electronics

Article

Efficient 3D Indoor Visible Light Positioning via Stabilized Regularized Least Squares

Yang Wang and Xiaona Liu *

The Institute of IoT, Shenzhen Polytechnic University, Shenzhen 518055, China; wyang@szpu.edu.cn
* Correspondence: liuxiaona@szpu.edu.cn

Abstract

We study indoor visible light positioning (VLP) in a fully three-dimensional setting under thermal noise and signal-dependent shot noise. Existing shot-noise-aware least-squares (LS) VLP methods largely rely on a fixed receiver height, so their key simplifications break down when the vertical coordinate is unknown. We derive a 3D likelihood-inspired surrogate, lift the nonlinear geometry into a weighted LS form, and add a geometry-aware regularizer that restores approximate consistency among redundant lifted variables. To make the iterative solver reliable in sparse layouts, we introduce room-feasible projection, consistent pilot reconstruction, clipped frozen weights, and damped objective-accepted updates. Each inner step remains a closed-form 4×4 linear solve. In deterministic synthetic simulations of an $8 \times 8 \times 3$ m room with layouts containing 4, 9, and 16 light-emitting diodes (LEDs), high- and low-power regimes, and 400 test samples per scenario, the stabilized solver removes the catastrophic sparse-layout failures observed for the tested LS-family baselines and retains a 14×–23× runtime advantage over the nonlinear baseline. The nonlinear baseline often attains lower mean error in sparse and moderate layouts; the contribution here is the resulting speed–tail-control trade-off within the fixed-orientation line-of-sight model. The resulting estimator is intended for synthetic 3D VLP simulations; hardware validation is outside the present study.

Keywords: visible light positioning; indoor localization; received signal strength; least squares; shot noise; 3D positioning

1. Introduction

Visible light positioning (VLP) is an attractive indoor localization technology because ceiling light-emitting diodes (LEDs) are already deployed for illumination and can simultaneously provide communication and positioning signals [1–3]. Compared with radio-frequency (RF) indoor positioning systems such as WiFi, Bluetooth, and ultra-wideband (UWB) [4,5], VLP can offer high accuracy, electromagnetic compatibility, and reuse of lighting infrastructure [6]. Practical applications include autonomous mobile robots, indoor navigation, asset tracking, and augmented-reality systems requiring reliable spatial estimates [7]. Surveys and tutorials [3,8–11] show that received signal strength (RSS)-based VLP remains appealing when one seeks low-cost photodiode (PD) receivers and lightweight inference. However, many RSS pipelines still use simplified additive-noise models or rely on nonlinear/metaheuristic optimization, as illustrated by recent evolutionary RSS-based VLP work [12].

Check for updates

Academic Editors: Walter Daniel Leon-Salas and Juan Carlos Torres Zafra

Received: 4 June 2026
Revised: 13 July 2026
Accepted: 15 July 2026
Published: 19 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

 https://doi.org/10.3390/electronics15143172

Our recent 2D work [13] showed that shot noise materially changes the optimization landscape and proposed an efficient least-squares (LS) alternative to direct gradient-based maximum likelihood estimation. That formulation assumes a fixed vertical separation between the LEDs and the receiver, effectively constraining the positioning problem to a horizontal plane. In many practical indoor applications, this assumption is too restrictive: the receiver height varies significantly across different use cases [14,15]. Mobile robots operate at various heights depending on their design [7], drones fly at different altitudes, handheld devices are carried at user-dependent heights [16], and even fixed installations may be mounted at non-standard elevations. These variations make the positioning task intrinsically three-dimensional, requiring estimation of the vertical coordinate alongside the horizontal position [17].

The 3D extension is not a trivial coordinate augmentation of the 2D setting. Once the height is unknown, the received optical power depends on the horizontal position and the vertical separation in a more strongly coupled way. At the same time, shot noise induces heteroscedasticity that now varies with the unknown height. A naive LS lifting introduces redundant variables whose inconsistency can destabilize the estimate, especially in sparse LED layouts or low signal-to-noise ratio (SNR) regimes.

The link to the fixed-height 2D formulation can be made explicit without adding another layer of notation. When the vertical separation is fixed at $u = u_0$, the terms u_0^2, u_0^ρ, and the u_0-dependent part of the variance weight are constants. Once u is unknown, the residual contains both u^2 and u^ρ, while the shot-noise-aware weight also depends on $\tau_i^2(u)u^4$. The 3D lift must therefore carry $t = u^\rho$ alongside $s = x^2 + y^2 + u^2$. These four lifted coordinates are not independent: they must satisfy $s = x^2 + y^2 + t^q$. The geometry-aware regularizer is needed to keep the locally solved LS variables near that physical consistency relation; it is not introduced merely because the coordinate vector has one additional entry.

The practical difficulty is threefold: unknown height changes the RSS scale and hence the shot-noise-aware weights; lifted LS variables make the algebra cheap but need not correspond to a physical 3D point; and sparse ceiling layouts amplify both effects because several heights can explain similar RSS patterns. The proposed solver addresses these issues through pilot and iterative weights, a linearized consistency regularizer, and a stabilization layer that rejects physically implausible or objective-worsening updates.

This paper builds a 3D LS-style estimator for this setting. We first derive a likelihood-inspired 3D surrogate under thermal noise and shot noise. We then convert it into a lifted weighted LS problem, and add a geometry-aware regularizer obtained from a first-order linearization of the consistency manifold linking the lifted variables. Finally, we stabilize the iterative reweighting stage by projecting iterates to the feasible room volume, reconstructing consistent pilot variables before each linearization, clipping extreme weights, and accepting only damped updates that improve the original surrogate. The goal is not to replace nonlinear maximum-likelihood optimization in every accuracy-critical regime, but to keep the LS subproblems closed form while reducing the divergent estimates seen in sparse 3D layouts.

The main contributions are:

- We extend the shot-noise-aware LS methodology from 2D to true 3D localization, where the unknown height fundamentally alters both the surrogate objective and the heteroscedastic weight structure, requiring a new formulation that properly handles the coupled 3D geometry.

- We derive a 3D lifted weighted LS formulation that turns each frozen, locally linearized subproblem into a structured linear form, together with a closed-form geometry-aware regularized subproblem that enforces consistency among redundant lifted variables. Theorem 1 provides the exact closed-form solution for this regularized quadratic subproblem.
- We propose a stabilized iterative reweighted regularized LS solver that preserves the computational efficiency of a 4×4 linear solve per iteration while improving robustness through room-feasible projection, consistent pilot reconstruction, clipped frozen weights, and damped surrogate-accepted updates. We state the closed-form guarantee for each frozen quadratic subproblem separately from the heuristic stabilization layer, and prove monotonic non-increase of the evaluated surrogate under the acceptance rule.
- We evaluate the methods in deterministic synthetic 3D simulations with systematic parameter tuning across 4-, 9-, and 16-LED layouts. The tables report accuracy and runtime, while the layout-scaling and survival-curve figures show how LED geometry and stabilization affect large errors. The convergence, λ-sensitivity, and Cramér–Rao lower-bound (CRLB) figures provide additional diagnostics. The stabilized method eliminates the catastrophic outliers observed in the tested sparse 4-LED settings and preserves approximately $14\times$–$23\times$ speedup over nonlinear optimization, while the nonlinear baseline often has lower mean error in sparse and moderate layouts. We also report repeated-seed and LED-position-jitter analysis.

The proposed stabilized LS solver is designed for settings where closed-form updates, runtime, and tail behavior matter. The nonlinear baseline attains lower mean error in some regimes.

The rest of the paper reviews related work, presents the model and surrogate, develops the solver and CRLB benchmark, and reports the synthetic experiments.

2. Related Work

VLP has been widely studied; see [2,3,8–11,18] for surveys and tutorials. RSS-based methods dominate many low-cost deployments because they require only intensity measurements: RSS/lateration-style models and hybrid frameworks [1,19,20] often assume comparatively simple noise models and can achieve 10–30 cm accuracy in favorable settings [21,22]. Recent machine-learning approaches [23–25] improve empirical robustness to model mismatch at the cost of training data, hardware-specific calibration, and weaker analytical guarantees. Timing-based methods can be accurate but require tight synchronization [26]. Angle-of-arrival (AOA) and hybrid AOA–RSS designs improve 3D observability [27–29], but they generally require angular-sensitive receivers or additional hardware complexity [30].

Full 3D VLP remains more delicate than fixed-height 2D localization because the vertical coordinate changes both the optical path geometry and the received-power scale. Dedicated 3D systems include multi-orientation or accelerometer-aided receiver designs [31,32], Bayesian formulations [33], particle-swarm optimization [34], low-cost 3D validation [15], robot-navigation validation [35], and RSS-ratio methods designed to reduce height dependence [14]. These works address important geometric or systems issues, but most of them either do not model shot noise explicitly or rely on nonlinear, metaheuristic, or data-driven inference. Recent calibration work further shows that installation parameters such as LED tilt and gain can materially affect received-signal-strength-based weighted least-squares (WLS) performance [36], reinforcing the need for algorithms whose assumptions and failure modes are explicit.

Two directly comparable 3D RSS studies further delimit the present setting. Almadani et al. [37] estimate a 3D position without prior height information by combining a Cayley–Menger determinant construction with a cost function and report experimental results for an industrial application. Almadani et al. [38] then experimentally evaluate a 3D VLP system in an industrial environment under receiver tilt and multipath reflections. These studies report unknown-height RSS positioning beyond the ideal LOS model. Their determinant-based formulation and measurement conditions differ from the thermal-plus-shot-noise LS model studied here, so their reported errors are not used as a cross-study numerical benchmark. Instead, they clarify that the present contribution is an algorithmic fixed-orientation LOS result, not an experimental robustness claim.

Recent studies also consider simultaneous 3D position and receiver-orientation estimation. Xu et al. [39] study a reconfigurable-intelligent-surface (RIS)-assisted visible-light-communication (VLC) system and jointly estimate position and orientation with Cramér–Rao lower-bound (CRLB) analysis, while Zhou et al. [40] exploit multipath interference and random fading for simultaneous position and orientation detection. Both studies estimate a broader state than the present work and use channel or orientation effects that are deliberately held fixed in our benchmark. In contrast, we estimate 3D position under fixed LED–PD orientation in a thermal-plus-shot-noise, line-of-sight RSS model, so that each update remains an LS-family subproblem with a closed-form solution. This narrower formulation makes the LS derivation and the tail-failure analysis tractable, but it also limits the claim. The proposed method should therefore be read as complementary to position–orientation estimators, not as a replacement for them.

Shot noise is signal-dependent and can dominate in low-illumination optical systems [41]. Nevertheless, much of the VLP literature treats noise as additive and signal-independent. Recent CRLB and Bayesian-bound analyses under signal-dependent or orientation-dependent models [42–44] show that the variance model affects both achievable accuracy and estimator design. Our prior 2D work [13] incorporated shot-noise heteroscedasticity into an iterative LS solver, but it relied on a fixed receiver height. That assumption removes the height-dependent terms that become central in the present paper.

