\overline{\mathbf{RSS}}^k_L$, or optionally for k CSI matrices, denoted by $\overline{\mathbf{CSI}}^k_L$, where L denotes the *true* measurement location. Let $\mathcal{L}$ denote the set of all discrete candidate location centeroids defined on MAP at a certain spatial resolution ρ (e.g., $1\,\mathrm{m} \times 1\,\mathrm{m}$), depending on specific application requirements and training data availability. For an input of k aggregated measurements, the model produces a probability distribution $p^k_{L_i}$ over all candidate locations $L_i \in \mathcal{L}$. The predicted location is then obtained as,

$$\hat{L}^k = \mathcal{M}_{\text{loc}}\big(\overline{\mathbf{RSS}}^k_L \text{ or } \overline{\mathbf{CSI}}^k_L\big) = \mathop{}_{L_i \in \mathcal{L}} p^k_{L_i} \cdot L_i. \quad (1)$$

As is common in the literature, the localization error decreases with larger k, since more measurements reduce uncertainty. The model itself may be realized using a simple neural network or other standard classifiers. Importantly, the focus of this paper is not on designing a new localization model, but on optimizing how such models are used in deployment through energy-efficient scan orchestration.

Fading predictor ($\mathcal{M}_{\text{fade}}$). In addition to localization, our framework employs a model to predict fading coefficients within a given zone. The intuition is that spatial fluctuations of CSI across nearby points encode the degree of environmental heterogeneity, which in turn governs temporal fading on each AP link. Formally, for a candidate zone $\mathcal{Z}$ around location L, we collect CSI measurements $\{\mathbf{CSI}_{L'}\}$ for all $L' \in \mathcal{Z}$. For each anchor (AP) ℓ, the regression model $\mathcal{M}_{\text{fade}}$ maps the spatial variation of $\{\mathbf{CSI}^{\ell}_{L'}\}$ to a temporal fading coefficient $\hat{\sigma}_\ell$, i.e.,

$$\hat{\boldsymbol{\sigma}_{\mathcal{Z}}} = \mathcal{M}_{\text{fade}}\big(\{\mathbf{CSI}_{L'} \mid L' \in \mathcal{Z}\}\big),$$

where $\hat{\boldsymbol{\sigma}_{\mathcal{Z}}} = \hat{\sigma}_1, \ldots, \hat{\sigma}_N$ and each $\hat{\sigma}_\ell$ represents the predicted fading severity on link ℓ. While in principle the model could output parameters of a fading distribution (e.g., Nakagami or Rayleigh), for simplicity in this paper we use the standard deviation of received signal amplitude as the fading coefficient.

B. Functional Pipeline of the LowFI System

Before deployment, LowFI is bootstrapped with the necessary database and pre-trained models. These models are periodically updated and do not directly influence the real-time operation of the system. Post-deployment, the framework operates in three stages. It begins by generating a coarse location estimate from a single WiFi scan, followed by a neighborhood search and ray tracing to capture multipath effects. Using this information, the system estimates the fading characteristics in the environment, which influences the uncertainty of the RF measurements. Finally, the framework optimizes the number of scans and selects the most informative access points (APs) based on their contribution to reducing uncertainty. This process progressively refines the location estimate while minimizing energy consumption and redundant measurements.

(i) Coarse-grained localization ($\hat{L}$). Given that the original device location $L \in \mathcal{L}$, the process begins with a single WiFi scan that records the $\mathbf{RSS}_L \in \mathbb{R}^N$, where N is the number of access points. This measurement is passed to the localization model $\mathcal{M}_{\text{loc}}$, which provides a coarse-grained estimate, denoted as $\hat{L}$, of the device's location. Depending on the amount of multipath, the error $|L - \hat{L}|$ can vary. WiFi Fine Time Measurement (FTM) can also generate this coarse estimate, though it may be less accurate at 20 or 40 MHz bandwidths, especially in environments with interference.

(ii) Estimating number of scans. In order to decide how many Wi-Fi scans are necessary for reliable localization, we first need to quantify the uncertainty of the received signals. After obtaining the coarse location estimate $\hat{L}$, we define a zone L_{zone} centered at $\hat{L}$ with radius R, where R corresponds to the median localization error observed in historical data. The choice and selection of R is thoroughly discussed in Section IV. Within this zone, a lightweight ray-tracing module (e.g., Sionna RT) is used to generate CSI-like representations for candidate points. Although material properties are not modeled explicitly, including them could further improve amplitude predictions in future extensions.

