mathematics

MDPI

Article

FR3 Path Loss in Outdoor Corridors: Physics-Guided Two-Ray Residual Learning

Jorge Celades-Martínez [1,*], Jorge Rojas-Vivanco [2], Melissa Diago-Mosquera [3], Alvaro Peña [2] and Jose García [2,*]

[1] Doctorado en Industria Inteligente, Pontificia Universidad Católica de Valparaíso, Valparaíso 2362804, Chile
[2] Escuela de Ingeniería de Construcción y Transporte, Pontificia Universidad Católica de Valparaíso, Valparaíso 2362804, Chile; jorge.rojas.v@pucv.cl (J.R.-V.); alvaro.pena@pucv.cl (A.P.)
[3] Departamento de Electrónica, Universidad Técnica Federico Santa María, Valparaíso 2390123, Chile; melissa.diago@usm.cl
* Correspondence: jorge.celades.m@mail.pucv.cl (J.C.-M.); jose.garcia@pucv.cl (J.G.)

Abstract

Accurate path-loss characterization in the upper mid-band is critical for 5G/6G outdoor planning, yet classical deterministic expressions lose fidelity at 18 GHz, and purely data-driven regressors offer limited physical insight. We present a physics-guided residual learner that couples a calibrated two-ray model with an XGBoost regressor trained on the deterministic residuals. To enlarge the feature space without promoting overfitting, synthetic samples obtained by perturbing antenna height and ground permittivity within realistic bounds are introduced with a weight of $w = 0.3$. The methodology is validated with narrowband measurements collected along two straight 25 m corridors. Under cross-corridor transfer, the hybrid predictor attains 0.59–0.62 dB RMSE and $R^2 \geq 0.996$, reducing the error of a pure-ML baseline by half and surpassing deterministic formulas by a factor of four. Small-scale analysis yields decorrelation lengths of 0.23 m and 0.41 m; a cross-correlation peak of unity at $\Delta = 0.10$ m confirms the physical coherence of both corridors. We achieve <1 dB error using a small set of field measurements plus simple synthetic data. The method keeps a clear mathematical core and can be extended to other priors, NLOS cases, and semi-open hotspots.

Keywords: upper mid-band; FR3; path-loss modeling; physics-informed machine learning; two-ray model; residual learning; outdoor corridors; 18 GHz; synthetic data augmentation

MSC: 78M50; 68T05

check for updates

Academic Editors: Haolan Zhang, Xingsen Li and Yunfa Fu

Received: 14 July 2025
Revised: 19 August 2025
Accepted: 21 August 2025
Published: 23 August 2025

Citation: Celades-Martínez, J.; Rojas-Vivanco, J.; Diago-Mosquera, M.; Peña, A.; García, J. FR3 Path Loss in Outdoor Corridors: Physics-Guided Two-Ray Residual Learning. *Mathematics* **2025**, *13*, 2713. https://doi.org/10.3390/math13172713

Copyright: © 2025 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://creativecommons.org/licenses/by/4.0/).

1. Introduction

A need for new spectral resources has been driven by emerging 6G services, such as AI-native communications, extended reality, and ultra-reliable machine-type applications [1]. The upper mid-band between 7 and 24 GHz (FR3) has been identified as a promising complement to the existing bands [2]. Congestion in sub-6 GHz (FR1) has been reported, while FR3 offers wider bandwidths but exhibits stronger propagation losses and reduced coverage [3–6]. Owing to this balance, FR3 has been positioned to enable enhanced mobile broadband while leveraging the current deployments.

At 18 GHz, limitations in traditional path-loss modeling have been observed [7]. Deterministic models (e.g., ray tracing) provide detailed site-specific predictions but require

extensive geometric and material data and substantial computation [8,9]. Empirical formulas are easier to apply but often lack generality across sites or frequencies and require large measurement campaigns [10–13]. Deterministic models are accurate but costly. Empirical models are simpler but often less accurate or reliable. This trade-off worsens at high frequencies, where small scene changes can cause large path-loss shifts.

In this work, a physics-guided residual learner based on a calibrated two-ray model and an XGBoost corrector is proposed for outdoor corridors at 18 GHz. The two-ray term is calibrated to measured data to capture distance-dependent loss and LOS–ground interference. Residual errors are then learned by XGBoost to correct systematic deviations. Light synthetic augmentation is included during training with a down-weight $w = 0.3$ to enlarge coverage without overfitting.

Adding synthetic samples exposes XGBoost to more scenarios. This makes it more robust and helps it to learn a smoother more general function

$$L_{\text{pred}}(d) \;=\; L_{\text{Two-Ray}}(d) \;+\; \Delta_{\text{XGB}}(d),$$

such that known propagation physics are enforced by $L_{\text{Two-Ray}}$, while site-specific effects are absorbed by Δ_{XGB}. Consistency with the physics-enhanced residual learning (PERL) paradigm has been maintained, as adopted in other domains, including computational fluid dynamics [14,15], biomedical applications [16,17], and climate modeling [18].

The contributions of this study are as follows. A physics-guided residual predictor for outdoor FR3 corridors is formulated and evaluated under strict cross-corridor transfer. A simple calibration and error-surface analysis of the two-ray prior is provided to quantify sensitivity to antenna height and permittivity. An adaptive hybrid variant is defined in which lightly weighted synthetic residuals are used to improve robustness without overfitting. A reproducible pipeline is documented, including fixed feature definitions, frozen hyperparameters across splits, dataset fields and sample counts, and supporting small-scale statistics to aid replication.

Validation has been conducted with 18 GHz measurements along two 25 m outdoor corridors under LOS. The hybrid predictor has been benchmarked against a pure-ML XGBoost baseline and deterministic formulas (two-ray, free space, and single-slope). Sub-1 dB RMSE and $R^2 \approx 1.0$ have been achieved under cross-corridor transfer, with residuals showing low autocorrelation and high spectral entropy. It is therefore indicated that near-deterministic accuracy can be reached while preserving interpretability for FR3 planning.

The remainder of the article is organized as follows. Section 2 reviews the related work on learning-based path-loss modeling across sub-7 GHz, FR3, and mmWave. Section 3 describes the measurement campaign, preprocessing, the deterministic two-ray model and its calibration, the machine-learning pipeline, and the small-scale statistics. Section 4 presents the results and discussion for large-scale modeling, cross-corridor transfer, residual learning, and the adaptive hybrid with synthetic augmentation. Section 5 concludes and outlines future directions.

2. Related Work

Considerable activity has been reported on learning-based path-loss modeling for FR3 and adjacent bands. A surge of ML applications has been motivated by the need to learn propagation behavior directly from data [19]. Work has been concentrated in indoor offices [20,21] and corridors [5,22]. Concerns regarding the black-box nature of purely data-driven models have been expressed, which has motivated physics-informed approaches [23–25]. A gap has remained for outdoor corridor validation at the upper mid-band, where dominant mechanisms differ from indoor cases. A consolidated overview is provided in Table 1 for sub-7 GHz (FR1) and in Table 2 for FR2/FR3.

ML techniques have been adopted across a wide range of scenarios. Supervised models have included SVR, random forests, MLPs, LSTMs, and boosted ensembles. In indoor corridors at 14/18/22 GHz, sub-decibel RMSE has been reported with MLP/LSTM [11], while empirical single/double-slope formulas have underperformed [22]. In complex layouts generated via full-wave simulation, CatBoost has outperformed deep MLPs [21]. In measurement-only environments, SVR and random forests have yielded sub-2 dB within aircraft cabins [26], and KNN bagging has shown competitiveness in rural macro-cells [27]. These trends are consistent with the summaries in Tables 1 and 2, which group representative studies by paradigm, data source, band, scenario, and reported RMSE.

Table 1. Summary of supervised learning methods for path-loss prediction in Frequency Range 1 (sub-7 GHz). Each entry lists the architecture, key idea, data source, frequency band, scenario, and RMSE (dB).

ML Method	Architecture/Type	Key Feature	Reference	Channel Model	Frequency [GHz]	Scenario	RMSE [dB]
Convolutional Neural Networks (CNNs)							
CNN	Convolutional layers for spatial modeling	Captures spatial dependencies	[28]	Ray Tracing	5	Small offices	–
Multi-Layer Perceptron Networks (MLPs)							
MLP	Multi-layer feed-forward network	Captures complex mappings	[29]	Ray Tracing	0.9	Urban	44
			[30,31]	Ray Tracing	0.9	Urban	–
			[32,33]	Ray Tracing	0.9, 3.5	Urban	–
			[34]	Only measurements	7	Urban	4.5
			[35]	Only measurements	sub-6	Coastal/Vegetative	1.9/4.3
Tree-based Ensemble Methods							
XGBoost	Gradient boosting with regularization	Handles sparse data	[36]	Only measurements	0.449–5.85	Urban/Suburban (UK)	7.44
Random Forest	Bootstrap aggregated trees	Reduces variance, interpretable	[26]	Only measurements	2.4	Aircraft cabin	1.76
AdaBoost	Adaptive boosting of weak learners	Focuses on hard examples	[26]	Only measurements	2.4	Aircraft cabin	2.12
Kernel- and Distance-based Methods							
SVR	Kernel-based regression	Works with small data; margin-based	[26]	Only measurements	2.4	Aircraft cabin	2.20
Bagging-KNN	KNN with bootstrap aggregation	Robust to noise	[27]	Only measurements	3.7	Rural Greece	4.3

Table 2. Summary of supervised learning methods for path-loss prediction in Frequency Ranges 2 and 3 (above 7 GHz). Studies marked with an asterisk (*) also include lower-frequency data.

ML Method	Architecture/Type	Key Feature	Reference	Channel Model	Frequency [GHz]	Scenario	RMSE [dB]
Convolutional Neural Networks (CNN)							
CNN	Convolutional layers for spatial modeling	Captures spatial dependencies	[37] *	Ray Tracing	0.8–60	Urban	22
			[38]	Single slope model	28	Suburban	7.2
			[39]	Only measurements	28	Suburban	8.6
			[40]	Only measurements	mmWave	Urban	6

Table 2. *Cont.*

ML Method	Architecture/Type	Key Feature	Reference	Channel Model	Frequency [GHz]	Scenario	RMSE [dB]
Multi-Layer Perceptron Networks (MLP) and RNN Variants							
MLP	Multi-layer feed-forward network	Captures complex mappings	[41] *	–	0.5, 28	Railway station	0.8
			[42] *	–	0.8–70	Urban/Suburban	8.3
			[43]	Ray Tracing	30	Urban	5.8
MLP	Multi-layer feed-forward on measured data	Captures non-linear PL behavior	[11]	Only measurements	14, 18, 22	Enclosed corridor	0.036
RNN-LSTM	Recurrent net with memory cells	Learns temporal dependencies	[11]	Only measurements	14, 18, 22	Enclosed corridor	0.042
MLP	8-layer trained on CST sim.	Indoor path loss learning	[21]	CST-EM simulation	28	Indoor office (LOS/NLOS)	6.70
Tree-based Ensemble Methods							
CatBoost	Gradient-boosted decision trees	Lowest RMSE in study	[21]	CST-EM simulation	28	Indoor office (LOS/NLOS)	4.68

Despite strong accuracy, purely ML models have been characterized as black boxes with limited physical transparency [23,24]. Generalization can degrade when deployment conditions depart from training data [19], and large datasets are often required at FR3, which raises cost. These factors have motivated hybrid modeling, where physics-based priors are combined with data-driven residuals [25]. The present study follows this direction for outdoor corridors at 18 GHz.

3. Materials and Methods

3.1. Measurement Campaign and Raw Data Acquisition

To create the dataset for propagation analysis and subsequent ML model training, a comprehensive measurement campaign was executed. The characterization was performed at a frequency of $f = 18$ GHz, utilizing a continuous-wave (CW) signal transmitted at a nominal power of $Ptx = 0$ dBm. The equipment setup included a dedicated transmitter (Tx) and receiver (Rx) chain using the compact signal generator J0GSAG1311 and the spectrum analyzer J0SSAP53, both manufactured by SAF Tehnika (Riga, Latvia). Both terminals were equipped with conical horn antennas (model J0AA1724HG01), each providing a gain of $G_{tx} = 21.1$ dBi and $G_{rx} = 21.1$ dBi, with a half-power beamwidth (HPBW) of $\approx 13.5^\circ$ in E-plane and $\approx 16.5^\circ$ in H-plane. To ensure measurement integrity and verify antenna performance, the entire system was verified in an anechoic chamber prior to field deployment, as seen in Figure 1. The total measured losses attributed to cables and connectors in the system were $L_{trx} = 3.4$ dB.