On the optimization side, bisection-based estimators [45], convex RSS relaxations [46], and closed-form WLS designs [47] trade modeling tightness, speed, and robustness in different ways. Robust RSS localization with unknown model parameters [48], integrated positioning–communication power allocation [49], intelligent-reflecting-surface (IRS)-assisted VLP under orientation mismatch [50], and evolutionary RSS optimization under line-of-sight/non-line-of-sight (LOS/NLOS) conditions [12,51] broaden the design space. In particular, the evolutionary RSS optimizer of Pathak et al. [12] is the closest metaheuristic comparison point: it searches an RSS objective under LOS/NLOS conditions rather than exploiting a frozen closed-form LS subproblem. Its forward model, NLOS setting, and evaluation protocol differ from ours, so a numerical comparison across papers would conflate channel mismatch with optimization behavior. For this reason, the experimental baseline in this paper is a nonlinear optimizer applied to the same surrogate and room constraints; Pathak et al. are discussed as a complementary metaheuristic alternative rather than reproduced as an unmatched benchmark. The present paper is complementary: it focuses on a shot-noise-aware, true-3D RSS model and asks how far one can push an LS-family solver before resorting to generic nonlinear optimization. The main technical differences from [13] are the lifted 3D surrogate, the geometry-aware consistency regularizer, and the stabilization layer that controls sparse-layout failures.

Thus, the closest prior threads each cover only part of our target setting: fixed-height shot-noise LS provides the statistical motivation but not 3D geometry; existing 3D RSS systems provide the height-varying application setting but usually not shot-noise-aware LS structure; and CRLB or calibration studies clarify fundamental limits and model sensitivity but do not produce a closed-form iterative estimator. This paper sits at their intersection.

3. System Model and Likelihood-Inspired Surrogate

We consider an indoor VLP system with n LEDs mounted on the ceiling at a common height h [1,41]. The i-th LED is located at $\boldsymbol{p}_i = [x_i, y_i, h]^\top$. The PD receiver is located at $\boldsymbol{x} = [x, y, z]^\top$, where the height $z < h$ is unknown. We use the vertical separation $u := h - z > 0$ and the 3D parameter vector $\boldsymbol{\theta} := [x, y, u]^\top$ so that the optical distance and received-power terms can be written without repeatedly carrying the ceiling height. Figure 1 is placed before the channel equation to define the link distance d_i, the vertical separation u, and the common angle ϕ_i before these quantities enter the algebra.

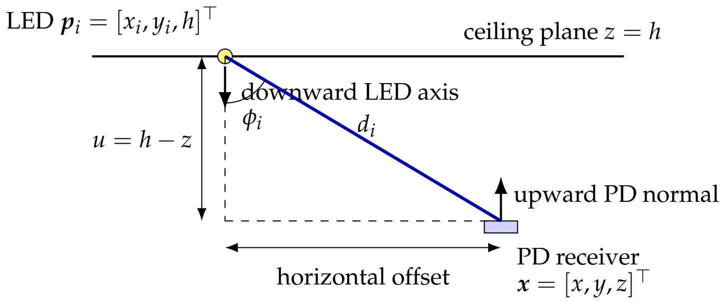

Figure 1. Fixed-orientation single-link geometry. The LED is mounted at ceiling height h, the receiver is at $[x, y, z]^\top$, the vertical separation is $u = h - z$, the link distance is d_i, and ϕ_i is the common irradiance/incidence angle under the vertically aligned LED–PD orientation assumption.

For the LED at position $\boldsymbol{p}_i$, the Lambertian channel gain can be written as

$$\gamma_i = \frac{\Gamma}{u^2} \cos^{m+3}(\phi_i), \qquad \cos(\phi_i) = \frac{u}{\sqrt{u^2 + (x_i - x)^2 + (y_i - y)^2}}, \tag{1}$$

where m is the Lambertian order and Γ is a location-independent optical-gain constant.

This channel model deliberately defines the fixed-orientation LOS framework used for the surrogate, the lifted LS subproblems, and the numerical Cramér–Rao lower-bound reference. Each LED points vertically downward, the PD normal points vertically upward, and all simulated links lie inside the PD field of view. Under these parallel vertical optical axes, the LED irradiance angle and the PD incidence angle have the same cosine. Thus, the standard LOS Lambertian factor $\cos^m(\phi_i)\cos(\psi_i)/d_i^2$ reduces to $\cos^{m+3}(\phi_i)/u^2$ after using $d_i = u/\cos(\phi_i)$. The constant Γ absorbs the optical area, filter gain, concentrator gain, Lambertian normalization, and fixed calibration gain. Receiver tilt changes the incidence-angle and visibility terms; multipath adds reflected optical paths and changes the signal mean, the signal-dependent variance, and the Fisher-information model. These effects are therefore not represented by simply adding another noise term to the present benchmark. A credible extension would require a re-derived forward model, likelihood-inspired surrogate, weights, and information bound, together with calibrated orientation, field-of-view, and reflection parameters. Such a setting is outside the fixed-orientation LS result reported here. Hence, the received signal from the i-th LED takes the form

$$r_i = P\gamma_i + n_i = \frac{P\Gamma}{u^2} \cos^{m+3}(\phi_i) + n_i, \tag{2}$$

$$\mathbb{E}[n_i] = 0, \qquad \mathbb{V}[n_i] = \sigma^2 + \zeta^2 P \gamma_i. \tag{3}$$

Here, r_i is the received RSS sample, P is the transmit power, n_i is the sum of thermal noise and shot noise, σ is the thermal-noise standard deviation, and ζ is the shot-noise coefficient.

Likelihood-Inspired 3D Surrogate

By rearranging the signal model, we obtain

$$\cos^{m+3}(\phi_i) = \frac{r_i u^2}{P\Gamma} - \frac{n_i u^2}{P\Gamma}. \tag{4}$$

Under a high-SNR regime, the ratio n_i/r_i is small and a first-order Taylor expansion yields

$$\cos^{-2}(\phi_i) = \left(\frac{r_i u^2}{P\Gamma}\right)^{-\frac{2}{m+3}} + \left(\frac{u^2}{P\Gamma}\right)^{-\frac{2}{m+3}} \frac{2}{m+3} r_i^{-\frac{m+5}{m+3}} n_i + o\left(\frac{n_i}{r_i}\right). \tag{5}$$

Define

$$\widetilde{r}_i := \max\{r_i, \epsilon_r\}, \qquad b_i := \left(\frac{\widetilde{r}_i}{P\Gamma}\right)^{-\frac{2}{m+3}}, \tag{6}$$

where $\epsilon_r > 0$ is a positive RSS floor. In all reported experiments $\epsilon_r = 10^{-6}$. The Taylor expansion above is a positive-RSS, small-noise derivation; replacing r_i by $\widetilde{r}_i$ should therefore be read as a floor-regularized extension of the likelihood-inspired surrogate, not as an exact likelihood transformation. This distinction matters in low-power regimes, where negative Gaussian RSS samples can occur. Ignoring the higher-order term, the first-order approximation treats $\cos^{-2}(\phi_i)$ as approximately Gaussian with mean $u^{-4/(m+3)} b_i$ and variance τ_i^2, where

$$\begin{aligned} \tau_i^2 &:= \eta_i(u, \widetilde{r}_i)^2 \left(\sigma^2 + \zeta^2 P \gamma_i\right), \\ \eta_i(u, \widetilde{r}_i) &:= \left(\frac{u^2}{P\Gamma}\right)^{-\frac{2}{m+3}} \frac{2}{m+3} \widetilde{r}_i^{-\frac{m+5}{m+3}}. \end{aligned} \tag{7}$$

Therefore, a likelihood-inspired surrogate is

$$\mathrm{LLR}(\boldsymbol{\theta}) := \sum_{i=1}^{n} \frac{\left(\cos^{-2}(\phi_i) - u^{-\frac{4}{m+3}} b_i\right)^2}{\tau_i^2}. \tag{8}$$

Using Equation (1), this can be rewritten as

$$\min_{x,y,u>0} \quad \mathrm{LLR}(\boldsymbol{\theta}) = \sum_{i=1}^{n} \frac{\left[u^2 + (x_i - x)^2 + (y_i - y)^2 - b_i u^{\rho}\right]^2}{\tau_i^2 u^4}, \tag{9}$$

where $\rho := (2m+2)/(m+3) = 2(m+1)/(m+3)$. The surrogate in Equation (9) remains nonconvex in (x, y, u), and the dependence of τ_i^2 on the unknown location makes direct optimization inconvenient. All low-power results below therefore evaluate the proposed algorithms under this floor-regularized surrogate; the floor is part of the estimator definition, not an unreported post-processing step.

4. Proposed 3D Solver

The solver follows a three-step progression. First, we lift the 3D RSS geometry into a weighted LS problem that can be solved in closed form. Second, because the lifted variables are redundant, we add a linearized consistency regularizer that keeps the estimate close to a physically meaningful 3D point. Third, we stabilize the outer reweighting loop with room projection, clipped weights, and damped acceptance so that sparse layouts do not produce catastrophic updates.

The lifting should not be read as a global convexification of the original 3D VLP problem. The surrogate in Equation (9) remains nonconvex because range geometry, unknown height, and shot-noise-dependent weights are coupled. The proposed formulation bypasses this difficulty only locally. At a fixed outer iteration, the current pilot point freezes the variance weights and linearizes the nonlinear consistency relation among the lifted variables; only after these two choices does the update become a four-variable quadratic LS subproblem with a closed-form solution. The outer loop then decides whether the projected, damped candidate should be accepted under the original surrogate. Thus, the mathematical guarantee is deliberately limited to exact solution of each frozen quadratic subproblem and accepted-surrogate non-increase. We do not claim global optimality of the original nonconvex objective.

4.1. Pilot Weighted Least-Squares Formulation

The exact surrogate in Equation (9) suggests the location-dependent weights $\omega_i(u) := 1/(\tau_i^2 u^4)$, which cannot be used directly because u is unknown. We therefore start from a pilot LS problem that keeps only the observable cross-LED scaling in $\widetilde{r}_i$, and later refine the weights using the current 3D estimate. When thermal noise dominates, or when the γ_i's are comparable across LEDs, the dominant observable RSS scaling in Equation (7) is $\tau_i^2 \propto \widetilde{r}_i^{-2(m+5)/(m+3)}$. The corresponding pilot weights therefore use the inverse scaling $w_i := \widetilde{r}_i^{2(m+5)/(m+3)}$.

Then we form the pilot weighted least-squares objective

$$\min_{x,y,u>0} \quad L_{\mathrm{WLS}}(x,y,u) := \sum_{i=1}^{n} w_i \left[u^2 + (x_i - x)^2 + (y_i - y)^2 - b_i u^{\rho}\right]^2. \tag{10}$$

This pilot objective is a heuristic initialization surrogate rather than an exact reformulation of Equation (9).

