2024 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN) | 979-8-3503-1764-0/24/$31.00 ©2024 IEEE | DOI: 10.1109/DySPAN60163.2024.10632815

Less is More: Improved Path Loss Prediction Using Simple Interpolation Models

Frost Mitchell
Aditya Bhaskara
Kahlert School of Computing
University of Utah
Salt Lake City, Utah, USA
frost.mitchell@utah.edu

Jie Wang
Neal Patwari
McKelvey School of Engineering
Washington University in St. Louis
St. Louis, Missouri, USA

***Abstract*—Accurately predicting path loss in wireless channels is essential for effective network planning, particularly in dynamic spectrum scenarios where users must minimize interference to other devices. We evaluate multiple path loss prediction methods on three large city-scale datasets, testing out-of-distribution predictions to assess the robustness of each technique.**

Surprisingly, interpolation methods outperform computationally intensive physics-based techniques and complex machine learning (ML) models. In our evaluation, a radial basis function interpolator featuring a linear kernel emerges as the most accurate method due to utilizing local information while producing a smooth interpolation surface for OOD regions. Even a simple linear interpolation model is more accurate on average than a sophisticated physics+ML method, demonstrating the efficacy of simpler approaches in practical wireless channel prediction.

***Index Terms*—path loss models, path loss estimation, received signal strength prediction, interpolation**

I. Introduction

Predicting the path loss for a wireless channel is a crucial problem in network planning, especially in a dynamic spectrum setting where the safe coexistence of multiple users of a wireless channel depends on accurate estimations of the expected interference to and from other users. For example, Tadik et al. [1] propose a Digital Spectrum Twin which can estimate received signal strength (RSS) from sources to determine if wireless users can transmit while meeting minimum interference thresholds which protect other users.

Recent works have proposed novel techniques for predicting path loss on a channel [2]–[5], but these methods have not been extensively evaluated in the challenging setting of out-of-distribution (OOD) prediction, where test inputs are not drawn from the same distribution as training data inputs; e.g., when test samples are drawn exclusively from regions with no training measurements. Specifically, [3], [4] evaluate only on simulated data in extremely simplified environments, while [2], [5] do not consider OOD prediction. In this paper, we present a thorough evaluation of several RSS prediction methods. We compare basic interpolation methods that use exclusively local information, a technique that predicts RSS based on the estimated shadowing loss [2], and one technique that combines physics-based estimations with training data, adapted from [5].

We consider a problem setting where the transmitter or receiver of the channel is in a fixed location and the other device is mobile. We predict RSS values at the location of this mobile device, which is not near any training samples. Ultimately, we find that a radial basis function (RBF) interpolator with a linear kernel is more accurate on average than all other techniques. This includes a complex physics+ML model adapted from [1] which uses features based on the geometry of the environment. These results indicate that as with many interpolation problems, the smoothness of the interpolation surface is essential for accuracy in sparse regions.

This work was supported by the National Science Foundation as part of awards CNS-1827940 and CNS-1564287, as well as the PAWR Project Office grant 10046930.

II. Problem Setting

We consider the problem of predicting RSS between a stationary device (base station, access point, or fixed sensor) and a secondary device at a known location. Given a *digital surface model* (DSM), a stationary device locations, and RSS measurements for many locations relative to the stationary device, the RSS estimation problem is to determine the RSS at unseen locations. We also assume we have a high-fidelity ($\leq$ 1 m resolution) DSM map of the region of interest, which includes height for both natural and artificial features, such as ground elevation, buildings, and trees. Challenges in this setting include shadow fading introduced by obstacles and the time-varying nature of wireless channels. Due to the random variations in channel conditions, RSS predictions that match a particular set of training data may not resemble the average channel conditions or measurements taken at a different time.

A. Demonstrating Robustness through Test Set Selection

In many machine learning problems, training and test data are randomly sampled from the same distribution, which is appropriate if future samples are also known to come from this distribution. But when conditions are known to change, such as in wireless channels, this random selection of data leads to model bias and a failure to generalize to unfamiliar inputs, as described in [6]–[8].

Fig. 1. Plots of Rx and Tx locations with building heights. Left to right: SLC1+Grid2, SLC2+Grid5, ANTW+Grid10.

