Accepted for presentation in: IEEE Future Networks World Forum (FNWF), Bangalore, India, November 2025.
– Best Paper Award (Second Place) –

Taming mmWave Connectivity Prediction with DRaGon: AI Propagation Modeling for Cluttered Industrial Environments

Melina Geis, Simon Häger, and Christian Wietfeld
Communication Networks Institute (CNI), TU Dortmund University, 44227 Dortmund, Germany
E-mail: {Melina.Geis, Simon.Haeger, Christian.Wietfeld}@tu-dortmund.de

***Abstract*—Future 6G networks are expected to enable a wide range of industrial applications that demand ultra-high data rates. To meet these requirements, communication systems will increasingly rely on millimeter-wave (mmWave) frequencies, which offer large available bandwidths but exhibit complex and sensitive propagation behavior. Accurate network planning for such environments requires reliable Radio Environmental Maps (REMs), which in turn depend on precise channel modeling. Machine learning (ML) has shown promise in bridging the gap between computation time and prediction accuracy left by traditional modeling methods, but it has not yet been transferred to large-scale, complex Indoor Factory (InF) environments. In this paper, we present an ML-based propagation model for taming mmWave connectivity prediction in a densely cluttered real-world industrial scenario. The model utilizes features derived from privacy-preserving level-of-detail environmental representations, such as top and side-view image projections. We evaluate two ML model types – LGBM and ANN – on a dataset collected in a complex industrial environment. Both models accurately capture the underlying propagation characteristics and outperform traditional empirical models by 8.18 dB in prediction RMSE and by being approximately 30-times faster than ray tracing. We demonstrate the potential of ML-based models to support fast and reliable connectivity map generation, making them suitable for efficient mmWave network planning for industrial scenarios.**

***Index Terms*—Machine Learning, mmWave Communication, Radio Propagation Modeling, Indoor Factory**

I. Introduction

In modern industrial environments, wireless communication plays a central role in enabling flexibility, scalability, and real-time coordination across production and logistics systems. With the advent of Industry 4.0, a wide range of participants – including Automated Guided Vehicles (AGVs), mobile robots, production machines, edge devices, and sensors – require reliable connectivity to operate efficiently and collaboratively. To meet these requirements, private 5G networks are increasingly deployed within factories and warehouses, with future 6G expected to further extend capabilities [1–3]. To support the high data rates demanded by many of these applications, the broad available bandwidth at millimeter-wave (mmWave) frequencies is often used - specifically, 5G Frequency Range 2 (FR2) spanning from 24 GHz to 71 GHz [4]. However, at such high frequencies the propagation characteristics differ compared to the traditional communications within sub-6 GHz (i.e., Frequency Range 1 (FR1)). mmWave signals suffer from higher path loss, limited diffraction, and increased sensitivity to environmental factors such as heavily material-dependent object penetration capabilities [5]. As a result, even minor changes in the environment can lead to considerable variations in the radio field.

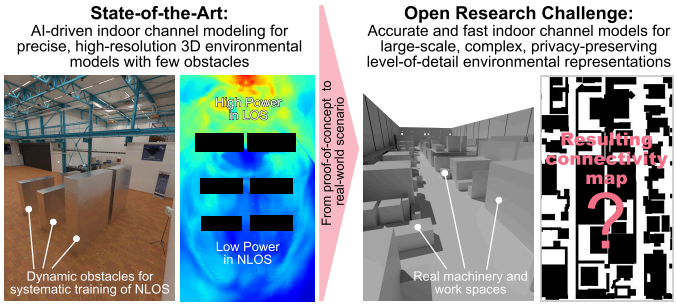

Fig. 1. From modeling radio propagation using highly detailed environmental models in controlled environments to accurate mmWave modeling relying on low-detail representations of complex real-world indoor scenarios

In industrial environments, where dynamic environmental changes are common, channel modeling must not only be precise but also computationally efficient to allow for adaptive network planning to ensure robust and reliable wireless communication, especially in critical applications. This requires an accurate, up-to-date 3D environmental model. While public data sources, such as *OpenStreetMap* are available for creating 3D models of outdoor environments, it is a major challenge with indoor scenarios, where environmental models typically need to be created manually and seldom made public for reasons of strategic privacy. The quality of environmental indoor models can vary significantly (see Fig. 1): In controlled laboratory settings, detailed 3D models with accurate geometry and material properties are more likely to be available. In contrast, models of real-world industry environments are often generated from sensor data such as LiDAR scans for preserving privacy [6], but may contain simplifications (abstraction of specific object details), modeling errors, and lack information about material properties [7]. Consequently, accurate channel modeling must be achieved quickly based on low-detail and easily obtainable environmental representations.

Machine Learning (ML)-driven propagation models are well-suited for this task, as they can implicitly learn environmental characteristics from the training data, allowing accurate predictions even when using simple environmental models without explicit material information. Inspired by our

previously presented ML-driven indoor propagation model, IndoorDRaGon [8], designed for sub-6 GHz communication, we present an ML-based mmWave radio propagation model for complex industry environments in this work. As IndoorDRaGon was trained on data from environments with only few obstacles, the entire feature generation process is redesigned to address the complexity of real-world industry scenarios.

The remainder of the paper is structured as follows. After discussing related works on propagation modeling in Sec. II, Sec. III introduces the new DRaGon [1] approach for for prognosis of mmWave network connectivity. In Sec. IV, detailed results on the ML regressors are presented and followed up by a discussion of the impact of extracted features on computation time and prediction accuracy. Last, we summarize our results in Sec. V and provide insights into our ongoing research.