Define the lifted variable

$$\boldsymbol{\vartheta} := \left[s, x, y, t\right]^{\top} := \left[x^2 + y^2 + u^2,\ x,\ y,\ u^{\rho}\right]^{\top}. \tag{11}$$

Also define

$$\boldsymbol{\phi}_i := \left[1, -2x_i, -2y_i, -b_i\right]^{\top}, \qquad a_i := -x_i^2 - y_i^2. \tag{12}$$

Then, Equation (10) becomes

$$L_{\mathrm{WLS}}(\boldsymbol{\vartheta}) = \sum_{i=1}^{n} w_i \left(\boldsymbol{\vartheta}^{\top} \boldsymbol{\phi}_i - a_i\right)^2. \tag{13}$$

Let $\mathbf{M} \in \mathbb{R}^{n \times 4}$ and $\mathbf{a} \in \mathbb{R}^n$ be defined row-wise by $\mathbf{M}_{i,:} := \sqrt{w_i}\, \boldsymbol{\phi}_i^{\top}$ and $\mathbf{a}_i := \sqrt{w_i}\, a_i$. Then, Equation (13) is exactly

$$L_{\mathrm{WLS}}(\boldsymbol{\vartheta}) = \|\mathbf{M}\boldsymbol{\vartheta} - \mathbf{a}\|_2^2. \tag{14}$$

Proposition 1. *If* $\mathbf{M}$ *has full column rank, then the weighted least-squares problem in Equation (14) has the unique minimizer*

$$\boldsymbol{\vartheta}_{\mathrm{LS}} = (\mathbf{M}^{\top}\mathbf{M})^{-1}\mathbf{M}^{\top}\mathbf{a}. \tag{15}$$

Proof. The objective in Equation (14) is a strictly convex quadratic function with Hessian $2\mathbf{M}^\top\mathbf{M}$. The full-column-rank assumption implies that $\mathbf{M}^\top\mathbf{M}$ is positive definite, so the minimizer is unique. Setting the gradient

$$\nabla_{\boldsymbol{\vartheta}}\|\mathbf{M}\boldsymbol{\vartheta}-\mathbf{a}\|_2^2 = 2\mathbf{M}^\top(\mathbf{M}\boldsymbol{\vartheta}-\mathbf{a})$$

to zero gives the normal equation

$$\mathbf{M}^\top\mathbf{M}\,\boldsymbol{\vartheta} = \mathbf{M}^\top\mathbf{a},$$

whose solution is Equation (15). □

The full-column-rank condition can fail or become numerically fragile in practical layouts. It requires at least four informative LED equations and sufficient geometric diversity in the rows $\boldsymbol{\phi}_i^\top = [1, -2x_i, -2y_i, -b_i]$. Degenerate cases include too few visible LEDs, nearly collinear or highly symmetric LED placements combined with similar RSS values, saturated or invalid RSS readings, and positions for which the lifted equations become nearly dependent. Even when exact rank is four, the normal matrix can be badly conditioned, so algebraic identifiability is not the same as numerical reliability. The proposition states the mathematical full-rank solution; the implementation uses a double-precision least-squares routine for the raw WLS step rather than explicitly forming $(\mathbf{M}^\top\mathbf{M})^{-1}$. Regularized updates solve the corresponding 4×4 linear systems, recovered points are projected to the room in the stabilized variant, and any linear-solver failure or non-finite prediction is recorded separately as a solver failure; continuous-error summaries use finite estimates only. This handling is important for interpreting the sparse-layout results, where the design matrices can be full rank but still ill-conditioned.

The lifted parameterization is redundant, because the four components of $\boldsymbol{\vartheta}$ are induced by only three physical variables. If $\boldsymbol{\vartheta}$ comes from some (x, y, u) through Equation (11), then

$$\vartheta_1 = \vartheta_2^2 + \vartheta_3^2 + \vartheta_4^q, \qquad q := \frac{m+3}{m+1} = \frac{2}{\rho}. \tag{16}$$

Moreover, the 3D location can be recovered from a consistent $\boldsymbol{\vartheta}$ via

$$x = \vartheta_2, \qquad y = \vartheta_3, \qquad z = h - \vartheta_4^{1/\rho}. \tag{17}$$

The raw solver in Equation (15) does not enforce Equation (16), which is the main source of geometric mismatch.

For a general lifted vector $\boldsymbol{\vartheta}$, we use the safeguarded recovery map

$$\mathcal{R}_\varepsilon(\boldsymbol{\vartheta}) := \left[\vartheta_2,\ \vartheta_3,\ h - (\max\{\vartheta_4, \varepsilon\})^{1/\rho}\right]^\top, \tag{18}$$

where $\varepsilon > 0$ is a fixed numerical safeguard; the experiments use $\varepsilon = 10^{-4}$. Equation (17) is the special case of Equation (18) on the consistency manifold in Equation (16).

4.2. Geometry-Aware Regularized Least Squares

The exact constraint in Equation (16) is nonlinear because of the term ϑ_4^q. Instead of directly regularizing the nonlinear manifold, we linearize it at a lifted pilot point $p = [s_p, x_p, y_p, t_p]^\top$ with $t_p > 0$. Here, t_p is the pilot value of the lifted height variable $t = u^\rho$, not time. For $g(s, x, y, t) := x^2 + y^2 + t^q - s$, the first-order approximation at p can be written as

$$\ell_p(\boldsymbol{\vartheta}) = \boldsymbol{c}_p^\top\boldsymbol{\vartheta} - d_p, \qquad \boldsymbol{c}_p = [-1, 2x_p, 2y_p, qt_p^{q-1}]^\top, \quad d_p = x_p^2 + y_p^2 + (q-1)t_p^q. \tag{19}$$

We then define the geometry-aware regularized least-squares subproblem

$$\min_{\boldsymbol{\vartheta}\in\mathbb{R}^4} \quad L_R(\boldsymbol{\vartheta};\lambda,p) := \|\mathbf{M}\boldsymbol{\vartheta}-\mathbf{a}\|_2^2 + \lambda\left(\boldsymbol{c}_p^\top\boldsymbol{\vartheta}-d_p\right)^2, \tag{20}$$

where $\lambda > 0$ is the regularization parameter.

Theorem 1. *Assume that* $\mathbf{M}$ *has full column rank and let* $\lambda > 0$*. Then the regularized problem in Equation (20) has a unique minimizer given by*

$$\boldsymbol{\vartheta}_\lambda(p) = \left(\mathbf{M}^\top\mathbf{M}+\lambda\boldsymbol{c}_p\boldsymbol{c}_p^\top\right)^{-1}\left(\mathbf{M}^\top\mathbf{a}+\lambda d_p\boldsymbol{c}_p\right). \tag{21}$$

Proof. Let $\mathbf{Q}_p := \mathbf{M}^\top\mathbf{M}+\lambda\boldsymbol{c}_p\boldsymbol{c}_p^\top$ and $\mathbf{r}_p := \mathbf{M}^\top\mathbf{a}+\lambda d_p\boldsymbol{c}_p$. Up to constants independent of $\boldsymbol{\vartheta}$, the objective is $\boldsymbol{\vartheta}^\top\mathbf{Q}_p\boldsymbol{\vartheta}-2\mathbf{r}_p^\top\boldsymbol{\vartheta}$. Since $\mathbf{M}$ has full column rank, $\mathbf{M}^\top\mathbf{M}$ is positive definite; adding $\lambda\boldsymbol{c}_p\boldsymbol{c}_p^\top$ preserves positive definiteness. Hence, the objective is strictly convex, and the first-order condition gives

$$\mathbf{Q}_p\boldsymbol{\vartheta} = \mathbf{r}_p$$

which immediately implies Equation (21). □

The parameter λ therefore controls a data-fit–consistency tradeoff: small values keep the solution close to the raw LS fit, whereas larger values penalize violation of the local consistency relation more strongly and may introduce bias. The experimental sensitivity analysis in Section 5 uses this interpretation when discussing the empirical bias–stability tradeoff.

Theorem 1 gives an exact closed-form solution for every fixed pilot point p under the stated full-column-rank condition. In the first regularized solve we use $p = [(\boldsymbol{\vartheta}_{\mathrm{LS}})_1,(\boldsymbol{\vartheta}_{\mathrm{LS}})_2,(\boldsymbol{\vartheta}_{\mathrm{LS}})_3,\max\{(\boldsymbol{\vartheta}_{\mathrm{LS}})_4,\varepsilon\}]^\top$, which guarantees that the coefficients t_p^{q-1} and t_p^q are well-defined.

4.3. Reweighted and Stabilized Iterative Solver

The pilot weights above give a cheap first estimate, but the shot-noise variance still depends on the current 3D position. The final solver therefore has three steps per outer iteration: update the frozen variance weights from a room-feasible pilot, solve the same 4×4 regularized LS subproblem, and accept only a projected damped update that reduces the original surrogate. A direct reweighted version without the last step is kept as an ablation baseline.

Given the current lifted iterate $\boldsymbol{\vartheta}^{(k)}$, let $p^{(k)}$ denote the pilot used for both variance freezing and consistency linearization, with components $s_p^{(k)}, x_p^{(k)}, y_p^{(k)}, t_p^{(k)}$. The direct reweighted baseline sets $p^{(k)}$ by keeping the first three components of $\boldsymbol{\vartheta}^{(k)}$ and flooring its fourth component at ε, whereas the stabilized solver below uses a room-feasible pilot. From $p^{(k)}$, define $u_p^{(k)} := (t_p^{(k)})^{1/\rho}$. The channel gain estimate used for freezing the variance is

$$d_{i,p}^{(k)} := \sqrt{(u_p^{(k)})^2+(x_i-x_p^{(k)})^2+(y_i-y_p^{(k)})^2}, \qquad \gamma_{i,p}^{(k)} := \frac{\Gamma}{(u_p^{(k)})^2}\left(\frac{u_p^{(k)}}{d_{i,p}^{(k)}}\right)^{m+3}. \tag{22}$$

The frozen variance uses the same first-order noise scale as Equation (7), evaluated at the pilot:

$$\tau_{i,p}^{2,(k)} := (\eta_{i,p}^{(k)})^2(\sigma^2+\zeta^2P\gamma_{i,p}^{(k)}), \qquad \eta_{i,p}^{(k)} := \left((u_p^{(k)})^2/(P\Gamma)\right)^{-\frac{2}{m+3}}\frac{2}{m+3}\tilde{r}_i^{-\frac{m+5}{m+3}}. \tag{23}$$

Using the full frozen weights suggested by Equation (9), $\omega_i^{(k)} := 1/(\tau_{i,p}^{2,(k)}(u_p^{(k)})^4)$. We set $\mathbf{M}_{i,:}^{(k)} := \sqrt{\omega_i^{(k)}}\,\boldsymbol{\phi}_i^\top$ and $\mathbf{a}_i^{(k)} := \sqrt{\omega_i^{(k)}}\,a_i$, linearize the consistency relation at $p = p^{(k)}$, and solve

$$\boldsymbol{\vartheta}^{(k+1)} = \arg\min_{\boldsymbol{\vartheta}\in\mathbb{R}^4}\left\{\left\|\mathbf{M}^{(k)}\boldsymbol{\vartheta} - \mathbf{a}^{(k)}\right\|_2^2 + \lambda\left(\boldsymbol{c}_{p^{(k)}}^\top\boldsymbol{\vartheta} - d_{p^{(k)}}\right)^2\right\}. \tag{24}$$

When the frozen weighted design has the required full column rank, Theorem 1 gives the closed-form update

$$\boldsymbol{\vartheta}^{(k+1)} = \left(\mathbf{H}^{(k)} + \lambda\boldsymbol{c}_k\boldsymbol{c}_k^\top\right)^{-1}\left(\mathbf{f}^{(k)} + \lambda d_k\boldsymbol{c}_k\right), \tag{25}$$

where $\mathbf{H}^{(k)} := (\mathbf{M}^{(k)})^\top\mathbf{M}^{(k)}$, $\mathbf{f}^{(k)} := (\mathbf{M}^{(k)})^\top\mathbf{a}^{(k)}$, $\boldsymbol{c}_k := \boldsymbol{c}_{p^{(k)}}$, and $d_k := d_{p^{(k)}}$.

