



Article History:

Received:

15/06/2026

Accepted:

20/08/2026

Published:

02/09/2026

**A Controlled Comparison of Classical and Neural
Models for Indoor Visible Light Positioning**

Mustafa M. Abuali ^{1*}, Alhusein Almahjoub ²

¹ Department of Computer and Information Technology, College
of Electronic Technology, Bani Walid, Libya

² Department of Communications Engineering, College of
Electronic Technology, Bani Walid, Libya.

**مقارنة محكمة بين النماذج الكلاسيكية والعصبية لتحديد
المواقع باستخدام الضوء المرئي داخل المباني**

مصطفى مفتاح ابو علي^{1*}، الحسين المحجوب²

¹ قسم هندسة الاتصالات، كلية التقنية الإلكترونية، بني وليد، ليبيا

² قسم الحاسوب وتقنية المعلومات، كلية التقنية الإلكترونية، بني وليد، ليبيا

*Corresponding author: maboali2050@gmail.com



Copyright: © 2026 by the authors. This article is an open-access article
distributed under the terms and conditions of the Creative Commons
Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract

Indoor visible light positioning (VLP) studies often use simplified channel and noise assumptions. In particular, signals from several LEDs may be combined at each photodiode, while receiver noise is sometimes represented as a fixed percentage of the received signal. In this work, we use a simulation setup in which the LEDs are treated as individually identifiable sources, representing an idealised time-, frequency-, or code-multiplexed system. Single-bounce NLOS reflections from the four vertical walls are included using an adaptive mesh, and receiver noise is modelled using shot and thermal noise. Six learning models are evaluated: an MLP, a CRNN, an RBF-SVM, a random forest, a k -nearest-neighbour regressor, and a decision tree. Their results are compared with an Oracle maximum-likelihood reference and a Centre baseline. Two room geometries, four receiver heights, and LED counts from 1 to 16 are considered. The RBF-SVM gives the lowest mean RMSE for $N_{LED} \leq 8$, whereas the MLP gives the lowest mean RMSE at 16 LEDs. For $N_{LED} \geq 4$, both models give mean RMSE values between 0.10 and 0.22 m. The CRNN does not improve on the MLP in the tested cases, while the tree-based and instance-based methods give higher errors. In the $10 \times 10 \times 5$ m room, the Oracle result remains close to zero at 16 LEDs. This suggests that the higher error of the learned models is related to the learning setup rather than a loss of positioning information in the channel.

Keywords: visible light positioning, optical wireless, LED identification, machine learning, channel modelling, physical receiver noise.

الملخص

غالباً ما تعتمد دراسات تحديد المواقع باستخدام الضوء المرئي (VLP) داخل المباني على فرضيات مبسطة للقنوات والضوضاء. وتحديداً، قد يتم دمج الإشارات القادمة من عدة صمامات ثنائية باعثة للضوء (LEDs) عند كل ثنائي ضوئي (Photodiode)، في حين يتم تمثيل ضوضاء المستقبل أحياناً كنسبة مئوية ثابتة من الإشارة المستقبلية. في هذا العمل، نستخدم إعداداً للمحاكاة تُعامل فيه مصابيح الـ LED كمصادر يمكن التعرف عليها بشكل فردي، مما يمثل نظاماً مثالياً متعدد الإرسال القائم على الوقت أو التردد أو الشفرة. تم تضمين انعكاسات الإرسال غير المباشر (NLOS) ذات الارتداد الواحد من الجدران الأربعة باستخدام شبكة متكيفة، كما تم نمذجة ضوضاء المستقبل باستخدام اللقطة (Shot noise) والضوضاء الحرارية (Thermal noise). تم تقييم ستة نماذج تعلم آلي: الإدراك المرتد متعدد الطبقات (MLP)، والشبكة العصبية المتكررة التلافيفية (CRNN)، وآلة ناقلات الدعم ذات دالة