II. Toward Propagation Modeling for the 6G Era

A wide range of channel models is available, partly designed especially for indoor scenarios, but only a few specialize on the complex propagation characteristics at mmWave frequencies. The highly dynamic and complex channel conditions push classic *empirical channel models* to their limits, especially because they typically rely on oversimplified assumptions that do not truly take the environment into account. For example, 3GPP presents a family of empirical channel models for 5G including mmWave frequencies [9]. They distinguish between several scenarios based on the environmental setting. For Indoor Factory (InF) scenarios, the factory halls are further categorized based on their size and clutter density. However, these models differentiate only between Line-of-Sight (LOS) and Non-LOS (NLOS) conditions, without explicitly considering the environment. ITU provides parameters for various construction materials, which can be used to derive corresponding obstacle shadowing effects [10].

Physics-based *ray tracing* simulations offer a more accurate alternative by explicitly modeling Electromagnetic Waves (EM) wave interactions with the environment [11]. These simulations can account for reflections, diffraction, and scattering based on detailed 3D models. However, the quality of the predictions is strongly related to the quality of the environment model [12]. Moreover, the calculations are computationally expensive, making ray tracing simulations less practical for large-scale, densely cluttered, or real-time applications.

In other domains, such as sub-6 GHz for outdoor and indoor cellular communication, *ML-based propagation models* have proven themselves, as they can implicitly learn environmental information, material properties, and propagation behavior from the measurement data [13]. In many cases, ML methods have been shown to outperform ray tracing simulations, such as in [14]. While some works rely on a high number of measurements as provided in [15], others rely solely on synthetic data [16]. In our previous work we introduced IndoorDRaGon [8] as an ML-based path loss prediction method for sub-6 GHz in a lab like environment with commercial-of-the-shelf 5G equipment given a precise environment model (cf. Sec. III-B). The authors in [17], apply various lightweight ML models trained on measurements at 28 GHz conducted in a manufacturing lab by using only a limited number of features, e.g., the transceiver positions and the number of obstructions in the direct path based on the floor plan. In [18], the authors avoid using environmental information by solely relying on features based on channel state information to train an Aritifical Neural Network (ANN) with measurements at 28 GHz in a small-scale indoor corridor environment. The authors in [19] rely on synthetic data obtained from ray tracing simulations in industry-inspired pseudo-environments. They train an ANN for different frequencies, including 28 GHz, using both channel-related features and features classifying the environment.

In this work, we now take up the idea of [17] to generate area-covering data for, compared to [18], large-scale indoor environments and thus derive an ML method for accurate Radio Environmental Map (REM) generation similar to our validated method in [8]. As the authors in [19], we utilize ANNs trained on a wider range of features than in [17].

III. Proposed ML-Based mmWave Channel Model for Industrial Environments

Our goal is to retrieve a model to accurately predict the Received Signal Strength (RSS) given a 3D receiver and transmitter position and a 3D model of a complex indoor industrial environment for mmWave frequencies.

A. Complex Industry Scenario under Investigation

In this work, we consider a real-world large-scale production scenario from *Fraunhofer IPT* [20] in Aachen, Germany. The factory hall has a length of 97.0 m, a width of 28.5 m, and

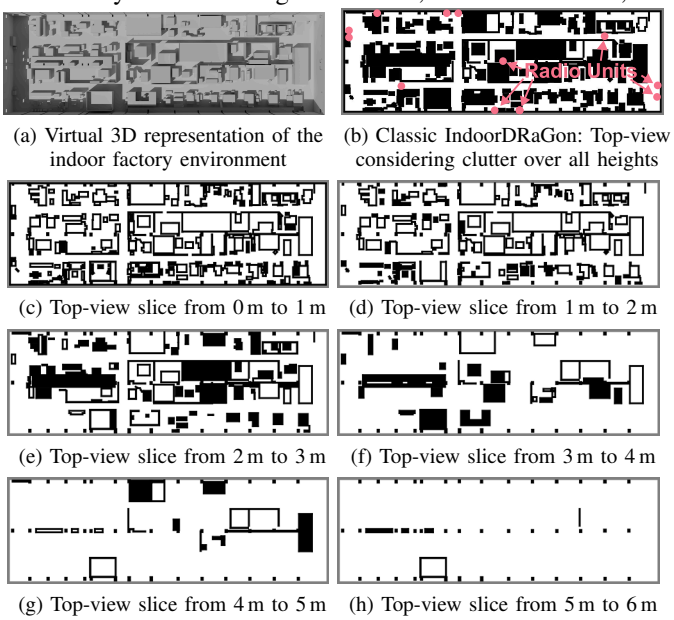

(a) Virtual 3D representation of the indoor factory environment

(b) Classic IndoorDRaGon: Top-view considering clutter over all heights

(c) Top-view slice from 0 m to 1 m

(d) Top-view slice from 1 m to 2 m

(e) Top-view slice from 2 m to 3 m

(f) Top-view slice from 3 m to 4 m

(g) Top-view slice from 4 m to 5 m

(h) Top-view slice from 5 m to 6 m

Fig. 2. Top-view representations of complex industry scenario: As one single top-view (b) cannot reflect the clutter's complexity, multiple top-view slices (c)-(h) are employed

[1]The acronym DRaGon originates from our initial ML-enabled channel prediction method, Deep RAdio channel modeling from GeOinformatioN, but today refers generally to our ML-enabled channel prediction methods.

a height of roughly 10 m, resulting in a shop floor area of more than 2,700 m^2. A 3D top-view of the environment can be seen in Fig. 2a. The environment contains more than 50 machine tools and workspaces used across various stages of production. These elements, along with their surrounding structures, are largely composed of metallic materials, which significantly attenuate or block the penetration of electromagnetic waves. In previous works, we investigated the mmWave coverage at *Fraunhofer IPT* by collecting measurements for a single cell [21] and for two cells [1]. According to the 3GPP definition, the scenario corresponds to InF DH (dense clutter, high Radio Unit (RU)) with the RUs being mounted at 6 m height.