The spatial variance of CSI across L_{zone} is then fed into the fading predictor $\mathcal{M}_{\text{fade}}$, which outputs, $\hat{\boldsymbol{\sigma}}_{\mathcal{Z}} = \hat{\sigma}_1, \ldots, \hat{\sigma}_N$, where each $\hat{\sigma}_\ell$ represents the expected temporal standard deviation (fading coefficient) of the signal from AP$_\ell$. Larger values of $\hat{\sigma}_\ell$ indicate stronger multipath and greater variability in the received signal.

Number of scans. We now connect these fading estimates to the number of scans required. For the true location L and anchor ℓ, let the averaged RSS measurement from k scans be denoted as $\overline{\mathbf{RSS}}^{k}_{L,\ell}$. Assuming that RSS (in dB) is approximately Gaussian, we have $\overline{\mathbf{RSS}}^{k}_{L,\ell} \sim \mathcal{N}\big(\mu_{L,\ell}, \frac{\hat{\sigma}^2_{L,\ell}}{k}\big)$ where $\mu_{L,\ell}$ is the true mean RSS from AP_ℓ at location L. To bound the estimation error of the mean within ε_ℓ with confidence level $1-\alpha$, we require,

$$\Pr\big(|\,\overline{\mathbf{RSS}}^{k}_{L,\ell} - \mu_{L,\ell}\,| \;\le\; \varepsilon_\ell\big) \;\ge\; 1-\alpha, \tag{2}$$

which leads to the condition

$$k_{L,\ell} \;\ge\; \big(\frac{z_{1-\alpha 2}\,\hat{\sigma}_{L,\ell}}{\varepsilon_\ell}\big)^2 \tag{3}$$

where $z_{1-\alpha 2}$ is the $1-\alpha 2$ quantile of the standard normal distribution. Hence, for each AP we can estimate the expected scan requirement $k_{L,\ell}$. In practice, consecutive measurements may exhibit mild temporal correlation, which slightly relaxes the independence assumption in eqn. 3. Collectively, these form the scan-budget vector $\mathbf{N}^{L}_{\text{MAX-SCANS}} = k_{L,1}, k_{L,2}, \ldots, k_{L,N}$, which specifies the maximum scans (expected) needed per AP to achieve the target confidence level $1-\alpha$ at location L. In the next stage, this vector is used to discard APs that remain too noisy and to guide joint scan budgeting and AP selection.

(iii) Scan budgeting and AP selection. In practice, the vector $\mathbf{N}^{L}_{\text{MAX-SCANS}}$ often contains a wide spread of values: some entries are very small indicating stable links, while others are extremely large reflecting noisy or highly variable links. A naive hard-thresholding strategy, for instance, discarding all APs above a fixed $k_{L,\ell}$ or restricting to a fixed number of APs, performs poorly, as it either eliminates informative APs or wastes effort on redundant ones. Instead, the challenge is to strike a balance: we must minimize the total number of scans while retaining those APs that provide the most value for localization. This requires quantifying not only how reliable each AP is, but also how *independent* its contribution is from the others. Including many highly correlated APs adds little new information, whereas carefully chosen APs with complementary characteristics maximize the total information gain. Thus, the objective of this stage is to jointly select a subset of APs and allocate the scan budget so as to achieve the desired confidence with the fewest possible measurements.

We design a joint AP–scan selection procedure that balances informativeness against redundancy. For our zone of interest (coarse-grained location), Z, each AP ℓ is first assigned a Fisher score F^{Z}_{ℓ}, which measures the ratio of between-zone separability to within-zone fading variability:

$$F^{Z}_{\ell} \;=\; \frac{\text{Var}_{Z\in\mathcal{Z}}\,\mu_{Z,\ell}}{\frac{1}{|\mathcal{Z}|}\;{}_{Z\in\mathcal{Z}}\,\hat{\sigma}_{Z,\ell}{}^{2}} \tag{4}$$

where $\mu_{\ell,z}$ and $\hat{\sigma}_{Z,\ell}$ denote the mean and standard deviation of AP ℓ's RSS in zone Z. A large F^{Z}_{ℓ} indicates that the AP provides strong discrimination across zones while being relatively stable within a zone. Although other measures of independence could be used (e.g., mutual information or principal component analysis), we adopt the Fisher score because it is simple, interpretable and directly relates to separability under fading variability. APs are then ranked in descending order of F^{Z}_{ℓ} to form a candidate list π. Starting with an empty set $\mathcal{S}$, APs are added one by one from π according to their *net gain*:

$$\Delta^{Z}_{j} \;=\; F^{Z}_{j} - \gamma \;{}_{\ell\in\mathcal{S}}\,|\rho_{\ell j}|, \tag{5}$$

where $\rho_{\ell j}$ is the correlation between APs ℓ and j, and γ is a penalty parameter. Since the measurement database provides RSS/CSI fingerprints across all locations, $\rho_{\ell j}$ can be computed as the Pearson correlation between their signal vectors over $\mathcal{L}$. Alternatively, it can be estimated indirectly from $\mathcal{M}_{\text{loc}}$ by comparing how strongly the two APs co-influence the predicted location classes, thus capturing redundancy at the inference level. Both F^{Z}_{ℓ} and $\rho_{\ell j}$ are computed spatially across all reference locations.

Algorithm 1: Joint AP–Scan Selection (Zone Z)

Input:
Scan-budget: $\mathbf{N}^{Z}_{\text{MAX-SCANS}}$
Measurements: $\mathcal{D}_{\text{meas}}$
Gain threshold: τ
Max. scan constraint: $K_{\max}$
Output: Selected AP subset $\mathcal{S}^\star$; scan budget $k^\star$

Initialize: $\mathcal{S}^\star \leftarrow$; $\quad k^\star \leftarrow 0$; $\quad \mathcal{S} \leftarrow \{1, 2, \ldots, N\}$;
$\{F^{Z}_{\ell}\}, \{\rho_{\ell j}\} \leftarrow \mathcal{D}_{\text{meas}}$
$\pi \leftarrow$ AP indices sorted in descending order of F^{Z}_{ℓ};
for $j \in \pi$ **do**
  $\Delta^{Z}_{j} \leftarrow F^{Z}_{j} - \gamma \;{}_{\ell\in\mathcal{S}^\star}\,|\rho_{\ell j}|$
  if $\Delta^{Z}_{j} \ge \tau$ **then**
    $\mathcal{S}^\star \leftarrow \mathcal{S}^\star \cup \{j\}$ $k^\star \leftarrow \max\big(k^\star,\ k_{L,j}\big)$
    if $k^\star > K_{\max}$ **then**
      $\mathcal{S}^\star \leftarrow \mathcal{S}^\star \setminus \{j\}$ $k^\star \leftarrow \max_{\ell\in\mathcal{S}^\star} k_{L,\ell}$
return $\mathcal{S}^\star$, $k^\star$

Intuitively, the net gain Δ^{Z}_{j} balances two factors: the Fisher score F^{Z}_{j}, which rewards an AP for its ability to distinguish between zones, and the redundancy term, which penalizes j if it is highly correlated with the APs already included in the current set $\mathcal{S}$. The process begins with $\mathcal{S}$ = . At each step, the next candidate j from the ranked list π is considered. Its net gain Δ^{Z}_{j} is computed relative to the APs already in $\mathcal{S}$, i.e., $\ell \in \mathcal{S}$ (see eqn. 5). If $\Delta^{Z}_{j} \ge \tau$, the AP is accepted and added to $\mathcal{S}$; otherwise it is discarded. This greedy procedure repeats until either no more APs pass the threshold or the AP budget $m_{\max}$ is reached.

The outcome of this selection is a subset $\hat{\mathcal{S}}$ of $m^\star = |\hat{\mathcal{S}}|$ APs that maximize zone separability while minimizing redundancy. Once $\hat{\mathcal{S}}$ is determined, the fading coefficients $\{\hat{\sigma}_\ell\}$ for $\ell \in \hat{\mathcal{S}}$ are substituted into the confidence bound in eqn. 3, yielding per-AP scan requirements $k_{L,\ell}$. Since the system must guarantee a joint confidence of at least $1-\alpha$, the scan budget is set conservatively as the worst case across the selected APs, $k_L = \max_{\ell \in \hat{\mathcal{S}}} k_{L,\ell}$.

Thus, the final configuration $\mathcal{S}^\star, k^\star$ emerges naturally from the two stages: $\mathcal{S}^\star = \hat{\mathcal{S}}$ is the subset of APs chosen by the Fisher–redundancy tradeoff, and $k^\star = k_L$ is the scan budget derived from their fading coefficients. Together, $\mathcal{S}^\star, k^\star$ specifies both *which* anchors to listen to and *how many* measurements to collect, ensuring localization accuracy with confidence $1-\alpha$ while avoiding redundant scans and unnecessary energy expenditure. We summarize our scan-budgeting and AP-selection technique in Algorithm 1. In our implementation, the complete LowFI pipeline runs in approximately 0.16 s for an indoor setup with 10 APs on a standard CPU. This confirms that the framework is suitable for real-time deployment in typical indoor environments.