Each full-rank frozen subproblem is solved exactly by Theorem 1. The outer loop is the fragile part: in sparse 3D layouts, relinearization away from the consistency manifold, highly uneven weights, and unconstrained updates can move the next iterate outside the physical room. The stabilization layer below addresses those failure modes while leaving the closed-form subproblem unchanged.

Let the room dimensions be $L_x \times L_y \times h$, and let $[z_{\min}, z_{\max}]$ be the admissible receiver-height range. Define $\mathcal{P}_{\text{room}}$ as component-wise clipping of $[x, y, z]^\top$ to $[0, L_x] \times [0, L_y] \times [z_{\min}, z_{\max}]$. We also define the consistent reconstruction map $\mathcal{C}([x, y, z]^\top) := [x^2 + y^2 + (h-z)^2, x, y, (h-z)^\rho]^\top$. Combining Equation (18) with these operators yields the consistent room-feasible projection

$$\mathcal{T}(\boldsymbol{\vartheta}) := \mathcal{C}\left(\mathcal{P}_{\text{room}}(\mathcal{R}_\varepsilon(\boldsymbol{\vartheta}))\right). \tag{26}$$

This map sends any lifted estimate to a room-feasible point satisfying the nonlinear consistency relation in Equation (16).

At iteration k, the stabilized solver sets $p^{(k)} = \mathcal{T}(\boldsymbol{\vartheta}^{(k)})$ and computes the frozen weights $\omega_i^{(k)}$. In sparse or low-power layouts, a single very large or very small frozen weight can make the quadratic subproblem dominated by one LED and can amplify the lifted-variable inconsistency. We therefore clip the weights around their median $m_\omega^{(k)} := \operatorname{median}_{i=1,\ldots,n}\omega_i^{(k)}$:

$$\widetilde{\omega}_i^{(k)} := \operatorname{clip}\left(\omega_i^{(k)}, \frac{m_\omega^{(k)}}{c_\omega}, m_\omega^{(k)}c_\omega\right), \qquad c_\omega > 1. \tag{27}$$

This median-centered clipping preserves the relative scale of typical LED measurements while preventing outlying frozen variances from controlling the LS update. It then assembles the quadratic subproblem using $\widetilde{\omega}_i^{(k)}$ in place of $\omega_i^{(k)}$, maps the candidate back through $\mathcal{T}$, and tests damped versions of that candidate against

$$\widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}) := \mathrm{LLR}\left(\left[\vartheta_2, \vartheta_3, (\max\{\vartheta_4, \varepsilon\})^{1/\rho}\right]^\top\right), \tag{28}$$

i.e., the original surrogate in Equation (9) evaluated at the recovered (x, y, u) point. For a finite damping set $\mathcal{A} \subset (0, 1]$, each candidate is $\boldsymbol{\vartheta}_\alpha^{(k+1)} := \mathcal{T}((1-\alpha)\boldsymbol{\vartheta}^{(k)} + \alpha\boldsymbol{\vartheta}_{\text{cand}}^{(k+1)})$. The best improving candidate is accepted; if no α yields a sufficient decrease, the loop stops. Thus, Theorem 1 covers the frozen quadratic solve, while projection, clipping, damping, and acceptance control the outer-loop failure modes.

https://doi.org/10.3390/electronics15143172

Proposition 2 (Monotonic surrogate decrease)**.** *Let $\{\boldsymbol{\vartheta}^{(k)}\}_{k\geq 0}$ be the sequence produced by Algorithm 1. Then, the original surrogate evaluated at the corresponding 3D points is non-increasing:*

$$\widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}^{(k+1)}) \;\leq\; \widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}^{(k)}), \qquad \forall\, k \geq 0. \tag{29}$$

Moreover, the sequence terminates in at most K iterations. If the algorithm terminates early at iteration k, then every tested damped candidate at that iteration satisfies $\widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}_\alpha^{(k+1)}) \geq \widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}^{(k)}) - \epsilon$, i.e., no further ϵ-improving step exists within the damping set for the current frozen candidate.

Proof. At each iteration, Algorithm 1 accepts an update only if some damped candidate $\boldsymbol{\vartheta}_\alpha$ satisfies $\widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}_\alpha) < \widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}^{(k)}) - \epsilon$. If no such candidate exists, the iterate is left unchanged and the loop terminates. Therefore every accepted update strictly decreases $\widetilde{\mathrm{LLR}}$ by at least ϵ, and a rejected iteration leaves it unchanged. The iteration bound K is explicit in the loop. □

Algorithm 1: Stabilized iterative reweighted regularized LS.

Input: RSS data, LED layout, room parameters, λ, $\mathcal{A}$, c_ω, and K.
Output: Room-feasible 3D estimate $\hat{\boldsymbol{x}}$.
Set $\tilde{r}_i \leftarrow \max\{r_i, \epsilon_r\}$, compute $w_i = \tilde{r}_i^{2(m+5)/(m+3)}$, and form $\mathbf{M}, \mathbf{a}$.
Solve raw WLS $\boldsymbol{\vartheta}_{\mathrm{LS}} \leftarrow \arg\min_{\boldsymbol{\vartheta}} \|\mathbf{M}\boldsymbol{\vartheta} - \mathbf{a}\|_2^2$ using a numerical least-squares solve.
Initialize $\boldsymbol{\vartheta}^{(0)} \leftarrow \mathcal{T}(\boldsymbol{\vartheta}_{\mathrm{LS}})$ and $J^{(0)} \leftarrow \widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}^{(0)})$.
for $k = 0, \ldots, K-1$ **do**
  Set $p^{(k)} \leftarrow \mathcal{T}(\boldsymbol{\vartheta}^{(k)})$.
  Compute $\tau_{i,p}^{2,(k)}$ and $\omega_i^{(k)} = 1/[\tau_{i,p}^{2,(k)}(u_p^{(k)})^4]$.
  Set $m_\omega^{(k)} \leftarrow \mathrm{median}_{i=1,\ldots,n}\, \omega_i^{(k)}$ and clip $\omega_i^{(k)}$ around $m_\omega^{(k)}$ by factor c_ω.
  Build $\mathbf{M}^{(k)}, \mathbf{a}^{(k)}$, linearize at $p^{(k)}$, and solve the regularized LS subproblem.
  Project the solution through $\mathcal{T}$ to obtain $\boldsymbol{\vartheta}_{\mathrm{cand}}^{(k+1)}$.
  For each $\alpha \in \mathcal{A}$, compute $\boldsymbol{\vartheta}_\alpha^{(k+1)} \leftarrow \mathcal{T}((1-\alpha)\boldsymbol{\vartheta}^{(k)} + \alpha\boldsymbol{\vartheta}_{\mathrm{cand}}^{(k+1)})$.
  Set $J_\alpha \leftarrow \widetilde{\mathrm{LLR}}(\boldsymbol{\vartheta}_\alpha^{(k+1)})$ and choose $\alpha^\star \in \arg\min_{\alpha\in\mathcal{A}} J_\alpha$.
  if $J_{\alpha^\star} < J^{(k)} - \epsilon$ **then**
    $\boldsymbol{\vartheta}^{(k+1)} \leftarrow \boldsymbol{\vartheta}_{\alpha^\star}^{(k+1)}$ and $J^{(k+1)} \leftarrow J_{\alpha^\star}$.
  else
break
  end if
end for
return $\hat{\boldsymbol{x}} \leftarrow \mathcal{R}_\varepsilon(\boldsymbol{\vartheta}^{(k_{\mathrm{last}})})$.

Proposition 2 is deliberately modest: it guarantees accepted-surrogate non-increase, not global optimality. This is the intended division of labor in the solver: closed-form quadratic subproblems provide speed, and the stabilization layer prevents the outer reweighting loop from worsening the original surrogate.

Baselines for Comparison

For ablation, we also evaluate simpler LS-style variants: (i) raw weighted LS, solving only Equation (15); (ii) one-step regularized LS, applying a single geometry-aware correction via Equation (21); (iii) iterative reweighted regularized LS, updating variances and the linearized regularizer without stabilization; and (iv) a nonlinear SciPy baseline that attempts local minimization of the floor-regularized surrogate in Equation (9) using the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with bound constraints (L-BFGS-B) with multiple restarts. Other estimator families are also reasonable in this setting. Direct maximum-likelihood or nonlinear least-squares solvers can often fit the original surrogate more tightly in practice; particle-swarm or other metaheuristic searches can reduce local-minimum sensitivity at higher computational cost; Bayesian filters are natural for tracking; and fingerprinting or machine-learning methods can absorb hardware mismatch when calibration data are available. The nonlinear SciPy optimizer is included as a strong direct nonlinear baseline in this benchmark because it uses the same forward model and objective, but it remains a finite-budget local optimizer rather than a guaranteed global solver.

4.4. Computational Complexity

Each iteration of the stabilized solver costs $O(n)$ in the number of LEDs: computing the n weights is $O(n)$, assembling the $n \times 4$ design matrix $\mathbf{M}^{(k)}$ is $O(n)$, forming the 4×4 Gram matrix $(\mathbf{M}^{(k)})^\top \mathbf{M}^{(k)}$ is $O(n)$, and inverting the 4×4 system is $O(1)$. The damped line search evaluates the surrogate at $|\mathcal{A}|$ candidate points, each costing $O(n)$. With at most K iterations and $|\mathcal{A}| = 5$, the overall complexity is $O(Kn)$ with a small constant factor, explaining the measured speed advantage up to the tested 16-LED setting.