TABLE I
DATASET INFORMATION

	SLC1	SLC2	ANTW
Frequency [MHz]	462.7	3543	868
Area [km^2]	2.2×2.4	1.4×1.4	6.1×5.6
Stationary Devices	23 (Rx)	6 (Tx)	15 (Rx)
Num. RSS Samples	95,517	77,639	57,777
Meas. Start/End	Apr25 - Nov23	June 27-29	Nov16 - Feb5
Meas. Days	5	3	65

When attempting to deploy RSS prediction over some region of interest, it is prohibitively expensive and difficult to sample RSS at high density in every area of the region. We explore a setting where RSS samples are available to train a prediction model, but RSS samples cannot be obtained from some portion(s) of the region of interest. We use the approach used in our previous work [6], where we divide the region of interest into an $N \times N$ grid and each cell in the grid is assigned randomly to the training or test sets with probability 0.8 and 0.2, respectively. This ensures some separation between measurements used to fit or train an RSS prediction model, and the samples used to evaluate the model.

We consider grid sizes of N=2, 5, and 10. Examples of grid separation for each dataset are shown in Fig. 1. A smaller grid size represents a large distribution shift, where the OOD test data differs significantly from the training distribution, with a larger grid size implying the opposite.

III. Evaluation Datasets and OOD RSS Prediction

We use three large datasets of RSS measurements, with two captured on the University of Utah campus in Salt Lake City, USA (SLC1, SLC2), and one public set of measurements from Antwerp, Belgium (ANTW). Here we describe each of the datasets and the specific details that are relevant to RSS prediction, with details provided in Table I and plots of the measurement locations and environment shown in Fig. 1.

A. Dataset Descriptions

1) SLC1: The first dataset at 462.7 MHz in the FRS band was taken from [9] and is described in detail in [6]. In these measurements, a 1 W mobile transmitter moved through campus, with the RSS values recorded at stationary receivers. There are three different types of receivers with differing hardware and antenna placement, with placement on cell towers, rooftops, and at ground level on the side of buildings. The mobile sensor was carried while walking, on a bicycle, and while in a car, so the spatial density of measurements varies throughout. Some measurements in this dataset are known to have interference from third-party sources.

Regarding the RSS prediction problem, the three different types of sensors have dramatically different noise floors and sensitivities, due to the different antenna placements, hardware types, and gain settings. For example, rooftop radios had an RSS above the noise floor at distances approaching and sometimes exceeding 1 km, but the shorter cell towers had a more limited range of 300-400 m. Ground-level receivers had a range between 400-600 m.

2) SLC2: In the same environment as SLC1, a single mobile receiver was carried on foot and in a car while capturing continuous-wave transmissions from 8 m tall cell towers. Transmission was at 3.54 MHz. There were 4 or 5 active transmitters at a time, with a total of 6 different transmitters used. Due to the limited power of these transmitters, we only considered locations in a smaller area of 2 km^2, compared with the 5 km^2 of SLC1. Compared with SLC1, the method of collection guaranteed minimal third-party interference.

3) ANTW: The final dataset is a subset of the data from [10]. A set of RSS measurements were taken in Antwerp, Belgium. LoRaWAN transmitters were attacked to mail trucks and location data was transmitted to nearby cell towers, which recorded the RSS.

Unlike the SLC datasets, researchers could not control the movement of mobile devices, so the samples are highly skewed in terms of density, with samples concentrated in certain areas, with little to no coverage for much of the region of interest.

IV. Path Loss Prediction Techniques

In this section, we introduce the propagation models evaluated in this work. We explore two standard interpolation methods, Linear and RBF interpolation. Then we introduce a loss field technique that estimates the amount of loss due to

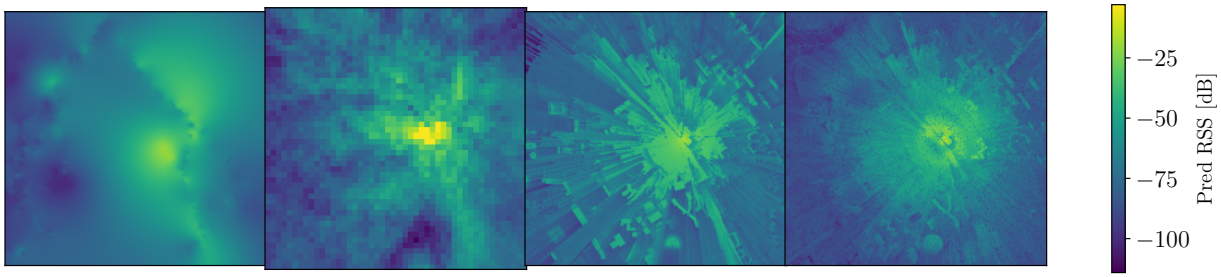

Fig. 2. Predicted RSS values over SLC1, Grid10 for RSS prediction methods. Left to right: Linear, RBF, CELF, TIREM++.