The experiments were conducted in an outdoor courtyard at the Universidad Técnica Federico Santa María campus in Valparaíso, Chile. This environment is a representative scenario featuring open areas bordered by building facades of concrete and glass, thus introducing a rich multipath environment with specular components. Two distinct pedestrian corridors within this courtyard, labeled as C1 and C2, as seen in Figure 2, were selected as measurement routes to capture channel variations under LOS conditions. A dynamic sampling methodology was employed to capture the spatial variability of the channel. The Tx was mounted on a mobile platform and moved at a slow constant velocity along a linear trajectory, while the Rx remained stationary at one end of the corridor. Throughout all measurements, both Tx and Rx antennas were maintained at a fixed height of $h_{tx} = 1.2$ m and $h_{rx} = 1.2$ m

above the ground to ensure consistent ground–reflection geometry. This procedure was repeated for both corridors, covering a total distance range, d, from 3.15 m to 25 m.

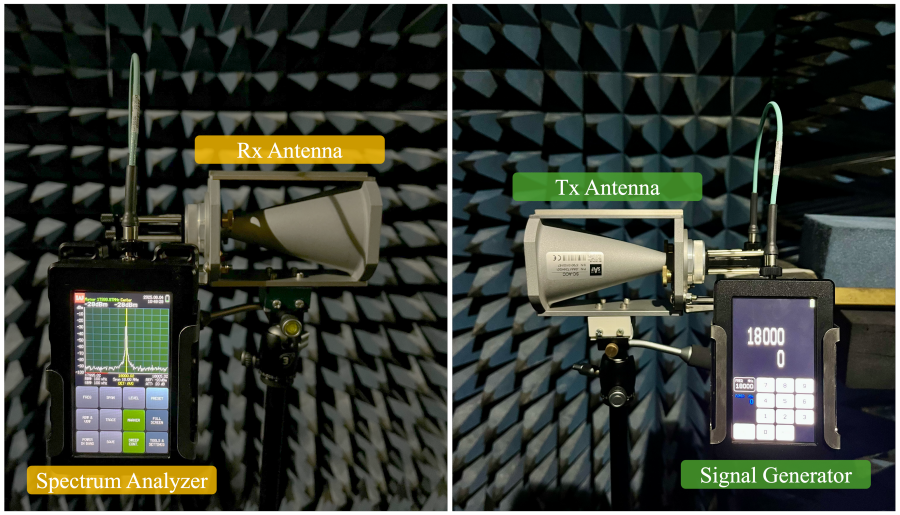


Figure 1. Verification of equipment in anechoic chamber.



Figure 2. Measurement scenario and equipment.

At the receiving end, a high-sensitivity spectrum analyzer, configured with an internal attenuation of 16 dB to optimize its dynamic range, was used to continuously log the incoming signal power. This continuous data acquisition method provides high-density spatial sampling, which is crucial for accurately characterizing large-scale path loss trends. The campaign yielded a comprehensive dataset of approximately 500 received power samples, providing a statistical foundation for the analysis presented in the subsequent sections.

3.2. Dataset Summary and Splits

Two outdoor-corridor datasets were used (C1: $N = 186$, C2: $N = 334$), totaling $N = 520$ large-scale samples after the 30λ Lee filter. For each sample, the fields distance, PL_mean, Prx_mean, Pr_Raw_dBm, and PL_Raw_dB were recorded.

Headline results were obtained under strict cross-corridor transfer. Models were trained on all samples from C1 and tested on all samples from C2, and vice versa, with no intermixing between train and test sets. A random 80/20 split on the combined set ($N = 520$; random_state = 42) was employed only as a diagnostic baseline and was not used for the reported cross-corridor metrics.

For pure ML baselines, the target was PL_{mean} and the features were $\{\log_{10} d,\ \text{corridor_id}\}$. For residual/hybrid models, the target was $\varepsilon(d) = PL_{\text{mean}} - PL_{\text{2-ray}}(d)$ and the features were $\{\log_{10} d,\ PL_{\text{2-ray}}(d),\ h_t(=h_r),\ \varepsilon_r,\ \text{corridor_id}\}$.

3.3. Preprocessing

Accurate characterization of the channel requires isolating the large-scale path-loss trend from the rapid fluctuations caused by small-scale multipath fading. The raw power trace $P_{\text{raw}}(d)$ acquired while the receiver traverses the corridor is therefore subjected to a two–step preprocessing pipeline.

(i) Large-scale averaging.

A sliding-window average of the received power is computed to filter the small-scale fading, following Lee's classical procedure [44]. A window length of 30λ was chosen, a value that falls well within the 10λ to 40λ range recommended in the literature for effectively separating short-term and long-term signal variations [45,46]. At the operating frequency of 18 GHz ($\lambda = 16.7$ mm, hence $30\lambda \simeq 0.51$ m), this length provides a robust trade-off: it is sufficiently large to average over the rapid constructive/destructive interference that occurs on centimeter scales (i.e., on the order of $\lambda/2$), yet it is physically small enough to avoid over-smoothing the data and masking the genuine gradual distance-dependent decay of signal power. Throughout, powers are expressed in dBm, antenna gains in dBi, and losses in dB, and $P_{\text{raw}}(d)$ denotes the instantaneous received-power sample prior to averaging. Therefore, the path loss is computed using the standard link–budget formula:

$$PL(d) \;=\; P_t + G_t + G_r - L_{trx} \;-\; \underbrace{\big\langle P_{\text{raw}}(d') \big\rangle_{d' \in \mathcal{W}(d)}}_{\text{Lee filter}}, \tag{1}$$

where $\mathcal{W}(d)$ denotes the spatial window centered at distance d with width 30λ and $\langle \cdot \rangle$ the arithmetic mean inside that window. Here, P_t is the transmit power at the source output; G_t and G_r are the transmit and receive antenna gains, respectively; L_{trx} denotes the fixed RF front-end loss (cables/connectors and any fixed analyzer attenuation, if not corrected separately in P_{raw}); and $P_{\text{raw}}(d)$ is the instantaneous received-power sample before spatial averaging. These estimates isolate the large-scale path-loss component and serve as the basis for all subsequent analyses and modeling steps.

(ii) Excess path loss (EPL).

To analyze the residual variability of the channel, the deterministic free-space loss predicted by

$$PL_{\text{FSPL}}(d) \;=\; 20 \log_{10}\!\left(\frac{4\pi d}{\lambda}\right) \tag{2}$$

is subtracted from $PL(d)$.

The resulting excess path loss

$$\text{EPL}(d) \;=\; \overline{PL}(d) \;-\; PL_{\text{FSPL}}(d) \tag{3}$$

is therefore referenced to an ideal line-of-sight scenario and retains only those deviations induced by the environment (e.g., floor reflections or wall scattering). This quantity is subsequently employed for all small-scale statistical analyses (decorrelation length, spectral entropy, etc.) presented in Section 3.7.

To decouple large-scale attenuation (i.e., path loss and shadowing) from small-scale fluctuations (multipath fading), we extracted the local mean power from the measurement data using a 30λ sliding window. This process yields an empirical estimate of the combined path loss and shadowing, which enables a meaningful comparison between our results and

the generative statistical models defined in the IEEE 802.11n standard [47]. Crucially, this separation of scales is essential to prevent bias in performance metrics, such as the root mean squared error (RMSE), and in the residuals learned during the ML stage.

3.4. Deterministic Reference Model

Two-ray formulation: To capture the dominant large-scale behavior of the outdoor corridor at 18 GHz, the two-ray ground–reflection model is adopted. The received complex field is represented as the coherent sum of an LOS component and a floor-reflected component [48], namely

$$\overline{P_{rx}}(d) = P_{tx}\left(\frac{\lambda}{4\pi}\right)^2 \left| \frac{\sqrt{G_0}}{d_0} + \frac{\rho_s R \sqrt{G_1} e^{-j\Delta\phi}}{d_1} \right|^2 \tag{4}$$

where d_0 and d_1 are the lengths of the direct and reflected paths. $\Delta\phi$ is the phase difference between the two paths resulting from the difference in path length given by $\Delta\phi = \frac{2\pi}{\lambda}\left(\sqrt{(h_t+h_r)^2+d^2} - \sqrt{(h_t-h_r)^2+d^2}\right)$. G_0 and G_1 are the products of the antenna gains of the transmitter and receiver for the direct and reflected paths.

The factor ρ_s (roughness attenuation factor) attenuates the specular reflection according to the surface roughness σ_h

$$\rho_s = \exp\left[-8\pi^2\sigma_h^2 \cos^2\theta_g / \lambda^2\right], \tag{5}$$

nominal electromagnetic parameters for concrete are adopted from the literature: relative permittivity $\varepsilon_r = 7$, conductivity $\sigma = 0.05\,\mathrm{S/m}$, and roughness $\sigma_h = 8.5\,\mathrm{mm}$. Substituting $E_{\mathrm{tot}}(d)$ into Friis' equation yields the closed-form path-loss prediction $PL_{2\text{-ray}}(d)$ used throughout this study [49–51].

R is the ground reflection coefficient, which was modeled for vertical polarization

$$R = \frac{\sin\theta - Z}{\sin\theta + Z}, \quad Z = \frac{\sqrt{\epsilon_r - \cos^2\theta}}{\epsilon_r}, \tag{6}$$

The value of R is determined by the dielectric properties of the concrete floor, encapsulated in the complex permittivity model [52] $\epsilon_r = \epsilon - j60\sigma\lambda$ using the specific values of relative permittivity ϵ_0 and conductivity σ for concrete.

Calibration sweep: Although antenna heights were nominally fixed at $h_t = h_r = 1.20\,\mathrm{m}$ during the measurement campaign, small installation tolerances ($\pm 0.1\,\mathrm{m}$) and dielectric uncertainty prompted an in silico calibration. A full-factorial sweep was therefore carried out over the grid

$$h_t = h_r \in \{0.8, 1.0, 1.2, 1.4, 1.6\}\mathrm{m}, \qquad \varepsilon_r \in \{4.0, 4.5, \ldots, 7.0\},$$

and the corresponding $PL_{2\text{-ray}}(d)$ values were simulated for each pair. The same 30λ averaging window described in Section 3.3 was applied, and the results were interpolated to match the exact distance samples of corridors C1 and C2. The RMSE with respect to the measured $\overline{PL}(d)$ was computed separately for each corridor and summed as

$$\mathrm{RMSE}_{\mathrm{tot}}(h, \varepsilon_r) = \mathrm{RMSE}_{\mathrm{C1}} + \mathrm{RMSE}_{\mathrm{C2}},$$

producing an error surface whose minimum was used to identify the optimal parameter pair $(h^\star, \varepsilon_r^\star)$. In the present dataset, the minimum was found at $h^\star \approx 1.2\,\mathrm{m}$, $\varepsilon_r^\star \approx 7$, confirming the correctness of the installation and providing a physically grounded baseline (total RMSE $\approx 1.9\,\mathrm{dB}$) against which all machine-learning models were subsequently compared.

3.5. Machine-Learning Pipeline

The pipeline (Figure 3) follows a physics-informed data-driven workflow. Raw logs for C1 and C2 are first converted from timestamps to distance and averaged with a 30λ Lee filter to obtain large-scale path loss. A deterministic two-ray model then provides a baseline $PL_{\text{2-ray}}(d)$ under nominal physical settings, which is used both as a prior and as an input feature together with $\log_{10} d$ and a corridor identifier. Global Bayesian hyperparameter optimization with `Optuna` selects the best configurations for XGBoost, random forest, and MLP using a cross-validated objective that trades off validation RMSE against training time. Model fitting proceeds through three schemes aligned with the experiments: (a) cross-corridor transfer to test spatial generalization, (b) residual learning that predicts $\varepsilon = PL_{\text{meas}} - PL_{\text{2-ray}}$ and reconstructs $\widehat{PL}$ by superposition, and (c) an adaptive hybrid variant that augments training with lightly weighted synthetic residuals generated from the two-ray model under moderate perturbations of antenna height and floor permittivity. At inference, the final predictor returns $\widehat{PL}(d) = PL_{\text{2-ray}}(d) + \hat{\varepsilon}(d)$.