الأساس الشعاعي (RBF-SVM)، والغابة العشوائية (Random Forest)، والانحدار بأقرب جيران (k-NN regressor)، وشجرة القرار (Decision Tree). وقورنت نتائجها مع مرجع الاحتمال الأقصى (Oracle) وخط الأساس المركزي (Centre baseline). تم أخذ هندستين للغرف في الاعتبار، وأربعة ارتفاعات للمستقبل، وأعداد من مصابيح الـ LED تتراوح من 1 إلى 16. أعطى نموذج RBF-SVM أدنى متوسط لجذر متوسط مربع الخطأ (RMSE) عندما كان عدد مصابيح الـ LED أقل من أو يساوي 8، في حين سجل نموذج MLP أدنى متوسط لـ RMSE عند 16 مصباح. LED بالنسبة لعدد مصابيح LED أكبر من أو يساوي 4، أعطى كلا النموذجين قيم متوسط RMSE تتراوح بين 0.10 و 0.22 متر. لم يظهر نموذج CRNN أي تحسن مقارنة بـ MLP في الحالات المختبرة، بينما سجلت الطرق القائمة على الأشجار والطرق القائمة على الحالات (Instance-based) أخطاءً أعلى. في غرفة بأبعاد $5 \times 10 \times 10$ أمتار، ظلت نتيجة مرجع (Oracle) قريبة من الصفر عند 16 مصباح LED، وهذا يشير إلى أن الخطأ الأعلى في النماذج المتعلمة يرتبط بإعدادات التعلم نفسها وليس بفقدان معلومات تحديد المواقع داخل القناة.

الكلمات المفتاحية: تحديد المواقع بالضوء المرئي، الاتصالات اللاسلكية الضوئية، تحديد هوية مصابيح الـ LED، التعلم الآلي، نمذجة القنوات، الضوضاء الفيزيائية للمستقبل.

1. Introduction

Visible light communication (VLC) has matured over the past two decades from a niche laboratory demonstration into a viable complement to radio-frequency wireless access [1]. Because the same luminaries that provide illumination can also carry an identifying signal, VLC naturally enables *visible light positioning* (VLP): a receiver equipped with photodiodes (PDs) estimates its position from the received optical power pattern [2]. VLP is attractive for indoor location-based services because it reuses existing lighting infrastructure, does not interfere with licensed radio spectrum, and can in principle deliver decimetre accuracy [3]. Previous work has considered several aspects of VLP. Saadi *et al.* [2] review geometric, fingerprinting, and learning-based positioning methods. Pugliese *et al.* [4] study a single-LED system assisted by distributed optical intelligent reflecting surfaces and show that useful positioning accuracy can also be obtained by changing the transmitter-side configuration. Alijani *et al.* [5] consider visible-light sensing for motion and presence detection, while Dratnal *et al.* [6] review bio-inspired methods for optimising LED placement and receiver geometry. Machine-learning methods have also received increasing attention in indoor VLP. Yu *et al.* [7] review these approaches for edge-centric systems and discuss the computational cost of different model families.

Two modelling choices are particularly relevant to this study. In some simulations, each PD is assumed to measure the sum of the powers received from all LEDs. With this assumption, adding LEDs can make the total received-power field more uniform and may increase the positioning error. This is different from a multiplexed system in which the contribution from each LED can be identified separately. Receiver noise is also sometimes modelled as a fixed fraction of the signal. Such a model keeps the relative noise level approximately constant and does not represent how SNR changes with received optical power. Saadi *et al.* [8] showed that LED placement affects the trade-off between coverage uniformity and throughput, but the combined effect of LED count and receiver noise on positioning was not considered.

This paper investigates these two modelling issues using a controlled simulation study. The main contributions are:

- We treat the LEDs as individually identifiable sources at each PD, so that each measurement corresponds to one LED. This is an *idealised abstraction* of a physical multiplexing scheme (TDM, FDM or CDM); we do not model the duty-cycle, orthogonality, or code-length overhead of any specific implementation. See Section II.E.
- We integrate single-bounce NLOS reflections from the four vertical walls with an adaptive mesh that subdivides elements close to a receiver into ~ 2 cm sub-elements. A convergence study confirms $< 1\%$ 99th-percentile agreement against the finest mesh.
- We replace relative noise with a physical shot-plus-thermal receiver model due to Komine and Nakagawa [9], so that weak signals are relatively noisier and the total transmit power matters.

