



**Visible Light Positioning with TinyML:
Improving Data Quality and Reducing Data Collection Effort**

Jakub Adam Trzykowski¹

Supervisor(s): Qing Wang¹, Ran Zhu¹

¹EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 23, 2024

Name of the student: Jakub Adam Trzykowski
Final project course: CSE3000 Research Project
Thesis committee: Qing Wang, Ran Zhu, Johan Pouwelse

Abstract

Visible light positioning (VLP) systems enable indoor positioning through a deployment of light-emitting diodes (LEDs) as transmitters and photodiodes (PDs) as receivers. A promising approach in VLP involves recording the received signal strength (RSS) to construct fingerprint samples for later use in positioning. However, achieving high accuracy demands a labor-intensive data collection process.

In this study, we propose improvements to a data cleaning and augmentation pipeline. Our improvements focus on preserving more source data during cleaning and data-based LED position estimation for more reliable data augmentation. Experimental results show that our approach maintains comparable positioning accuracy while reducing data collection efforts by over 99%. Furthermore, we conduct experiments to investigate the impact of spatially irregular data collection strategies on positioning accuracy. Finally, we deploy a machine learning model on a microcontroller to demonstrate the practical feasibility of our proposed methods.

1 Introduction

Outdoor localization systems, such as the Global Positioning System (GPS), have completely transformed how we determine our location in the world. However, with the rise of Internet of Things (IoT) devices envisioned for the future, there is also a growing need for accurate indoor positioning. This technology is crucial for industries such as smart manufacturing, smart cities, and smart homes. Unfortunately, outdoor systems like GPS perform poorly inside buildings and cannot provide the high levels of precision often required in these environments.

To address this demand, several potential solutions have already been proposed. These range from using wireless technologies such as Bluetooth [1], Wi-Fi [2], ultra-wideband [3], or Zigbee [4], to using camera imagery [5]. Each of these solutions suffers from several issues. Radio frequency-based solutions are prone to electromagnetic interference, tend to achieve lower accuracy, and compete for crowded spectrum resources [6]. Moreover, their deployment is often expensive and complex [7]. Meanwhile, image-based localization requires significant computational power, making it suboptimal for IoT devices.

The widespread adoption of light-emitting diode (LED) technology, alongside the advancements in visible light communication (VLC) technology, has made visible light positioning (VLP) a promising candidate in the field [8]. This technology offers several advantages, including better accuracy compared to other technologies [9], use of widely available lighting infrastructure, and high security characteristics. Additionally, LED light operates in the license-free spectrum and is safe for humans.

Visible light positioning is an umbrella term for various techniques, including angle-of-arrival (AOA) [10], time-of-arrival (TOA), time-difference-of-arrival (TDOA) [11], and

received signal strength (RSS) [12]. These methods rely on a similar setup comprising LEDs and photodiodes to capture signals. While AOA offers high accuracy, it demands significant computational resources and costly equipment. TOA and TDOA require precise synchronization and sensitive hardware. In contrast, RSS stands out as the most cost-effective and readily deployable option.

RSS VLP operates on a straightforward principle: an array of LEDs serves as transmitters, while a photodiode (PD) acts as the receiver, capturing signal strengths from each light source. These intensity readings form a distinctive fingerprint for the location. Maintaining a comprehensive database of such fingerprints is impractical, particularly for IoT devices. However, leveraging machine learning enables extrapolation from existing data points, enabling position estimation.

The main challenge associated with the RSS method lies in the labor and resource-intensive data collection process. Achieving sufficient accuracy requires a dense fingerprint dataset. Zhu et al. [13] proposed data cleaning and augmentation techniques to enable sparser data acquisition, using an 8 cm resolution instead of the original 1 cm, while maintaining comparable accuracy. However, we believe there are opportunities to improve their methods and enhance the accuracy of the augmentation.

Moreover, their approach required collecting data points in a rigid grid pattern. In this study, we aim to additionally assess the impact of utilizing spatially irregular sampling on the accuracy of the positioning system. We conduct experiments simulating different data collection strategies to allow more flexibility in data acquisition, thereby reducing the effort required even further.

With the ultimate aim of enabling indoor positioning on IoT devices, our focus remains on showcasing real-world applicability. To this end, we deploy the developed solutions on an Arduino Nano 33 BLE Sense microcontroller. This device is readily available off-the-shelf, making it an ideal candidate for widespread deployment.

In Section 2, we discuss previous work on data cleaning and augmentation and describe the dataset used for the experiments. Section 3 outlines our approach to improving the pipeline and details further experiments on data collection. The experimental setup is briefly discussed in Section 4, while Section 5 presents and analyzes the obtained results. In Section 6, we reflect on the ethical aspects and reproducibility of our work. Finally, Section 7 places our experimental results in a broader context, and Sections 8 and 9 summarize conclusions and suggest possible future developments.

2 Background

This section discusses the dataset used in this work, as well as, summarizes the previous contributions of Zhu et al. [13] on improving data quality and reducing the data collection effort.

2.1 DenseVLC RSS Dataset

To perform our experiments, we use RSS dataset from the DenseVLC testbed [14] as opposed to simulation. This grounds our research in the real world, closely mimicking the

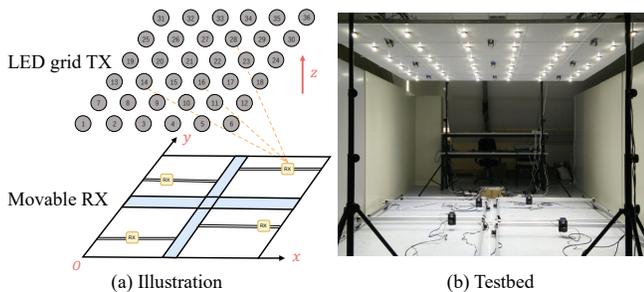


Figure 1: The setup used to obtain the RSS dataset for VLP, an illustration in (a) and the actual testbed in (b).

potential difficulties with imperfect data when deploying a VLP system. Additionally, Zhu et al. used this exact data set in their works, allowing us to directly compare the obtained results. Figure 1 showcases the testbed used for acquiring the data.

