

LocaGen: Low-Overhead Indoor Localization Through Spatial Augmentation

Abdelrahman Abdelmotlb*

es-abdulrahmanm.mohammed2025@alexu.edu.eg

Alexandria University

Alexandria, Egypt

Sherif Mostafa

sherif_mostafa@aucegypt.edu

The American University in Cairo

Cairo, Egypt

Abdallah Taman*

es-abdullahm.abdullah2025@alexu.edu.eg

Alexandria University

Alexandria, Egypt

Moustafa Youssef

moustafa-youssef@aucegypt.edu

The American University in Cairo

Cairo, Egypt

Abstract

Indoor localization systems commonly rely on fingerprinting, which requires extensive survey efforts to obtain location-tagged signal data, limiting their real-world deployability. Recent approaches that attempt to reduce this overhead either suffer from low representation ability, mode collapse issues, or require the effort of collecting data at all target locations.

We present *LocaGen*, a novel spatial augmentation framework that significantly reduces fingerprinting overhead by generating high-quality synthetic data at completely unseen locations. *LocaGen* leverages a conditional diffusion model guided by a novel spatially-aware optimization strategy to synthesize realistic fingerprints at unseen locations using only a subset of seen locations. To further improve our diffusion model performance, *LocaGen* augments seen location data based on domain-specific heuristics and strategically selects the seen and unseen locations using a novel density-based approach that ensures robust coverage. Our extensive evaluation on a real-world WiFi fingerprinting dataset shows that *LocaGen* maintains the same localization accuracy even with 30% of the locations unseen and achieves up to 28% improvement in accuracy over state-of-the-art augmentation methods.

ACM Reference Format:

Abdelrahman Abdelmotlb, Abdallah Taman, Sherif Mostafa, and Moustafa Youssef. 2025. LocaGen: Low-Overhead Indoor Localization Through Spatial Augmentation. In *4th ACM SIGSPATIAL International Workshop on Spatial Big Data and AI for Industrial Applications (GeoIndustry '25)*, November 3, 2025, Minneapolis, MN, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

With the growing adoption of location-based services, indoor localization systems have recently gained increasing attention like

healthcare, emergency response, logistics, and military [1, 3]. Currently, the most prominent indoor localization approach is to rely on signal fingerprinting [3]. Fingerprinting depends on collecting **large amounts** of location-tagged signal data in the environment (e.g., from WiFi access points) to create a “fingerprint map” and train a localization model. It then relies on learning location-dependent signal attenuation effects (e.g., due to changing distance from the transmitter and passing through obstructions such as furniture) to allow the localization model to extract location-distinguishing patterns from real-time signal data and infer the location. *Although fingerprinting approaches provide state-of-the-art localization accuracy [3], their practical deployability is severely hindered by the data collection overhead needed to train the localization model [4, 5].*

Recently, several data augmentation-based approaches have been proposed to reduce the overhead of fingerprint data collection [7, 10–12]. However, these approaches either rely on classical machine learning methods [10–12], which have low representational ability; or employ the commonly used generative adversarial networks (GANs) [7] that suffer from mode collapse and training instability problems [6]. This causes the available methods to provide synthetic samples from a limited distribution, reducing localization accuracy. Moreover, some of these approaches require collecting data at all the fingerprint locations, necessitating the effort of surveying the whole environment. *This raises the need for an augmentation method capable of providing realistic synthetic samples to ensure high localization accuracy while mitigating the need to survey the entire environment.*

In this paper, we propose *LocaGen*, a novel spatial augmentation framework designed to generate fingerprint data at locations completely unseen during training. By doing so, *LocaGen* can be plugged into any fingerprinting-based localization system to significantly reduce its data collection overhead, since it mitigates the need to survey the entire environment. The main idea behind *LocaGen* is to choose a set of *seen locations* where fingerprint data are collected, use these seen location data to train a conditional deep generative model to synthesize fingerprint data conditioned on the location, and then use it to generate data at *unseen locations*. *LocaGen* then passes the collected seen and generated unseen location data to the localization system to train its location inference model.