- We compare six learning models and two reference baselines under identical train/validation/test splits, and report mean $\pm$ std over three repeated spatial splits, following the quality-of-service perspective of Bastiaens *et al.* [10].

The results show that the RBF-SVM gives the lowest mean RMSE for $N_{\text{LED}} \leq 8$, while the MLP gives the lowest mean at 16 LEDs. The CRNN does not improve on the MLP for the repeated measurements used here. In the square room, the Oracle error remains very small at 16 LEDs, which suggests that the poorer performance of the learned models at this LED count is not caused only by the channel. Possible reasons include the larger feature dimension, model capacity, and optimisation, but these effects are not separated in the present experiments.

II. System Model

A. Rooms and Transmitters

We consider two rooms of equal area but different aspect ratio: S1 ($10 \times 10 \times 5$ m) and S2 ($5 \times 20 \times 5$ m). LEDs are mounted on the ceiling at $z = L_z$ and point straight down. They are arranged in a uniform grid: 1 at the centre, 4 in a 2×2 arrangement, 8 in a 4×2 grid aligned with the long axis of the room, and 16 in a 4×4 grid. The emission follows a Lambertian pattern with half-power semi-angle $\phi_{1/2} = 60^\circ$, giving order $m = 1$. A total optical power $P_{\text{tot}} = 30$ W is shared equally among the $N_{\text{LED}} \in \{1, 4, 8, 16\}$ LEDs, so the per-LED power is $P_{\text{tot}}/N_{\text{LED}}$.

B. Receiver

The receiver consists of four PDs. PD 0 is horizontal with normal $\mathbf{n}_0 = (0, 0, 1)$. PDs 1–3 are tilted by $\beta = 30^\circ$ from vertical, with normals

$$\mathbf{n}_k = (\sin\beta \cos\alpha_k, \sin\beta \sin\alpha_k, \cos\beta), \quad (1)$$

where $\alpha_k \in \{0^\circ, 120^\circ, 240^\circ\}$ for $k = 1, 2, 3$. Their centres are offset from the reference point by $l_{\text{pd}} = 5$ cm in the direction of their azimuth. Each PD has active area $A_p = 1$ cm², field of view (FOV) 70° , concentrator refractive index $n = 1.5$, and optical filter gain $T_s = 1$. The horizontal grid spacing is 0.2 m; receiver heights are $h \in \{0, 1, 2, 3\}$ m.

C. Line-of-Sight Channel

The DC LOS gain from a single LED to a PD follows the standard Lambertian model

$$H_{\text{LOS}} = \frac{(m+1)A_p}{2\pi d^2} \cos^m \phi T_s g(\psi) \cos\psi, \quad (2)$$

where d is the LED-to-PD distance, ϕ is the emission angle (with respect to the LED's downward normal), ψ is the incidence angle at the PD (with respect to the PD normal), and $g(\psi)$ is the ideal non-imaging concentrator gain

$$g(\psi) = \begin{cases} n^2/\sin^2(\text{FOV}), & 0 \leq \psi \leq \text{FOV}, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

D. Non-Line-of-Sight Channel

We model single-bounce reflections from the four vertical walls, with reflectance $\rho = 0.8$; the floor and ceiling are treated as non-reflecting. Each wall is discretised into planar elements of area A_w . For a Lambertian reflector, the received optical power at a PD via a single wall element is

$$H_{\text{NLOS}} = \sum_w \frac{(m+1)A_p T_s \rho A_w}{2\pi^2 d_1^2 d_2^2} \cos^m \phi_1 \cos\alpha_1 g(\psi_2) \cos\psi_2 \cos\beta_2, \quad (4)$$

where, for wall element w with centre $\mathbf{p}_w$ and inward unit normal $\mathbf{n}_w$:

- $d_1 = \|\mathbf{p}_w - \mathbf{p}_{\text{LED}}\|$ is the LED-to-wall distance,
- $d_2 = \|\mathbf{p}_{\text{PD}} - \mathbf{p}_w\|$ is the wall-to-PD distance,
- ϕ_1 is the emission angle at the LED (relative to the LED's downward normal),
- α_1 is the incidence angle at the wall (between the LED-to-wall vector and $\mathbf{n}_w$),
- ψ_2 is the incidence angle at the PD (between the wall-to-PD vector and the PD normal $\mathbf{n}_{\text{PD}}$),

- β_2 is the emission angle at the wall (between the wall-to-PD vector and $\mathbf{n}_w$),
- the factor π in the denominator arises from the Lambertian reflector radiance model,
- the concentrator term $g(\psi_2)$ applies the same FOV rejection as in the LOS case.

To evaluate (4) numerically, each wall element is treated as a point reflector at its centre.

The default element size is $h_e = 0.125$ m. When a PD is close to a wall element, the point approximation becomes less accurate. For receivers within $3h_e + l_{pd}$ of a wall, nearby elements are therefore divided into ~ 2 cm sub-elements, and the corresponding coarse contribution is removed to avoid double-counting. Table 1 gives the convergence results for the NLOS component at $h = 0$ m. With refinement, the 0.125 m mesh has a 99th-percentile relative difference of 0.95% compared with the 0.0625 m reference mesh. Without refinement, the corresponding value is 26%, with a maximum difference above 55%. For this reason, the refined mesh is used in the remaining simulations.

Table 1. NLOS convergence, $h = 0$ m, SI, 4 LED. Reference: 0.0625 m mesh.

Element	Refined	p99 rel. diff.	Max rel. diff.
(m)		(%)	(%)
0.500	yes	2.33	3.46
0.250	yes	1.35	1.95
0.125	yes	0.95	1.31
0.125	no	26.4	54.8
0.0625	yes	0 (ref.)	0 (ref.)

E. Receiver Noise and Multiplexing Abstraction

For physical realism we use the Komine–Nakagawa receiver model [9]. The shot-noise contribution is $\sigma_{\text{shot}}^2 = 2qRPB + 2qI_{bg}I_2B$, and the thermal-noise contribution is

$$\sigma_{\text{th}}^2 = \frac{8\pi kT\eta A I_2 B^2}{G} + \frac{16\pi^2 kT\Gamma\eta^2 A^2 I_3 B^3}{g_m}, \quad (5)$$

with $R = 0.53$ A/W, $B = 100$ MHz, $I_{bg} = 5.1$ mA, $I_2 = 0.562$, $I_3 = 0.0868$, $T = 295$ K, $G = 10$, $\eta = 1.12 \times 10^{-6}$ F/m², $\Gamma = 1.5$, and $g_m = 30$ mS. Each position is measured $N_r = 10$ times with independent noise realisations.

Multiplexing abstraction. We assume that the received contribution from each LED can be identified separately. Each per-LED measurement is taken over the full symbol duration, so the receiver-noise model is applied without an additional multiplexing penalty. A practical TDM system would reduce the available duty cycle for each LED, while FDM and CDM would introduce other losses associated with bandwidth allocation or code properties. These implementation effects are not included here. The reported positioning accuracy therefore represents an idealised case for a multiplexed VLC system with the same total optical power.

F. Features

From the raw power array of shape (4 PDs, N_{LED}) we build two features per PD and LED:

- a compressed level $\text{arcsinh}(P/s)$, with $s = 10^{-6}$ W, which tolerates small negative noisy values and compresses the dynamic range of P ;
- a *directional power ratio* $P/(\sum_{\text{PD}} |P| + s)$, which measures the share of one LED's light captured by each tilted PD.

The second feature describes how the received power from one LED is distributed across the four PD orientations. It is motivated by angle-of-arrival behaviour, but it is not used as an explicit AoA estimate. After concatenating the two features over all PDs and LEDs, each position is represented by $8N_{\text{LED}}$ input features.