B. Preliminary Work on Industrial Indoor Channel Modeling

As mentioned in Sec. II, we introduced IndoorDRaGon in our previous work [8] as an ML-based path loss prediction method. Given an environmental model of the indoor scenario, a transmitter configuration, and receiver position, two synthetic black-white images of the propagation environment are generated. One showing the Direct Path (DP)'s side-view and one showing the corresponding top-view. Based on these images, numerical features are extracted, indicating the number of obstacle pixels in five vertical and three horizontal split sections. These image-based features are combined with additional numerical inputs such as characteristics of the DP, the relative positions of transmitter and receiver, and communication parameters to form the feature vector $\mathbf{x}$. This vector is fed into the already trained Random Forest (RF) model, which then predicts a correction term to refine an underlying empirical channel model. In a subsequent study [22], we replaced the RF model with an ANN, enabling faster model adaptation through Transfer Learning (TL).

We have shown that IndoorDRaGon can outperform ray tracing simulations and empirical models. However, it was trained on sub-6 GHz measurement data and in a less complex, rather lab-like environment with only a few obstacles and a detailed environmental model. As demonstrated in [22], fine-tuning of IndoorDRaGon using TL is a suitable approach to adapt the IndoorDRaGon model to new scenarios. The authors in [23] have demonstrated that this procedure can also be used for rapid adaptation to new frequency ranges. For the particular use case in this work, a 3D environmental model was created based on 3D LiDAR points. The model consists of roughly 6,000 polygonal surfaces representing the sides and tops of obstacles without standardized geometric structure or material annotations. Due to this increased complexity and reduced model regularity, the original IndoorDRaGon approach, which relied on simplified cuboid representations for feature extraction, is no longer directly applicable.

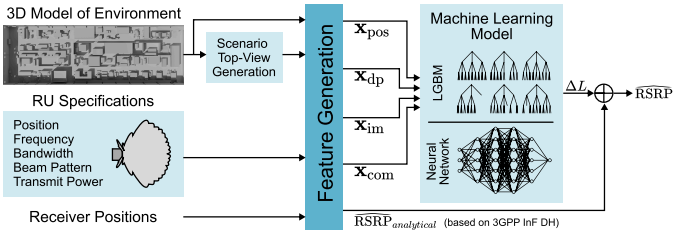

Fig. 4. Overall architecture of the proposed ML-based channel model

C. Towards mmWave-IndoorDRaGon

We modified the process chain to address the aforementioned problems. Fig. 4 provides an architectural overview of the proposed ML-based mmWave channel model for industrial environments. Based on a 3D model of the environment, the RU specifications and a 3D receiver position, various numerical features are extracted. The latter are fed into an ML regressor that predicts a correction term ΔL used to correct the analytical Reference Signal Received Power (RSRP) prediction based on an empirical path loss model. More detailed explanations follow below.

1) Feature Generation Adaptations: Although a side-view image of the DP (cf. Fig. 3) – capturing object heights – is generated in [8], the full spatial complexity of the environment in such densely cluttered scenarios cannot be mapped sufficiently, leading to significant information loss. To address the increased scenario's complexity – particularly the fact that obstacles do not uniformly start at ground level – we extend the original image-based representation by using not just a single black-white top-view image (cf. Fig. 2b), but seven distinct top-view slices at different heights. Each one shows a 1 m-high horizontal slice, with the exception of the last one covering the upper heights from 6 to 10 m. Scenario covering top-views can be seen in Fig. 2c to h.

Fig. 3 depicts the image feature extraction for one of the seven top-views and side-view for one transmitter-receiver pair. The scenario top-views can be generated in advance, while the DP's side-view needs to be created ad-hoc. Therefore, intersections of the DP plane orthogonal to the xy-plane with the 3D environment model are determined geometrically. Since identical operations are performed here on many elements in parallel, we leverage `CuPy` [24] for GPU-accelerated computation. Afterwards, the side- and top-views are tailored

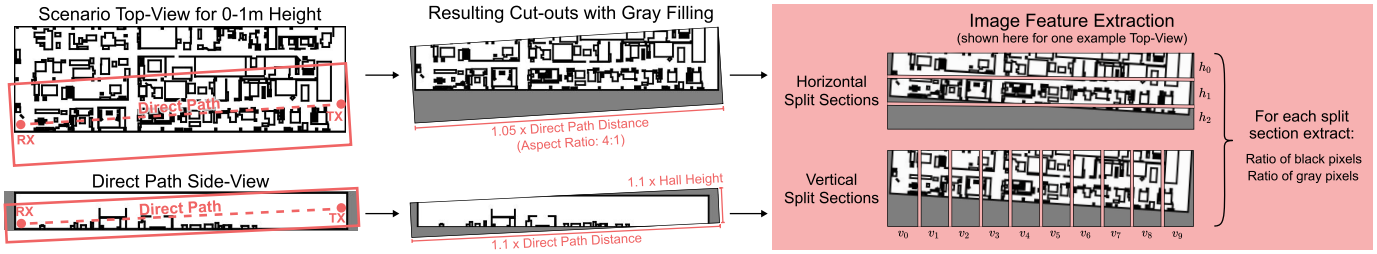

Fig. 3. Example image feature extraction for top-view slice 0 m to 1 m and one transmitter-receiver pair

according to a DP bounding box. The latter shows $1.05 \cdot d_{\text{dp}}$ width and is fixed to a $4:1$ aspect ratio for top-views and $1.1 \cdot d_{\text{dp}}$ width and $1.1 \cdot h_{\text{hall}}$ height for side-views. As these bounding boxes might include areas outside of the scenario, we use a gray color to distinguish between obstacles and the outside area. Since rotating and cropping images is rather computationally intensive, we apply `CuPy` for the tailoring process. This allows us to extract the ratio of black and gray pixels easily. As the distances between transmitter and receiver are comparable large and we lose a lot of detail with only a few image split sections, we use three horizontal and ten vertical split sections (cf. Fig. 3) for every side- and top-view image. This results in 208 numerical features in the $\mathbf{x}_{\text{im}}$ subvector.