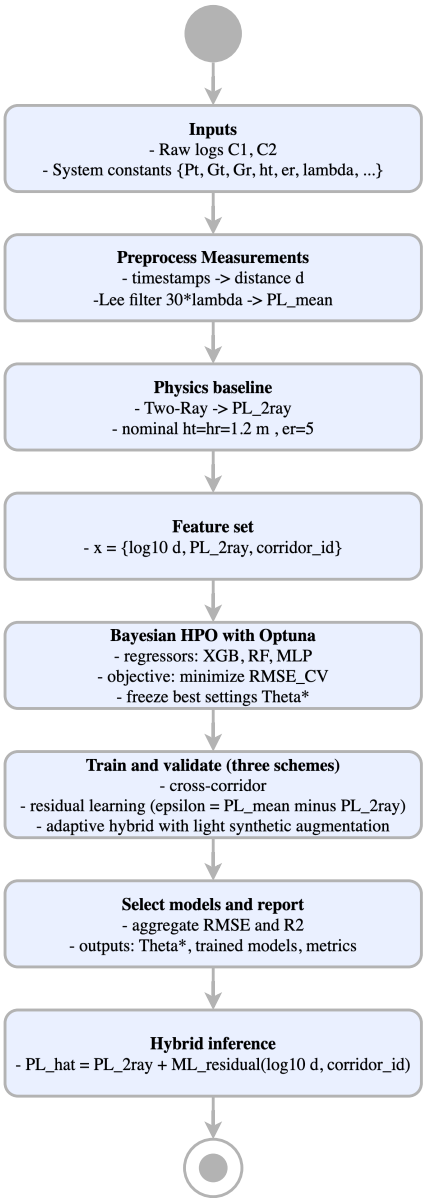


Figure 3. High-level physics-informed ML pipeline used in this study. Note that *Theta** denotes the optimal hyperparameter configuration found by Bayesian optimization.

The data-driven component of this study was structured into three successive configurations, each mapped to a specific experimental scenario (see Table 3). All models were implemented using the scikit-learn [53] and xgboost [54] libraries.

Table 3. Mapping between experiments and ML training schemes.

Experiment	Training Set	Prediction Target
a—Cross-corridor	Real data only	PL_{meas}
b—Residual model	Real data only	$\varepsilon = PL_{\text{meas}} - PL_{\text{2-ray}}$
c—Adaptive hybrid	Real + synthetic ($w = 0.3$)	ε (as above)

Feature engineering.

For each distance sample d, the feature vector was constructed as

$$\mathbf{x}(d) = \{\log_{10} d,\ PL_{\text{2-ray}}(d),\ h_t,\ \varepsilon_r,\ \text{corridor_id}\},$$

where $\log_{10} d$ encodes distance decay, $PL_{\text{2-ray}}$ provides a physics-based prior, (h_t, ε_r) represent antenna height and floor permittivity, and corridor_id $\in \{0, 1\}$ enables the model to capture corridor-specific patterns.

Three types of regressors were considered:

- XGBoost (gradient boosting trees): 200 trees of depth 6 and learning rate 0.10 with subsampling of 0.9 were used in the cross-corridor case; 50 trees of depth 4 were applied in the residual-learning configuration; and 150 trees of depth 5 were employed in the adaptive hybrid configuration.
- Random forest: 300 trees with unlimited depth and bootstrap sampling.
- Multi-layer perceptron neural network: two hidden layers (64 and 32 neurons), with ReLU activation and the Adam optimizer, trained for 1500 epochs.

All hyperparameters were tuned using Bayesian optimization via the Optuna framework [55] with a cross-validated objective by which validation RMSE was minimized under a runtime penalty. Both RMSE and training time were included in the objective function to balance accuracy with computational cost. To ensure a fair search, the same data partitions, preprocessing pipeline, and scoring metric were used across model families. Once selected, the optimal hyperparameters were held fixed across all experiments to ensure consistency and comparability. Per-experiment retuning was avoided, and reproducible head-to-head comparisons were enabled by freezing the settings.

Three distinct training strategies were followed, each aligned with one of the experimental configurations:

(a) Cross-corridor transfer. The complete dataset from corridor C1 was used for training, the model was validated on corridor C2, and then the roles were reversed. This setup assessed spatial generalization without overlap between training and test data.

(b) Residual learning. The target variable was defined as the residual $\varepsilon(d) = PL_{\text{meas}}(d) - PL_{\text{2-ray}}(d)$. The model was trained to predict the fine-scale deviations from the deterministic baseline. The final prediction was computed as the sum $PL_{\text{2-ray}} + \hat{\varepsilon}$.

(c) Adaptive hybrid with synthetic augmentation. To enhance robustness against moderate variations in physical parameters, synthetic residuals were generated using the two-ray model across four parameter combinations: $(h_t, \varepsilon_r) \in \{1.0, 1.4\} \times \{5, 7\}$. These synthetic instances were added to the training set and down-weighted ($w = 0.3$) to influence, but not dominate, the learning process. Evaluation was conducted exclusively on real measurement data.

All models were evaluated using the RMSE in dB and the coefficient of determination (R^2). Algorithm 1 summarizes the complete data-driven workflow as a single reproducible sequence. It begins with large-scale averaging of the raw measurements and the construction of a physics-based two-ray baseline. A unified feature table is then assembled, followed by the execution of the three learning schemes: (a) cross-corridor transfer to assess spatial generalization, (b) residual learning to capture deviations from the deterministic model, and (c) an adaptive hybrid strategy that combines real and augmented data to improve robustness. The algorithm also specifies where hyperparameter tuning and post hoc interpretability analyses are applied. This end-to-end representation complements Table 3 by illustrating how each experimental setting is embedded within the unified processing pipeline.

Algorithm 1: End-to-end workflow for the hybrid ML path-loss study.

Input: Raw logs $P_{\text{raw}}(t)$ for corridors C1, C2; system constants $\{P_t, G_t, G_r, h_t, \varepsilon_r, \lambda, \dots\}$.
Output: Best hyperparameters $\Theta^\star$, trained models $\mathcal{M}_{\{a,b,c\}\times\{\text{XGB,RF,MLP}\}}$, and performance metrics.

`// 1. Measurement preprocessing`
foreach *corridor file* **do**
  convert timestamps $\rightarrow$ distance $d(t)$;
  apply Lee filter $(30\lambda) \Rightarrow PL_{\text{mean}}(d)$;

`// 2. Physics-based baseline`
compute $PL_{\text{2-ray}}(d)$ with nominal $(h_t = h_r = 1.2\text{ m},\ \varepsilon_r = 5)$;

`// 3. Feature engineering`
foreach *sample d* **do**
  $\mathbf{x}(d) = \{\log_{10} d,\ PL_{\text{2-ray}}(d),\ h_t,\ \varepsilon_r,\ \texttt{corridor_id}\}$;

`// 4. Global Bayesian hyperparameter search (Optuna)`
foreach *regressor* $r \in \{XGB, RF, MLP\}$ **do**
  run Optuna to minimise $\text{RMSE}_{\text{CV}} + \alpha\text{time}_{\text{train}}$;
  store best setting $\Theta_r^\star$;

freeze $\Theta_r^\star$ for all subsequent fits;

`// 5. Training strategy (a) Cross-corridor`
foreach *(train C1 $\rightarrow$ test C2, train C2 $\rightarrow$ test C1)* **do**
  $y \leftarrow PL_{\text{mean}}$;
  foreach *r* **do**
    train $r(\Theta_r^\star)$ on train corridor; evaluate on test corridor;

`// 6. Training strategy (b) Residual learning`
compute residuals $\varepsilon(d) = PL_{\text{mean}}(d) - PL_{\text{2-ray}}(d)$;
$y \leftarrow \varepsilon(d)$;
foreach *r* **do**
  train on mixed corridors; predict $\hat{\varepsilon}$;
  output $\hat{PL} = PL_{\text{2-ray}} + \hat{\varepsilon}$;

`// 7. Training strategy (c) Adaptive hybrid with synthetic augmentation`
generate synthetic residuals using $(h_t, \varepsilon_r) \in \{1.0, 1.4\} \times \{5, 7\}$;
weight synthetic rows by $w = 0.3$; set $y \leftarrow \varepsilon$;
foreach *r* **do**
  train on real + synthetic; evaluate on real test data;

`// 8. Evaluation and reporting`
collect RMSE and R^2 for every (scheme, r) pair;

3.6. Complexity of the Hybrid vs. Pure ML and Pure Deterministic Models (5G/6G Scale)

The residual is learned in our approach by a hybrid in which a small XGBoost is stacked on a calibrated two-ray prior. Concretely, XGBoost with 50 trees of depth 4 is employed for residual learning and 150 trees of depth 5 are applied in the adaptive hybrid; when only real measurement data are used, XGBoost with 200 trees of depth 6, learning rate 0.10, and subsampling 0.9 is trained. Random forest is configured with 300 trees with unlimited depth and bootstrap, while the MLP is implemented with two hidden layers (64 and 32), ReLU activation, the Adam optimizer, and 1500 training epochs. These hyperparameters were selected by Bayesian optimization with Optuna [55], using a cross-validated objective in which validation RMSE was minimized under a runtime penalty. Once determined, the optimal hyperparameters were fixed across experiments so that strict head-to-head comparability was ensured.

Deterministic (two-ray):

Training: none. Inference per query point: closed-form evaluation with a handful of arithmetic operations $\rightarrow \mathcal{O}(1)$; negligible memory. This serves as the fastest and most scalable baseline, although it is less accurate in complex scenes.

Pure ML (RF/XGBoost/MLP):

- XGBoost (data-real only, cross-corridor): training $\approx \mathcal{O}(Tnd\text{depth})$; inference $\approx \mathcal{O}(T\text{depth})$ node tests per point. With $T = 200$, depth $= 6$, this amounts to $\sim$1200 node tests per query; model size grows with T and depth.
- Random forest: similar training/inference scaling but typically larger T (here 300) and unconstrained depth, yielding heavier inference (often 2–3$\times$ the XGBoost above) and a larger memory footprint.
- MLP (64–32): training $\approx \mathcal{O}(EnP)$ and inference $\approx \mathcal{O}(P)$ MACs per point (a few thousand operations), with regularization/early-stopping required to match ensemble accuracy.

Hybrid (two-ray + residual XGBoost):

Training $\approx \mathcal{O}(n)$ to compute the two-ray prior + boosted-trees on the residual $\varepsilon(d) = PL_{\text{meas}}(d) - PL_{\text{2-ray}}(d)$ with a compact XGBoost. Inference per point = two-ray $\mathcal{O}(1)$ + a short tree walk with $T = 150$ and depth $= 5$, which corresponds to $\sim$750 node tests per query. In practice, this is $\sim$1.5$\times$ lighter than the XGBoost trained only on real data (200×6) and $\gg$ lighter than RF, with a correspondingly smaller model (few thousand nodes). The synthetic augmentation used in this configuration (parameter pairs $(h_t, \varepsilon_r) \in \{1.0, 1.4\} \times \{5, 7\}$, weighted $w = 0.3$) adds negligible inference cost and only modest training overhead.

For large-scale evaluation grids, the deterministic two-ray model establishes the reference limit in terms of computational efficiency. The hybrid approach remains close to this limit since only a lightweight tree ensemble is added on top of the two-ray prediction while still reaching data-driven accuracy. In contrast, pure ML ensembles increase per-point latency and memory requirements.

3.7. Small-Scale Channel Statistics

To characterize the fine-grained variability left after large-scale filtering, three complementary analyses were performed, each tied to its corresponding experiment.

- Autocorrelation of EPL.
  The empirical spatial autocorrelation function $\rho(\Delta) = \text{corr}\big(\text{EPL}(x), \text{EPL}(x + \Delta)\big)$ was evaluated for lags $\Delta \leq 5$m ($\approx 300\lambda$). An exponential model $\rho(\Delta) = \exp(-\Delta/L_c)$ was fitted via non-linear least squares, yielding the decorrelation length L_c for each

corridor (Section 4.3.1). This metric determines the minimum spacing required for statistically independent samples.