III. Experimental Protocol

Positions are grouped into 1×1 m spatial blocks. Complete blocks are then assigned randomly to the training (70%), validation (15%), and test (15%) sets. Both rooms contain 100 blocks: a 10×10 arrangement for S1 and a 5×20 arrangement for S2. With a grid spacing of 0.2 m, each block contains about 25 receiver positions, giving approximately 1750 training, 375 validation, and 375 test positions in each repetition. Within a repetition, the same spatial split is used for every LED count, receiver height, and model. Three repetitions are performed. Performance is measured using the RMSE of the 2D horizontal positioning error. Median error, P95, maximum error, and the fraction of errors below 0.5 m are also recorded. Since only three repetitions are used, small differences between models should be interpreted with caution. Across the tested configurations, the standard deviation between repetitions is approximately 0.01–0.10 m.

A. Models

Oracle. Maximum-likelihood matching is performed using the true noiseless channel map and the true noise standard deviations on the same grid as the test positions. This is used only as a reference because it assumes exact knowledge of the simulated channel. It is not a Cramér–Rao or MMSE bound, but it helps show whether the simulated channel contains enough information for accurate position estimation.

Centre. Predicts the mean position of the training set and is used as a simple no-information baseline.

MLP. Fully connected network on averaged features. Two hidden-layer sizes and two regularisation settings are tuned on validation.

CRNN. A Conv1D–Conv1D–LSTM–FC network is applied to the N_r repeated measurements. These repeats are independent measurements rather than a time sequence. The CRNN is included to examine whether processing the repeats jointly provides any advantage over using averaged features.

RBF-SVM, RF, KNN, DT. Hyperparameters are selected on the same validation blocks using GridSearchCV with a predefined split. The choice of models is motivated by the survey of Yu *et al.* [7], which identifies MLPs and SVM regressors as the most commonly deployed families in edge-centric VLP deployments.

IV. Results

A. Model Comparison Across LED Counts

Fig. 1 plots the test RMSE against the LED count for all models in S1, at each receiver height. Fig. 2 shows the same for the elongated room S2. Table 2 gives the numerical summary, averaged over heights and the three repeats.




Figure 1. Test RMSE (mean $\pm$ std over 3 repeats) vs. LED count for the 10 $\times$ 10 $\times$ 5 m room. One panel per receiver height. Dashed curves are the Oracle and Centre references.




Figure 2. Test RMSE (mean $\pm$ std over 3 repeats) vs. LED count for the 5 $\times$ 20 $\times$ 5 m room.

The elongated geometry has a wider spread between the best and worst models, and the Centre baseline is systematically higher than in S1.

Table 2. Test RMSE (m), mean over heights and three repeats. Best learning model per row in bold. The standard deviation across repeats is on the order of 0.01–0.10 m and is omitted for compactness.

Room	N	MLP	CRNN	SVM	RF	KNN	DT
S1	1	0.311	0.348	0.305	0.506	0.589	0.796
S1	4	0.138	0.199	0.101	0.298	0.555	0.685
S1	8	0.157	0.212	0.146	0.306	0.507	0.776
S1	16	0.176	0.227	0.246	0.338	0.467	0.858
S2	1	0.469	0.469	0.461	0.774	0.667	0.892
S2	4	0.187	0.254	0.151	0.463	0.673	0.749
S2	8	0.199	0.247	0.169	0.425	0.651	0.798
S2	16	0.222	0.250	0.280	0.438	0.623	0.850

The RBF-SVM gives the lowest mean RMSE for 1, 4, and 8 LEDs in both rooms. For the 4- and 8-LED cases, its mean RMSE is between 0.10 and 0.17 m. This is in line with the observations summarised by Yu *et al.* [7], who reported that kernel-based approaches can remain competitive with neural networks in indoor VLP. At 16 LEDs, the MLP gives the lowest mean RMSE in both rooms. For $N_{\text{LED}} \geq 4$, its mean RMSE varies from 0.14 to 0.22 m. The CRNN gives errors about 0.03–0.07 m higher than the MLP in the corresponding configurations and also requires more fitting time in our implementation (15–60 s for the MLP and 35–150 s for the CRNN). The RF, KNN, and DT give larger errors than the MLP and RBF-SVM in most of the tested cases.

Only three spatial splits were used, so small differences in the reported mean values should not be interpreted as statistically significant differences between the models.