Whereas most features of our previous work [8] are reused, we now include the distance to the last intersection instead of the obstructed-LOS (OLOS) distance. In addition, the receiver height is omitted, as the considered receivers operate on the same height of 1 m. The resulting feature vector contains 218 features that can be classified into four subvectors (see Tab. I).

As pointed out in Section III-A, the scenario corresponds to 3GPP's InF-DH definition. Therefore, we rely on the corresponding NLOS path loss prediction model introduced in [9, Tab. 7.4.1-1] for estimating L_{PL} and $\hat{\text{RSRP}}_{\text{analytical}}$.

TABLE I. LOGICAL SUBVECTORS FORMING THE FEATURE VECTOR

Subvector	Features
$\mathbf{x}_{\text{pos}}$	Delta positions $(\Delta x, \Delta y, \Delta z)$, main antenna beam deviation $(\Delta\varphi, \Delta\theta)$, 3D distance d_{3D}
$\mathbf{x}_{\text{dp}}$	Distance to first intersection, distance to last intersection, number of intersections
$\mathbf{x}_{\text{im}}$	Ratio of black and grey pixels for three horizontal $(h_0 ... h_2)$ and ten vertical $(v_0 ... v_9)$ split sections for side- $(\mathbf{x}_{\text{sv}})$ and top-views $(\mathbf{x}_{\text{tv}})$ of the DP
$\mathbf{x}_{\text{com}}$	Empirical path loss L_{PL}

2) Synthetic Data Basis: Thousands of data points are required to train and validate an ML model. We rely on synthetic data generated with the commercial ray tracer *Wireless InSite* [25], which is calibrated based on our real-world measurements from [21] that would be insufficient on their own as they only cover certain pathways and were conducted with mobile devices. Ray tracing simulations are performed for twelve distinct transmitter positions, located at the walls and ceiling (cf. Fig. 2b) of the hall using two distinct antenna configurations and a REM resolution of 0.5 m. The resulting receiver grid contains 9,936 positions at a fixed height of 1 m. The obstacle distribution at this height results in a clutter density of 25.05 % with these positions being eliminated from the data set. The beam patterns are learned implicitly with one value in the feature vector indicating the utilized beam pattern and given the main antenna beam direction. The feature vector is extended by a numerical value indicating whether beam pattern one or two is utilized. While antenna configuration one has a field of view of roughly 90° [1], the other antenna radiates omnidirectionally. Both antennas are using bandwidth of 100 MHz, a subcarrier spacing of 60 kHz, and a center frequency of 27.1 GHz. Overall, the resulting dataset consists of 88,820 RSRP samples. See references [21] and [26] for detailed information on the ray tracing specifications.

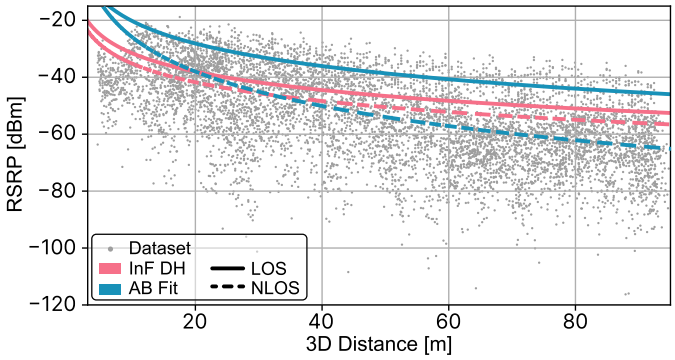

Fig. 5. RSRP scattered over the distance for one RU together with 3GPP InF DH and fitted AB models

IV. DATA ANALYSIS, ML REGRESSOR DESIGN, AND PERFORMANCE EVALUATION

A. Fitting and Evaluating Empirical Path Loss Models

Fig. 5 shows the determind RSRP over the distance for one RU together with the 3GPP InF DH LOS and NLOS models. As it is evident, the empirical model cannot capture the complexity of the environment and greatly simplifies. It achieves an Root Mean Square Error (RMSE) of 49.24 dBm. The 3GPP models belong to the class of the *floating intercept*, or *alpha-beta* (AB) models, following the equation below:

$$\text{PL} = \alpha + 10 \cdot \beta \log_{10}(d) + X_\sigma \tag{1}$$

where α is the intercept in dB, β is the slope, and X_σ is a zero mean Gaussian random variable with standard deviation σ in dB [5]. We used the measurements for one RU near the main antenna beam ($-5° < \Delta\theta < 5°$ and $-5° < \Delta\varphi < 5°$) and fitted α and β for LOS and NLOS situations. The fitted parameters in the LOS case are $\alpha = 42.91$, $\beta = 2.63$ and in NLOS case $\alpha = 34.75$, $\beta = 4.00$. Although the fitted model appears to follow the data trend better, it achieves an RMSE of 49.67 dB, which is as poor as with the 3GPP model. In [10], the effect of building materials is provided by the ITU. The attenuation rate per meter can be derived by the following equation, with f being the frequency in GHz:

$$\text{A} = 1{,}636 \cdot \frac{c \cdot f^d}{\sqrt{a \cdot f^b}} \tag{2}$$

However, the parameters (a, b, c, d) given for metallic surfaces lead to extremely high obstacle shadowing values. Therefore, we fitted a parameter set for the obstacles in our environment with $a = 1.7541$, $b = -0.0536$, $c = 0.0076$, and $d = 1.0784$. Using the fitted *alpha-beta* model together with the obstacle shadowing reduces the RMSE to 33.30 dB. While significantly better, this enormous RMSE value nonetheless demonstrates that more accurate channel models are required.