The setup comprises 36 LED transmitters (TXs) and 4 receivers (RXs). The TXs are arranged in a 6×6 array on a height-adjustable ceiling with an inter-TX distance of about 0.5 m. The system utilizes high-performance CREE XT-E LEDs, each equipped with a lens TINA FA10645 to limit the field of view to 30° . The RXs, controlled by OpenBuilds ACRO Systems, are positioned on the floor, each equipped with a photodiode S5971.

The floor area, measuring $3 \text{ m} \times 3 \text{ m}$, is divided into 4 square grids with a cross-like gap in the middle. Each grid covers an accessible area of approximately $1.2 \text{ m} \times 1.2 \text{ m}$, enabling movement in both the x and y directions.

Each sample fingerprint in the dataset consists of the measured RSS values from the 36 LEDs and the corresponding x, y coordinates of the measurement position. Specifically, the measurements are conducted at 1 cm intervals in both x and y directions. For each RX at a specific height (172 cm and 196 cm), there are 121 steps in both x and y directions in each quadrant. Additionally, at every sampling position, the measurement is repeated three times. This results in $121 \times 121 \times 4 \times 3 \times 2 = 351,384$ data points per each LED.

2.2 RSS Data Cleaning and Augmentation

The performance of RSS-based VLP is highly dependent on the density and quality of collected fingerprint data, which traditionally requires labor-intensive data collection efforts. To mitigate this, recent research in [13] proposed data preprocessing techniques, including data cleaning and augmentation, to construct reliable and dense fingerprint datasets.

The proposed data cleaning strategy addresses the inherent noise in RSS measurements, which can arise from device thermal noise, sampling errors, and varying ambient light conditions. They employed a two-stage cleaning process that retains the most consistent measurement for each location and estimates missing or erroneous values. More precisely, the first stage begins by computing ‘‘continuity scores’’ for each sample at each location. The sample with the best score is selected, or if none meets the required threshold, all samples are discarded. In such cases, the second stage of the data cleaning pipeline is activated. This stage uses the surrounding clean

samples to estimate the true light intensity value based on the Lambertian radiation model, described by the equation:

$$I^r = \begin{cases} I^t A(d) \cos^m(\phi) \cos(\psi), & 0 \leq \psi \leq \psi_c \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

with I^t being the lamp emission power, ψ – the incident angle, ψ_c – the FoV of the LED lamp, ϕ – the irradiation angle, m – the LED’s Lambertian order, $A(d)$ – a propagation loss function over the distance d between RX and TX.

To further increase the density of samples, the researchers proposed a data augmentation step. For example, if the data was collected at 8 cm intervals, this method allows increasing the resolution to 1 cm, enabling higher positioning accuracy in machine learning systems. The generation of new samples follows a similar approach to the second stage of the cleaning pipeline – it uses surrounding data samples to approximate the expected values using the Lambertian model.

Experimental results validate the effectiveness of these methods, demonstrating an average positioning error of around 2 cm while using an 8 cm resolution dataset. Overall, these methods reduced the data collection effort by 98% while maintaining accuracy.

3 Approach

Although the aforementioned work has achieved promising performance through data preprocessing methods, there is still room to explore more careful designs for data cleaning and data augmentation to further improve data quality. In this section, we outline our observations and propose an improved approach to the pipeline.

3.1 Improving Data Cleaning

In looking for possible areas of improvement, we started by reproducing the data cleaning and augmentation methods. In the cleaning stage, one measures how much a sample point differs from its immediate neighborhood. This involves calculating the average received signal strength (RSS) from nearby points, and then finding the absolute value of the difference between this average and the sample’s RSS. Points with the smallest differences are considered to conform to their neighborhood the best.

Formally, the scoring metric is described by the equation:

$$S_{x,y} = \left| I_{x,y} - \frac{1}{|\mathbf{I}_N|} \sum_{I \in \mathbf{I}_N} I \right|, \quad (2)$$

where \mathbf{I}_N are the samples located in the neighborhood around (x, y) , usually a circular region. After calculating this score, for each location (x, y) the sample that minimized $S_{x,y}$ is chosen. That was unless none of the samples managed to land under a predefined threshold, in which case the sample point is discarded. While this works relatively well for LEDs whose measurements consisted of within-standard values, like LED13 (Figures 2a and 2c), for excessively noisy ones, like LED7 (Figure 2b), it means that many accurate measurements are discarded due to surrounding noise (Figure 2b).

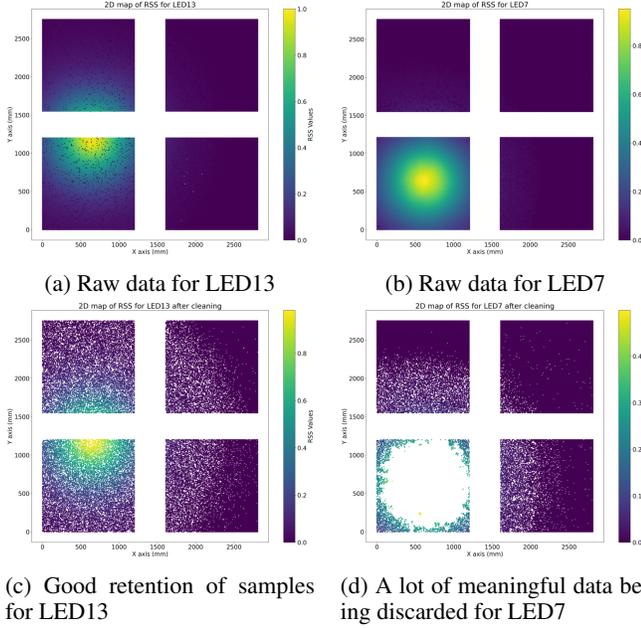


Figure 2: Comparison between overly noisy measurements and comparatively clean ones, along with the implications on the data cleaning process.

An obvious observation when looking at the structure of the noise is that the vast majority of erroneous values result in close-to-zero (dark) measurements. We can exploit this by modifying the scoring function to reward points that exhibit brightness above environment average. The revised scoring metric looks as follows:

$$\widetilde{\mathbf{I}}_N := \frac{1}{|\mathbf{I}_N|} \sum_{I \in \mathbf{I}_N} I \quad (3)$$

$$S_{x,y} = \begin{cases} (I_{x,y} - \widetilde{\mathbf{I}}_N) \cdot r, & \text{if } I_{x,y} \geq \widetilde{\mathbf{I}}_N \\ (\mathbf{I}_N - I_{x,y}), & \text{otherwise,} \end{cases} \quad (4)$$

where r is the brightness reward that lowers the score for bright points.