- Spectral entropy.
  Using sliding windows of 25 cm with 50% overlap (N samples per window), the one-sided FFT power spectrum $\{P_k\}_{k=1}^{N/2}$ of the EPL was obtained. Local unpredictability was quantified as Shannon entropy

$$H = -\sum_k P_k \log P_k,$$

  with spectrum normalized so that $\sum_k P_k = 1$. The resulting entropy curve $H(x)$ highlights sections of the corridor with highly irregular fading versus smoother regions (Section 4.3.2).
- Cross-correlation between corridors.
  After interpolating both EPL series onto a common 5 cm grid, the normalized cross-correlation $\rho_\times(\Delta)$ was calculated via FFT convolution. The peak value $\rho_{\max}$ and its lag $\Delta_{\max}$ reveal the global similarity and any horizontal shift between the two "fading fingerprints" (Section 4.3.3).

4. Results and Discussion

4.1. Large-Scale Path-Loss Modeling

This section establishes the foundation for all the subsequent analyses by quantifying the large-scale path-loss behavior. A systematic calibration of the classical two-ray model was carried out in which antenna height and floor permittivity were treated as free parameters. The resulting fit and associated error surface define a physics-based reference that data-driven and hybrid approaches are later expected to outperform.

Two-Ray Calibration

In Figure 4, the total calibration error $\mathrm{RMSE}_{\mathrm{C1}} + \mathrm{RMSE}_{\mathrm{C2}}$ is shown as obtained from the full-factorial sweep of the two key physical parameters in the classical two-ray model: the effective antenna height ($h_t = h_r \in [0.8, 1.6]$ m) and the real part of the concrete permittivity ($\varepsilon_r \in [4, 7]$). The surface reveals a narrow well-defined valley whose floor extends approximately along $h_t = h_r \approx 1.2$ m for $\varepsilon_r \gtrsim 6.5$. An increase in error is observed when deviating from this ridge in either direction, confirming the strong sensitivity of the two-ray interference pattern to the electrical characteristics of the floor. The global minimum (marked in black) is found at $h^\star = 1.20$ m and $\varepsilon_r^\star = 7$, with a combined RMSE of 1.90 dB.

The corridor-specific results reported in Table 4 indicate that the deterministic fit was slightly better in corridor C2 (0.88 dB) than in corridor C1 (1.02 dB). This difference is attributed to a minor variation in the local floor surface, resulting in a smoother concrete finish that improved the strength of the specular reflection, as seen in Figure 2. Despite this detail, the optimal antenna height obtained in both corridors matches the nominal installation ($h_t = h_r = 1.20$ m), thus validating the experimental setup.

Figure 5 presents a comparison between the measured large-scale path loss, the calibrated two-ray prediction, and the ideal free-space loss (FSPL). Over the full 25 m range, the two-ray prediction remains within the shaded ± 1 dB corridor relative to the measured data, although clear oscillations persist. These residual ripples correspond to the constructive and destructive interference minima predicted by the analytical model; however, their amplitudes and positions show slight discrepancies due to floor roughness and secondary wall reflections—factors that are not captured by the simplified two-path formulation.

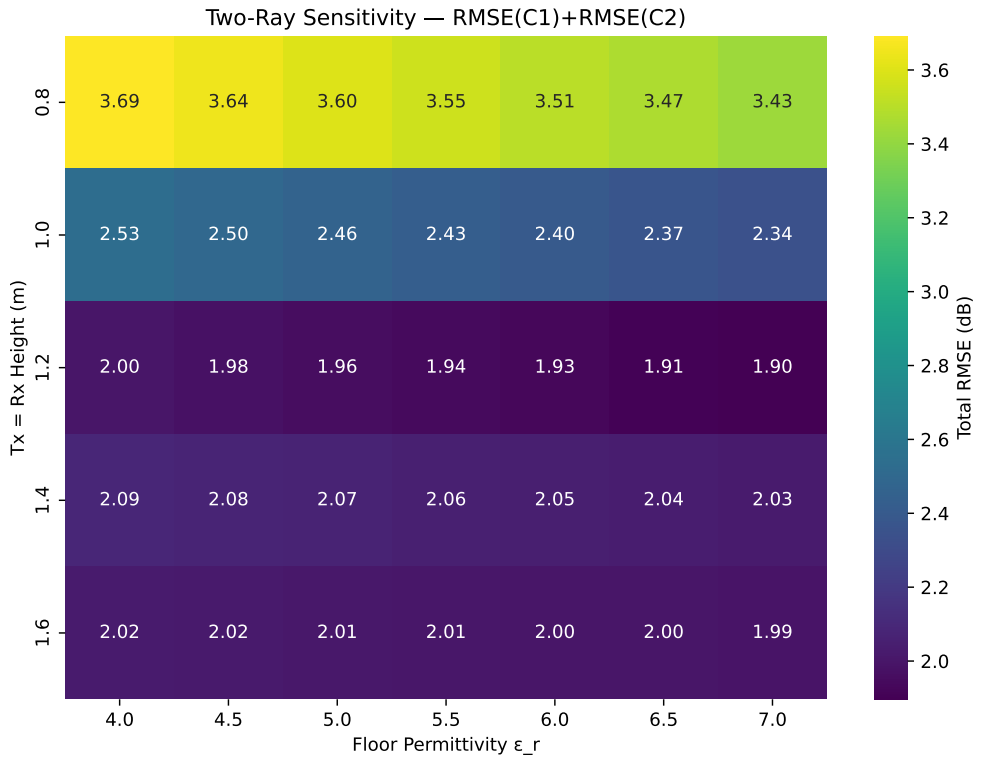


Figure 4. Total calibration error $\mathrm{RMSE}_{C1} + \mathrm{RMSE}_{C2}$ as a function of co-located antenna height $h_t = h_r$ and floor permittivity ε_r. The global minimum (black dot) occurs at $h^\star = 1.20$ m, $\varepsilon_r^\star = 7$, with $\mathrm{RMSE}_{\mathrm{tot}} = 1.90$ dB.

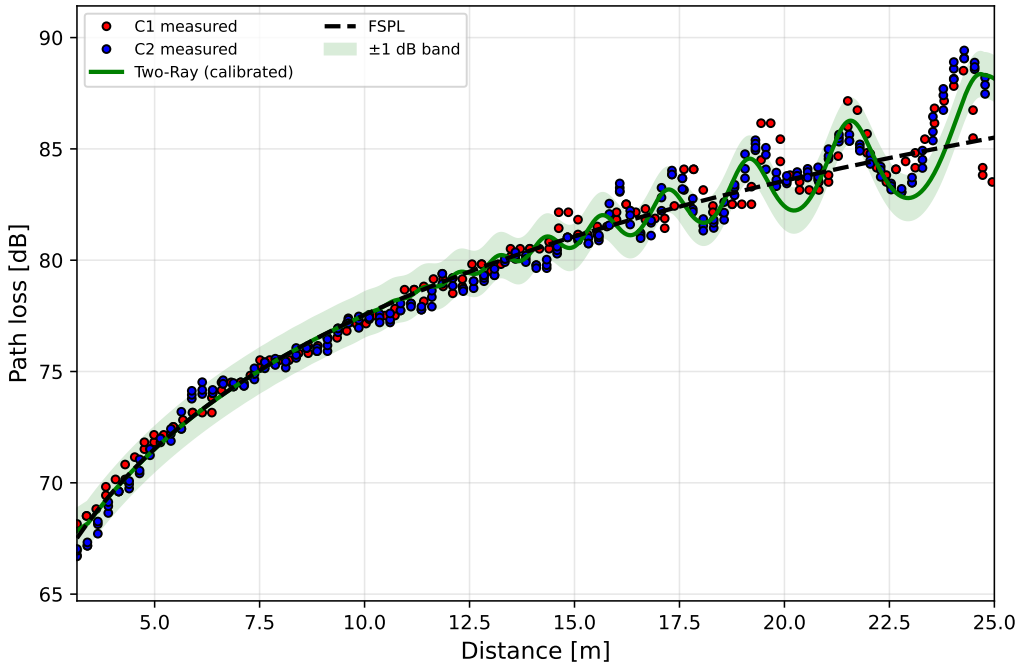


Figure 5. Measured large-scale path loss (points), calibrated two-ray prediction (solid green), and free-space loss (dashed). The shaded band denotes a ±1 dB corridor around the two-ray curve, illustrating that most measured samples remain within this margin except for residual oscillations due to floor roughness and higher-order reflections.

Table 4. Optimal physical parameters and deterministic fit quality (after 30λ averaging).

Corridor	$h^\star$ [m]	RMSE [dB]
C1	1.20	1.02
C2	1.20	0.88
Total	–	1.90

The calibration exercise serves two purposes. First, it provides a physically interpretable baseline whose error (1.90 dB total) is already lower than that reported for many empirical one-slope models at similar frequencies ($\sim$3–4 dB). The error surface shows the closed-form model is fragile: shifting antennas by $\pm$0.2 m or mis-setting ε_r by 10% can raise RMSE by $>$0.5 dB. That can negate the benefit of the physics-based model. These findings motivate the hybrid strategy adopted in the remainder of the paper. By keeping the analytically calibrated two-ray term fixed as a strong prior and allowing the ML residual to compensate for the systematic discrepancies, the interpretability and parameter efficiency of the deterministic model are expected to be preserved. At the same time, the fine-scale deviations introduced by multipath clutter and material uncertainty can be absorbed. This strategy ultimately aims to reduce the RMSE below the 1 dB threshold reported in Section 4.2.

4.2. Machine-Learning Performance

In this section, the predictive performance and generalization capability of several ML models are evaluated in the context of outdoor FR3 upper mid-band path-loss estimation. A purely data-driven scenario is first considered, in which models are trained on one corridor and tested on another to assess their ability to handle domain shift. Next, a residual-learning scheme is introduced, where physics-based priors are incorporated to enhance robustness and interpretability. Finally, an adaptive hybrid model is presented in which real measurements are complemented with synthetic data to improve generalization across moderate variations in geometry and material properties. Through these experiments, the impact of physics-informed learning and data augmentation is quantified in the development of reliable low-error outdoor propagation models.

4.2.1. Cross-Corridor Transfer

The aim of this experiment is to quantify how well a purely data-driven regressor copes with a domain-shift scenario in which the model is trained on all samples from one corridor and then evaluated on the other. Only two explanatory variables are provided—$\log_{10}$-distance and a binary corridor-id flag—so that any transfer ability must arise from the model's intrinsic bias rather than explicit knowledge of the test environment. This set-up emulates the practical case where a path-loss predictor calibrated in one hallway is re-used in a similar but previously unseen outdoor layout.

The tree ensembles generalize almost symmetrically. XGBoost attains an RMSE of 0.97 dB when transferred from C1 $\rightarrow$ C2 and 1.11 dB in the reverse direction, whereas random forest is marginally better at 0.93 dB and 1.03 dB, respectively. In both directions, the coefficient of determination remains above 0.95 and the scatter points cluster tightly around the 45° line (Figure 6). The MLP now performs on par with the ensembles, yielding 0.98 dB RMSE in both splits and maintaining $R^2 > 0.96$; see Figure 7.

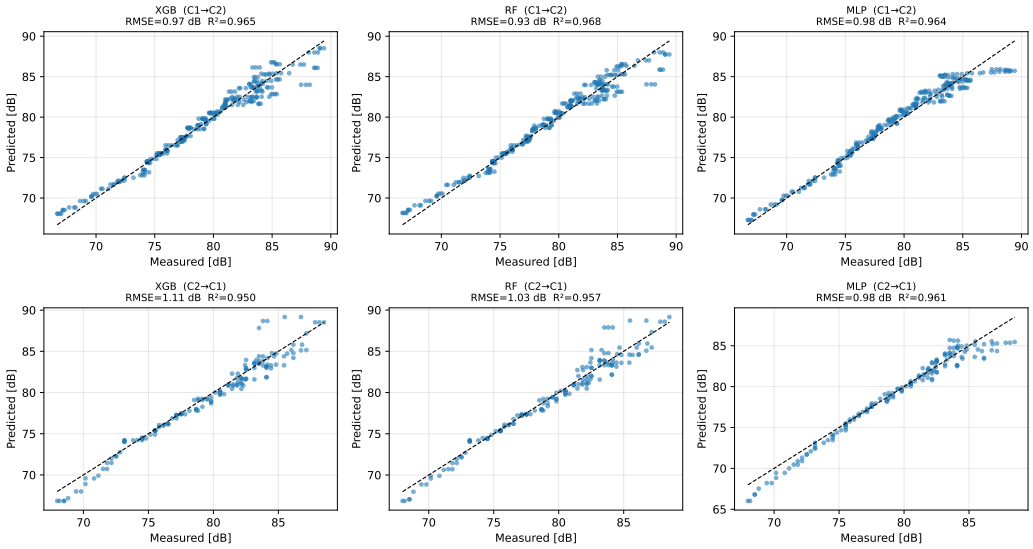


Figure 6. Cross-corridor scatter plots of predicted versus measured path loss. Each panel corresponds to a specific model (XGB, RF, or MLP) and training–testing direction (C1 → C2 in the first row; C2 → C1 in the second). The dashed diagonal marks the ideal $y = x$ line.