B. Machine Learning Model Selection and Optimization

Due to the dataset's high dimensionality, we focus our model selection on two powerful, non-linear regressors: a gradient boosted RF – namely Light Gradient Boosting Machine

(LGBM) – and an ANN, both capable of capturing complex relationships in high-dimensional spaces. We build on the Python libraries `scikit-learn` and `PyTorch`.

For both models, we performed a Bayesian optimization to find a suitable hyperparameter set with roughly 1,000 combinations tested for each. The upper part of Tab. II gives an overview of the candidate hyperparameter values together with the found optimal combinations for the considered ANN. The blue highlights indicate the found optimal combination. Similar to [22], the hidden layers double in size until the maximum number of neurons is reached and then halve in size. The considered hyperparameters of the LGBM are provided in the lower part of Tab. II.

TABLE II. CANDIDATE AND SELECTED HYPERPARAMETERS

Hyperparameter ANN	**Candidate Values**
# Neurons in First Layer	[64, 128, 256, **512**]
Maximum # Neurons	[128, 256, **512** , 1024]
# Hidden Layers	[4, 5, 6, 7, **8** , 9, 10]
Batch Size	[64, 128, 256, 512, **1024** , 2048]
Learning Rate	[**1e-3** , 5e-4, 1e-4]
Weight Decay Factor	[**1e-3** , 5e-4, 1e-4]
Hyperparameter LGBM	**Candidate Values**
# Trees	[50, 100, 150, 200, 250, **300** , 500, 1000]
Max. depth	[5, 10, 15, 20, 25, 30, **40** , 50, 70, 100]
Max. leaf number	[20, 50, 100, 200, **300** , 500, 1000]
Learning rate	[**1e-1** , 1e-2, 1e-3, 1e-4]
Bagging frequency	[1, 2, 3, 4, 5, **6** , 7, 8, 9, 10]
Min. Number of data in one leaf	[1, 5, 10, 20, **50** , 100, 200, 500, 1000]
Subsample ratio of the training instances	[0.5, 0.6, 0.7, 0.8, 0.9, **1.0**]
Subsample ratio of columns when construction each tree	[0.5, 0.6, **0.7** , 0.8, 0.9, 1.0]

selected hyperparameter

C. Performance Comparison and Evaluation

Based on the hyperparameter combinations found, the final models can be trained. In order to examine the generalization behavior, not only a *global*, but also *cross* models are trained:

- **Global**: The model is trained on the basis of 80% of the entire data set.
- **Cross**: For each RU position, the ML models are trained based on all remaining RU positions, but without taking the RU position of interest into account.

Fig. 6 shows the resulting Empirical Cumulative Distribution Function (ECDF) on the remaining 20% that were not considered during *global* training. For *cross* evaluation, the test samples were evaluated on their respective *cross* trained models. As it can be observed, the *global* LGBM model achieves the best performance with an RMSE of 7.16 dB. It outperforms the corresponding ANN variant by 0.43 dB in RMSE. However, ANN allows for concepts like TL, which may be relevant for future model adjustments. The *cross* variants achieve poorer RMSE values of 9.77 and 10.75 dB, respectively. Nevertheless, this is significantly more accurate than the predictions of the empirical models. The 3GPP InF DH and fitted AB model achieve similar RMSEs of only 15.34 and 15.96 dB, respectively. The fitted AB model combined with fitted obstacle shadowing yields the worst prediction accuracy with 27.12 dB in RMSE as it mainly overpredicts the path loss in NLOS case.

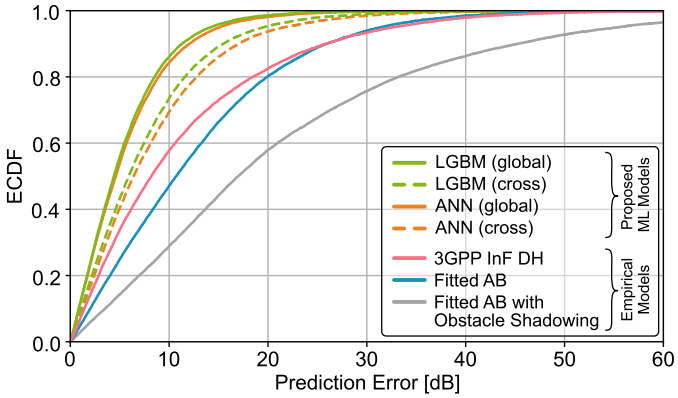

Fig. 6. Comparison of ML-based models *global* as well as *cross* trained together with empirical benchmark models

D. Spatial Prediction Fidelity

To gain more insight into the differences between the specific models, Fig. 7 shows REMs generated by various models for a specific transmitter position. The ground truth based on the ray tracing simulation can be seen on the left, followed by the REMs generated by the *global* and *cross* trained LGBM and ANN models. This is complemented by two empirical benchmarks: the 3GPP InF DH model and our fitted AB model with obstacle shadowing (cf. Sec. IV-A). While the *global* trained LGBM model provides the most accurate REM with only 4.28 dB in RMSE, the worst REM is achieved by the fitted AB model with obstacle shadowing with an RMSE of 28.47 dB. Overall, it can be seen that all ML variants can replicate the ground truth very well, even if the transmitter position was not seen during training. The ML-generated REMs look somewhat smoother than the ground truth, as the models can mainly reproduce the large-scale effects, and local small-scale effects cannot be captured with certainty. It can be observed that the LGBM can partly map these local effects when trained in a *global* manner due to its decision tree structure, which leads to a noticeably lower RMSE than achieved by the ANN.