4.5. Cramér–Rao Lower Bound

To benchmark any estimator, we derive the Fisher information matrix (FIM) for the 3D position under the combined thermal-plus-shot-noise model [20,42,44]. Let $\sigma_i^2(\boldsymbol{\theta}) := \sigma^2 + \zeta^2 P\gamma_i(\boldsymbol{\theta})$ denote the position-dependent noise variance. The observations $\{r_i\}$ are conditionally independent with $r_i \mid \boldsymbol{\theta} \sim \mathcal{N}(P\gamma_i(\boldsymbol{\theta}), \sigma_i^2(\boldsymbol{\theta}))$, so the FIM takes the heteroscedastic Gaussian form

$$[\mathbf{F}(\boldsymbol{\theta})]_{jk} = \sum_{i=1}^{n} \frac{1}{\sigma_i^2(\boldsymbol{\theta})} \frac{\partial(P\gamma_i)}{\partial\theta_j} \frac{\partial(P\gamma_i)}{\partial\theta_k} + \frac{1}{2\sigma_i^4(\boldsymbol{\theta})} \frac{\partial\sigma_i^2}{\partial\theta_j} \frac{\partial\sigma_i^2}{\partial\theta_k}, \tag{30}$$

where $\boldsymbol{\theta} = [x, y, u]^\top$. The second term accounts for the information contained in the variance itself, which is non-negligible under shot noise.

Let $d_i := \sqrt{u^2 + (x_i - x)^2 + (y_i - y)^2}$. From the Lambertian model,

$$P\gamma_i = P\Gamma \frac{u^{m+1}}{d_i^{m+3}}. \tag{31}$$

The required partial derivatives are

$$\frac{\partial(P\gamma_i)}{\partial x} = (m+3)\, P\gamma_i \frac{x_i - x}{d_i^2}, \tag{32}$$

$$\frac{\partial(P\gamma_i)}{\partial y} = (m+3)\, P\gamma_i \frac{y_i - y}{d_i^2}, \tag{33}$$

$$\frac{\partial(P\gamma_i)}{\partial u} = P\gamma_i \left(\frac{m+1}{u} - \frac{(m+3)u}{d_i^2} \right). \tag{34}$$

The variance derivatives follow from $\partial\sigma_i^2 / \partial\theta_j = \zeta^2 P\, \partial\gamma_i / \partial\theta_j$ for $\theta_j \in \{x, y, u\}$.

The Cramér–Rao lower bound (CRLB) for the 3D position error is then

$$\mathrm{CRLB}(\boldsymbol{\theta}) = \mathrm{tr}\big(\mathbf{F}(\boldsymbol{\theta})^{-1}\big), \tag{35}$$

and $\sqrt{\mathrm{CRLB}(\boldsymbol{\theta})}$ lower-bounds the root-mean-square 3D error of any unbiased estimator at position $\boldsymbol{\theta}$ [43]. Since the proposed regularized and projected estimators are generally biased, the CRLB is used as a diagnostic reference rather than as a strict lower bound on their observed mean-squared error. In practice, the second term in Equation (30) (the variance-information contribution) is small compared to the first when $\zeta^2 P\gamma_i \ll \sigma^2$, i.e., when thermal noise dominates; under strong shot noise, the two terms are comparable, and omitting the second term would understate the available information [44]. We evaluate CRLB numerically at each test position in Section 5 and compare it with the achieved errors of the stabilized solver and the nonlinear baseline.

5. Experiments

5.1. Experimental Setup

We evaluate the proposed solvers in a synthetic 3D simulation that follows the model assumptions and emphasizes sparse layouts, where 3D ambiguity is strongest. Following established evaluation methodology in VLP research [10,20], we compare against both LS-style variants and a nonlinear optimization baseline. The simulations assume LOS Lambertian RSS propagation with all links inside the PD field of view, known LED coordinates and ceiling height, fixed downward LED axes, a fixed upward PD normal, normalized optical gain, and no wall reflections, diffuse multipath, shadowing, or NLOS component. This is a controlled LOS study; it does not include orientation mismatch, multipath, installation error, or calibration uncertainty. Tilt and multipath are not represented as arbitrary RSS perturbations because they change the observation model and would make the surrogate, LS weights, and CRLB diagnostic internally inconsistent. The numerical settings use the RSS floor $\widetilde{r}_i = \max\{r_i, 10^{-6}\}$, room-box projection in the stabilized solver, six outer iterations, damping set $\{1, 0.5, 0.25, 0.1, 0.05\}$, median-centered weight clipping with $c_\omega = 50$, three deterministic L-BFGS-B starts for the nonlinear baseline, and separate recording of non-finite predictions as solver failures; continuous-error statistics use finite estimates only.

The room size is $8 \times 8 \times 3$ m, with all LEDs mounted at height $h = 3$ m. Receiver positions are sampled uniformly from

$$x, y \in [0.25, 7.75], \qquad z \in [0.4, 2.4].$$

We use Lambertian order $m = 2$ and set $\Gamma = 1$ for normalization. The layout sweep uses three ceiling grids: a sparse 4-LED layout at $(2, 2)$, $(2, 6)$, $(6, 2)$, $(6, 6)$; a denser 9-LED grid $\{2, 4, 6\} \times \{2, 4, 6\}$; and a 16-LED grid $\{1, 3, 5, 7\} \times \{1, 3, 5, 7\}$.

The noise and gain parameters in the benchmark are simulation parameters, not measurements from a hardware campaign. In a physical system, these quantities would be estimated before positioning is run. The thermal/electronic noise level σ would normally come from LED-off or dark-current samples, because those measurements isolate receiver noise from the optical signal. The shot-noise coefficient ζ would be obtained from repeated RSS measurements at known illumination levels by fitting the variance-versus-mean relation. The gain term Γ, together with any per-LED gain factors, would be calibrated using receiver positions with known coordinates and known orientation, then checked on held-out calibration positions. This matters for the present method because gain and noise errors change both the RSS scale and the frozen weights in the LS subproblems. The present study assumes these quantities are known after calibration and focuses on estimator behavior under that model; calibration uncertainty is not modeled in the reported synthetic results.

For each layout we consider a high-power regime ($P = 60$, $\sigma = 0.015$, $\zeta = 0.05$) and a low-power regime ($P = 25$, $\sigma = 0.03$, $\zeta = 0.08$), giving six scenarios in total. We use deterministic seed `20260327`, generate 80 tuning samples and 400 test samples per scenario, and tune λ over $\{10^{-2}, 10^{-1}, 1, 10, 100\}$ on the tuning split by minimizing mean error (median as tie-breaker). The selected regularization parameters $(\lambda_{\text{reg}}, \lambda_{\text{iter}}, \lambda_{\text{stab}})$ are $(100, 100, 100)$ for 4-LED high power, $(1, 0.01, 100)$ for 4-LED low power, $(100, 100, 100)$ for 9-LED high power, $(0.1, 100, 100)$ for 9-LED low power, $(100, 100, 100)$ for 16-LED high power, and $(1, 100, 100)$ for 16-LED low power. Figure 2 makes the geometry change explicit.

Tables 1 and 2 and Figure 3 use this same fully tuned six-scenario protocol. The convergence and regularization-sensitivity diagnostics include representative sparse, moderate, and dense layouts. The CRLB diagnostic focuses on the 9- and 16-LED layouts, where empirical RMSE can be compared with the information bound without sparse-layout bias effects dominating the figure.

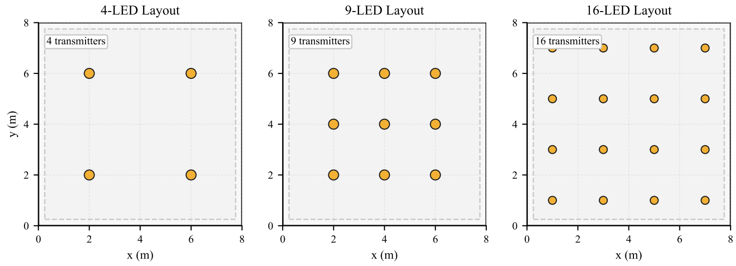

Figure 2. Top-view LED layouts used in the six-scenario benchmark. The dashed inner square marks the common receiver sampling region $[0.25, 7.75] \times [0.25, 7.75]$ m^2; all LEDs are mounted at ceiling height $h = 3$ m.

Table 1. Results in sparse 4-LED scenarios. Bold = best LS-style.

Method	Mean (m)	P95 (m)	Max (m)	Time (s)
4-LED high power				
Raw WLS	3.402	4.783	5.754	3×10^{-5}
Reg. LS	28,217	3.663	1.1×10^{7}	4×10^{-5}
Orig. iter.	66,833	3.561	2.7×10^{7}	1.6×10^{-4}
Stab. (ours)	**1.247**	**3.144**	**4.223**	1.7×10^{-3}
SciPy NL	0.891	2.895	5.564	2.8×10^{-2}
4-LED low power				
Raw WLS	3.422	4.872	5.508	3×10^{-5}
Reg. LS	4.0×10^{5}	4.555	1.6×10^{8}	4×10^{-5}
Orig. iter.	620.4	8.376	2.4×10^{5}	1.7×10^{-4}
Stab. (ours)	**1.555**	**3.395**	**4.404**	1.5×10^{-3}
SciPy NL	1.283	3.372	6.097	2.6×10^{-2}

Table 2. Results in denser 9- and 16-LED scenarios. Bold = best LS-style.

Method	9-LED				16-LED			
	Mean	P95	Max	Time	Mean	P95	Max	Time
High power								
Raw WLS	0.429	**1.919**	6.921	3×10^{-5}	0.140	0.256	7.212	3×10^{-5}
Reg. LS	0.343	2.535	8.749	4×10^{-5}	0.121	**0.173**	10.706	4×10^{-5}
Orig. iter.	**0.304**	2.206	6.919	2.1×10^{-4}	0.128	0.183	7.211	2.1×10^{-4}
Stab. (ours)	0.307	2.536	**3.981**	1.0×10^{-3}	**0.108**	0.192	**3.910**	7.1×10^{-4}
SciPy NL	0.225	1.741	3.810	1.7×10^{-2}	0.122	0.833	2.618	1.5×10^{-2}
Low power								
Raw WLS	1.259	4.985	9.477	3×10^{-5}	0.461	2.062	8.599	3×10^{-5}
Reg. LS	1539.7	3.831	6.2×10^{5}	4×10^{-5}	**0.321**	1.628	7.077	4×10^{-5}
Orig. iter.	89.66	4.450	3.5×10^{4}	1.7×10^{-4}	0.422	**1.447**	8.598	1.9×10^{-4}
Stab. (ours)	**0.985**	**3.411**	**4.676**	9.1×10^{-4}	0.335	1.600	**4.903**	6.4×10^{-4}
SciPy NL	0.631	2.910	5.509	1.7×10^{-2}	0.425	1.693	2.559	1.5×10^{-2}

Units are meters for Mean/P95/Max and seconds per sample for Time. In Tables 1 and 2, Mean, P95, and maximum error are computed over finite estimates; non-finite outputs are recorded separately as solver failures.