LocaGen includes modules that address several domain-specific design challenges to ensure realistic data generation at unseen locations for accurate localization with minimal data collection

*Equal contribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GeoIndustry '25, November 3, 2025, Minneapolis, MN, USA

© 2025 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>




Figure 1: LocaGen Framework Architecture.

overhead. Specifically, *LocaGen* optimizes the division of fingerprint locations into seen and unseen locations based on a novel density-based approach that ensures that the distribution of the seen location data provides enough coverage for realistic generation of unseen location data. Moreover, *LocaGen* employs heuristic-based domain-specific augmentation methods on seen location data, e.g., through Gaussian noise injection, to increase the amount of training data available to the generative model, avoiding overfitting. Furthermore, in contrast to available augmentation methods that rely on classical machine learning techniques and GANs [7, 10–12], *LocaGen* adapts a state-of-the-art conditional diffusion model [2] to perform data generation, improving the quality of generated data. Finally, *LocaGen* introduces spatially-aware optimization in training its diffusion model to ensure that the model prioritizes learning the data distribution that allows generating realistic fingerprint data at the target unseen locations.

We implement and evaluate *LocaGen* using a real-world WiFi fingerprinting dataset composed of location-tagged received signal strength (RSS) measurements from 520 WiFi access points collected by three users [13]. Our results show that plugging *LocaGen* into a localization system results in the localization accuracy remaining unchanged even with 30% of the locations being unseen. Furthermore, it provides an improvement of up to 28% compared to adapting the state-of-the-art augmentation methods to generate data in unseen locations [7, 10–12]. This highlights the effectiveness of *LocaGen* in significantly reducing data collection overhead in fingerprinting-based systems with minimal reduction in localization accuracy.

The remainder of this paper is structured as follows. Section 2 provides an overview of *LocaGen*'s framework architecture and discusses the details of the modules that handle different design challenges. Section 3 presents the evaluation of *LocaGen*. Finally, Section 4 concludes our paper.

2 The LocaGen Framework

In this section, we present an overview of the *LocaGen* framework, outlining our approach. We then discuss the design challenges faced by *LocaGen* and present the details of the modules addressing them.

2.1 Framework Overview

Figure 1 illustrates the architecture of our *LocaGen* framework, which is designed to be integrated into the data collection process during the construction of the fingerprinting localization model to

reduce the collection overhead. *LocaGen* starts with our *Fingerprint Initializer* module dividing the target fingerprint locations into two groups: **seen locations**, in which fingerprint data are manually collected and constitute the data collection overhead, and **unseen locations**, whose fingerprint data will be automatically generated by *LocaGen* through spatial augmentation based on modeling the seen location data distribution. To ensure the generation of realistic samples in the unseen locations while requiring a minimal number of seen locations, our *Fingerprint Initializer* proposes a density-based approach for its locations' division. After the localization system collects data at the seen locations, our *Data Synthesizer* module employs heuristic-based domain-specific data augmentation to create synthetic training data at the seen locations, enhancing *LocaGen*'s generative model ability to learn the fingerprint distribution and generate diverse, realistic samples at the unseen locations. Finally, our *Spatial Augmenter* module trains a conditional generative diffusion model [2] using data from the seen locations, proposing a novel spatially-aware optimization approach to generate realistic synthetic fingerprints conditioned on the coordinates of the unseen location. Employing a state-of-the-art diffusion model ensures the generation of realistic and diverse data as it mitigates limitations in common generative models, including the mode collapse issues faced by the commonly used GANs [6]. The augmented seen and generated unseen location data form the full fingerprint map, which is then used to train the localization model.

2.2 Design Challenges and Framework Details

The main challenges *LocaGen* faces include optimizing the selection of seen and unseen locations, the risk of generative model overfitting, and the lack of generative model spatial understanding.

2.2.1 Optimizing Seen Locations Selection. The reliance on data collected at only the seen locations restricts the coverage of the collected signals. This can cause some areas in the environment to be underrepresented and might result in some signal transmitters being completely unseen during training data collection (e.g., when relying on WiFi fingerprinting with limited access point coverage). This makes the selection of seen and unseen locations challenging, as we need to maximize the coverage of collected data to ensure realistic data generation at the unseen locations while minimizing the seen location data collection overhead.