IV. Results

A. Experimental Setup and Dataset:

We evaluate our framework in a 20×12 m indoor arena instrumented with ten static access points (APs) at fixed locations. A grid of measurement locations is shown in Fig. 2. In total, we collect data at *150+* grid points, covering both high- and low-multipath regions for RF scanning.

Devices and Data Collection. We employ three device classes: (i) Nexmon-enabled Wi-Fi devices [21](Raspberry Pi), (ii) ESP32-S2 modules [22]enabled with CSI toolkit [23], and (iii) Android smartphones. Nexmon and ESP32-CSI toolkit operate in the 2.4 GHz band with 20 MHz channels and 64 subcarriers, enabling CSI and RSSI logging at the rate of 1000Hz and 100Hz respectively; Android phones collect periodic RSSI scans at rate of 1 scan per 5 seconds via the native Wi-Fi framework. We use all three devices for RSSI measurements and power analysis, while CSI is obtained only from Raspberry Pi(s) and ESP32(s). At each grid location, we record continuously for 60 s. This corresponds to roughly $60{,}000$ CSI entries from Nexmon, $6{,}000$ CSI entries from ESP32, and about 30 RSSI scans from Android smartphones per location. Together, these synchronized traces yield stacked CSI snapshots (64 subcarriers) and RSSI vectors that serve as input for training the baseline localizer $\mathcal{M}_{loc}$. [1]

LoS versus nLoS classification. We distinguish LoS and nLoS using the measured fading deviation: points with $\sigma < 2.5$ dB are labeled LoS (dominant direct paths), while $\sigma \geq 2.5$ dB indicates nLoS (multipath-rich). This threshold was set via ROC analysis to balance errors. In our arena (Fig. 2), about 40% of locations are LoS and 60% nLoS, covering both corridor-facing points and cluttered, wall-shielded rooms.

3D Simulation using Ray Tracing. Accurate evaluation of our approach requires modeling multipath propagation effects. To capture this, we recreate the arena geometry as a 3D replica in Blender [24] and import it into Sionna-RT [25], to further compute the RF ray paths. This scene is integrated into our pipeline to synthesize CSI at locations within a radius R around the zone of interest $\mathcal{Z}$. CSI synthesized from this neighborhood is used by $\mathcal{M}_{\text{fade}}$ to predict the zone fading $\hat{\sigma}_{\mathcal{Z}}$. The model $\mathcal{M}_{fade}$ is trained on the same samples collected accross 150 grid points in the same arena. The performance is discussed in later parts of this section.

[1]Android RSSI scans are captured at the device's native scan cadence; ESP32/Raspberry Pi capture CSI at device-native packet rates (hundreds to thousands of frames/s), ensuring sufficient temporal diversity without over-constraining to a fixed rate.

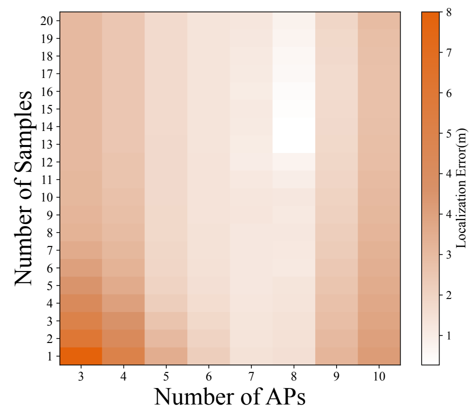

Fig. 4. Trade-off between links(number of APs) and samples. Heatmap of median localization error (m) as a function of the number of APs (x-axis) and the number of RSSI samples averaged (y-axis). Error drops sharply when moving from 3 to 4 APs to ≥ 5 APs, while increasing samples beyond ≈ 10 yields diminishing returns—useful for choosing the minimum AP-sample budget that meets a target error.

Battery Profiling. We report end-to-end battery costs that include not only active scanning but also CSI extraction and processing overheads (dominant in CSI modes). For Android devices, power statistics are obtained using the native `dumpsys` tool, which provides fine-grained logs of system resource usage. For Nexmon and ESP32 devices, we use an external power profiler to record current over time and convert it to milliampere-hours (mAh). This unified methodology allows homogeneous comparison of RSSI, CSI, and fused localization modalities across heterogeneous devices.