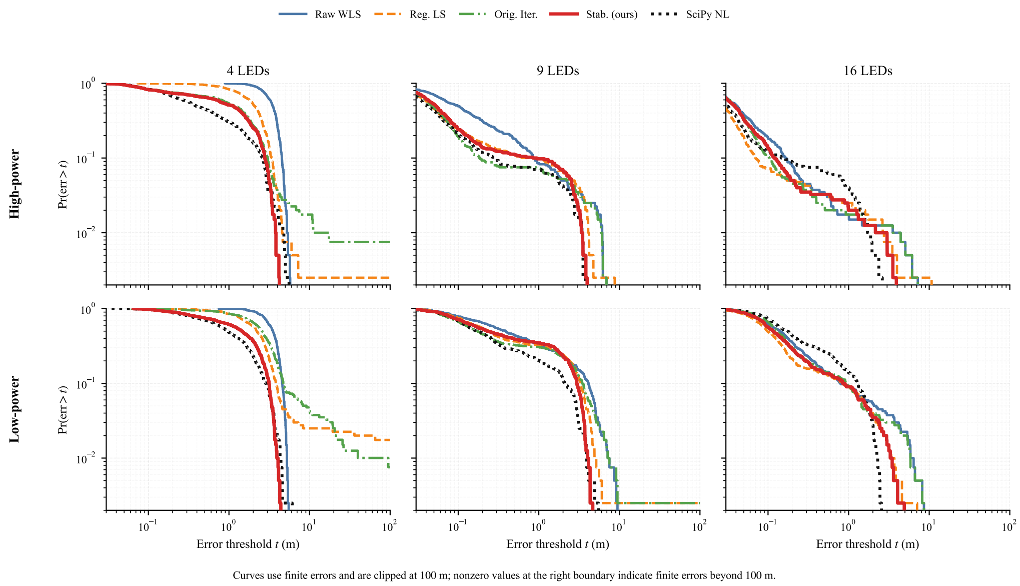

Figure 3. Error exceedance curves for finite estimates on the full benchmark. The horizontal axis is clipped at 100 m; nonzero values at the right boundary indicate finite errors exceeding that range. Non-finite outputs are recorded separately as solver failures.

- Reproducibility.

All reported results were obtained using the same deterministic simulation procedure. Within each scenario, the room geometry, LED coordinates, receiver sampling region, random seed, tuning/test split, λ grid, damping set, clipping factor, and nonlinear-baseline restart budget were fixed. The six scenarios use the stated LED layouts and power regimes. The regularization parameter was selected on the tuning split and then held fixed on the test split. This ensures that differences in mean error across layouts reflect estimator behavior under the stated layouts rather than changes in tuning. The simulations are reproducible within the synthetic LOS Lambertian model, but they do not replace hardware validation.

We compare five solvers: raw weighted LS, one-step regularized LS, original iterative reweighted regularized LS, the proposed stabilized iterative reweighted regularized LS, and a nonlinear SciPy baseline that attempts local minimization of the floor-regularized surrogate in Equation (9) using L-BFGS-B with three restarts. The first four methods separate the effect of the LS lifting, the geometry-aware regularizer, iterative reweighting, and the stabilization layer. The SciPy baseline is different in purpose: it is a direct nonlinear reference using the same forward model and objective, so it tests how much accuracy is lost or gained by replacing nonlinear search with closed-form LS-family updates.

Raw weighted LS and the SciPy baseline have no tuned regularization parameter; only the three regularized LS-style methods use the λ grid. For the stabilized solver, we use six outer iterations, damping set $\mathcal{A} = \{1, 0.5, 0.25, 0.1, 0.05\}$, clipping factor $c_\omega = 50$, and tolerance 10^{-9}. The SciPy baseline uses L-BFGS-B with room-box constraints, three deterministic starts (raw-WLS, room center, strongest-LED horizontal location), at most 90 iterations per restart, and default SciPy convergence tolerances. Because this baseline can move directly in the nonlinear 3D variables, it is a strong accuracy comparator in sparse layouts. Its finite restart budget does not guarantee a global optimum. Runtime was measured on the same workstation using the Python 3.9.6 implementation for all methods. The tables report mean, P95, maximum error, and runtime; figures add median and exceedance views. Because the primary six-scenario sweep uses one seed and the nonlinear baseline uses a finite restart budget, small dense-layout gaps are interpreted conservatively.

For numerical linear algebra, the raw WLS step uses double-precision least squares rather than an explicit inverse; Proposition 1 states the mathematical full-rank solution.

Regularized systems are solved as 4×4 linear systems, recovered positions are projected to the room box in the stabilized variant, and linear-solver errors or non-finite predictions are recorded separately as solver failures. Mean, P95, and maximum errors are computed over finite estimates. No stabilized-solver failures occur in the reported benchmark.

- Rank and conditioning tests.

We also ran a robustness analysis using five seeds, 200 test samples per seed, and fixed tuned λ values. All design matrices retained rank four, so the main failure mode was not exact algebraic rank loss. The difficulty was numerical conditioning: in the 4-LED high/low cases, the median condition numbers of $\mathbf{M}^\top\mathbf{M}$ were 5.0×10^9 and 1.2×10^{10}, 95th percentiles were 7.3×10^{16} and 2.1×10^{17}, and worst cases reached 10^{23}–10^{33}. After clipping and regularization in the first stabilized subproblem, the corresponding average median/p95 condition numbers were reduced to $4.9 \times 10^4/1.1 \times 10^6$ and $3.8 \times 10^4/3.4 \times 10^6$, although rare maxima still reached 10^{16}–10^{17}. This explains why Proposition 1 is not enough by itself: a matrix can have rank four and still be too ill-conditioned for an unstabilized lifted LS estimate to be reliable.

- RSS-floor and repeated-seed tests.

Nonpositive RSS readings occurred in 3.6–15.2% of individual LED readings, affecting 14.8–83.9% of samples, so the low-power results should be interpreted as performance under the stated floor-regularized heuristic rather than as exact likelihood behavior in a high-SNR regime. A floor-sensitivity analysis over $\epsilon_r \in \{10^{-8}, 10^{-6}, 10^{-4}\}$ changed low-power stabilized mean errors only from 1.51–1.56, 0.903–0.927, and 0.388–0.399 m for 4-, 9-, and 16-LED layouts. Across the five seeds, stabilized mean errors, ordered as 4-/9-/16-LED high/low power, were 1.28 [1.16, 1.41], 1.56 [1.43, 1.70], 0.385 [0.318, 0.452], 0.924 [0.826, 1.02], 0.078 [0.050, 0.106], and 0.398 [0.338, 0.458] m; brackets give approximate 95% confidence intervals over seed means. With 5 cm LED-position perturbations during data generation, the corresponding means were 1.30, 1.56, 0.425, 0.970, 0.124, and 0.410 m. These checks support the LS-family robustness and geometry-conditioning claims, while preserving the limitation that the study remains synthetic and does not cover orientation or multipath mismatch.

- Ablations.

Table 3 and Figure 2 define the geometry test under different noise levels. Tables 1 and 2 report the main accuracy and runtime results, and Figure 3 shows tail failures that are hidden by central summaries. Figures 4 and 5 provide convergence and information-bound diagnostics.

Table 3. System parameters used in experiments.

Parameter	Symbol	Value
Room size	$L_x \times L_y \times h$	$8 \times 8 \times 3$ m
4-LED positions	(x_i, y_i)	$(2,2), (2,6), (6,2), (6,6)$
9-LED positions	(x_i, y_i)	$\{2,4,6\} \times \{2,4,6\}$
16-LED positions	(x_i, y_i)	$\{1,3,5,7\} \times \{1,3,5,7\}$
Lambertian order	m	2
Transmit power (hi/lo)	P	60/25 W
Thermal noise (hi/lo)	σ	0.015/0.03
Shot noise (hi/lo)	ζ	0.05/0.08
Optical gain	Γ	1 (norm.)
Receiver height range	z	$[0.4, 2.4]$ m

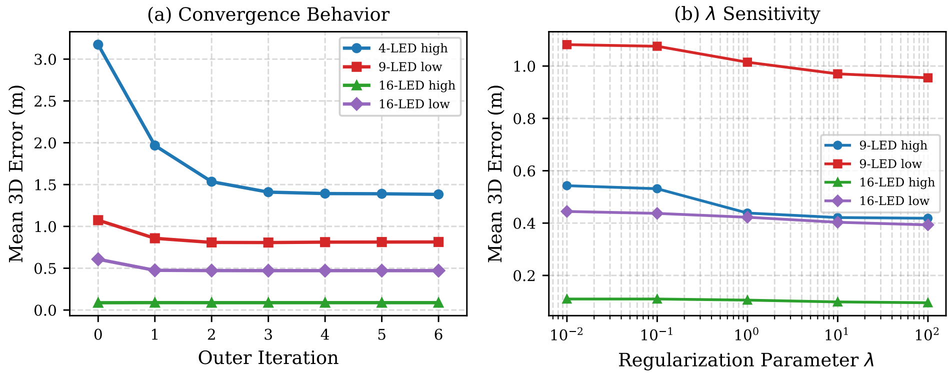

Figure 4. Convergence and regularization diagnostics. Panel (**a**) shows mean 3D error over outer iterations; panel (**b**) shows tuning-split mean error versus λ.

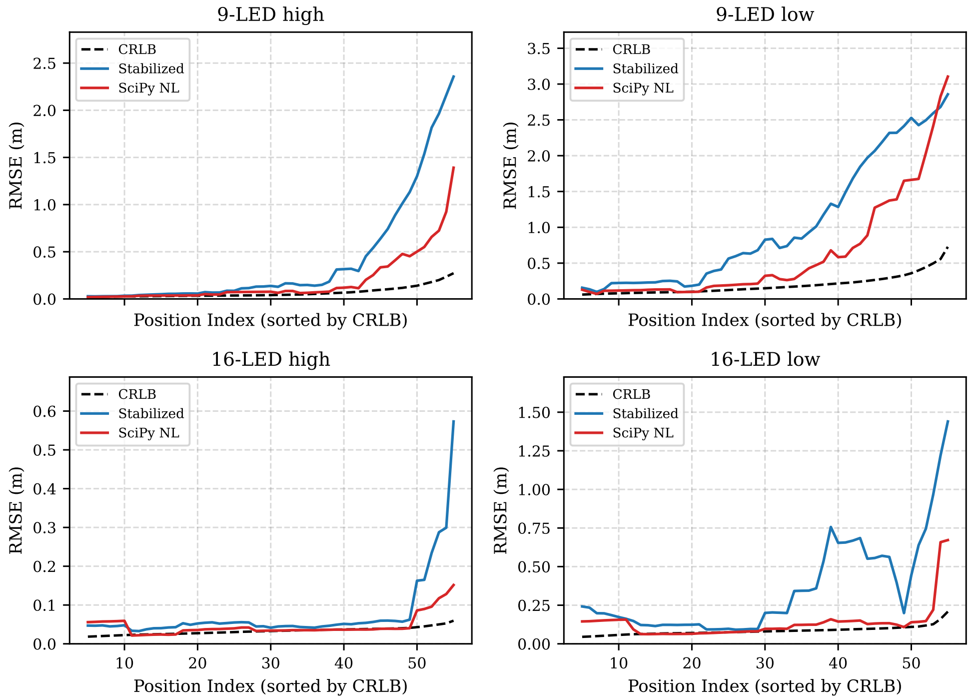

Figure 5. Per-position root-mean-square error (RMSE) and numerical CRLB, sorted by CRLB, for 9- and 16-LED scenarios. Dashed black curves denote analytical CRLB values; solid colored curves denote simulated estimator RMSE.