This method allows us to retain more valuable data points needed to reconstruct missing values. Figure 3 shows what effects the new scoring function had on the cleaning pipeline. The reward value was chosen based on how well the data looked visually, and we admit that a more rigorous method of determining it would be more desirable. However, we deem this outside the scope of this paper.

3.2 Improving Data Reconstruction

The data augmentation process proposed by Zhu et al. [13] is based on the Lambertian radiation model. However, since the exact emission power I^t of the LED is unknown, the researchers used neighboring RSS measurements to predict the value from a simple proportion equation:

$$\frac{I^r_{x_1, y_1}}{I^r_{x_2, y_2}} = \frac{A(d_1)}{A(d_2)} \left[\frac{\cos(\phi_1)}{\cos(\phi_2)} \right]^{m+1}, \quad (5)$$

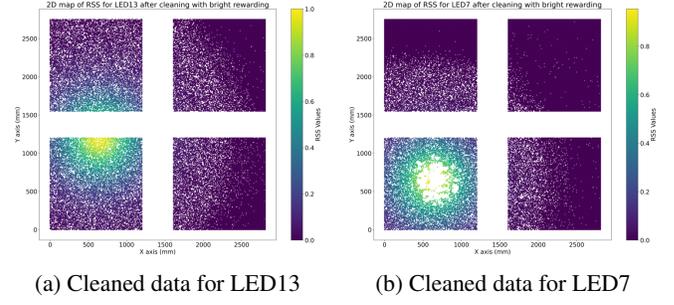


Figure 3: The improved scoring function did not have a significant effect on the already clean-looking LED13, while at the same time, it dramatically improves data retention for LED7.

where the I^r 's are the RSS values, d_i is the distance between the LED and the corresponding (x, y) point, and ϕ_i is the angle between the light ray and the normal to the horizontal plane.

This model, however, presumes that the LED positions are both known and precise, which is not always the case. Figure 4 illustrates the discrepancy between the provided LED positions and their actual footprints. Consequently, the quality of the augmentation is impeded. Figure 4b highlights the negative effects of inaccurate LED positioning on data sampled at an 8 cm resolution, which was subsequently augmented to a 1 cm resolution. This results in skewed data and an overall lack of smoothness.

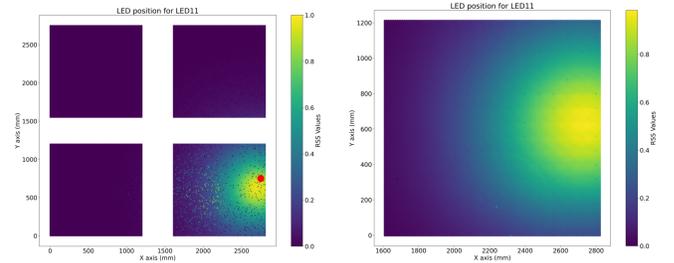


Figure 4: Example of how inaccurate LED position estimates affect data augmentation.

Figure 4: Example of how inaccurate LED position estimates affect data augmentation.

To address this issue, we propose a reliable and automated method to estimate the (x, y) positions of the light-emitting diodes. The core concept is to robustly fit a circle to the LED footprint, with the center of the circle providing a more accurate approximation of the source coordinates. The process involves the following steps:

1. **Initial Data Cleaning:** As described in Section 3.1. Removing anomalous data helps the accuracy of this estimation.
2. **Circular Strip Extraction:** Create a circular strip of RSS values by filtering positions that fall within specific RSS thresholds (Figure 5a).
3. **Kasa [15] circle fitting:** Apply Kasa's robust circle-fitting and retrieve the new position estimate.

The Kasa method’s robustness assures that the method is effective even in cases where a significant portion of the circle is missing. Figure 5 illustrates the circle-fitting process for LED11 with Subfigure 5d showing how the 8-to-1 cm augmented data, refer to Figure 4b to compare how it would look before.

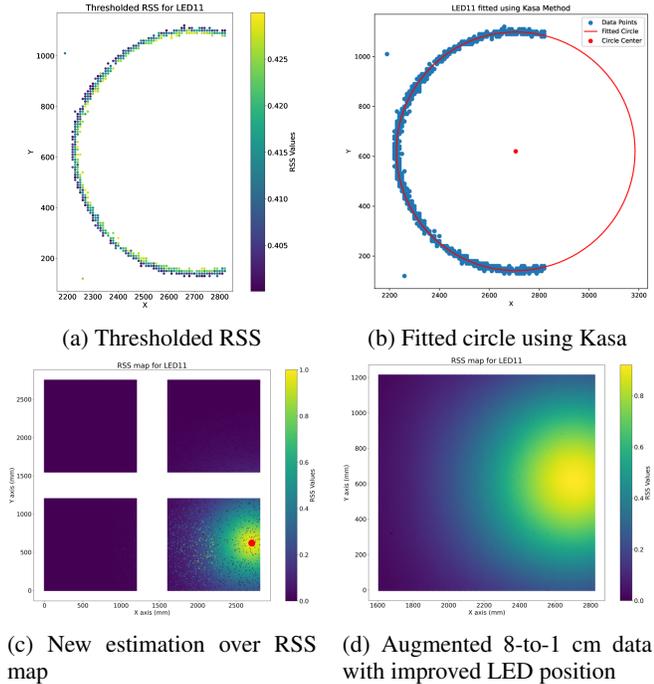


Figure 5: LED center position estimation processes for LED11.

Using the newly estimated LED positions we can better clean the data, and later if desired, use it to increase the data resolution. The entirety of the revised process is summarized in Figure 6.