B. Error Distribution

Fig. 3 shows the empirical CDF of the 2D error for S1 with 4 LEDs, pooling all heights and repeats. Fig. 4 shows the same for S2 with 16 LEDs, a lower-SNR regime.




Figure 3. Empirical CDF of the 2D error. S1, 4 LEDs, all heights and repeats pooled. The MLP, SVM, and CRNN cluster tightly on the left; the Centre baseline occupies the far right.




Figure 4. Empirical CDF of the 2D error. S2, 16 LEDs, all heights and repeats pooled. The elongated room at high LED count is the harder configuration; the Oracle curve develops a visible tail at low SNR.

The CDFs also show that RMSE does not describe the full error distribution. This is relevant to the quality-of-service discussion in Bastiaens *et al.* [10]. For example, for the CRNN in S1 with 16 LEDs and $h = 0$, the median error is 0.189 m, the RMSE is 0.254 m, and the maximum error is 0.86 m. Reporting median and P95 errors together with RMSE is therefore useful when the application has a maximum acceptable positioning error.

C. Role of Receiver Height and SNR

Table 3 shows the mean RMSE at each receiver height for the MLP, SVM, DT, and Oracle in S1 with 16 LEDs. The MLP changes only slightly with height. The SVM error is lower at $h = 3$ m than at $h = 0$ m in this case, although the present experiments do not identify the reason for this trend. The DT shows a less consistent response and gives its largest error at $h = 3$ m.

Table 3. RMSE (m) vs. receiver height. S1, 16 LEDs.

h (m)	MLP	SVM	DT	Oracle
0	0.194	0.281	0.889	0.052
1	0.164	0.249	0.811	0.016
2	0.166	0.236	0.765	~ 0
3	0.181	0.217	0.969	~ 0

Table 4 gives the median per-measurement SNR and NLOS-to-LOS ratio for four representative cases. The SNR values include LEDs that fall outside the receiver FOV. These measurements have zero signal, so the reported median combines out-of-FOV measurements with in-FOV measurements that have low SNR. The values should therefore be interpreted as an overall summary rather than an in-FOV SNR measure. With the receiver parameters used here, the median SNR in S1 falls by 27.4 dB between the 1-LED and 16-LED cases. A similar change is observed in S2. At 16 LEDs and $h = 3$ m, the median SNR is -1.3 dB in S1 and -11.3 dB in S2; more than 60% of the individual measurements in the latter case are below 0 dB.

Table 4. Channel statistics. Median per-measurement SNR and median NLOS/LOS ratio. The median SNR includes measurements from LEDs outside the receiver FOV and is therefore a conservative summary.

Config.	N	h (m)	SNR (dB)	NLOS/LOS
S1	1	0	35.6	0.21
S1	16	0	8.2	0.29
S1	16	3	-1.3	0.11
S2	16	3	-11.3	0.17

The Oracle results are different in these two cases. In S1 with 16 LEDs, the mean Oracle RMSE remains close to zero at all receiver heights (Table 3), showing that accurate position recovery is still possible with the simulated channel. In S2 with one LED at $h = 3$ m, the Oracle mean RMSE varies from 0.39 to 1.24 m across repetitions, with some errors above 4 m. Inspection of the likelihood surface shows several similar minima at symmetric positions in the long room, causing the maximum-likelihood estimate to move between them. The poor result in this case is therefore associated with ambiguity in the channel, whereas this behaviour is not observed for S1 with 16 LEDs.

D. Role of the LED Count

For both the MLP and SVM, the lowest mean RMSE among the tested LED counts occurs at $N_{\text{LED}} = 4$. For the MLP in S1, the height-averaged mean RMSE increases from 0.138 m at 4 LEDs to 0.157 m at 8 LEDs and 0.176 m at 16 LEDs (Table 2).

Several effects change simultaneously as N_{LED} increases. The per-LED power scales as $1/N_{\text{LED}}$, so under the simulated receiver-noise parameters the median per-measurement SNR falls by approximately 27.4 dB from 1 to 16 LEDs (Table 4). At the same time the input feature dimension grows linearly with N_{LED} ($8N_{\text{LED}}$ features), so each LED count is a separate learning

problem with a different input dimensionality, and the model capacity and optimisation landscape also change.