B. Performance Evaluation.

As discussed, localization accuracy is jointly governed by two factors: the number of scans aggregated at each location and the number of anchors (APs) contributing measurements. Both factors come with costs: higher battery drain for scans and potential redundancy when too many APs are included. To better understand this trade-off, we evaluate the combined effect of scans and APs on accuracy.

Fig. 4 shows a heatmap of localization error for the high multipath(nLoS) setting, plotted as a function of the number of APs (horizontal axis) and the number of scans (vertical axis). Each cell in the heatmap corresponds to the median localization error observed under the respective configuration. Two key observations emerge.

First, the localization error decreases as the number of APs or the number of scans increases. This confirms that both additional links and scan averaging (larger k) contribute to mitigating severe multipath effects in nLoS settings. However, beyond 6–7 APs or 7–10 scans, the error reduction saturates, with only marginal improvements observed thereafter.

Second, the heatmap highlights the diversity–redundancy trade-off that motivates our joint AP–scan selection strategy. The lighter region in the plot shows that near-optimal accuracy can be achieved with a modest budget of $\approx$ 5 APs and $\approx$ 7 scans, avoiding the cost of excessive scanning or over-provisioning APs. This motivates the adaptive design of our framework, which explicitly targets this low-error, low-cost region of the design space.

Overall, these results demonstrate that in challenging nLoS environments, careful orchestration of scans and APs is crucial to balance accuracy against power consumption, rather than indiscriminately increasing both dimensions.

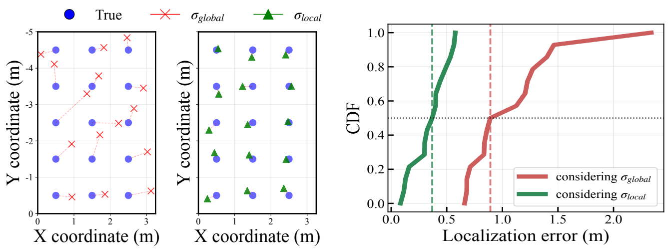

Fig. 5. Effect of zone radius R on localization. **Left**: Using a large R yields a global fading measure (σ_{global}) that ignores local spatial effects, while setting R to the median error radius produces a local fading measure (σ_{local}) that captures neighborhood-specific multipath. **Right**: CDF of localization error under both definitions; σ_{local} provides finer-grained estimates and yields consistently lower prediction error compared to σ_{global}.

Effect of Zone Radius R on Fading Estimation. An important design choice in our framework is the radius R used to define the neighborhood zone $\mathcal{Z}$ around the initially predicted location $\hat{L}$. This radius of $\mathcal{Z}$ directly determines the set of CSI samples aggregated to compute the standard deviation of the zone as $\hat{\sigma_{\mathcal{Z}}}$ as discussed in section III.

If R is chosen to be very large, the neighborhood spans the entire arena, and $\hat{\sigma}_{\mathcal{Z}}$ degenerates to a global indicator, which we denote as σ_{global}. In this case, the predicted fading coefficient reflects aggregate multipath behavior across the whole deployment, disregarding local spatial effects. In contrast, if R is chosen to be very small, the neighborhood collapses to a single point, making $\hat{\sigma}_{\mathcal{Z}}$ equivalent to the fading at that location $\hat{L}$ alone.

In our experiments, we exploit the availability of the baseline localizer $\mathcal{M}_{loc}$ to set R equal to the median localization error of $\mathcal{M}_{loc}$. This choice provides a principled compromise: the resulting standard deviation, denoted as σ_{local}, captures the multipath variability within the typical error radius of the baseline system.

Fig. 5 (top) compare the impact of these two design choices. Fig. 5 (bottom) illustrates how the predicted fading distribution differs when σ_{global} and σ_{local} are employed. Fig. 5 (right) presents the cumulative distribution function (CDF) of localization error under both configurations. The results show that using σ_{global} leads to conservative but less discriminative predictions, while σ_{local} provides finer-grained fading estimates that align more closely with the actual error distribution. This demonstrates that adopting σ_{local} yields more representative confidence modeling for zone-specific inference.

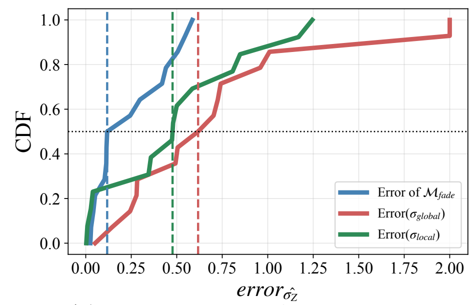

Fig. 6. Error of $\mathcal{M}_{fade}$ for predicting $\hat{\sigma_{\mathcal{Z}}}$ with σ_{global}, and with σ_{local}.