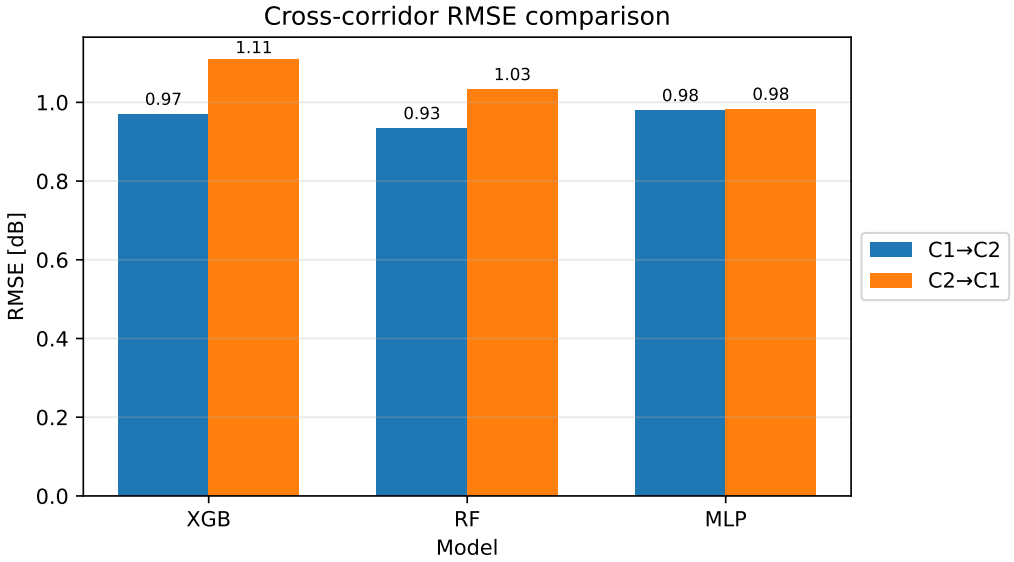


Figure 7. RMSE for the three learning models in the cross-corridor experiment. Each cluster corresponds to a training–testing direction (C1 → C2 and C2 → C1). Random forest attains the lowest RMSE; XGBoost is a close second, while the MLP now matches the ensembles in both directions.

The feature-importance scores in Figure 8 confirm that both trees rely almost exclusively on distance, effectively learning a corridor-agnostic propagation law. Two immediate conclusions follow. First, shallow ensembles (and, under careful tuning, a modest MLP) provide a robust bias–variance trade-off for small outdoor path-loss datasets, sustaining sub-1 dB accuracy under domain shift. Second, the negligible contribution of corridor-id indicates that, for corridors of identical geometry, distance alone captures the dominant deterministic trend; adding further categorical flags would inflate complexity without tangible gains. The residual ∼1 dB error band therefore motivates the physics-guided residual-learning strategy developed in the next section.

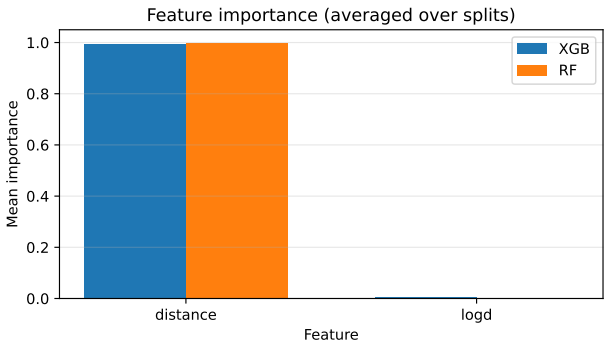


Figure 8. Average feature importance (over both cross-corridor splits) for the tree-based models. Both XGBoost and random forest place virtually all predictive weight on the logarithmic distance term, assigning the `corridor-id` flag a mean importance < 0.01.

The learning curves in Figure 9 show that, for XGB, the minimum test MSE occurs near 150 trees, while, for the MLP, the optimal point lies near 1300 epochs. In both cases, the training curves (dashed) continue to decrease beyond the optimum, indicating mild overfitting past this point, whereas the test curves (solid) flatten or begin to rise. These trends confirm that the selected hyperparameters correspond to near-optimal complexity under the cross-corridor setting, balancing bias and variance without excessive training.

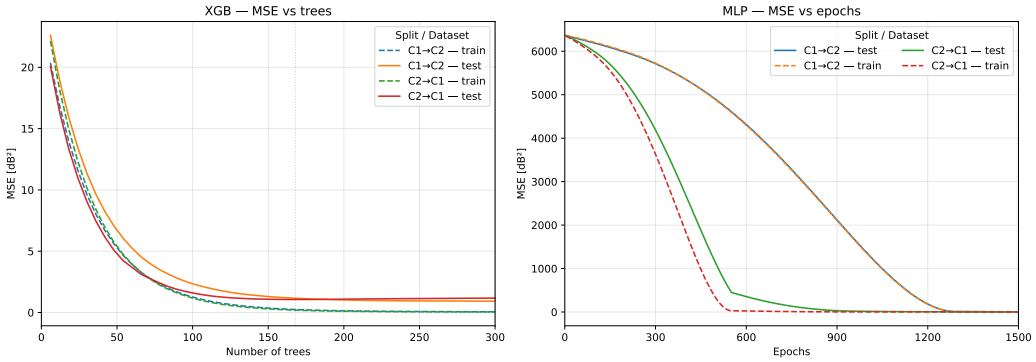
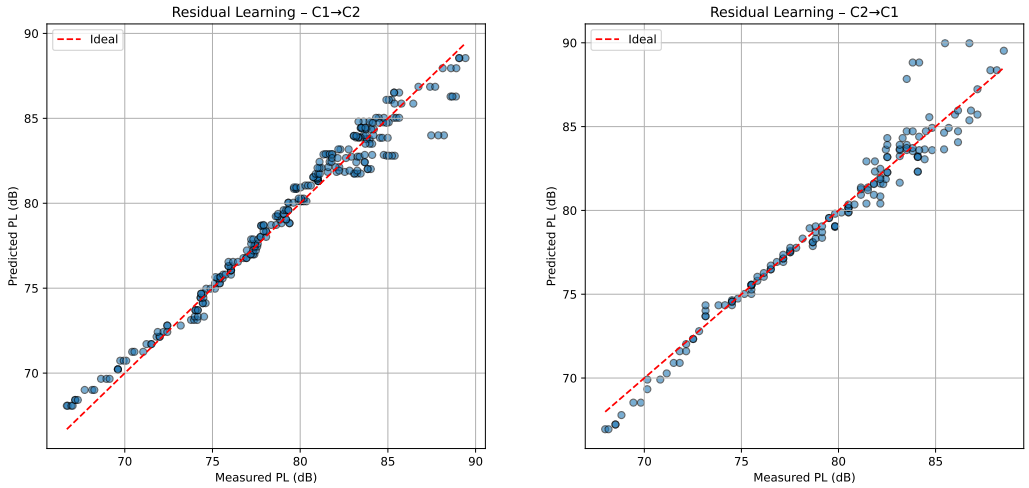


Figure 9. Learning curves for XGB (**left**) and MLP (**right**) models, showing mean squared error (MSE) for training (dashed) and test (solid) sets. Horizontal axes are rescaled to 0–300 trees (XGB) and 0–1500 epochs (MLP). Vertical dotted lines mark the minimum test MSE.

4.2.2. Residual-Learning Cross-Transfer

To visualize how the residual-learning scheme performs under the same cross-corridor split as in the previous subsection, predictions are juxtaposed for the cases where the model is (i) trained on Corridor 1 and tested on Corridor 2, and (ii) trained on Corridor 2 and tested on Corridor 1. Figure 10 presents both scatter plots side by side to enable direct comparison.

For the C1 $\rightarrow$ C2 transfer, the hybrid model attains 0.59 dB RMSE and an R^2 of 0.997; in the opposite direction, the error remains essentially identical (0.62 dB, $R^2 = 0.996$). These values represent almost a 40% reduction with respect to the best pure-ML ensemble and a three-fold improvement over the calibrated two-ray baseline. Because the residual learner is trained on the difference $PL_{\text{meas}} - PL_{\text{2-ray}}$, it only needs to model the fine-scale discrepancies that the deterministic term cannot capture—chiefly the frequency of constructive/destructive fringes and second-order wall reflections. As these effects are weakly dependent on which corridor is used, the residual function generalizes almost perfectly across the two layouts.



(**a**) (**b**)

Figure 10. Hybrid residual-learning performance under cross-corridor transfer. In both directions, the predictions lie close to the $y = x$ bisector, removing nearly all systematic bias. (**a**) Model trained on Corridor 1, tested on Corridor 2. (**b**) Model trained on Corridor 2, tested on Corridor 1.

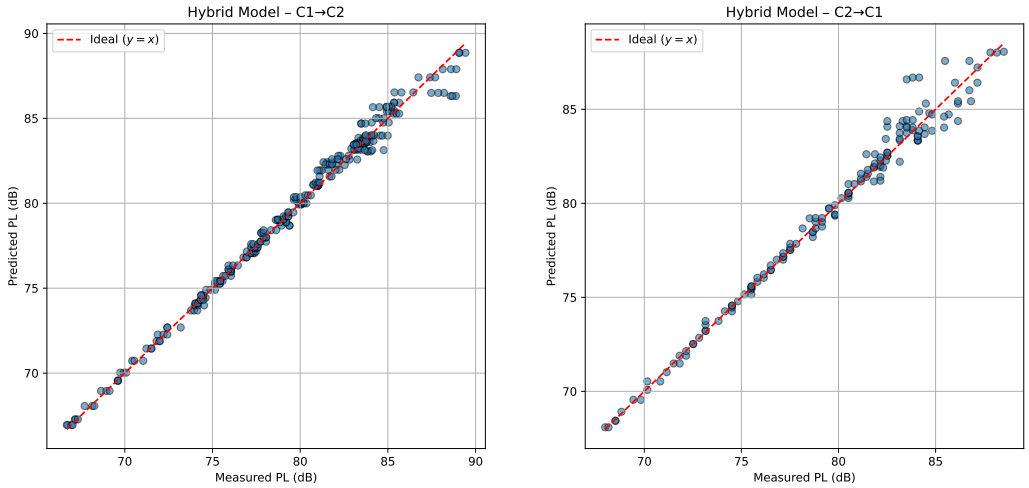
The experiment confirms two advantages of the physics-guided residual approach. First, it inherits the inherent interpretability of the two-ray term: large-scale attenuation and the LOS–floor interaction remain explicit. Second, by constraining the data-driven component to small corrections, the method achieves higher data-efficiency and robustness than a free-standing regressor—evidenced by the near-identical RMSE in both transfer directions. The residual hybrid thus provides a reliable sub-dB predictor that is portable across corridors without retraining, a desirable property for rapid outdoor network design.

4.2.3. Adaptive Hybrid with Synthetic Augmentation

This experiment evaluates a physics-guided residual learner whose input features combine (i) the calibrated two-ray prediction $PL_{\text{2-ray}}$ and (ii) a small set of explanatory variables ($\log_{10} d$, antenna height $h_t = h_r$, concrete permittivity ε_r, and a corridor flag). The model is trained to predict the residual $\varepsilon(d) = PL_{\text{meas}}(d) - PL_{\text{2-ray}}(d)$. To improve robustness, the real measurements are complemented with $\sim 4\times$ as many synthetic samples generated by perturbing h_t and ε_r within realistic bounds (± 20 cm, 5–7). Synthetic rows receive a weight of 0.3 so that they inform but do not dominate the fit.

When trained on Corridor 1 and tested on Corridor 2, the hybrid model achieves 0.59 dB RMSE ($R^2 = 0.997$); the reverse split yields 0.62 dB RMSE ($R^2 = 0.996$). Five-fold cross-validation on the combined real dataset confirms the stability of the approach, with an average error of 0.49 ± 0.11 dB. Figure 11 shows that the predicted points lie almost perfectly on the $y = x$ bisector in both transfer directions.