shadow-fading between any two points in the region [2], and finally, we explore *TIREM++*, an ML model based on [5] that uses predictions from the TIREM propagation model as well as other environmental features for RSS estimation.

An example of each of the RSS prediction methods (Linear, RBF, CELF, TIREM++) is shown in Fig. 2. The Linear interpolation surface consists of triangles between points in the training data, while the RBF model appears very similar, though with a considerably smoother interpolation surface, without the sharp discontinuities of the Linear model. The CELF predictions have a generally similar appearance and are reasonably smooth, though there are large curves throughout the predictions, which are an artifact of the method used to estimate the loss field. TIREM++ predictions seem to have a much lower noise floor than all other predictions. There are rays visibly radiating out from the stationary device, indicating line-of-sight paths in different directions. On very close inspection, the prediction has a distinct graininess which is likely due to noise in the input features.

a) Interpolation Methods: We applied two prepackaged interpolation methods from SciPy, LinearNDInterpolator, and RBFInterpolator [11]. The linear interpolation method creates triangles between the input data and then applies barycentric interpolation on each triangle. The other method is *radial basis function interpolation*, which constructs an interpolant based on a linear combination of radial basis functions. We used a linear kernel with smoothing parameter $S = 1$. Typically an RBF kernel is nonlinear, allowing the interpolation function to handle nonlinearities in the data, but our experiments found that the linear kernel was more accurate for our validation sets.

Interpolation can only be done within the convex hull of the training data. To provide full interpolation of the region of interest, we first fit a simple log-distance path loss model to the training set, then use this model to predict the RSS at the four corners of the region of interest and add these four predicted RSS values to the training data. This allows us to interpolate over the entire region.

b) Channel Estimation via Loss Field (CELF): Next, we explored a loss field estimation method, CELF, from [2]. This technique proposes a site-trained loss field model to compensate for the shadowing loss error in the log-distance path loss model. First, a training dataset is applied to fit the log-distance path loss model. Then CELF utilizes Bayesian linear regression to learn the underlying loss field from the modeling error. The loss field explains the propagation loss due to the shadowing effect from the site's obstructions, e.g., buildings and terrain. At prediction time, the path loss is calculated by summing the field values between a transmit/receive pair on the map. The primary advantage offered by CELF is that it provides physics-based loss estimates without using any environmental information from a DSM.

c) TIREM-based Predictions: The Terrain Integrated Rough Earth Model (TIREM) [12] is a propagation model that uses the elevation of a path between a transmit/receive pair to estimate the path loss. Along this 1D path, a simple diffraction model is used to calculate the expected signal strength.

In [5], Tadik et al. propose a simple neural network which uses environmental features to learn a correction factor for the TIREM prediction. These features include the elevation angle and line-of-sight between Tx and Rx, distance to the closest obstacle, path length, diffraction angles, shadowing angles, and the original TIREM prediction. We do not provide the (x, y) coordinates of each location to the model. This is done to produce a generalized model that predicts RSS solely based on environmental features, rather than from local information. The intent is that for OOD regions where no local information is available, the model can use environmental features to accurately predict RSS.

Using these features, we train a simple fully connected neural network to predict path loss. Unlike in [5], where the neural network is extremely small and simple to provide interpretability, we use a model with 3 hidden layers of 200 neurons each. We use a learning rate of 10^{-3}, with L2 regularization with weight decay of 10^{-6} and dropout rates of 0.1 for input features and 0.01 for hidden layers, and train for 40,000 epochs. For each dataset, we combine the RSS values from all stationary locations into a single training set and train a single model to predict RSS values (with three models used for SLC1, where there are three types of stationary devices).