We next evaluate the accuracy of our fading prediction model $\mathcal{M}_{fade}$. The model is trained to predict the per-link standard deviation of incoming measurements from CSI-derived features, and its accuracy is measured by the absolute error between predicted and ground-truth deviation vectors.

Fig. 6 presents the cumulative distribution function (CDF) of the prediction error of $\hat{\sigma_{\mathcal{Z}}}$ under two alternative definitions of the standard deviations reference: σ_{global} and σ_{local}. The results demonstrate that adopting σ_{local} yields consistently lower prediction errors compared to σ_{global}. This confirms that σ_{local} provides a more representative and discriminative notion of fading, enabling $\mathcal{M}_{fade}$ to learn and generalize more effectively. In contrast, using σ_{global} dilutes the fading representation by averaging across the entire arena, thereby introducing noise and reducing prediction fidelity. These findings corroborate our earlier analysis: σ_{local}, defined with respect to the median localization error radius of $\mathcal{M}_{loc}$, offers the most appropriate fading target for training and evaluating $\mathcal{M}_{fade}$.

Effect of Number of Scans on Accuracy and Power Consumption. To quantify the impact of scan averaging on robustness, Fig. 7 (top) plots the CDF of the difference in localization error between nLoS and LoS predictions across increasing numbers of scans. With a single scan, the distribution is wide, and the median error gap between nLoS and LoS exceeds 5 m. As additional scans are averaged, the distributions shift leftward and tighten, with the median gap dropping to $\approx 2\ m$ at $k = 7$. In other words, taking multiple scans smooths out the scan fluctuations, hence, accuracy improves and the gap between LoS and nLoS errors becomes much smaller.

In contrast, the power consumption grows with k, starting at about 0.008 mAh for a single scan and reaching approximately 0.232 mAh at $k = 20$. This highlights that while larger k improves accuracy, the additional energy cost must be considered, with an optimal operating point lying in the range of 5–7 scans where accuracy is substantially improved while the battery budget remains practical.

Overall, these results emphasize two points: (i) only a limited number of scans and APs are needed to capture most of the accuracy improvement, depending on LoS and nLoS conditions, and (ii) σ_{local} yields more precise predictions than σ_{global}, thereby enabling efficient operation at lower scan counts.

Error-Battery Trade-off Analysis. To further evaluate the efficiency of scan usage, Fig. 8 presents a heatmap of localization error normalized by the corresponding battery drain

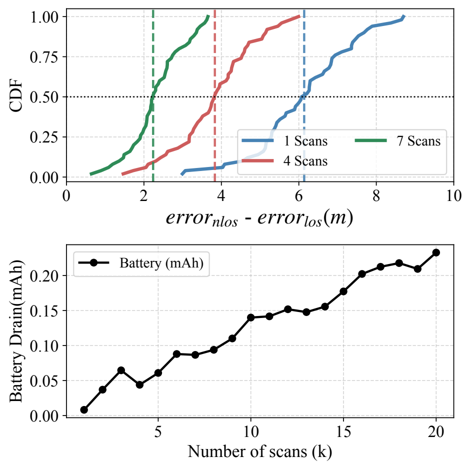

Fig. 7. **Top**: CDF of the difference in localization error for increasing numbers of cumulative RSSI scans. Temporal averaging progressively reduces the median error gap (from > 5 m with a single scan to ≈ 2 m at $k = 7$). **Bottom**: Battery consumption as a function of scan count. The battery cost rises nearly linearly (from 0.008 mAh at $k = 1$ to 0.245 mAh at $k = 20$).

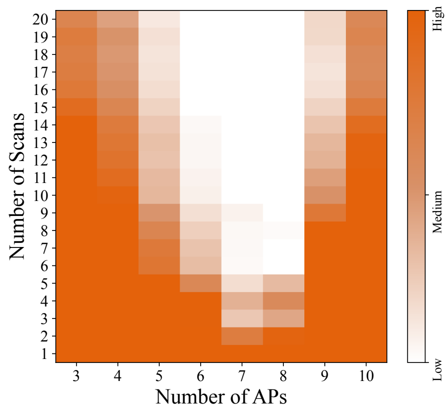

Fig. 8. Spectrogram of localization error per unit scan cost (m per mAh) as a function of the number of APs and scans. The trade-off zone is visible: few scans yield low energy cost but high error, while higher scans reduce error at the expense of energy. The mid-range (5–7 scans, 4–6 APs) achieves the best balance between accuracy and efficiency.

per scan. This highlights the effective "error per unit energy" across varying numbers of APs and scan counts. We observe a clear trade-off region: with very few scans the error is high, while with more scans the error drops but the gains soon level off even as energy cost rises. The middle band (5-7 scans with 4-6 APs) represents the most balanced operating zone, where the error is substantially reduced compared to single-scan cases, yet the scan cost remains moderate. Overall, the figure shows that adding scans beyond this range brings little accuracy benefit for the extra energy spent.