A final enhancement to the data augmentation process involves applying an alternative augmentation technique for data points that fall outside the LED’s light cone. The Lambertian model fails to predict these points. We suggest employing a simple inverse-distance weighting (IDW) to estimate values for such points. The predicted RSS values for a point p can be computed according to the formula:

$$\tilde{I}_p^r = \frac{\sum_{p' \in \mathbf{I}_N} I_{p'}^r \frac{1}{d(p,p')}}{\sum_{p' \in \mathbf{I}_N} \frac{1}{d(p,p')}} \quad (6)$$

where \mathbf{I}_N is some neighborhood around p and $d(p,p')$ denotes the Euclidean distance between the two locations.

3.3 Experimentation on Spatially Irregular Data

The goal of these experiments is to evaluate the impact of alternative sample collection strategies on the quality of the model. Specifically, we aim to determine whether focusing data acquisition efforts on certain areas in the testbed can enhance model accuracy. This could involve prioritizing locations where higher accuracy is required or targeting areas that maximize information gain for the system.

We propose the following three experimental scenarios:

1. **Uniformly Distributed Data Points:** This scenario involves data points that are uniformly distributed. The aim is to evaluate whether the augmentation benefits from structured grid data samples as opposed to unorganized structure. Illustrated in Figure 7b.
2. **Normally Distributed, LED Centered Data Points:** This scenario involves data points that are normally distributed around each LED. The objective is to focus data collection efforts on areas with the highest potential information gain and measure its influence. Shown in Figure 7d.
3. **Normally Distributed, Globally Centered Data Points:** In this scenario, data points are normally distributed with a global center. The goal is to simulate a more extensive data collection effort concentrated in the center, where high positioning accuracy could be more important. Example in Figure 7c.

4 Experimental Setup

It is important to design careful and transparent experiments to measure whether the proposed methods actually reflect in tangible improvements. To this end, we will detail the process of obtaining the test data, while specifying the parameters used. Additionally, we will describe the machine learning model employed to evaluate our experiments.

4.1 RSS Dataset for Experimentation

To conduct our experiments, we utilized a part of the dataset with samples collected at a height of 176 cm. We developed a script designed to perform data cleaning and data augmentation with modifiable parameters. Initially, all measurements were normalized to the range of 0 to 1 using MinMax normalization¹.

In the first stage of creating clean data from the 1 cm resolution dataset, we employed a 15 mm radius around each position to compute the continuity score. Using the revised scoring method, we discarded samples that differed from their neighborhood average by more than 0.08, applying a brightness reward of $\frac{1}{3}$ (Equation (4)). In the baseline version, we discarded samples if their score exceeded 0.13. We intentionally set different thresholds in the baseline to retain more samples at the expense of less effective noise filtering.

To estimate improved LED positions, we thresholded the RSS values to the range of 0.4 to 0.43 of the clean data before any sample reconstructions. We then used an unmodified implementation of the Kasa algorithm to obtain improved estimates.

In the second stage of data cleaning, missing points were reconstructed using the proportion from Equation (5). We took the average of surrounding points within a 100 mm radius, provided at least three samples were present. We employed the standard propagation loss function $A(d) = \frac{1}{d^2}$.

¹MinMax normalizarion – https://en.wikipedia.org/wiki/Feature_scaling

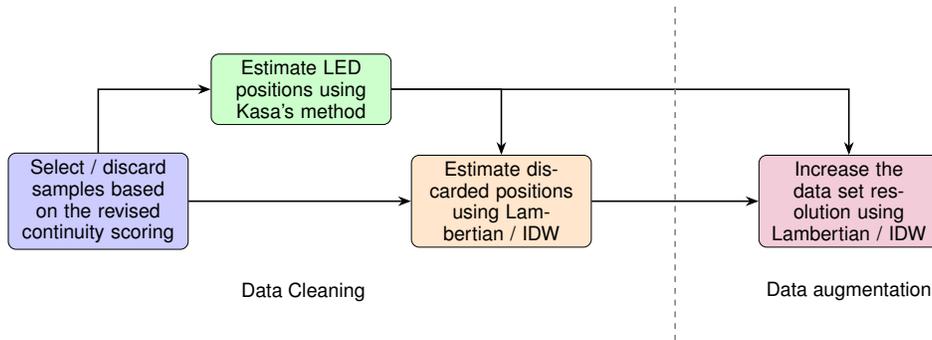


Figure 6: Summarized pipeline of data cleaning and augmentation. The estimated LED positions are used both to reconstruct missing points and increase the density of the dataset.

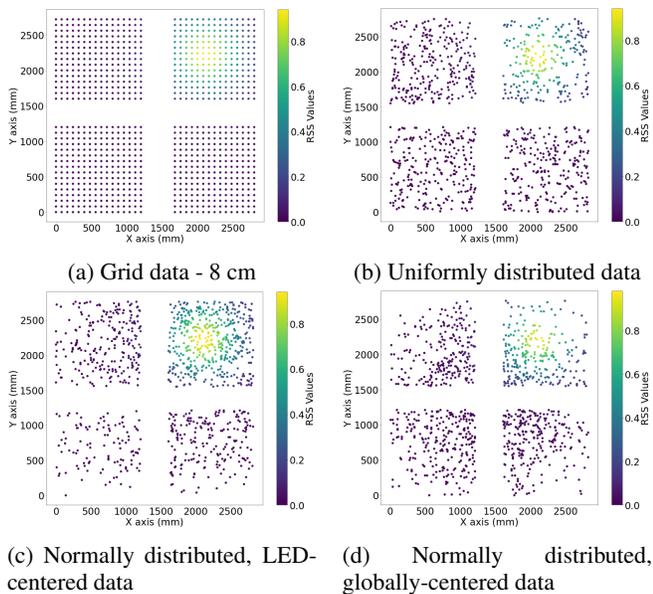


Figure 7: Showcase of different sampling methods imitating different data collection strategies. All approaches sample around 1000 points.

The revised version differed from the baseline only in the improved LED positions and the use of inverse distance weighting (IDW) estimation for points outside the LEDs’ cone of light.

To simulate the collection of lower-density fingerprint sets, we employed various strategies. For grid sampling, samples were taken every 8 or 16 cm. For uniform sampling, random x and y coordinates were generated until their count matched the number of samples in the 8 and 16 cm grids, respectively. A similar approach was used for LED-centered normal samples, with a standard deviation (σ) of 1250 mm for each dimension. For globally-centered normal data, we used $\sigma = 750$ mm.