Two observations stand out. First, augmenting the training set with lightly weighted synthetic variations eliminates the small performance gap that remained after pure residual learning: accuracy is now consistently sub-0.6 dB—roughly a 40% improvement over the best corridor–transfer ensemble and a three-fold gain relative to the calibrated two-ray. Second, the near-identical error in both directions shows that the residual function generalizes across modest changes in height and permittivity, indicating that the model has captured the underlying physics rather than memorizing corridor-specific artifacts. These results underscore the practical benefit of the hybrid strategy: it retains interpretability, requires only a few dozen real measurements, and delivers record-level accuracy that is portable across similar outdoor environments without retraining.



(**a**) (**b**)

Figure 11. Hybrid two-ray + XGBoost performance under cross-corridor transfer. In both directions, the residual learner removes nearly all bias left by the deterministic baseline, keeping the scatter within a ± 1 dB corridor around the ideal $y = x$ line. (**a**) Model trained on Corridor 1, tested on Corridor 2. (**b**) Model trained on Corridor 2, tested on Corridor 1.

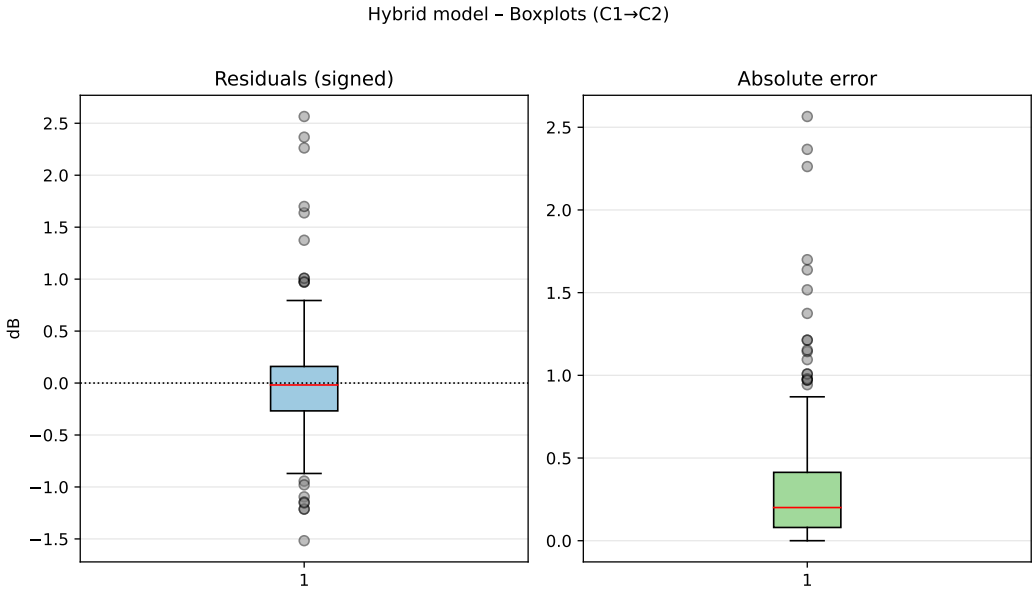
To complement these accuracy metrics, we also quantify prediction uncertainty. The idea is simple: residuals are gathered from the held-out set, their spread is summarized by the RMSE, and a two-sided 95% prediction interval is obtained as $\hat{y} \pm 1.96\text{RMSE}$. This approach assumes residuals are approximately Gaussian and unbiased, which is consistent with the empirical distributions observed. Figure 12 presents the residual and absolute-error boxplots for the two transfer directions. For C1 $\rightarrow$ C2, the RMSE of 0.59 dB implies a confidence band of ± 1.16 dB, while, for C2 $\rightarrow$ C1, the 0.62 dB RMSE translates to ± 1.22 dB. In both cases, the interquartile ranges remain narrow (0.43 and 0.41 dB), and whiskers seldom exceed ± 1 dB. These results show that the hybrid model not only achieves sub-dB accuracy on average but also maintains tight error bounds, providing practical confidence margins for real-world deployment.

4.3. Small-Scale Channel Statistics

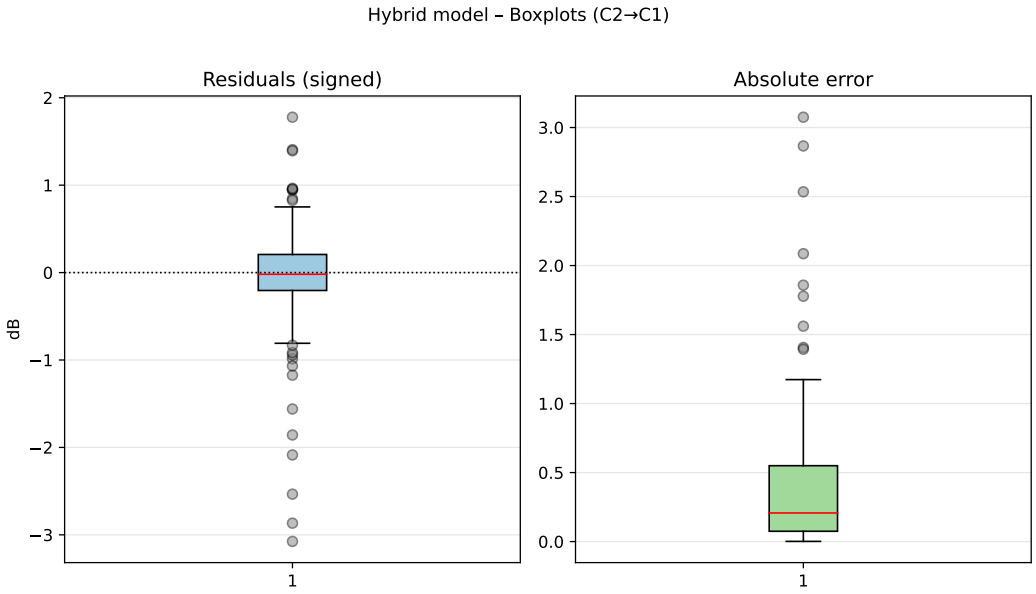
In this section, small-scale fading characteristics of the FR3 upper mid-band outdoor channel are examined through spatial correlation analysis. The decorrelation length is first estimated by fitting an exponential model to the empirical autocorrelation of the excess path loss, providing insight into the coherence distance relevant for measurement design and stochastic simulation. Subsequently, the spatial cross-correlation between the two corridor environments is analyzed to quantify the structural similarity of their fading profiles. These analyses are intended to support the interpretation of residual model behavior and to guide the construction of statistically independent training datasets.

4.3.1. Spatial Autocorrelation of Excess Path Loss

While the large-scale trend is captured by the two-ray baseline, a reliable estimate of the coherence distance—defined as the spacing at which local fading samples become statistically independent—is also required for link–budget design. For this purpose, the empirical spatial autocorrelation function (ACF) of the excess-path-loss series ($\text{EPL}(d) = \overline{PL}(d) - PL_{\text{FSPL}}(d)$) is computed, and the classical exponential model $\rho(\Delta) = \exp(-\Delta / L_c)$ is fitted, where L_c denotes the decorrelation length.



(**a**)



(**b**)

Figure 12. Residual and absolute-error boxplots for cross-corridor transfer. Approximate 95% prediction bounds are computed as $\hat{y} \pm 1.96\text{RMSE}$, while the empirical spread is summarized with the interquartile range (box) and whiskers. (**a**) Adaptive Hybrid C1 $\rightarrow$ C2. RMSE = 0.59 dB ($R^2 = 0.997$). 95% prediction interval (Normal approx.): $\pm 1.96 \times 0.59 = \pm 1.16$ dB around $\hat{y}$. Residual IQR $[-0.27, 0.16]$ dB (IQR $= 0.43$); whiskers $[-0.91, 0.80]$ dB. (**b**) Adaptive Hybrid C2 $\rightarrow$ C1. RMSE = 0.62 dB ($R^2 = 0.996$). 95% prediction interval (Normal approx.): $\pm 1.96 \times 0.62 = \pm 1.22$ dB around $\hat{y}$. Residual IQR $[-0.20, 0.21]$ dB (IQR $= 0.41$); whiskers $[-0.82, 0.83]$ dB.

Figure 13 superimposes the measured ACFs for corridors C1 and C2 together with their exponential fits. The estimated decorrelation lengths are $L_c^{\text{C1}} \approx 0.23$ m and $L_c^{\text{C2}} \approx 0.41$ m. Both curves cross the $1/e$ reference line within 0.5 m, indicating that samples separated by more than three to four wavelengths at 18 GHz can be considered quasi-independent.

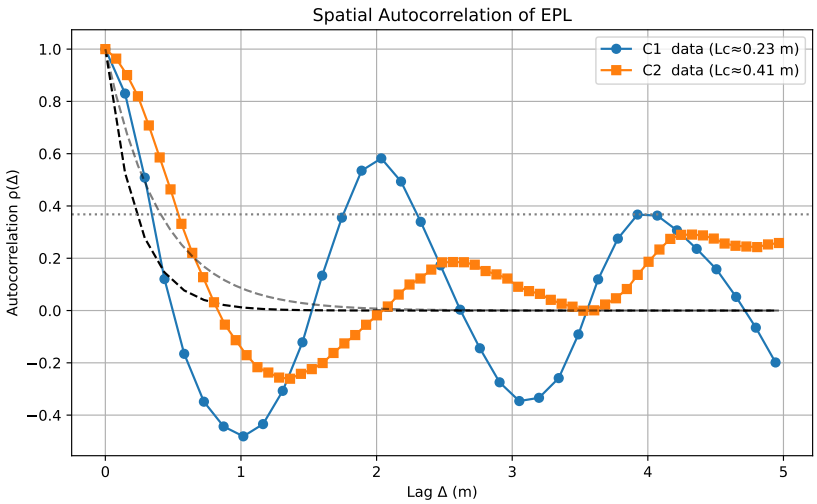


Figure 13. Empirical spatial autocorrelation of the excess path loss (symbols); black dashed curves show the exponential fits for each corridor. The gray dotted horizontal line marks the $1/e$ reference; the intersections with the fits yield decorrelation lengths of $L_c^{\mathrm{C1}} \approx 0.23\,\mathrm{m}$ and $L_c^{\mathrm{C2}} \approx 0.41\,\mathrm{m}$.

Corridor C1 decorrelates almost twice as fast as C2, a difference that aligns with its coarser wall finish and larger number of open doorways, both of which increase multipath richness at centimeter scales. In contrast, C2's smoother concrete surfaces produce a more persistent dominant specular component, extending L_c to roughly 0.4 m. Practically, these values set the minimum antenna-movement step for independent field measurements and can be used to parameterize stochastic channel simulators. The short coherence distance also helps to explain why the residual learner in this experiment benefited from dense synthetic sampling: even modest spatial perturbations provided new weakly correlated information to the model.

4.3.2. Local Spectral Entropy of Excess Path Loss

While autocorrelation quantifies coherence distance, spectral entropy provides a complementary view of how rich or structured the local multipath spectrum is along the corridor. A low-entropy window indicates a narrowband quasi-periodic interference pattern dominated by a few spatial frequencies, whereas high entropy signals a broadband noise-like residual with no obvious dominant tones. We compute the Shannon entropy of the one-sided FFT magnitude in 25-centimeter sliding windows with 50% overlap—sufficient to encompass several wavelengths at 18 GHz while retaining fine spatial resolution.

Figure 14 contrasts the spectral entropy before (raw EPL) and after (hybrid residual) the hybrid correction for each corridor (top: C1; bottom: C2). Across both layouts the two curves largely overlap—indicating that most EPL energy is already spread over a wide band—yet the hybrid trace slightly raises the lowest entropy valleys, which is consistent with a modest whitening associated with removing residual quasi-periodic structure from the two-ray fit. These observations agree with the strong similarity revealed by the cross-correlation analysis (Figure 15) and with the short decorrelation lengths in Figure 13. Taken together, once deterministic bias is removed, the residual channel is effectively broadband and statistically homogeneous, which is favorable for robust machine-learning prediction and for stochastic channel simulators.