Trade-off between RSSI, CSI, and Fused Models. We primarily focus on RSSI-based localization in this work. Nevertheless, the framework can naturally be extended to incorporate CSI or fused (RSSI+CSI) inputs. To illustrate this trade-off, Fig. 9 compares accuracy and energy across modalities in a controlled LoS setting. As expected, CSI and fusion yield lower localization error than RSSI, since they capture richer channel information. However, these gains come at a steep energy cost: on ESP32-S2, CSI logging (including extraction and processing) consumes about $10\times$ more power

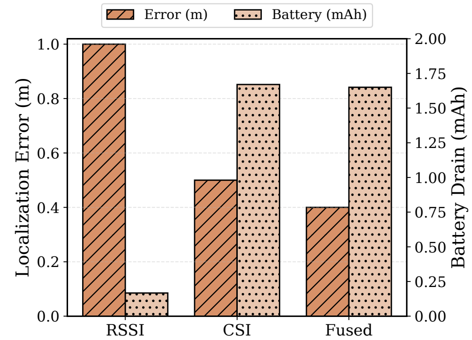

Fig. 9. Trade-off between accuracy and energy across modalities. While CSI and fusion reduce localization error compared to RSSI, they draw almost ten times more power. RSSI, though less accurate in comparison, is far more energy-efficient.

($\approx$1.67 mAh/min) than RSSI ($\approx$0.167 mAh/min), with fusion exhibiting similar drain. Thus, while CSI and fusion can provide higher accuracy in principle, RSSI strikes a far more practical balance between accuracy and device energy, which aligns with the goals of our study.

V. Conclusion

Our study highlights the design principles for building energy-aware WiFi localization systems on resource-constrained devices. Instead of treating accuracy as the sole objective or prescribing a fixed number of scans irrespective of location/layout, we show that careful orchestration of scans and anchor diversity is key to balancing accuracy and battery drain. By identifying the 'sweet spot"-roughly 5-7 scans and 4–6 APs for our setup-most of the accuracy benefits can be captured without draining device batteries.

Limitations and generalizability. Our evaluation is limited to controlled indoor spaces with static AP placement and known layouts; real-world dynamics such as network traffic, mobility, and layout changes may alter the optimal scan/AP balance. Although battery profiling across devices (Android, ESP32-S2, Raspberry Pi) may vary in absolute values, the observed trends remain consistent. Extending LowFI to dynamic, mobile, and multi-user settings—potentially by incorporating IMU cues to suppress scans during motion-is an important direction for future work.

Summary of contributions. We introduced LowFI, a layout-aware, low-energy WiFi localization framework that adaptively chooses scan count and AP subset. Across multiple devices, LowFI cuts nLoS error from $\approx$ 8 m to $\approx$ 2 m with 5 scans and achieves sub-meter accuracy in LoS, while maintaining low energy by operating in the $5-7$-scan range. Although CSI/fusion can reach $<$ 0.10 m error, they cost $\approx$ $10\times$ more energy. Thus, LowFI demonstrates that careful scan and AP orchestration enables accurate and battery-efficient wireless localization, with implications for other RF sensing modalities.

References

[1] P. E. Numan, H. Park, C. Laoudias, S. Horsmanheimo, and S. Kim, "Dnn-based indoor fingerprinting localization with wifi ftm," in *2022 23rd IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 2022, pp. 367–371.