Finally, the augmentation of sampled data followed the same steps and parameters as the second stage of data cleaning. This process was repeated until the entire testbed area was recreated.

4.2 Experimental Environment

To assess the accuracy of the proposed methods, we conducted experiments using a simple 2-layer Multi-Layer Perceptron (MLP) with 64 and 32 neurons, respectively. This network configuration was chosen to demonstrate its sufficiency in extracting relevant information from the RSS dataset, while being more suitable for embedded environments compared to the larger 2.5k-neuron network utilized by Zhu et al. [13].

The experiments were performed using the *tensorflow*² library. The ‘Sequential’ MLP used ‘relu6’ (ReLU limited to value of 6) activation function for the first hidden layer and ‘selu’ (Scaled Exponential Linear Unit) for the second hidden layer, as we have observed this combination to converge the fastest. We compiled the model to use the Adam optimizer with default parameters and employed the Euclidean distance metric as the loss function. Additionally, data normalization was performed using the ‘StandardScaler’ from *sklearn*³ prior to training.

The network was trained using a 70-20-10% training-test-validation split to adhere to best machine learning practices. For each experiment, the network was trained for 50 epochs with a batch size of 32.

To simulate different environments, we tested various LED configurations, ranging from very dense to sparse setups. Figure 8 illustrates the different LED configurations tested.

5 Evaluation

After running the experiments, we detail the results on how the modified pipeline performed in comparison to the original. On top of that, we examine how alternative data collection approaches compare to the standard rigid grid data.

5.1 Improvements on Data Preprocessing

Table 1 summarizes the errors obtained across various LED configurations. Initially, we establish a baseline using a model trained on raw data (column ‘Raw data’), providing an initial assessment of the challenges posed by each scenario.

Confs 1, 2, 4, and 5 yield reasonably accurate results, all achieving average errors under 10 cm. In contrast, Confs 3

²TensorFlow – <https://tensorflow.google.cn/>

³scikit-learn – <https://scikit-learn.org/stable/index.html>

Table 1: Summary of average errors (in cm) arising from the original and revised cleaning-augmentation methods. The table also shows error variations when measured against raw, clean, and source datasets (the dataset on which the network was trained). Improvements of the revised pipeline compared to the original are highlighted.

Error against	Raw data		Clean data				Augmented data (from 16 cm)					
	Raw	Original		Revised		Original			Revised			
		Raw	Raw	Clean	Raw	Clean	Raw	Clean	Aug.	Raw	Clean	Aug.
Conf 1	1.64	5.84	0.76	6.06	0.79	7.86	1.70	0.61	6.68	1.38	↓18.8%	0.60
Conf 2	4.55	13.66	2.59	13.52	2.91	14.39	4.66	2.03	14.01	4.09	↓12.2%	2.04
Conf 3	40.53	50.12	25.04	49.65	24.20	51.49	27.13	22.40	51.20	26.93	↓0.7%	22.41
Conf 4	2.55	7.14	1.43	7.61	1.51	9.84	3.33	1.52	8.71	3.08	↓7.5%	1.44
Conf 5	6.61	11.02	5.78	11.32	5.93	15.96	9.23	3.52	16.30	9.44	↑2.3%	3.43
Conf 6	12.28	17.62	10.68	17.26	10.17	25.14	16.08	6.10	25.21	15.31	↓4.8%	6.03

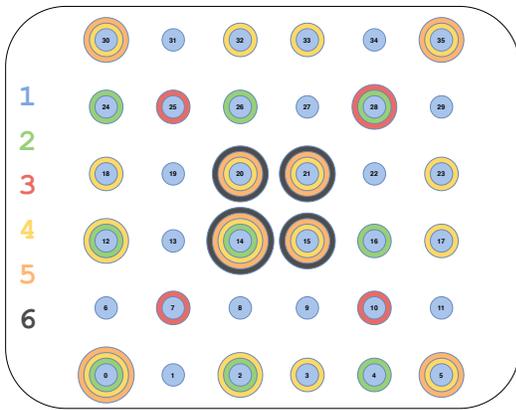


Figure 8: Evaluated LED configurations. The colors denote the subsets of LEDs used for each experiment. For example, Conf 1 (light blue) used all LEDs, while Conf 6 (black) only used the four middle ones.

and 6 prove to be the most challenging, with average errors of 40 cm and 12 cm, respectively. This increased error can be attributed to the limited visibility of LEDs from many positions within the testbed. For instance, predicting the exact measurement position for LED 7 in Conf 3 (see Figure 8) on an arbitrary ring of equal RSS values around it is inherently difficult.

Next, we compare the original data cleaning pipeline to the revised one (column "Clean Data"), which incorporates more accurate LED positioning and additional inverse distance weighting augmentation. Overall, the differences between the two methods are minimal. In some scenarios, one method achieves slightly better results, but the differences are generally insignificant. This minimal impact is likely because only a few points required reconstruction, limiting the overall effect of the modifications.

The final column presents the errors for the augmented data, with resolutions refined from 16 cm to 1 cm. We present the results at 16 cm granularity to highlight the most significant differences, as denser sample sets (≤ 8 cm) show only negligible error differences. We expect these differences to become more pronounced with even sparser datasets.

Errors against clean data are highlighted for both imple-

mentations to illustrate the accuracy of the pipeline in reconstructing the data. The revised method reduces errors by up to 20% compared to the original. This improvement is most noticeable in the denser configurations: Confs 1, 2, and 4. However, the improvement is much less significant for Confs 3 and 6, and for Conf 5, the error actually increased by 2.3%.

The likely reason for the less significant improvement in Conf 3 is that this scenario is particularly challenging, making the behavior of the multi-layer perceptron (MLP) more unpredictable. For Confs 5 and 6, the poorer performance can probably be attributed to the use of the four central LEDs. These LEDs were more difficult to reconstruct because they are located very close to the central non-accessible area, having fewer reference data points.

Several predictable trends follow from the results. Firstly, errors against the training datasets decrease from raw to clean to augmented data, which is expected: cleaning the data removes noise, making it more structured and easier for the MLP to learn its properties. When the data is sampled and augmented, it becomes even simpler. Conversely, errors against raw data are smallest when the MLP is trained on raw data, and they increase when the data is cleaned and further when it is sampled and augmented. This trend is expected, as each step moves further away from the originally taken RSS measurements.