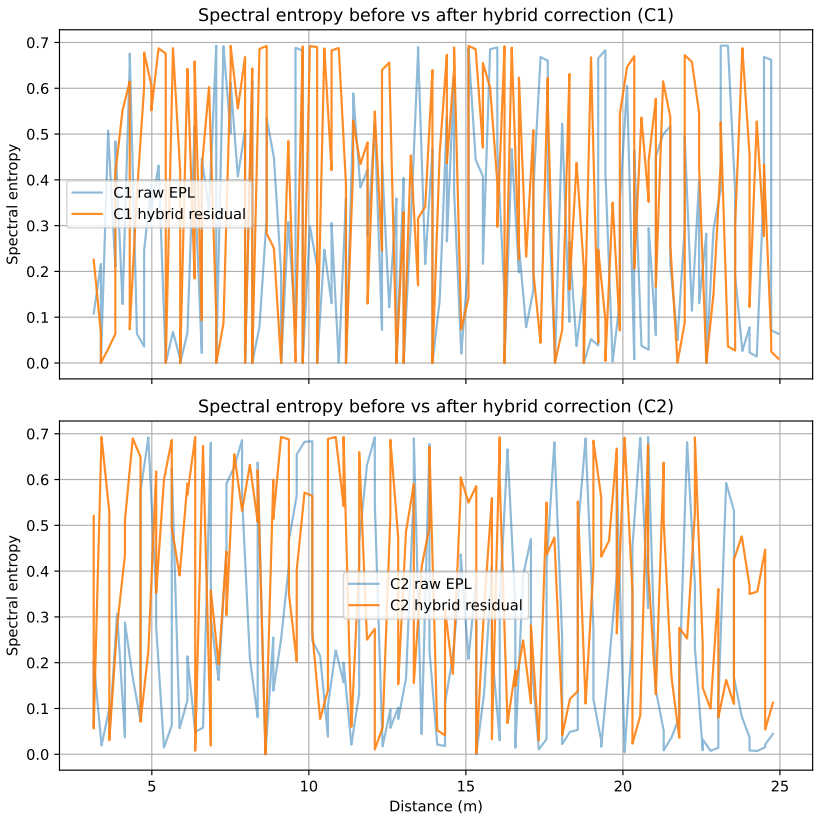


Figure 14. Local spectral entropy before (raw EPL) and after hybrid correction (hybrid residual) in 25 cm windows (50% overlap). Top: C1; bottom: C2. A slight lifting of the lowest entropy valleys after the hybrid step indicates modest whitening, while sustained plateaus denote broadband multipath regions.

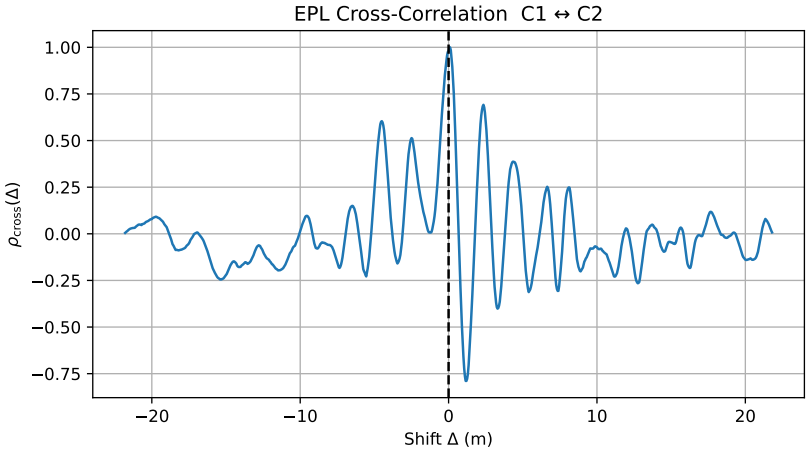


Figure 15. Spatial cross-correlation between the excess-path-loss sequences of corridors C1 and C2. The peak $\rho_{max} = 1.00$ occurs at a shift of $\Delta_{max} = 0.10$ m.

4.3.3. Cross-Correlation Between Corridors

After characterizing the autocorrelation within each corridor, we examine how similar the two fading fingerprints are to one another. The excess-path-loss traces of C1 and C2 were first resampled on a common 5 cm grid and then cross-correlated in the spatial domain. Shifting one trace against the other reveals whether major fading features coincide and quantifies any horizontal displacement.

The normalized cross-correlation curve $\rho_{\times}(\Delta)$ is shown in Figure 15. A single sharp peak reaches the maximum possible value $\rho_{max} = 1.00$ at a shift of $\Delta_{max} = +0.10$ m. No

secondary lobes exceed 0.25, indicating that the two profiles align almost perfectly once the 10 cm offset is applied.

A unit-height peak implies that the large-scale fading structures—including the periodic minima introduced by two-ray interference—are effectively identical in both corridors. The small 10 cm lag can be traced to the manual alignment of the measurement starting points and falls below the decorrelation lengths previously estimated. From a modeling perspective, this result justifies pooling the two datasets when training the residual learner: once the traces are aligned, they provide redundant information rather than conflicting patterns.

5. Conclusions

This study shows that a physics-enhanced residual-learning scheme can push outdoor 18 GHz path-loss prediction below the one-decibel barrier without sacrificing interpretability or requiring labor-intensive ray tracing. By fusing a calibrated two-ray baseline with a light XGBoost residual learner and augmenting the training set with low-weight synthetic variations, the hybrid model (i) inherits the deterministic term's physical transparency, (ii) captures corridor-specific fine-scale effects, and (iii) delivers cross-corridor generalization that pure ML or pure physics cannot match. Consistent sub-0.6 dB RMSE, unit-height cross-correlation peaks, and decorrelation lengths on the order of a few wavelengths collectively indicate that the residual learner has absorbed virtually all the systematic variability while leaving only white-noise residuals.

A practical implication is that network planners can achieve near-deterministic accuracy with a few tens of measurement points rather than hundreds, cutting survey time and enabling rapid what-if analyses (e.g., antenna-height changes) via the synthetic-augmentation trick. The approach also offers a clear upgrade path for legacy empirical models: the two-ray prior can be replaced with any site-specific deterministic engine, and the residual learner will adapt accordingly.

Regarding the applicability of the model to other FR3 environments, its performance depends on the dominant propagation physics. In open squares, the two-ray prior remains relevant, but the residual learner must handle increased scattering. For more complex scenarios such as urban canyons, the two-ray prior could be insufficient and should be replaced by a more suitable one, such as a simplified ray-tracing engine, while preserving the hybrid framework. Similarly, semi-indoor NLOS conditions would require physical priors that account for diffraction and penetration.

The current campaign was confined to straight LOS corridors with identical cross-sections, so the model's robustness against sharp bends, NLOS segments, or composite materials remains untested. Likewise, the two-ray assumption may break down in large atria or industrial halls where ceiling reflections dominate. Our next step is therefore two-fold:

- Diverse outdoor geometries. We will repeat the measurement-and-training pipeline in L-shaped corridors, open-plan offices, and multi-story atria to quantify how residual complexity grows with geometric diversity and to verify whether a single residual learner can be shared across multiple deterministic priors.
- Semi-open and outdoor transitions. Preliminary simulations suggest that the same residual framework, coupled with a single-slope free-space prior, could handle short-range outdoor hotspots (e.g., stadium concourses and campus walkways) where ground reflections dominate but clutter is sparse. We plan to collect 18 GHz data in an open field with sparse obstacles and test whether a hybrid model can maintain sub-2 dB RMSE while requiring far fewer rays than full-blown outdoor ray tracing.

Author Contributions: Conceptualization, J.G.; methodology, J.G.; software, J.C.-M. and J.R.-V.; validation, J.C.-M., J.R.-V. and M.D.-M.; formal analysis, J.C.-M.; investigation (data acquisition), J.C.-M., J.R.-V. and M.D.-M.; resources, J.G.; data curation, J.C.-M., J.R.-V. and M.D.-M.; writing—original draft preparation, J.G. and M.D.-M.; writing—review and editing, A.P. and J.G.; visualization, J.C.-M.; supervision, J.G.; project administration, J.G.; funding acquisition, J.G. and A.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the VINCI-DI Initiative of the Pontificia Universidad Católica de Valparaíso (PUCV) under project No. 039.706/2025. Additional support was provided by the National Agency for Research and Development (ANID) through FONDECYT, Grant No. 11240070, and the ANID Basal Project, AFB240002 (AC3E).

Data Availability Statement: The measurement campaign datasets and the Python code used for model training and evaluation are available from the corresponding authors (J.C.M. and J.G.) on request.

Acknowledgments: The authors would like to express their sincere gratitude to the Pontificia Universidad Católica de Valparaíso (PUCV) for the support provided through the VINCI-DI Initiative (Project No. 039.706/2025). The authors also acknowledge the support of the National Agency for Research and Development (ANID) through FONDECYT Grant No. 11240070 and the ANID Basal Project AFB240002 (AC3E). In addition, the authors acknowledge the support of the Doctorate in Smart Industry program of the Pontificia Universidad Católica de Valparaíso, Chile.

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

Abbreviations

The following acronyms are used in this manuscript:

FR3	Frequency Range 3 (upper mid-band, ∼7–24 GHz)
LOS	Line-of-Sight
NLOS	Non-Line-of-Sight
PL	Path Loss
FSPL	Free-Space Path Loss
EPL	Excess Path Loss
RMSE	Root Mean Squared Error
R^2	Coefficient of Determination
ML	Machine Learning
XGB	Extreme Gradient Boosting (XGBoost)
RF	Random Forest
MLP	Multi-Layer Perceptron
HPO	Hyperparameter Optimization
CV	Cross-Validation
CW	Continuous Wave
HPBW	Half-Power Beamwidth
Tx	Transmitter
Rx	Receiver
FFT	Fast Fourier Transform
ACF	Autocorrelation Function
dB	Decibel (power ratio)
dBi	Antenna gain referenced to isotropic
dBm	Power referenced to 1 mW
PERL	Physics-Enhanced Residual Learning