V. Results and Analysis

In this section, we present the results for different RSS predictors. Each RSS method was trained/fit on the same set of data, and evaluated on an *OOD Test* set, which consists of RSS samples in the OOD regions. We apply our grid-separation to each dataset to select the OOD Test set with grid sizes of 2×2,

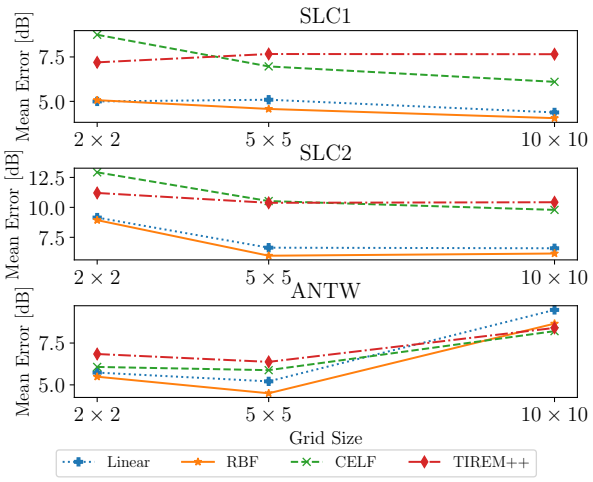

Fig. 3. Average absolute error on OOD Test samples across different grid sizes.

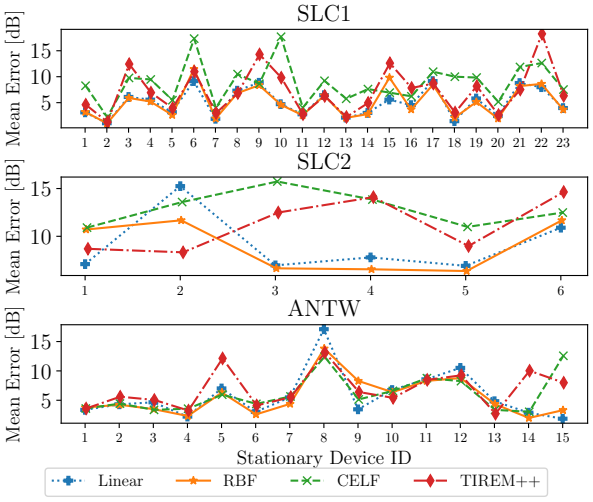

Fig. 4. Average absolute error on OOD Test samples for each stationary device for Grid2.

5×5, and 10×10, in order from **largest distribution shift to smallest**. The large cells on a 2×2 grid represent an entire quarter of the region without training data.

In Fig. 3, we display the average error on the OOD Test samples across different grid sizes for each dataset. We see that on average, the RBF interpolator is consistently more accurate than other methods. This holds across different grid sizes for all of our three datasets. This is a surprising but crucial result: **simple data-based interpolation methods are more accurate than complex physics and ML-informed models, even on OOD samples.** We discuss some possible explanations for this behavior later in this section.

To provide more insight into the accuracy of each method, Fig. 4 shows the average error for each stationary device individually for grid size of 2. We can see that Linear and RBF methods typically track together, with more variation for the CELF and TIREM++ methods. Error for all methods

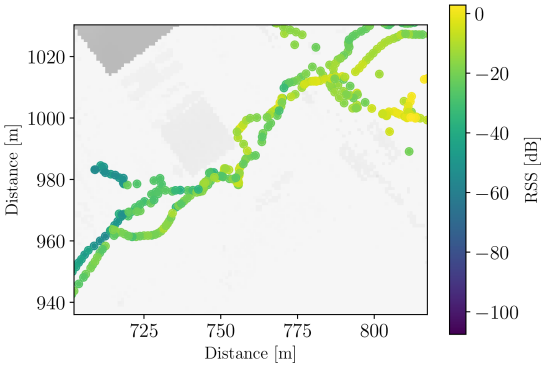

Fig. 5. RSS values in the SLC2 dataset, showing high variance for samples taken within a small area.

varies dramatically from device to device, which we found depends on the device's proximity to an OOD boundary and the diversity of training data available.

One major difference between each of the three datasets is the tightness of the error averages across all four methods. For ANTW, all four methods have extremely similar average error, with only a few exceptions. This trend is less notable but still present in the SLC1 results. The existence of this trend implies that achieving accurate RSS predictions may depend largely on the configuration of training data, more than on the prediction method used.