Overall, we uphold the belief that augmentation is a viable approach to achieve high-granularity datasets. The errors were within 1-2 cm of the model trained on the entirety of the clean dataset, except for Confs 5 and 6, where the differences were slightly larger.

5.2 Alternative Sampling Strategies

We investigated how the augmentation of data from structured grid-like data (Figure 7a) compares to the augmentation of data from uniformly distributed samples (Figure 7b), as well as, created from samples that were normally distributed around each LED (Figure 7d). Specifically, we examined two scenarios: ~ 1000 samples ≈ 8 cm resolution, and ~ 250 samples ≈ 16 cm resolution. Table 2 presents the results obtained.

When comparing the sampling methods at around 1000 samples, the results are quite inconclusive. Random uniform sampling performed comparably in Confs 1, 3 and 6. Conf 5 offered a minor improvement, while Conf 2 and especially

Table 2: Comparison of accuracies between models trained on augmented datasets constructed from structured, grid-like and from uniformly distributed samples.

	Grid	Uniform	Normal
~1000 samples			
Conf 1	1.16	1.17 $\uparrow 0.9\%$	1.13 $\downarrow 2.6\%$
Conf 2	3.65	3.90 $\uparrow 6.8\%$	3.76 $\uparrow 3\%$
Conf 3	26.12	25.60 $\downarrow 2\%$	25.64 $\downarrow 1.8\%$
Conf 4	2.18	2.57 $\uparrow 17.9\%$	2.48 $\uparrow 13.8\%$
Conf 5	8.62	8.21 $\downarrow 4.8\%$	8.38 $\downarrow 2.8\%$
Conf 6	15.47	15.70 $\uparrow 1.5\%$	15.01 $\downarrow 3\%$
~250 samples			
Conf 1	1.38	1.54 $\uparrow 11.6\%$	1.42 $\uparrow 2.9\%$
Conf 2	4.09	4.87 $\uparrow 19.1\%$	5.02 $\uparrow 22.7\%$
Conf 3	26.93	27.19 $\uparrow 1\%$	27.42 $\uparrow 1.8\%$
Conf 4	3.08	3.97 $\uparrow 28.9\%$	3.69 $\uparrow 19.8\%$
Conf 5	9.44	9.66 $\uparrow 2.3\%$	10.05 $\uparrow 6.5\%$
Conf 6	15.31	16.86 $\uparrow 10.1\%$	18.73 $\uparrow 22.3\%$

Conf 4 performed worse compared to the grid data. Normal LED-centered sampling performed all-around slightly better. Confs 1, 2, 3, 4 and 6 were all within a couple of percentage points of the baseline. Curiously enough, Conf 4 again achieved significantly worse performance.

Moving onto 250 samples, the results are more clear. Uniform samples performed slightly worse in Confs 3 and 5, and significantly worse for other scenarios. For normally distributed samples the situation is quite similar except for Conf 1 where it did not perform that bad.

Overall, the general trend is that grid samples tend to result in more accurate reconstruction of the dataset. When one is allowed a bit more data collection, making denser measurements around the LED position can offer improved accuracy. On the other hand, uniform sampling is consistently inferior to rigid grid data.

5.3 Concentrated data collection

In this experiment, we sampled the data points using the normal distribution positioned at the center of the testbed. The goal was to evaluate whether this approach boosts the accuracy of the model in some area of more interest. Figure 9 presents a comparison of accuracies over the testbed area between normally distributed and the baseline grid data.

The model trained on the new dataset achieved worse average errors across all scenarios. However, Confs 1, 2 and 4 (denser TX topologies) seem to be confirming our hypothesis – the area closer to the center testbed actually exhibit superior accuracies in comparison to grid data, while the areas towards the edges and corners of the testbed achieve higher average errors. Confs 5 and 6 tend to also adhere to a similar trend. However, the errors near the center of the testbed are quite similar. Lastly, it seems that Conf 3 is too unstable to be able to draw any meaningful conclusions.

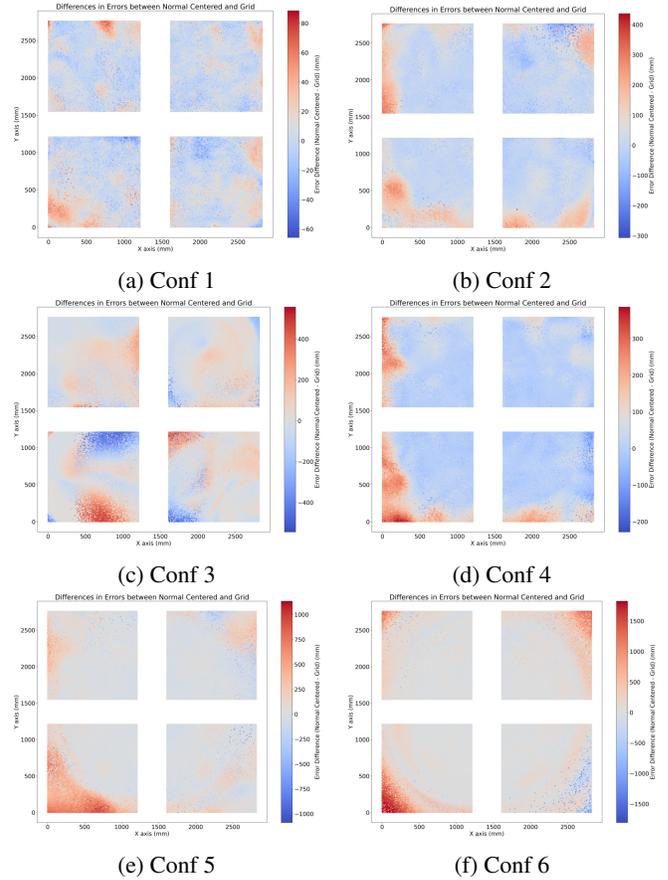


Figure 9: Comparison of error differences between two sampling methods: grid sampling at 16 cm resolution and normal distribution sampling (approximately 250 samples) centered at the testbed's center. Blue regions indicate areas where the normal distribution sampling method outperformed grid sampling, while red regions indicate areas where grid sampling performed better.