5.2. Main Results

Tables 1 and 2 summarize the final benchmark. The results separate the solver story into two regimes. In the sparse 4-LED cases, stabilization is the difference between usable LS-style estimates and occasional catastrophic failures. In the denser 9- and 16-LED cases, the LS-style family becomes much more competitive, and the role of stabilization shifts from uniform mean-error dominance to tail control and a predictable runtime–accuracy tradeoff.

The table-level pattern changes with geometry. In the 4-LED scenarios, stabilization is decisive: the proposed method reduces the mean error from 3.40 m for raw WLS to 1.25–1.55 m and avoids the catastrophic blow-ups observed for the unstabilized regularized variants. The meter-level and extremely large mean errors are not a consequence of room size alone. They arise mainly in sparse or low-power cases where height–RSS ambiguity, low-SNR measurements, the positive RSS floor applied after noisy nonpositive readings, and severe conditioning of the lifted design matrix can produce rare but very large outliers. This is visible in the table: for some unstabilized LS variants, the P95 error remains within a few meters, while the maximum error is several orders of magnitude larger, causing the mean to be dominated by a small number of divergent estimates.

In the 9-LED layouts, no LS-style method wins every metric, but the stabilized solver is strongest in the harder low-power case and gives the lowest LS-style worst-case error in both power regimes. With 16 LEDs, all LS-style methods return to sub-meter mean error even at low power; the stabilized solver is best on high-power mean error, while the 16-LED low-power case is a close runtime–accuracy tradeoff across the LS family. The nonlinear SciPy baseline remains strongest in sparse and moderate layouts because it attempts finite-budget local minimization of the nonlinear surrogate with bound constraints and multiple starts instead of relying on a frozen linearized LS subproblem. The proposed method should therefore be interpreted as a fast LS-family estimator with improved tail control, not as a uniformly more accurate replacement for nonlinear optimization.

5.3. Distributional Robustness

Figure 3 explains the main failure mode behind these table-level trends. The one-step and original iterative regularized variants can look reasonable through much of the practical error range, yet retain long right tails in the 4-LED and 9-LED low-power regimes. The layout sweep, exceedance curves, and conditioning analysis should be read together when assessing geometry sensitivity. The 4-LED layouts expose the strongest height–RSS ambiguity and the heaviest conditioning tails, while the denser 9- and 16-LED layouts improve the information profile and reduce the frequency of catastrophic LS-family failures. The stabilization layer targets exactly this behavior: room-feasible projection prevents physically invalid lifted estimates, weight clipping limits single-LED dominance, and surrogate-accepted damping rejects outer steps that worsen the original objective. By 16 LEDs, the curves cluster in the sub-meter regime and the question shifts from preventing breakdown to choosing among speed, mean error, P95 error, and maximum error.

To examine how errors vary with receiver geometry, we analyzed the per-sample simulation results. Receiver heights were split into three equal bins over $z \in [0.4, 2.4]$ m, and horizontal positions were grouped by a 3×3 partition of the sampled floor region into center, edge, and corner groups. In the 4-LED layouts, the stabilized mean error varies much more across horizontal regions than across height bins: at high power, the regional means span 0.08–1.92 m while the height-bin means span 1.17–1.37 m; at low power, the corresponding spans are 0.23–2.30 m and 1.50–1.63 m. The nonlinear baseline shows the same spatial dependence, with 4-LED regional means of 0.67–1.15 m at high power and 0.56–1.43 m at low power. With 16 LEDs, the stabilized regional span contracts to 0.03–0.18 m at high power and 0.09–0.52 m at low power. Thus the remaining errors are not spatially uniform; they are largest in geometrically less favorable regions, especially under sparse lighting. The broader LS-family error distribution is shown by the exceedance curves in Figure 3 and by the tail-rate analysis below.

The same analysis also distinguishes typical errors from rare divergent estimates. For 4-LED high power, regularized LS has median 2.06 m and P95 3.66 m but 0.2% of finite estimates exceed 100 m; the original iterative variant has median 1.18 m and P95 3.56 m, but 0.8% of finite estimates exceed 100 m. For 4-LED low power, the corresponding $>$100 m rates among finite estimates are 1.8% for one-step regularized LS and 0.8% for the original iterative variant. The original iterative variant additionally produced two non-finite outputs (0.5%) which are recorded separately as solver failures. The stabilized solver has no samples above 10 m or 100 m in any of the six main scenarios. This explains why some unstabilized LS means are extremely large even when their medians and P95 values appear moderate.

5.4. Convergence, Regularization, and CRLB

Figure 4a shows that most of the empirical error reduction of the stabilized solver occurs within 3–4 outer iterations. This supports the practical iteration budget used in the benchmark, while Proposition 2 separately guarantees that the accepted surrogate value does not increase. Figure 4b is consistent with the role of λ in the regularized subproblem: small λ emphasizes data fit, large λ enforces the linearized consistency relation more strongly, and the best predictive error depends on layout conditioning and noise. Figure 5 compares per-position RMSE with the CRLB derived in Section 4.5 on the 9- and 16-LED high- and low-power scenarios. Both estimators generally track the CRLB ordering, supporting the interpretation that denser geometry improves the information profile and that the remaining error variation is tied to position-dependent Fisher information.

6. Conclusions

We developed a shot-noise-aware 3D RSS-based VLP solver that keeps the LS subproblems closed form while addressing the instability introduced by unknown receiver height. The method combines a floor-regularized 3D surrogate, a lifted weighted LS reformulation, a geometry-aware regularizer, and a stabilized reweighting loop with room projection, weight clipping, damping, and accepted-surrogate monotonicity. In the synthetic LOS simulations, the stabilized solver removes the catastrophic sparse-layout failures observed for the tested LS-family baselines and retains a large runtime advantage over nonlinear optimization. The nonlinear baseline often attains lower mean error in sparse and moderate layouts. The proposed method is therefore a fast LS-family alternative for the fixed-orientation LOS setting when runtime and control of large errors are important.

- Limitations and future work.

The conclusions apply to a synthetic Lambertian LOS model with known LED geometry and a fixed-orientation photodiode receiver. The study does not evaluate hardware deployments, public datasets, receiver-orientation uncertainty, or multipath. A public-dataset evaluation for this model would require 3D RSS samples together with LED geometry, receiver-height ground truth, and orientation and calibration metadata. Receiver tilt requires calibrated PD orientation, field of view, and gain, whereas reflected paths require room geometry, surface reflectance, and receiver-response assumptions. These effects change the forward model, signal-dependent variance, and CRLB diagnostic; they cannot be represented by a generic additional noise term. The experimental studies of Almadani et al. [37,38] illustrate their relevance to unknown-height 3D VLP. Future deployment studies should first extend and validate the physical model, then assess the corresponding estimator. Installation errors [36], gain calibration, open-dataset evaluation, hardware validation, RSS–AOA fusion [29], IRS-assisted VLP [50], and time-varying mobile scenarios [25] are natural directions for further work.

Author Contributions: Conceptualization, Y.W. and X.L.; methodology, Y.W. and X.L.; software, Y.W.; validation, Y.W.; formal analysis, Y.W. and X.L.; investigation, Y.W.; data curation, Y.W.; writing—original draft preparation, Y.W.; writing—review and editing, Y.W. and X.L.; visualization, Y.W.; supervision, X.L.; funding acquisition, X.L. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Department of Education of Guangdong Province (Grant No.: 2023ZDZX1083).

Data Availability Statement: The simulation code and numerical results used in this study are available from the corresponding author upon reasonable request.

Acknowledgments: The authors have no additional acknowledgments to declare.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AOA	Angle of arrival
CRLB	Cramér–Rao lower bound
FIM	Fisher information matrix
LED	Light-emitting diode
LOS	Line of sight
LS	Least squares
NLOS	Non-line-of-sight
PD	Photodiode
RF	Radio frequency
RIS	Reconfigurable intelligent surface
RSS	Received signal strength
RMSE	Root-mean-square error
SNR	Signal-to-noise ratio
UWB	Ultra-wideband
VLP	Visible light positioning
VLC	Visible light communication
IRS	Intelligent reflecting surface
WLS	Weighted least squares

References

1. Komine, T.; Nakagawa, M. Fundamental Analysis for Visible-Light Communication System Using LED Lights. *IEEE Trans. Consum. Electron.* **2004**, *50*, 100–107. [CrossRef]
2. Gong, C. Visible Light Communication and Positioning: Present and Future. *Electronics* **2019**, *8*, 788. [CrossRef]
3. Keskin, M.F.; Sezer, A.D.; Gezici, S. Localization via Visible Light Systems. *Proc. IEEE* **2018**, *106*, 1063–1088. [CrossRef]
4. Zafari, F.; Gkelias, A.; Leung, K.K. A Survey of Indoor Localization Systems and Technologies. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 2568–2599. [CrossRef]
5. Roohi, K.; Fekr, A.R. A Comparative Analysis of Indoor Localization Technologies. *Comput. Netw.* **2025**, *270*, 111527. [CrossRef]
6. Maheepala, M.; Kouzani, A.Z.; Joordens, M.A. Light-Based Indoor Positioning Systems: A Review. *IEEE Sens. J.* **2020**, *20*, 3971–3995. [CrossRef]
7. Guan, W.; Chen, S.; Wen, S.; Tan, Z.; Song, H.; Hou, W. High-Accuracy Robot Indoor Localization Scheme Based on Robot Operating System Using Visible Light Positioning. *IEEE Photonics J.* **2020**, *12*, 7901716 . [CrossRef]
8. Zhuang, Y.; Hua, L.; Qi, L.; Yang, J.; Cao, P.; Cao, Y.; Wu, Y.; Thompson, J.; Haas, H. A Survey of Positioning Systems Using Visible LED Lights. *IEEE Commun. Surv. Tutor.* **2018**, *20*, 1963–1988. [CrossRef]
9. Luo, J.; Fan, L.; Li, H. Indoor Positioning Systems Based on Visible Light Communication: State of the Art. *IEEE Commun. Surv. Tutor.* **2017**, *19*, 2871–2893. [CrossRef]
10. Zhu, Z.; Yang, Y.; Chen, M.; Guo, C.; Cheng, J.; Cui, S. A Survey on Indoor Visible Light Positioning Systems: Fundamentals, Applications, and Challenges. *IEEE Commun. Surv. Tutor.* **2025**, *27*, 1656–1686. [CrossRef]
11. Wang, R.; Niu, G.; Cao, Q.; Chen, C.S.; Ho, S.W. A Survey of Visible-Light-Communication-Based Indoor Positioning Systems. *Sensors* **2024**, *24*, 5197. [CrossRef] [PubMed]
12. Pathak, P.; Sharavanan, S.; Karthikeyan, M.; Mubarakali, A.; Arun Kumar, U. Enhanced Optical Visible Light Positioning via RSS-Based Evolutionary Optimization under LOS and NLOS Conditions. *Sci. Rep.* **2025**, *15*, 42025. [CrossRef] [PubMed]
13. Liu, X.; Zou, D.; Huang, N.; Wang, Y. An Efficient Iterative Least Square Method for Indoor Visible Light Positioning Under Shot Noise. *IEEE Photonics J.* **2023**, *15*, 7300910. [CrossRef]
14. Xu, Y.; Hu, X.; Sun, Y.; Yang, Y.; Zhang, L.; Deng, X.; Chen, L. High-Accuracy Height-Independent 3D VLP Based on Received Signal Strength Ratio. *Sensors* **2022**, *22*, 7165. [CrossRef] [PubMed]
15. Khan, S.; Paradells, J.; Catalán, M. Low-Cost 3D Indoor Visible Light Positioning: Algorithms and Experimental Validation. *Photonics* **2024**, *11*, 626. [CrossRef]
16. Xu, J.; Gong, C.; Xu, Z. Experimental Indoor Visible Light Positioning Systems With Centimeter Accuracy Based on a Commercial Smartphone Camera. *IEEE Photonics J.* **2018**, *10*, 7908717. [CrossRef]
17. Saleh, S.S.; Kheirallah, H.N.; Aly, M.H. Efficient Three-Dimensional Indoor Dark Light Visible Light Positioning with Received Signal Strength Technique. *Opt. Quantum Electron.* **2024**, *56*, 952. [CrossRef]