In the lower part of Fig. 7 the REMs can be seen for a smaller region of $30\,\mathrm{m} \times 10\,\mathrm{m}$ allowing for a clearer identification of the differences. The *global* LGBM closely resembles the ground truth but with reduced noise. In the *cross* trained LGBM REM, it can be seen that the predictions for $(x, y) \approx (70\,\mathrm{m}, 10\,\mathrm{m})$ are too pessimistic and shading in the bottom left corner was not recorded correctly. The *global* trained ANN captures the overall trend pretty well, but appears to be too smooth compared to the ground truth and corresponding LGBM REM. The *cross* trained variant captures the shading effects in the lower left corner pretty well, especially compared to the LGBM model. In contrast to the latter, the model is too optimistic for $(x, y) \approx (70\,\mathrm{m}, 10\,\mathrm{m})$.

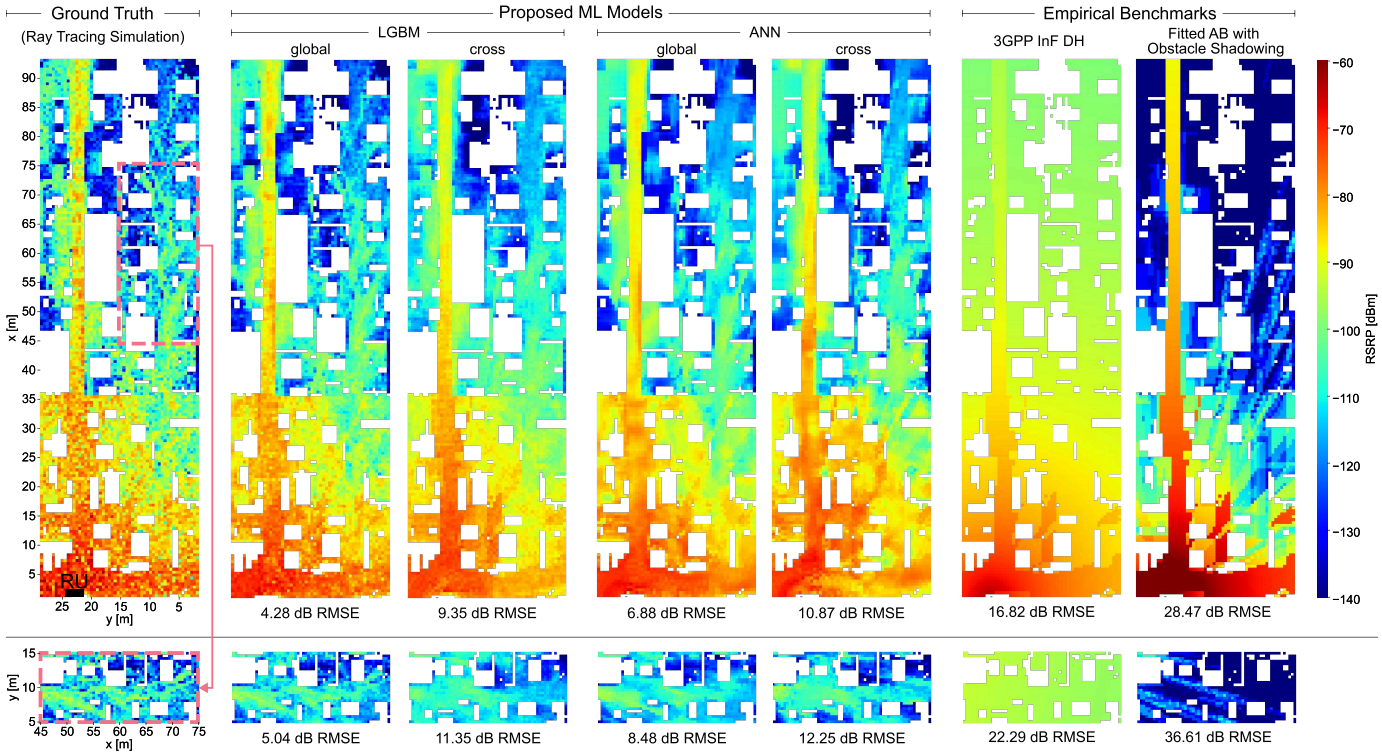

Fig. 7. Visual comparison of REMs for one RU with corresponding RMSE values: ground truth (ray tracing), proposed models (*global* and *cross* trained LGBM and ANN), and baseline models (3GPP InF DH and AB). A selected region is additionally enlarged to highlight local differences between the methods.

The 3GPP InF DH model is not able to cover the effects of the environment and is too optimistic overall. By contrast, the AB model with obstacle shadowing is overly pessimistic by default, and the captured obstacle shadowing does not align with the ground truth, as reflections, for example, are not modeled.

E. Analysis of Feature Contribution and Generation Overhead

Generating all features used by the ML models to predict the REM from Fig. 7 takes 198 s utilizing a *NVIDIA A100-PCIE-40GB* for GPU acceleration. Predicting those 7,285 data points using the ANN takes 1.11 s and using the LGBM only 0.05 s. For the sake of completeness: The training of the *global* ANN takes 214.62 s and of the LGBM model 9.47 s. Compared to the CPU-based ray tracing with ray interaction settings for high speed ($3\times$ reflections allowed, $0\times$ transmissions (no metal object penetration), $1\times$ diffraction, and paths with $< -150\,\mathrm{dBm}$ are dropped) that takes approximately 90 min for REM generation, the 3.5 min needed to perform the ML predictions appears to be very short. However, most of the time falls back on the feature generation. Therefore, we analyze the impact of the different feature classes and the time needed for generating those in more detail.