To address this challenge, our intuition is to prioritize surrounding the unseen locations with a high density of seen locations to capture the needed signal distribution for realistic synthetic data generation. To do so, our *Fingerprint Initializer* calculates the average distance from each location to its neighboring locations to quantify the neighbor density for each location in the entire set of target fingerprint locations. It then selects those with the highest neighbor density as unseen locations, removes them from the set of target locations, and iteratively repeats the process until the needed number of unseen locations is obtained. The remaining unselected points are designated as seen locations for data collection.

2.2.2 The Risk of Generative Model Overfitting. Since *LocaGen* relies on samples collected from the seen locations only to train its generative model, there is a risk of generative model overfitting. This results in the model generating synthetic data based on a




Figure 2: LocaGen’s spatial augmentation diffusion-based model. The model is optimized with our novel spatially-aware loss function $\mathcal{L}(\theta)$ to ensure synthesizing high-quality samples for the target unseen locations.

limited data distribution, which reduces the ability of the localization model to learn sufficient location-distinguishing features. This is because deep generative models require large amounts of data for training to ensure that they generate samples that reflect the real-world data distributions [5].

Our *Data Synthesizer* addresses this challenge by augmenting the seen location training data given to our generative model through domain-specific heuristic-based techniques that simulate real-world **temporal** signal variations. Specifically, it employs *Gaussian noise injection* to generate samples that mimic the effect of dynamic changes in the environment and the noise of the wireless channel on the signal [8]. Furthermore, it simulates real-world fluctuations in seen transmitters by dropping transmitters in the fingerprint sample whose signal strength falls below a certain threshold, setting their value to zero.

2.2.3 Our Generative Model.

Model Design. LocaGen introduces a deep generative model design based on the recently proposed continuous conditional diffusion model (CCDM) [2] to spatially augment the fingerprint data by generating samples in unseen locations (Figure 2). CCDMs rely on a forward diffusion process to iteratively add noise to the fingerprint signal data M_0 at a certain seen location in time steps t , generating a noisy version M_t . Then, the CCDM trains a U-Net model in a reverse diffusion process to reverse the noise and obtain a prediction of the original data, given a nearby unseen location as a conditional embedding. This allows the U-Net, after training, to generate synthetic data from noise, conditioned on a specific unseen location. To condition the U-Net generation, the location coordinates are embedded and concatenated with the input noisy signal measurements, and the training loss function is weighted based on the distance between the chosen unseen location condition and the ground truth seen location at each training step.

Integrating Spatial Awareness. The straightforward application of the CCDM samples unseen location condition labels randomly

around the ground truth seen location labels [2], aiming to cover the entire distribution of possible labels. However, this can cause the model to include impossible unseen location labels in our context of spatial augmentation, where our primary goal is to accurately model the distribution of fingerprint data at specific unseen locations. For example, the CCDM might sample locations completely outside of the building walls, which hinders its learning ability.

To this end, our *Spatial Augmenter* instead injects spatial awareness into our CCDM by supplying the unseen locations obtained by the *Fingerprint Initializer* directly into the U-Net optimization (Figure 2). This ensures that the surrounding seen locations are given high weights in the training process to learn their distribution, **which allows the model to generate high-quality realistic samples at the target unseen locations**, leading to the following spatially-aware loss function for optimizing our U-Net model parameters θ :

$$\mathcal{L}(\theta) = - \sum_{i=1}^{N_{\text{unseen}}} \sum_{j=1}^{N_{\text{seen}}} \left[w(\ell'^i, \ell^j) \cdot \|U_{\theta}(t, \ell'^i, M_t^j) - M_0^j\|^2 \right] \quad (1)$$

Where ℓ' and ℓ represent the embeddings of the unseen and seen locations respectively, the distance between which is calculated using $w(\cdot)$ and used as a weight for the loss function $\mathcal{L}(\theta)$. U_{θ} represents the U-Net model whose parameters are θ , and M_t and M_0 denote the noisy and non-noisy versions of the signal measurement respectively, obtained at the seen locations ℓ at time step t .