The Oracle does not show the same increase in error. In S1 with 16 LEDs, its mean RMSE remains close to zero at all heights (Table 3). This suggests that the channel still contains sufficient positioning information, while the learned models are affected by other parts of the modelling and training setup. As N_{LED} increases, the feature dimension, model fitting problem, and optimisation conditions all change at the same time. The current experiments do not separate these effects. The result therefore applies to the models and feature representation considered here and should not be taken to mean that four LEDs are generally optimal for VLP. A related effect was reported by Saadi *et al.* [8] for VLC system design, where adding more fixtures did not always improve performance when the power available to each fixture was constrained. The present positioning results show a similar trend for the models and features evaluated here.

E. Room Aspect Ratio

The elongated room S2 gives higher errors than S1 for most models. At 16 LEDs, the MLP mean RMSE is 0.176 m in S1 and 0.222 m in S2, while the Centre baseline is 3.81 m and 5.68 m, respectively. S2 also has a larger visible wall area and a higher NLOS contribution, and some positions between LEDs have relatively low SNR. The difference is about 30% for the MLP at 16 LEDs. This result is consistent with the point made by Bastiaens *et al.* [10] that positioning performance depends on the deployment geometry.

V. Discussion

The main observations from the experiments are discussed below.

Model performance changes with the LED count. The RBF-SVM has the lowest mean RMSE for $N_{\text{LED}} \leq 8$, whereas the MLP has the lowest mean at $N_{\text{LED}} = 16$. The RBF-SVM also has a lower fitting cost in our implementation. However, the difference between the RBF-SVM and MLP is within the variation between repetitions for several configurations. The results therefore do not establish one model as consistently better across all settings. They also show that increasing N_{LED} while keeping P_{tot} fixed changes both the channel SNR and the learning problem.

The MLP gives relatively stable results across the tested configurations. The CRNN does not provide an accuracy improvement despite its longer fitting time. One possible reason is that the repeated measurements are independent rather than temporally ordered, so averaging already reduces their measurement noise before the MLP is applied. For systems with limited computational resources, the RBF-SVM may also be useful, particularly for the smaller LED counts considered here [7].

The Oracle reference is useful for interpreting the model errors. In S2 with one LED at $h = 3$ m, the Oracle itself has a relatively large error, indicating ambiguity in the available channel information. In S1 with 16 LEDs, the Oracle error remains small even though the learned models perform less well. Without this reference, these two cases would appear similar if only the learned-model RMSE were reported. Pugliese *et al.* [4] also showed that changes on the transmitter side can improve VLP performance, which is relevant when the main limitation comes from the channel rather than the regression model.

This study has several limitations. Because the per-LED features are concatenated, the MLP input dimension increases with N_{LED} , and a separate model is trained for each LED count. A permutation-invariant architecture, such as DeepSets or an attention-based model, could be investigated in future work to handle different numbers of LEDs with one model. The experiments also use only three spatial splits, which limits the strength of comparisons between models. In addition, no sensitivity analysis was carried out for the receiver bandwidth B or background current I_{bg} . Both parameters affect the receiver noise and may influence the trend

with LED count. Finally, the LED-identification model is idealised and does not include the implementation losses of a practical multiplexing scheme. Future work could also examine device-free sensing [5] and joint sensing and positioning using the same channel model.

VI. Conclusion

This paper compared classical and neural learning models for indoor VLP using individually identifiable LED channels, single-bounce NLOS reflections, and a shot-plus-thermal receiver-noise model. The RBF-SVM gives the lowest mean RMSE for up to 8 LEDs, while the MLP gives the lowest mean RMSE at 16 LEDs. For 4–16 LEDs, both methods give mean RMSE values in the range 0.10–0.22 m. The CRNN does not improve on the MLP in the tested setup, and the RF, KNN, and DT give higher errors in most cases. When total optical power is fixed, increasing the number of LEDs beyond four also reduces the power available from each LED. The learned models show increasing error under these conditions, although the Oracle result in S1 remains close to zero at 16 LEDs. This indicates that the trend cannot be explained by channel information alone. In S2 with one LED at a receiver height of 3 m, the Oracle also gives a relatively large error because of ambiguity in the likelihood surface. Overall, the results show the importance of reporting channel conditions, including SNR and the NLOS/LOS ratio, when comparing VLP learning models.