5.4 Deployment on a microcontroller

To demonstrate the practical applicability of our work, we deployed the trained MLP model on an Arduino Nano 33 BLE Sense microcontroller. This deployment used a model trained with clean data under Conf 1, utilizing all 36 LEDs. Initially, the model size was 75 KB. After quantization, the model size was reduced to 21 KB. The quantized model efficiently provides position predictions in under 2000 μs , illustrating its suitability for real-time applications in IoT devices.

6 Responsible Research

6.1 Reproducibility

In order to make our research reproducible, we have provided all external software and configuration parameters that were part of the experiments. Using those, any knowledgeable person should be able to reproduce our methods. Additionally, the source code is available to view online⁴. However, due to

⁴<https://github.com/Trzyp0712/TinyML-VLP>

time constraints it is poorly maintained and not documented well.

6.2 Usage of Generative AI

Throughout the project, generative AI was used to get to know the basic syntax of the *tensorflow* and *matplotlib* libraries, as those were outside the area of our expertise. Additionally, large language models helped with \LaTeX , including writing small scripts and formatting tables.

Lastly, GPT models were used to improve the readability and flow of the text. They were tasked to revise paragraphs of the paper to avoid stylistic issues such as repetitions and poor sentence structure. The output of the models were then further tweaked manually to match our preferred style of writing and proofread to remove any inaccuracies.

6.3 Ethical Concerns

Visible light positioning (VLP) systems only record received signal strength (RSS) data, which is non-sensitive and does not allow for the identification of individuals. While there is a potential for surveillance, this would require individuals to carry a microcontroller equipped with a photodiode, and it would only be effective within the pre-mapped area of a building.

Moreover, since VLP systems utilize LEDs operating within the visible light spectrum, they pose no inherent health risks to humans. However, it is important to consider the safety of photosensitive individuals, as the rapid blinking of LEDs in a deployed system could trigger unwanted reactions.

7 Discussion

To fully interpret the results, it is essential to place them in a broader context and compare them with previously reported findings. Additionally, it is important to address potential shortcomings of the approach and offer explanations for any unexpected behaviors.

7.1 Data Preprocessing Pipeline

We successfully reproduced the results of Zhu et al. [13] for the MLP trained on raw data, with errors differing by at most 2 cm across all 6 TX configurations. However, when evaluating the results for clean data, Confs 1 and 4 slightly outperformed the solution proposed by Zhu et al. In contrast, Confs 2, 3, 5, and 6 exhibited errors up to three times worse than the previous work. This indicates that the smaller 96-perceptron network performs comparably well with the larger 2.5k-perceptron network in dense LED deployment scenarios. However, in sparser scenarios, the smaller network seemed to struggle to capture subtle fluctuations in the data, resulting in reduced accuracy.

As a general guideline, when designing a VLP system using RSS values for localization, it is especially important to ensure that at least three different LEDs are visible from every position. This allows the network to utilize principles of trilateration[16] to recognize these positions effectively. While denser TX deployments can achieve this, an alternative, more cost-effective approach is to employ lenses with a wider field of view (FoV).

Another consideration is the performance of a model trained on clean data in real-world deployments. As shown in Table 1, the model trained on raw data performs best on raw data as well. This could be due to the model either learning to disregard noise or overfitting to it. Deploying a model trained on clean data involves a tradeoff between speed and accuracy: either longer measurements are taken to suppress noise, or a single measurement is taken, accepting lower accuracy for reduced latency.

Finally, we recognize the need for a more challenging dataset to further evaluate and refine our methods. Although the testbed occasionally generated a substantial amount of noise, it had very little ambient light and no multipath reflections. The primary advantage of the RSS method is its ability to perform effectively in those more complex environments. Therefore, it would be worthwhile to experiment with the proposed pipeline on more varied datasets.

7.2 Alternative Sampling Methods

While there was no expectation for uniform sampling, contrary to our initial expectations, normal sampling did not outperform rigid grid data. The dataset constructed from LED-centered normally distributed data points exhibited notably poor performance. Specifically, using fewer samples resulted in disproportionately worse outcomes compared to the baseline.

It was anticipated that normally distributed data would at least match the performance of uniformly sampled data. However, this assumption did not hold, particularly at 250 samples, where normally distributed data led to significantly worse accuracies in some configurations. This indicates deeper underlying issues.

Visual inspection of the augmented datasets revealed estimation inaccuracies. It appeared that the LED’s light was spilling into the surrounding area. This phenomenon can be explained by the inverse-distance weighting (IDW) method used to estimate points outside the LED’s cone of light. IDW considers only the distance of samples, not their direction. Consequently, with a normal distribution, more samples are concentrated close to the light source, introducing bias in the estimation.

Several approaches can address this issue. One approach is to select samples in a more balanced manner. For instance, when determining the neighborhood, we can strive to balance the number of source samples from each quadrant around the point.

Another approach involves choosing an alternative estimation method, such as fitting some 2D function to the neighborhood around the sample point.

Finally, a hybrid approach could be employed. This involves using grid data collection with variable resolution. For example, within the LED’s neighborhood, samples could be collected at a 6 cm resolution, while using a 20 cm resolution elsewhere. This approach leverages the advantages of grid-structured data, which works effectively with IDW, while at the same time being more accurate where the data is more varied.

7.3 Usage of MLP for Evaluating Data Augmentation

There is a question to be asked about the viability of an MLP to assessing the quality of the augmentation. If the end goal is not necessarily to use a multi-layer perceptron network, but some other algorithm, this choice becomes questionable. MLPs are not easily explainable and quite unpredictable. They are somewhat resistant to noise and are able to generalize even when one would not necessarily expect it. It might be worth considering some other benchmark to compare augmentation pipelines.

8 Conclusions

We proposed a revised cleaning and augmentation pipeline, including an updated scoring function to retain more relevant data points and a method for estimating LED positions. These improvements, combined with our updated estimation algorithm, enabled more accurate RSS data reconstruction. Our benchmarks demonstrated that these modifications increase the accuracy of a multi-layer perceptron (MLP) by up to 20%. Furthermore, we successfully deployed the model on an Arduino Nano microcontroller.