18. Do, T.H.; Yoo, M. An In-depth Survey of Visible Light Communication Based Positioning Systems. *Sensors* **2016**, *16*, 678. [CrossRef] [PubMed]
19. Prince, G.B.; Little, T.D.C. Two-Phase Framework for Indoor Positioning Systems Using Visible Light. *Sensors* **2018**, *18*, 1917. [CrossRef] [PubMed]
20. Zhang, X.; Duan, J.; Fu, Y.; Shi, A. Theoretical Accuracy Analysis of Indoor Visible Light Communication Positioning System Based on Received Signal Strength Indicator. *J. Light. Technol.* **2014**, *32*, 4180–4186. [CrossRef]
21. Yin, L.; Wu, X.; Haas, H. Indoor Visible Light Positioning with Angle Diversity Transmitter. In *Proceedings of the 2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall), Boston, MA, USA, 6–9 September 2015*; IEEE: Piscataway, NJ, USA, 2015; pp. 1–5. [CrossRef]
22. Li, D.; Gong, C.; Xu, Z. A RSSI-Based Indoor Visible Light Positioning Approach. In *Proceedings of the 2016 10th International Symposium on Communication Systems, Networks and Digital Signal Processing (CSNDSP), Prague, Czech Republic, 20–22 July 2016*; IEEE: Piscataway, NJ, USA, 2016; pp. 1–6. [CrossRef]
23. Du, P.; Zhang, S.; Chen, C.; Yang, H.; Zhong, W.D.; Zhang, R.; Alphones, A.; Yang, Y. Experimental Demonstration of 3D Visible Light Positioning Using Received Signal Strength with Low-Complexity Trilateration Assisted by Deep Learning Technique. *IEEE Access* **2019**, *7*, 93986–93997. [CrossRef]
24. Tran, H.Q.; Ha, C. Machine Learning in Indoor Visible Light Positioning Systems: A Review. *Neurocomputing* **2022**, *491*, 117–131. [CrossRef]
25. Chen, Y.; Han, T.; Hu, J.; Liu, X. Dynamic Indoor Visible Light Positioning and Orientation Estimation Based on Spatiotemporal Feature Information Network. *Photonics* **2025**, *12*, 990. [CrossRef]
26. Keskin, M.F.; Gezici, S.; Arikan, O. Direct and Two-Step Positioning in Visible Light Systems. *IEEE Trans. Commun.* **2018**, *66*, 239–254. [CrossRef]
27. Sahin, A.; Eroglu, Y.S.; Guvenc, I.; Pala, N.; Yuksel, M. Accuracy of AOA-Based and RSS-Based 3D Localization for Visible Light Communications. In *Proceedings of the 2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall), Boston, MA, USA, 6–9 September 2015*; IEEE: Piscataway, NJ, USA, 2015; pp. 1–5. [CrossRef]
28. Zhu, B.; Cheng, J.; Wang, Y.; Yan, J.; Wang, J. Three-Dimensional VLC Positioning Based on Angle Difference of Arrival with Arbitrary Tilting Angle of Receiver. *IEEE J. Sel. Areas Commun.* **2018**, *36*, 8–22. [CrossRef]
29. Steendam, H. A 3-D Positioning Algorithm for AOA-Based VLP with an Aperture-Based Receiver. *IEEE J. Sel. Areas Commun.* **2018**, *36*, 23–33. [CrossRef]
30. Armstrong, J.; Sekercioglu, Y.A.; Neild, A. Visible Light Positioning: A Roadmap for International Standardization. *IEEE Commun. Mag.* **2013**, *51*, 68–73. [CrossRef]
31. Li, Q.L.; Wang, J.Y.; Huang, T.; Wang, Y. Three-Dimensional Indoor Visible Light Positioning System with a Single Transmitter and a Single Tilted Receiver. *Opt. Eng.* **2016**, *55*, 106103. [CrossRef]
32. Yasir, M.; Ho, S.W.; Vellambi, B.N. Indoor Positioning System Using Visible Light and Accelerometer. *J. Light. Technol.* **2014**, *32*, 3306–3316. [CrossRef]
33. Amsters, R.; Holm, D.; Joly, J.; Demeester, E.; Stevens, N.; Slaets, P. Visible Light Positioning Using Bayesian Filters. *J. Light. Technol.* **2020**, *38*, 5925–5936. [CrossRef]
34. Cai, Y.; Guan, W.; Wu, Y.; Xie, C.; Chen, Y.; Fang, L. Indoor High Precision Three-Dimensional Positioning System Based on Visible Light Communication Using Particle Swarm Optimization. *IEEE Photonics J.* **2017**, *9*, 7908120. [CrossRef]
35. Chew, M.T.; Alam, F.; Noble, F.K.; Legg, M.; Gupta, G.S. Visible Light Positioning-Based Robot Localization and Navigation. *Electronics* **2024**, *13*, 368. [CrossRef]
36. Wu, F.; Stevens, N.; De Strycker, L.; Rottenberg, F. Enhancing RSS-Based Visible Light Positioning by Optimal Calibration of LED Tilt and Gain. *IEEE Trans. Commun.* **2025**, *73*, 7771–7783. [CrossRef]
37. Almadani, Y.; Ijaz, M.; Joseph, W.; Bastiaens, S.; Rajbhandari, S.; Adebisi, B.; Plets, D. A Novel 3D Visible Light Positioning Method Using Received Signal Strength for Industrial Applications. *Electronics* **2019**, *8*, 1311. [CrossRef]
38. Almadani, Y.; Ijaz, M.; Adebisi, B.; Rajbhandari, S.; Bastiaens, S.; Joseph, W.; Plets, D. An Experimental Evaluation of a 3D Visible Light Positioning System in an Industrial Environment with Receiver Tilt and Multipath Reflections. *Opt. Commun.* **2021**, *483*, 126654. [CrossRef]
39. Xu, S.; Wu, Y.; Zhang, L.; Xing, S.; Wei, F. Simultaneous Position and Orientation Estimation for Reconfigurable Intelligent Surfaces-Assisted Visible Light Communications. *IEEE Trans. Veh. Technol.* **2025**, *74*, 861–876. [CrossRef]
40. Zhou, B.; Liu, A.; So, H.C. Visible Light Communication-Enabled Simultaneous Position and Orientation Detection for Harnessing Multipath Interference and Random Fading. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 1938–1952. [CrossRef]
41. Ghassemlooy, Z.; Popoola, W.; Rajbhandari, S. *Optical Wireless Communications: System and Channel Modelling with MATLAB*; CRC Press: Boca Raton, FL, USA, 2019. [CrossRef]
42. Deng, L.; Fan, Y.; Zhao, Q.; Wu, P. Indoor Visible Light Positioning Based on Channel Estimation and Cramér–Rao Low Bound Analysis with Random Receiving Orientation of User Equipment. *Photonics* **2023**, *10*, 812. [CrossRef]

43. ElFar, S.H.; Yaseen, M.; Ikki, S. Insights into Visible Light Positioning: Range Tracking and Bayesian Cramér–Rao Lower Bound Analysis. *IEEE Commun. Lett.* **2024**, *28*, 2056–2060. [CrossRef]
44. Yazar, O.; Gezici, S. Optimal Pulse Design for Visible Light Positioning Systems. *Signal Process.* **2022**, *191*, 108334. [CrossRef]
45. Tomic, S.; Beko, M.; Dinis, R.; Montezuma, P. A Robust Bisection-Based Estimator for TOA-Based Target Localization in NLOS Environments. *IEEE Commun. Lett.* **2017**, *21*, 2488–2491. [CrossRef]
46. Vaghefi, R.M.; Gholami, M.R.; Buehrer, R.M.; Strom, E.G. Cooperative Received Signal Strength-Based Sensor Localization with Unknown Transmit Powers. *IEEE Trans. Signal Process.* **2013**, *61*, 1389–1403. [CrossRef]
47. Zhu, X.; Mo, L.; Wu, X. Efficient Closed-Form Solutions for Visible Light Positioning in Low-Cost IoT Devices. *Electronics* **2024**, *13*, 614. [CrossRef]
48. Sun, Y.; Yang, S.; Wang, G.; Chen, H. Robust RSS-Based Source Localization with Unknown Model Parameters in Mixed LOS/NLOS Environments. *IEEE Trans. Veh. Technol.* **2021**, *70*, 3926–3931. [CrossRef]
49. Ma, S.; Yang, R.; Li, B.; Chen, Y.; Li, H.; Wu, Y.; Safari, M.; Li, S.; Al-Dhahir, N. Optimal Power Allocation for Integrated Visible Light Positioning and Communication System with a Single LED-Lamp. *IEEE Trans. Commun.* **2022**, *70*, 6734–6747. [CrossRef]
50. Iddrisu, I.; Gezici, S. Visible Light Positioning with Intelligent Reflecting Surfaces under Mismatched Orientations. *Signal Process.* **2025**, *230*, 109867. [CrossRef]
51. Cai, C.; Fu, M.; Meng, X.; Jia, C.; Pei, M. Indoor High-Precision Visible Light Positioning System Using Jaya Algorithm. *Math. Biosci. Eng.* **2023**, *20*, 10358–10375. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.