References

- [1] M. Saadi and L. Wuttisittikulkij, “Visible light communication--the journey so far,” *Journal of Optical Communications*, vol. 40, no. 4, pp. 447–453, 2019.
- [2] M. Saadi, A. Bajpai, and D. Z. Rodriguez, “Traditional and modern techniques for visible light positioning systems,” in *Applications of 5G and Beyond in Smart Cities*. CRC Press, 2023, p. 169.
- [3] J. Armstrong, Y. A. Sekercioglu, and A. Neild, “Visible light positioning: A roadmap for international standardization,” *IEEE Communications Magazine*, vol. 51, no. 12, pp. 68–73, 2013.
- [4] D. Pugliese, G. Iacovelli, A. Fascista, D. Striccoli, O. Romanov, L. A. Grieco, and G. Boggia, “Visible light indoor positioning with a single LED and distributed single-element OIRS: An iterative approach with adaptive beam steering,” *IEEE Transactions on Communications*, 2026.
- [5] M. Alijani, C. De Cock, W. Joseph, and D. Plets, “Device-free visible light sensing: A survey,” *IEEE Communications Surveys & Tutorials*, vol. 28, pp. 3791–3829, 2025.
- [6] M. Dratnal, L. Danys, and R. Martinek, “Bio-inspired optimization methods for visible light communication: A comprehensive review,” *Artificial Intelligence Review*, vol. 58, no. 8, p. 246, 2025.
- [7] Y. Yu, Y. Qu, D. Zhao, T. Zhong, T. H. Luan, and S. Yu, “Machine learning for edge-centric indoor visible light positioning: A comprehensive survey and future directions,” *IEEE Internet of Things Journal*, 2026.
- [8] M. Saadi, M. Khosravi, S. Mumtaz, and C. Tsimenidis, “Optimized LED placement for indoor visible light communication: Trade-offs between uniform coverage and high throughput,” in *2025 Horizons of Information Technology and Engineering (HITE)*. IEEE, 2025, pp. 1–5.

- [9] T. Komine and M. Nakagawa, "Fundamental analysis for visible-light communication system using LED lights," *IEEE Transactions on Consumer Electronics*, vol. 50, no. 1, pp. 100–107, 2004.
- [10] S. Bastiaens, M. Alijani, W. Joseph, and D. Plets, "A review of the quality of positioning service of VLP within the indoor positioning system landscape," *IEEE Transactions on Instrumentation and Measurement*, 2026.
- [11] *Robust Indoor Localization in VLC Systems Using Machine Learning: A Comparative Study under Environmental Variability*. (2026). *Albahit Journal of Applied Sciences*, 5(3), 114-124. <https://doi.org/10.65419/albahit.v5i3.159>
- [12] *Physics-Informed Neural Networks for Predictive Modeling of Energy Dissipation Pathways in IoT-Integrated Smart Solar Grids*. (2025). *Albahit Journal of Applied Sciences*, 4(1), 182-194. <https://doi.org/10.65419/albahit.v4i1.56>
- [13] *Digital Twin-Integrated Predictive Control for Intelligent Manufacturing Systems: A Simulation and AI Approach*. (2025). *Albahit Journal of Applied Sciences*, 4(1), 346-352. <https://doi.org/10.65419/albahit.v4i1.71>

Compliance with ethical standards**Conflict of Interest Disclosure**

The authors declare that they have no conflicts of interest related to the publication of this work.

Disclaimer / Publisher's Note

The statements, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of Al-Mutawassit Journal for Basic and Applied Sciences (MJBAS), its publisher, or its editors. Neither the journal, its publisher, nor its editors shall be held responsible for any injury to persons or damage to property resulting from any ideas, methods, instructions, or products referred to in the published content.