3 Framework Evaluation

3.1 Dataset and Experimental Setup

We use a public real-world indoor WiFi fingerprinting dataset [13] to evaluate LocaGen, which consists of RSS measurements giving a total of 1600 samples from 520 WiFi access points collected by three users at 70 locations along with the corresponding ground truth location coordinates. As our evaluation metric, we implement a localization model based on a fully connected neural network [9], integrating LocaGen into its data collection phase, and rely on the mean localization error of this localization model to quantify the impact of LocaGen. Unless otherwise specified, we use an unseen-to-seen locations ratio of 50% in our experiments.

3.2 Ablation Study

In this section, we verify the effectiveness of our *Fingerprint Initializer*’s seen location selection approach (Section 2.2.1) and the impact of our *Spatial Augmenter*’s spatially-aware CCDM training optimization (Section 2.2.3).

3.2.1 Verifying our Fingerprint Initializer. We compare our *Fingerprint Initializer*’s density-based seen location selection (Section 2.2.1) with two other methods: the commonly used grid center method, which divides the environment into regions and selects locations at each region’s center as seen locations [10] and to selecting the seen locations uniformly at random. Figure 3 shows that LocaGen’s density-based selection outperforms grid center selection and random selection **by over 31%** in terms of the mean location error.

3.2.2 Verifying our Spatially-Aware Optimization. Figure 4 shows that our *Spatial Augmenter*’s spatial awareness injection into the




Figure 3: Impact of our Finger-print Initializer.

loss function of our CCDM (Section 2.2.3) leads to a decrease in the localization error by **more than 15%** compared to directly employing the default CCDM optimization [2]. This is because our spatially-aware optimization enables our CCDM to better focus on optimizing data generation at the target unseen locations. This results in better representation ability for the signal distribution at these locations, leading to high-quality signal data generation.

3.3 Overall Framework Performance

In this section, we present a comprehensive evaluation of *LocaGen*'s impact on reducing the data collection overhead of localization systems. First, we examine the trade-off between localization accuracy and data collection overhead reduction when employing *LocaGen* with varying unseen-to-seen locations ratio. We then compare *LocaGen* to four state-of-the-art augmentation methods and to the case of using the collected data without spatial augmentation.

3.3.1 Localization Accuracy-Overhead Trade-Off. As seen in Figure 5, the data collection overhead decreases linearly with increasing the unseen-to-seen locations ratio, while the mean localization error stays relatively stable at first, *remaining nearly constant until a ratio of 30%* and showing only a marginal increase up to a ratio of 50%. Moreover, with a *threefold* reduction in data collection overhead (i.e., from 120 to 40 mins at a 70% unseen-to-seen locations ratio), the mean localization error increases by *2.24 meters only*. This highlights that *LocaGen*'s advantage in reducing the data collection overhead outweighs the increase in localization error. Localization system operators can select the suitable unseen-to-seen locations ratio based on the desired accuracy-overhead trade-off.

3.3.2 Comparison with State-of-the-Art. We compare *LocaGen* to employing state-of-the-art augmentation techniques to generate data at unseen locations, including a KNN-based spatial interpolator [10], a single-output Gaussian process (SOGP)-based approach [11], a multi-output Gaussian process (MOGP)-based method [12], and a selective generative adversarial network (SGAN) approach [7], and to using only the seen locations without augmentation. As seen in Figure 6, *LocaGen* results in superior localization accuracy compared to the state-of-the-art augmentation methods, improving the median location error by 28%, 27%, 18%, 17%, and 16%, compared to spatial interpolator [10], no augmentation, MOGP [12], SGAN [7], and SOGP [11], respectively. *LocaGen*'s superior performance stems from its use of a state-of-the-art conditional diffusion model [2] and from its spatially-aware model optimization (Section 2.2.3),




Figure 4: Impact of our spatially-aware optimization.

allowing the generation of high-quality and diverse unseen location data more effectively than other augmentation methods.