Additionally, we experimented with alternative data acquisition strategies. Our findings indicated that collecting data in a rigid grid remains the superior method with the current RSS interpolation algorithm. However, collecting data more extensively in chosen areas improved the local accuracy, particularly in denser LED deployments.

9 Future Work

Several opportunities for future research can be extrapolated from this study. Firstly, it is crucial to test these methods in more challenging environments where light reflection and multipath effects are present. Secondly, there is a recognized need for improved interpolation algorithms capable of handling non-grid data. A thorough analysis of different estimation methods and selection of the most promising one could substantially improve the accuracy of our experiments. Additionally, exploring hybrid approaches to data collection, such as variable density grids, would be valuable due to their relatively straightforward deployment and potential for improved accuracy.

References

- [1] D. Chen, K. Shin, Y. Jiang, and K.-H. Kim, "Locating and tracking BLE beacons with smartphones," presented at the CoNEXT 2017 - Proceedings of the 2017 13th International Conference on emerging Networking EXperiments and Technologies, 2017, pp. 263–275. DOI: 10.1145/3143361.3143385. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85040241402&doi=10.1145%2f3143361.3143385&partnerID=40&md5=4cdfcd75fd89c10902c601948b15ba8e>.
- [2] K. Chintalapudi, A. Iyer, and V. Padmanabhan, "Indoor localization without the pain," *ACM MobiCom 2010*, pp. 173–184, 2010.
- [3] L. Barbieri, M. Brambilla, A. Trabattoni, S. Mervic, and M. Nicoli, "UWB localization in a smart factory: Augmentation methods and experimental assessment," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, 2021. DOI: 10.1109/TIM.2021.3074403.
- [4] W.-H. Kuo, Y.-S. Chen, G.-T. Jen, and T.-W. Lu, "An intelligent positioning approach: RSSI-based indoor and outdoor localization scheme in zigbee networks," presented at the 2010 International Conference on Machine Learning and Cybernetics, ICMLC 2010, vol. 6, 2010, pp. 2754–2759. DOI: 10.1109/ICMLC.2010.5580783.
- [5] A. Kendall and R. Cipolla, "Geometric loss functions for camera pose regression with deep learning," presented at the Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, vol. 2017-January, 2017, pp. 6555–6564, ISBN: 978-1-5386-0457-1. DOI: 10.1109/CVPR.2017.694.
- [6] J. Torres-Sospedra, R. Montoliu, S. Trilles, Ó. Belmonte, and J. Huerta, "Comprehensive analysis of distance and similarity measures for wi-fi fingerprinting indoor positioning systems," *Expert Systems with Applications*, vol. 42, no. 23, pp. 9263–9278, 2015. DOI: 10.1016/j.eswa.2015.08.013.
- [7] H. Chan, C.-W. Chow, L.-S. Hsu, *et al.*, "Utilizing lighting design software for simulation and planning of machine learning based angle-of-arrival (AOA) visible light positioning (VLP) systems," *IEEE Photonics Journal*, vol. 14, no. 6, 2022, ISSN: 1943-0655. DOI: 10.1109/JPHOT.2022.3212628.
- [8] Z. Xu, C. Gong, and B. Bai, "Visible light positioning and communication," in *Visible Light Communication*, 2015, pp. 88–106. DOI: 10.1017/CBO9781107447981.005. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84952648450&doi=10.1017%2fCBO9781107447981.005&partnerID=40&md5=3131adff8a9a2dbbf5d33ee3ea1755b5>.
- [9] Y.-C. Wu, C.-W. Chow, Y. Liu, *et al.*, "Received-signal-strength (RSS) based 3d visible-light-positioning (VLP) system using kernel ridge regression machine learning algorithm with sigmoid function data preprocessing method," *IEEE Access*, vol. 8, pp. 214269–214281, 2020. DOI: 10.1109/ACCESS.2020.3041192.
- [10] A. Gradim, P. Fonseca, L. Alves, and R. Mohamed, "On the usage of machine learning techniques to improve position accuracy in visible light positioning systems," presented at the 2018 11th International Symposium on Communication Systems, Networks and Digital Signal Processing, CSNDSP 2018, 2018. DOI: 10.1109/CSNDSP.2018.8471773.
- [11] P. Du, S. Zhang, C. Chen, A. Alphones, and W.-D. Zhong, "Demonstration of a low-complexity indoor visible light positioning system using an enhanced tdoa scheme," *IEEE Photonics Journal*, vol. 10, no. 4, 2018. DOI: 10.1109/JPHOT.2018.2841831.
- [12] S. Zhang, P. Du, C. Chen, and W.-D. Zhong, "3d indoor visible light positioning system using RSS ratio with neural network," presented at the 23rd Opto-Electronics and Communications Conference, OECC 2018, 2018. DOI: 10.1109/OECC.2018.8729887.
- [13] R. Zhu, M. Van Den Abeele, J. Beysens, J. Yang, and Q. Wang, "Centimeter-level indoor visible light positioning," *IEEE Communications Magazine*, vol. 62, no. 3, pp. 48–53, 2024. DOI: 10.1109/MCOM.002.2300296. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85179779242&doi=10.1109%2fMCOM.002.2300296&partnerID=40&md5=3bf9e7e87a4d4188c7d24b6d45fc42f3>.
- [14] J. Beysens, A. Galisteo, Q. Wang, D. Juara, D. Giustiniano, and S. Pollin, "Densevlc: A cell-free massive mimo system with distributed leds," in *Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies*, 2018, pp. 320–332.
- [15] I. Kasa, "A circle fitting procedure and its error analysis," *IEEE Transactions on Instrumentation and Measurement*, vol. IM-25, no. 1, pp. 8–14, 1976, ISSN: 0018-9456. DOI: 10.1109/TIM.1976.6312298.
- [16] V. Rekkas, L. Iliadis, S. Sotiroudis, *et al.*, "Artificial intelligence in visible light positioning for indoor IoT: A methodological review," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 2838–2869, 2023. DOI: 10.1109/OJCOMS.2023.3327211. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85176335454&doi=10.1109%2fOJCOMS.2023.3327211&partnerID=40&md5=1d9d7b24ebcc828cc744030d9328ed91>.