References

1. Bazzi, A.; Bomfin, R.; Mezzavilla, M.; Rangan, S.; Rappaport, T.; Chafii, M. Upper mid-band spectrum for 6G: Vision, opportunity and challenges. *arXiv* **2025**, arXiv:2502.17914.
2. Cui, Z.; Zhang, P.; Pollin, S. 6G wireless communications in 7–24 GHz band: Opportunities, techniques, and challenges. *arXiv* **2023**, arXiv:2310.06425.
3. Shakya, D.; Ying, M.; Rappaport, T.S.; Poddar, H.; Ma, P.; Wang, Y.; Al-Wazani, I. Comprehensive FR1 (C) and FR3 lower and upper mid-band propagation and material penetration loss measurements and channel models in indoor environment for 5G and 6G. *IEEE Open J. Commun. Soc.* **2024**, *5*, 5192–5218. [CrossRef]
4. Deng, S.; Samimi, M.K.; Rappaport, T.S. 28 GHz and 73 GHz Millimeter-Wave Indoor Propagation Measurements and Path-Loss Models. In Proceedings of the IEEE International Conference on Communications Workshops (ICC Wkshp.), London, UK, 8–12 June 2015; pp. 1246–1250. [CrossRef]
5. Shen, Y.; Shao, Y.; Xi, L.; Zhang, H.; Zhang, J. Millimeter-Wave Propagation Measurement and Modeling in Indoor Corridor and Stairwell at 26 and 38 GHz. *IEEE Access* **2021**, *9*, 87792–87805. [CrossRef]
6. Oladimeji, T.T.; Kumar, P.; Elmezughi, M.K. Path-Loss Measurements and Model Analysis in an Indoor Corridor Environment at 28 GHz and 38 GHz. *Sensors* **2022**, *22*, 7642. [CrossRef] [PubMed]
7. Hu, Y.; Yin, M.; Mezzavilla, M.; Guo, H.; Rangan, S. Channel Modeling for FR3 Upper Mid-band via Generative Adversarial Networks. In Proceedings of the 2024 IEEE 25th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC), Lucca, Italy, 10–13 September 2024; IEEE: Lucca, Italy, 2024; pp. 776–780.
8. Cavalcanti, B.J.; Cavalcante, G.A.; Mendonça, L.M.d.; Cantanhede, G.M.; Oliveira, M.M.d.; D'Assunção, A.G. A hybrid path loss prediction model based on artificial neural networks using empirical models for LTE and LTE-A at 800 MHz and 2600 MHz. *J. Microwaves Optoelectron. Electromagn. Appl.* **2017**, *16*, 708–722. [CrossRef]
9. Lecci, M.; Testolina, P.; Giordani, M.; Polese, M.; Ropitault, T.; Gentile, C.; Varshney, N.; Bodi, A.; Zorzi, M. Simplified ray tracing for the millimeter wave channel: A performance evaluation. In Proceedings of the 2020 Information Theory and Applications Workshop (ITA), San Diego, CA, USA, 2–7 February 2020; IEEE: San Diego, CA, USA, 2020; pp. 1–6.
10. Kristem, V.; Bas, C.U.; Wang, R.; Molisch, A.F. Outdoor wideband channel measurements and modeling in the 3–18 GHz band. *IEEE Trans. Wirel. Commun.* **2018**, *17*, 4620–4633. [CrossRef]
11. Elmezughi, M.K.; Salih, O.; Afullo, T.J.; Duffy, K.J. Comparative analysis of major machine-learning-based path loss models for enclosed indoor channels. *Sensors* **2022**, *22*, 4967. [CrossRef]
12. Moraitis, N.; Rogaris, A.; Popescu, I.; Nikita, K.S. Measurements and Channel Characterization in Indoor Environments in the Upper Mid-Band. *IEEE Wirel. Commun. Lett.* **2025**, *14*, 1758–1762. [CrossRef]
13. Hu, J.; Al-Jzari, A.; He, Y.; Salous, S. FR3 Radio Propagation Channel Measurements and Modelling in Outdoor Environment for 5G and 6G Wireless Networks. In Proceedings of the 2025 19th European Conference on Antennas and Propagation (EuCAP), Stockholm, Sweden, 30 March–4 April 2025; IEEE: Stockholm, Sweden, 2025; pp. 1–5.
14. Beucler, T.; Pritchard, M.; Rasp, S.; Ott, J.; Baldi, P.; Gentine, P. Enforcing analytic constraints in neural networks emulating physical systems. *Phys. Rev. Lett.* **2021**, *126*, 098302. [CrossRef]
15. Wang, R.; Kashinath, K.; Mustafa, M.; Albert, A.; Yu, R. Towards physics-informed deep learning for turbulent flow prediction. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Virtually, 6–10 July 2020; pp. 1457–1466.
16. Kochkov, D.; Smith, J.A.; Alieva, A.; Wang, Q.; Brenner, M.P.; Hoyer, S. Machine learning–accelerated computational fluid dynamics. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, e2101784118. [CrossRef]
17. Jha, D.; Gupta, V.; Ward, L.; Yang, Z.; Wolverton, C.; Foster, I.; Liao, W.k.; Choudhary, A.; Agrawal, A. Enabling deeper learning on big data for materials informatics applications. *Sci. Rep.* **2021**, *11*, 4244. [CrossRef]
18. Zobeiry, N.; Joseph, A.; Eskandariyun, A.; Fu, H.; Brunton, S.; Alan, B.; Kabir, M.; Stere, A.; DePauw, T.C.; Fomin, S.; et al. Hybrid Physics and Machine Learning Modeling for Material Characterization and Failure Analysis. In Proceedings of the AIAA SCITECH 2024 Forum, Orlando, FL, USA, 8–12 January 2024; p. 1526.
19. Zhang, Y.; Wen, J.; Yang, G.; He, Z.; Wang, J. Path Loss Prediction Based on Machine Learning: Principle, Method, and Data Expansion. *Appl. Sci.* **2019**, *9*, 1908. [CrossRef]
20. MacCartney, G.; Rappaport, T.; Sun, S.; Deng, S. Indoor Office Wideband Millimeter-Wave Propagation Measurements and Channel Models at 28 and 73 GHz for Ultra-Dense 5G Wireless Networks. *IEEE Access* **2015**, *3*, 2388–2424. [CrossRef]
21. Zakeri, H.; Khoddami, P.; Moradi, G.; Alibakhshikenari, M.; Abd-Alhameed, R.; Koziel, S.; Dalarsson, M. Path Loss Model Estimation at Indoor Offices Environment by Using Deep Neural Network and CatBoost for Millimeter Wave 5G Wireless Application. *IEEE Access* **2024**, *12*, 159070–159085. [CrossRef]
22. Elmezughi, M.K.; Afullo, T.J.; Oyie, N.O. Performance Study of Path Loss Models at 14, 18, and 22 GHz in an Indoor Corridor Environment for Wireless Communications. *SAIEE Afr. Res. J.* **2021**, *112*, 32–45. [CrossRef]

23. Nuñez, Y.E.; Lovisolo, L.; da Silva Mello, L.A.R.; Orihuela, C. On the Interpretability of Machine Learning Regression for Path-Loss Prediction of Millimeter-Wave Links. *Expert Syst. Appl.* **2023**, *215*, 119324. [CrossRef]
24. Khalili, H.; Frey, H.; Wimmer, M.A. Balancing Prediction Accuracy and Explanation Power of Path Loss Modeling in a University Campus Environment via Explainable AI. *Future Internet* **2025**, *17*, 155. [CrossRef]
25. Limmer, S.; Alba, A.M.; Michailow, N. Physics-Informed Neural Networks for Pathloss Prediction. In Proceedings of the IEEE International Workshop on Machine Learning for Signal Processing (MLSP), Rome, Italy, 17–20 September 2023; pp. 1–6. [CrossRef]
26. Wen, J.; Zhang, Y.; Yang, G.; He, Z.; Zhang, W. Path Loss Prediction Based on Machine Learning Methods for Aircraft Cabin Environments. *IEEE Access* **2019**, *7*, 159251–159261. [CrossRef]
27. Moraitis, N.; Tsipi, L.; Vouyioukas, D.; Gkioni, A.; Louvros, S. Performance Evaluation of Machine Learning Methods for Path Loss Prediction in Rural Environment at 3.7 GHz. *Wirel. Netw.* **2021**, *27*, 1895–1910. [CrossRef]
28. Seretis, A.; Sarris, C.D. Toward Physics-Based Generalizable Convolutional Neural Network Models for Indoor Propagation. *IEEE Trans. Antennas Propag.* **2022**, *70*, 4112–4126. [CrossRef]
29. Marey, A.; Bal, M.; Ates, H.F.; Gunturk, B.K. PL-GAN: Path Loss Prediction Using Generative Adversarial Networks. *IEEE Access* **2022**, *10*, 90474–90480. [CrossRef]
30. Ates, H.F.; Hashir, S.M.; Baykas, T.; Gunturk, B.K. Path Loss Exponent and Shadowing Factor Prediction From Satellite Images Using Deep Learning. *IEEE Access* **2019**, *7*, 101366–101375. [CrossRef]
31. Bal, M.; Marey, A.; Ates, H.F.; Baykas, T.; Gunturk, B.K. Regression of Large-Scale Path Loss Parameters Using Deep Neural Networks. *IEEE Antennas Wirel. Propag. Lett.* **2022**, *21*, 1562–1566. [CrossRef]
32. Juang, R.T. Deep Learning-Based Path Loss Model in Urban Environments Using Image-to-Image Translation. *IEEE Trans. Antennas Propag.* **2022**, *70*, 12081–12091. [CrossRef]
33. Ahmadien, O.; Ates, H.F.; Baykas, T.; Gunturk, B.K. Predicting Path Loss Distribution of an Area From Satellite Images Using Deep Learning. *IEEE Access* **2020**, *8*, 64982–64991. [CrossRef]
34. Nguyen, T.T.; Yoza-Mitsuishi, N.; Caromi, R. Deep Learning for Path Loss Prediction at 7 GHz in Urban Environment. *IEEE Access* **2023**, *11*, 33498–33508. [CrossRef]
35. Kayaalp, K.; Metlek, S.; Genc, A. Prediction of path loss in coastal and vegetative environments with deep learning at 5G sub-6 GHz. *Wirel. Netw.* **2023**, *29*, 2471–2480. [CrossRef]
36. Ethier, J.; Châteauvert, M. Machine learning-based path loss modeling with simplified features. *IEEE Antennas Wirel. Propag. Lett.* **2024**, *23*, 3997–4001. [CrossRef]
37. Rafie, I.F.M.; Lim, S.Y.; Chung, M.J.H. Path Loss Prediction in Urban Areas: A Machine Learning Approach. *IEEE Antennas Wirel. Propag. Lett.* **2023**, *22*, 809–813. [CrossRef]
38. Cheng, H.; Ma, S.; Lee, H.; Cho, M. Millimeter Wave Path Loss Modeling for 5G Communications Using Deep Learning with Dilated Convolution and Attention. *IEEE Access* **2021**, *9*, 62867–62879. [CrossRef]
39. Cheng, H.; Ma, S.; Lee, H. CNN-Based mmWave Path Loss Modeling for Fixed Wireless Access in Suburban Scenarios. *IEEE Antennas Wirel. Propag. Lett.* **2020**, *19*, 1694–1698. [CrossRef]
40. Jin, W.; Kim, H.; Lee, H. A Novel Machine Learning Scheme for mmWave Path Loss Modeling for 5G Communications in Dense Urban Scenarios. *Electronics* **2022**, *11*, 1809. [CrossRef]
41. Fu, Z.; Du, F.; Zhao, X.; Geng, S.; Zhang, Y.; Qin, P. A Joint-Neural-Network-Based Channel Prediction for Millimeter-Wave Mobile Communications. *IEEE Antennas Wirel. Propag. Lett.* **2023**, *22*, 1064–1068. [CrossRef]
42. Nguyen, C.; Cheema, A.A. A Deep Neural Network-Based Multi-Frequency Path Loss Prediction Model from 0.8 GHz to 70 GHz. *Sensors* **2021**, *21*, 5100. [CrossRef]
43. Qiu, K.; Bakirtzis, S.; Song, H.; Zhang, J.; Wassell, I. Pseudo Ray-Tracing: Deep Leaning Assisted Outdoor mm-Wave Path Loss Prediction. *IEEE Wirel. Commun. Lett.* **2022**, *11*, 1699–1702. [CrossRef]
44. Lee, W.C. Estimate of local average power of a mobile radio signal. *IEEE Trans. Veh. Technol.* **2006**, *34*, 22–27. [CrossRef]
45. Espineira, P.M. *Modeling the Wireless Propagation Channel: A Simulation Approach with MATLAB*; Wiley: Hoboken, NJ, USA, 2008.
46. de la Vega, D.; LóPEZ, S.; Matias, J.M.; Gil, U.; Pena, I.; Velez, M.M.; Ordiales, J.L.; Angueira, P. Generalization of the Lee Method for the Analysis of the Signal Variability. *IEEE Trans. Veh. Technol.* **2008**, *58*, 506–516. [CrossRef]
47. *IEEE 802.11-03/940r4* ; TGn Channel Models, IEEE P802. 11 Wireless LANs Std. IEEE: Piscataway, NJ, USA, 2004.
48. Goldsmith, A. *Wireless Communications*; Cambridge University Press: Cambridge, UK, 2005.
49. Jamil, M.; Hassan, M.; Al-Mattarneh, H.; Zain, M. Concrete dielectric properties investigation using microwave nondestructive techniques. *Mater. Struct.* **2013**, *46*, 77–87. [CrossRef]
50. Chung, K.L.; Yuan, L.; Ji, S.; Sun, L.; Qu, C.; Zhang, C. Dielectric characterization of Chinese standard concrete for compressive strength evaluation. *Appl. Sci.* **2017**, *7*, 177. [CrossRef]
51. Santos, P.M.; Júlio, E.N. A state-of-the-art review on roughness quantification methods for concrete surfaces. *Constr. Build. Mater.* **2013**, *38*, 912–923. [CrossRef]

52. Jakes, W.C.; Cox, D.C. *Microwave Mobile Communications*; Wiley-IEEE Press: Hoboken, NJ, USA, 1994.
53. Kramer, O.; Kramer, O. Scikit-learn. In *Machine Learning for Evolution Strategies*; Springer: Cham, Switzerland, 2016; pp. 45–53.
54. Chen, T.; He, T.; Benesty, M.; Khotilovich, V.; Tang, Y.; Cho, H.; Chen, K.; Mitchell, R.; Cano, I.; Zhou, T.; et al. *Xgboost: Extreme Gradient Boosting*; R Package Version 0.4-2; 2015; Volume 1, pp. 1–4.
55. Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; Koyama, M. Optuna: A next-generation hyperparameter optimization framework. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Anchorage, AK, USA, 4–8 August 2019; pp. 2623–2631.

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.