4 Conclusion

We have presented *LocaGen*, a novel spatial augmentation framework, designed to be plugged into fingerprinting localization systems to significantly reduce data collection overhead. *LocaGen* adapts a state-of-the-art conditional diffusion model to generate realistic fingerprint data at unseen locations based on data from a few seen locations, proposing a novel spatially-aware optimization strategy to improve the quality of the generated fingerprints. Furthermore, *LocaGen* strategically divides the target fingerprint locations into seen and unseen locations to ensure robust coverage and employs domain-specific heuristic-based augmentation to increase the amount of seen location data, reducing overfitting. Our evaluation results on a public dataset demonstrate the effectiveness of *LocaGen* in reducing data collection overhead, maintaining the localization accuracy even with 30% of the locations unseen, and providing up to 28% accuracy improvements compared to state-of-the-art augmentation methods.

References

- [1] Toufiq Aziz and Insoo Koo. 2025. A Comprehensive Review of Indoor Localization Techniques and Applications in Various Sectors. *Applied Sciences* 15, 3 (2025).
- [2] Xin Ding, Yongwei Wang, Kao Zhang, and Z. Jane Wang. 2024. CCDM: Continuous Conditional Diffusion Models for Image Generation. *arXiv preprint arXiv:2405.03546* (2024), 19 pages.
- [3] Sherif Mostafa, Khaled Harras, and Moustafa Youssef. 2022. A Survey of Indoor Localization Systems in Multi-Floor Environments. <http://dx.doi.org/10.36227/techrxiv.20439648.v1>
- [4] Sherif Mostafa, Khaled A. Harras, and Moustafa Youssef. 2023. UniCellular: An Accurate and Ubiquitous Floor Identification System using Single Cell Tower Information. In *SIGSPATIAL*. ACM, New York, NY, USA, 10 pages.
- [5] Sherif Mostafa, Moustafa Youssef, and Khaled A. Harras. 2024. Accurate and Ubiquitous Floor Identification at the Edge Using a Single Cell Tower. In *IEEE/ACM SEC*. IEEE, Piscataway, NJ, USA, 14 pages.
- [6] Sherif Mostafa, Moustafa Youssef, and Khaled A. Harras. 2024. Ubiquitous and Low-Overhead Floor Identification with Limited Cellular Information. *ACM TSAS* (2024), 33 pages. Just Accepted.
- [7] Wafa Njima, Marwa Chafii, Arsenia Chorti, Raed M. Shubair, and H. Vincent Poor. 2021. Indoor Localization Using Data Augmentation via Selective Generative Adversarial Networks. *IEEE Access* 9 (2021), 98337–98347.
- [8] Hamada Rizk, Ahmed Shokry, and Moustafa Youssef. 2019. Effectiveness of Data Augmentation in Cellular-based Localization Using Deep Learning. In *IEEE WCNC*. IEEE, Piscataway, NJ, USA, 1–6.
- [9] Hamada Rizk, Marwan Torki, and Moustafa Youssef. 2019. CellinDeep: Robust and Accurate Cellular-Based Indoor Localization via Deep Learning. *IEEE Sensors Journal* 19, 6 (2019), 2305–2312.
- [10] Hamada Rizk and Moustafa Youssef. 2019. MonoDCell: A Ubiquitous and Low-Overhead Deep Learning-based Indoor Localization with Limited Cellular Information. In *SIGSPATIAL*. ACM, New York, NY, USA, 109–118.




Figure 6: Comparison to state-of-the-art.

- [11] Ahmed Shokry and Moustafa Youssef. 2024. DeepCell: A Ubiquitous Accurate Provider-side Cellular-based Localization. <http://arxiv.org/abs/2407.16927>
- [12] Zhe Tang, Sihao Li, Kyeong Soo Kim, and Jeremy S Smith. 2024. Multi-Dimensional Wi-Fi Received Signal Strength Indicator Data Augmentation Based on Multi-Output Gaussian Process for Large-Scale Indoor Localization. *Sensors* 24, 3 (2024), 1026.
- [13] Joaquín Torres-Sospedra, Raúl Montoliu, Adolfo Martínez-Usó, Joan P. Avariento, Tomás J. Arnau, Mauri Benedito-Bordonau, and Joaquín Huerta. 2014. UJIIndoor-Loc: A new multi-building and multi-floor database for WLAN fingerprint-based indoor localization problems. In *IPIN*. IEEE, Piscataway, NJ, USA, 261–270.