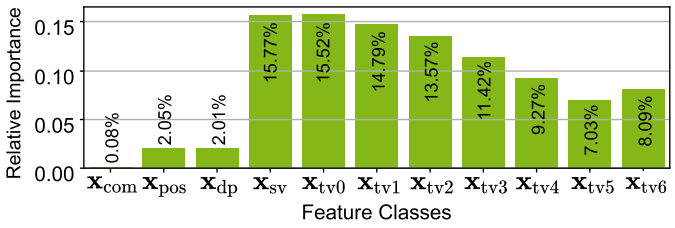

Fig. 8. Relative frequency importance of the feature subvectors (cf. Tab. I)

To analyze the importance of the features, we investigate the frequency importance of the *global* trained LGBM model. We grouped the features according to the classes introduced in Tab. I with additional division into the single images. The relative feature importances can be seen in Fig. 8. The features based on the side-view ($\mathbf{x}_{\mathrm{sv}}$) and lowest top-view image ($\mathbf{x}_{\mathrm{tv0}}$) are considered the most important with each having 15.8% importance. With increasing height of the top-view image, the importance of the features decreases, with the exception of the highest top-view image. This appears to be a reasonable outcome, given that the clutter density decreases with increasing height and that the receiver height is close to the ground level. Although the Direct Path feature class ($\mathbf{x}_{\mathrm{DP}}$) is not considered as important as the images, it holds the most important feature, which is found to be the distance to the first intersection.

Although the side-view features are recognized as the most important feature group, most of the time (42.45%) required to

TABLE III. Relationship among Feature Subvectors (cf. Tab. I), Feature Generation Time and Test RMSE

Feature Subvectors					Feature Generation Time [s]	Test RMSE [dB]	
$\mathbf{x}_{\mathrm{pos}}$	$\mathbf{x}_{\mathrm{com}}$	$\mathbf{x}_{\mathrm{dp}}$	$\mathbf{x}_{\mathrm{sv}}$	$\mathbf{x}_{\mathrm{tv}}$		global	cross
✓	✓	✓	✓	✓	173.35	7.16	9.35
✓	✓	✓	✓		132.41	7.32	9.92
✓	✓	✓		✓	73.58	7.21	9.38
✓	✓	✓			24.63	7.27	10.25
✓	✓				8.01	7.71	12.83

create the features falls into this category. Thus, we trained the LGBM without the side-view features to analyze the prediction accuracy for the *global* as well as one *cross* variant. The results can be seen in Tab. III. Eliminating the side-view images from the process reduces feature generation time, while the test RMSE only slightly deteriorates. In contrast, eliminating the top-view images saves less time, while the test RMSE increases noticeably. When using solely $\mathbf{x}_{\text{pos}}$ and $\mathbf{x}_{\text{com}}$ (cf. Tab. I), the feature generation takes less than ten seconds but at the cost of a high test RMSE in *cross* comparison.

V. Conclusions and Outlook

This work presented an ML-based radio propagation model for mmWave communication in complex industrial scenarios. The model relies on a vast range of numerical features, incorporating features extracted from multiple top-view images of the environment and a side-view image of the transmitter and receiver's direct path based on a privacy-preserving level-of-detail 3D representation of the scenario. Two ML regressors were investigated: an LGBM and an ANN. Both approaches are capable of accurately capturing the propagation characteristics and outperform traditional empirical modeling methods for unseen transmitter positions by 5.57 dB in RMSE.

In future work, the presented model will be further refined and validated using additional data from diverse and varying densely cluttered indoor industrial environments to improve its generalization capabilities. We also aim to develop a fine-tuning routine with real-world measurements to allow for efficient model adaptation even if only a small amount of measurement data is available. Furthermore, we aim to transfer the proposed mmWave propagation model into our existing network planning approach for active RUs [27] as well as for the network design with passive Intelligent Reflecting Surfaces (IRSs) [26].

Acknowledgment

This work was funded by the German Federal Ministry of Research, Technology and Space (BMFTR) in the course of the *6GEM* research hub under grant no. 16KISK038 and the *PANGOLIN Networks* project under grant no. 16KIS2357. The authors thank the Fraunhofer IPT for access to the shop floor model.

References

[1] M. Danger, C. Arendt, H. Schippers, S. Böcker, N. Beckmann, R. Schmitt, and C. Wietfeld, "6G Industrial Networks: Mobility-Centric Evaluation of Multi-cell mmWave Systems," in *Proc. IEEE VTC-Spring Conf.*, Jun. 2025.

[2] A. Ramírez-Arroyo, M. López, I. Rodríguez, T. B. Sørensen, S. Caporal del Barrio, P. Padilla, J. F. Valenzuela-Valdés, and P. Mogensen, "FR2 5G Networks for Industrial Scenarios: Experimental Characterization and Beam Management Procedures in Operational Conditions," *IEEE Trans. on Vehicular Technology*, vol. 73, no. 9, pp. 13 513–13 525, 2024.

[3] M. Ying, D. Shakya, T. S. Rappaport, P. Ma, Y. Wang, I. Al-Wazani, Y. Wu, and H. Poddar, "Upper Mid-Band Channel Measurements and Characterization at 6.75 GHz FR1(C) and 16.95 GHz FR3 in an Indoor Factory Scenario," in *Proc. IEEE ICC Conf.*, Jun. 2025.

[4] 5G-ACIA, "5G for Connected Industries and Automation," White Paper, 2019, accessed: 2025-07-16. [Online]. Available: https://www.5g-acia.org/publications/5g-for-connected-industries-and-automation/

[5] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.