[2] Y. Yu, R. Chen, L. Chen, S. Xu, W. Li, Y. Wu, and H. Zhou, "Precise 3-d indoor localization based on wi-fi ftm and built-in sensors," *IEEE Internet of Things Journal*, vol. 7, no. 12, pp. 11 753–11 765, 2020.
[3] C. Profentzas, M. Almgren, and O. Landsiedel, "Performance of deep neural networks on low-power iot devices," in *Proceedings of the workshop on benchmarking cyber-physical systems and Internet of things*, 2021, pp. 32–37.
[4] S. Heydari and Q. H. Mahmoud, "Tiny machine learning and on-device inference: A survey of applications, challenges, and future directions," *Sensors*, vol. 25, no. 10, p. 3191, 2025.
[5] B. Qiu, K. Feng, X. Li, H. Xiao, and Z. Zhang, "Mobility-aware user association and computation offloading in ultra-dense networks," *IEEE Transactions on Green Communications and Networking*, 2025.
[6] T. Choi, Y. Chon, and H. Cha, "Energy-efficient wifi scanning for localization," *Pervasive and Mobile Computing*, vol. 37, pp. 124–138, 2017.
[7] A. Fanariotis, T. Orphanoudakis, K. Kotrotsios, V. Fotopoulos, G. Keramidas, and P. Karkazis, "Power efficient machine learning models deployment on edge iot devices," *Sensors*, vol. 23, no. 3, p. 1595, 2023.
[8] K.-H. Kim, A. W. Min, D. Gupta, P. Mohapatra, and J. P. Singh, "Improving energy efficiency of wi-fi sensing on smartphones," in *2011 Proceedings IEEE INFOCOM*. IEEE, 2011, pp. 2930–2938.
[9] Android Developers, "Wi-fi scan overview," https://developer.android.com/topic/performance/vitals/bg-wifi, [Accessed: Apr. 13, 2025].
[10] Google Patents, "Wi-fi scan patent," https://patents.google.com/patent/WO2022148654A1/en, [Accessed: Apr. 13, 2025].
[11] X. Hu, L. Song, D. Van Bruggen, and A. Striegel, "Is there wifi yet? how aggressive probe requests deteriorate energy and throughput," in *Proceedings of the 2015 Internet Measurement Conference*, 2015, pp. 317–323.
[12] S. Wan, C. Gu, Y. Shu, and Z. Shi, "Last-seen time is critical: Revisiting rssi-based wifi indoor localization," *Signal Processing*, vol. 231, p. 109756, 2025.
[13] Y.-C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm, "Accuracy characterization for metropolitan-scale wi-fi localization," in *Proceedings of the 3rd international conference on Mobile systems, applications, and services*, 2005, pp. 233–245.
[14] M. R. Nowicki and P. Skrzypczyński, "Leveraging visual place recognition to improve indoor positioning with limited availability of wifi scans," *Sensors*, vol. 19, no. 17, p. 3657, 2019.
[15] Y. Liu, M. Dashti, and J. Zhang, "Indoor localization on mobile phone platforms using embedded inertial sensors," in *2013 10th Workshop on Positioning, Navigation and Communication (WPNC)*. IEEE, 2013, pp. 1–5.
[16] https://www.cognitivesystems.com/how-to-optimize-wi-fi-sensing-coverage/.
[17] Y. Tao, R. Yan, and L. Zhao, "An extreme value based algorithm for improving the accuracy of wifi localization," *Ad Hoc Networks*, vol. 143, p. 103131, 2023.
[18] J. Cheng, Y. Li, and Q. Xu, "An anchor node selection scheme for improving rss-based localization in wireless sensor network," *Mobile Information Systems*, vol. 2022, no. 1, p. 2611329, 2022.
[19] Y. Chen, J. Yang, W. Trappe, and R. P. Martin, "Impact of anchor placement and anchor selection on localization accuracy," *Handbook of Position Location: Theory, Practice, and Advances*, pp. 425–455, 2011.
[20] Q. Luo, C. Liu, X. Yan, Y. Shao, K. Yang, C. Wang, and Z. Zhou, "A distributed localization method for wireless sensor networks based on anchor node optimal selection and particle filter," *Sensors*, vol. 22, no. 3, p. 1003, 2022.
[21] "nexmon-csi," https://github.com/seemoo-lab/nexmon_csi, accessed: 2024-07-31.
[22] "Esp32 webpage:," https://www.espressif.com/en/products/socs/esp32, accessed: 2023-08-06.
[23] S. M. Hernandez and E. Bulut, "Lightweight and Standalone IoT Based WiFi Sensing for Active Repositioning and Mobility," in *21st International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM) (WoWMoM 2020)*, Cork, Ireland, Jun. 2020.
[24] Blender Online Community, "Blender – a 3d modelling and rendering package," https://www.blender.org/, 2024.
[25] J. Hoydis, S. Cammerer, F. Ait Aoudia, M. Nimier-David, L. Maggi, G. Marcus, A. Vem, and A. Keller, "Sionna," 2022, https://nvlabs.github.io/sionna/.