[6] B. Rodrigues, L. Müller, E. J. Scheid, M. F. Franco, C. Killer, and B. Stiller, "LaFlector: a Privacy-preserving LiDAR-based Approach for Accurate Indoor Tracking," in *Proc. IEEE LCN Conf.*, Oct. 2021.

[7] J. Chen, O. E. Mora, and K. C. Clarke, "Assessing the accuracy and precision of imperfect point clouds for 3d indoor mapping and modeling," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. IV-4/W6, pp. 3–10, 2018.

[8] M. Geis, H. Schippers, M. Danger, C. Krieger, S. Böcker, J. Freytag, I. Priyanta, M. Roidl, and C. Wietfeld, "IndoorDRaGon: Data-Driven 3D Radio Propagation Modeling for Highly Dynamic 6G Environments," in *Proc. European Wireless Conf.*, Oct 2023, IEEE.

[9] "Study on channel model for frequencies from 0.5 to 100 GHz, V 17.0.0," 3GPP, Tech. Rep. 38.901, Apr 2022.

[10] "Effects of building materials and structures on radiowave propagation above about 100 MHz," ITU, Tech. Rep. ITU-R P.2040-2, Sep 2021.

[11] Z. Yun and M. F. Iskander, "Ray Tracing for Radio Propagation Modeling: Principles and Applications," *IEEE Access*, vol. 3, pp. 1089–1100, 2015.

[12] A. Schott, A. Ichkov, P. Mähönen, and L. Simić, "Measurement Validation of Ray-Tracing Propagation Modeling for mm-Wave Networking Studies: How Detailed is Detailed Enough?" in *Proc. EuCAP Conf.*, Mar. 2023.

[13] A. Seretis and C. D. Sarris, "An Overview of Machine Learning Techniques for Radiowave Propagation Modeling," *IEEE Trans. on Antennas and Propagation*, vol. 70, no. 6, pp. 3970–3985, 2022.

[14] J. Thrane, B. Sliwa, C. Wietfeld, and H. L. Christiansen, "Deep Learning-based Signal Strength Prediction Using Geographical Images and Expert Knowledge," in *Proc. IEEE GLOBECOM Conf.*, Dec. 2020.

[15] H. Schippers, M. Geis, S. Böcker, and C. Wietfeld, "DoNext: An Open-Access Measurement Dataset for Machine Learning-Driven 5G Mobile Network Analysis," *IEEE Trans. on Machine Learning in Communications and Networking*, vol. 3, pp. 585–604, 2025.

[16] S. Liu, T. Onishi, M. Taki, and S. Watanabe, "A Generalizable Indoor Propagation Model Based on Graph Neural Networks," *IEEE Trans. on Antennas and Propagation*, vol. 71, no. 7, pp. 6098–6110, 2023.

[17] G. A. López-Ramírez and A. Aragón-Zavala, "Enhancing Indoor mmWave Communication With ML-Based Propagation Models," *IEEE Access*, vol. 13, pp. 13 748–13 769, 2025.

[18] F. D. Diba, M. A. Samad, and D.-Y. Choi, "Centimeter and Millimeter-Wave Propagation Characteristics for Indoor Corridors: Results From Measurements and Models," *IEEE Access*, vol. 9, pp. 158 726–158 737, 2021.

[19] M. H. Zadeh, M. Barbiroli, and F. Fuschini, "A Machine Learning Approach to Wireless Propagation Modeling in Industrial Environment," *IEEE Open Journal of Antennas and Propagation*, vol. 5, no. 3, pp. 727–738, 2024.

[20] Fraunhofer Institute for Production Technology (IPT). 5G-Industry Campus Europe. [Online]. Available: https://www.ipt.fraunhofer.de/content/dam/ipt/de/documents/Broschueren/5G-IndustryCampusEurope.pdf (Accessed 2025-07-07).

[21] M. Danger, C. Arendt, H. Schippers, S. Böcker, M. Mühleisen, P. Becker, J. B. Caro, G. Gjorgjievska, M. A. Latif, J. Ansari, N. Beckmann, N. König, R. Schmitt, and C. Wietfeld, "Performance Evaluation of IRS-enhanced mmWave Connectivity for 6G Industrial Networks," in *Proc. IEEE M&N Symp.*, Jul. 2024, Best Paper Young Author Award.

[22] M. Geis, N. Faust, H. Schippers, S. Böcker, and C. Wietfeld, "Leveraging Transfer Learning for Rapid Adaptation of ML-based Indoor Propagation Models," in *Proc. IEEE PIMRC Conf.*, Sep 2025, forthcoming. [Online]. Available: https://cni.etit.tu-dortmund.de/newsdetail/leveraging-transfer-learning-for-rapid-adaptation-of-ml-based-indoor-propagation-models-53469/

[23] T. Nagao and T. Hayashi, "Fine-Tuning for Propagation Modeling of Different Frequencies with Few Data," in *Proc. IEEE VTC-Fall*, Sep 2022.

[24] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, "CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations," in *Proc. LearningSys Workshop NIPS Conf.*, vol. 6, 2017.

[25] Remcom, Inc. Wireless InSite Propagation Software. [Online]. Available: https://www.remcom.com/wireless-insite (Accessed 2025-05-24).

[26] S. Häger, M. Geis, K.-I. Šabanović, P. Jörke, S. Böcker, and C. Wietfeld, "6G Network Design with Custom-tailored IRSs for Sustainable Millimeter-wave Connectivity Enhancements in Industrial Environments," in *Proc. IEEE FNWF Conf.*, Nov. 2025.

[27] M. Geis, C. Bektas, S. Böcker, and C. Wietfeld, "AI-Driven Planning of Private Networks for Shared Operator Models," in *Proc. IEEE LANMAN Symp.*, Jul. 2024.