The underlying cause of the higher spread for SLC2 is unclear, though it may be due to highly varying channels during data collection. An example of this high variance is shown in Fig. 5. Some RSS values taken during the collection of the SLC2 data are shown in this scatter plot. Particularly on the left side of the image, we see differences in RSS over 30 dB within just a few meters. This type of high variation along multiple journeys on the same path is common in this dataset, indicating highly varying channels.

a) CELF-Analysis: CELF performs poorly at a grid size of 2 on the SLC1 and SLC2 data, as shown in Fig. 3. This is expected behavior since an entire quarter of each region is OOD which has little intersection with the coverage area by the training dataset. As a result, CELF relies only on the loss field prior and thus cannot accurately observe the shadowing effect for that OOD area. In other words, CELF for Grid2 predictions falls back to a correlated log-normal shadowing model where the loss field for the OOD region is a pure zero-mean log-normal random field.

An example of this is shown in Fig. 6, which compares RBF and CELF predictions for one stationary sensor in SLC1+Grid2. RBF underpredicts the signal strength near the transmitter since pure interpolation fails to consider propagation physics whereas CELF predictions in the bottom left OOD region have a large positive error since the training data provides no shadowing information about the OOD region for CELF. Without the information about these loss-inducing obstacles, the RSS in this region is overestimated by CELF.

In Fig. 3, we see CELF's accuracy drastically improve with

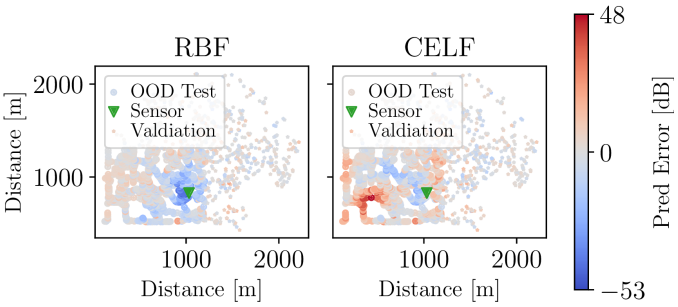

Fig. 6. RBF and CELF prediction error in SLC1+Grid2 for Sensor 9. Positive error indicates higher predicted RSS. RBF underpredicts the signal strength near the sensor, while CELF overpredicts values in the bottom left.

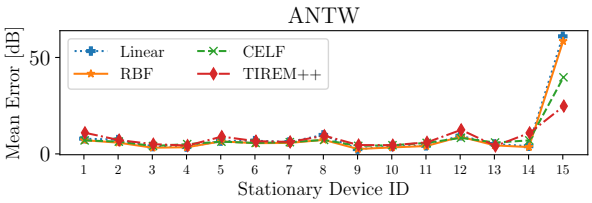

Fig. 7. Average error on OOD Test samples for each stationary device in ANTW+Grid10.

larger grid size (smaller distribution shift) for the SLC1 and SLC2 data, but not for ANTW data, due to the inconsistent density of samples and sparse coverage of the OOD region.

Like the more accurate interpolation methods, CELF does not use DSM information. The worse accuracy may be due to the highly varying channels demonstrated in Fig. 5, or due to a discretization process that introduces some loss of information. Due to memory limits the resolution could not be increased, but it's possible that with memory optimizations CELF could be improved to match the performance of the other interpolation methods.

b) TIREM++: Sometimes Effective with Limited Data: Although TIREM++ is sometimes competitive with the simpler interpolation methods, there are few cases where TIREM++ is significantly more accurate on average than any other techniques. One example is for SLC2 Device 2, where TIREM++ is more than 3 dB accurate on average, as shown in Fig. 4. Another is in the ANTW+Grid10 results shown in Fig. 3. Sensor 15 in this setting had 2-5× higher error than all other devices, which we show in context in Fig. 7. For Sensor 15, the average error for all prediction methods increases drastically. Error for the purely data-based methods (Linear, RBF) jumps to approximately 60 dB, while error for CELF and TIREM++ jump to 40 dB and 25 dB, respectively.

To explore this case further, in Fig. 8 we show a scatter plot of the actual prediction errors for RBF and TIREM++ predictions for this high-error sensor. The nearest training sample to Sensor 15 is over 500 m away, so the Linear and RBF interpolation models cannot predict a high RSS value without local information. Although the CELF model predicts a higher RSS value than the interpolation models, these

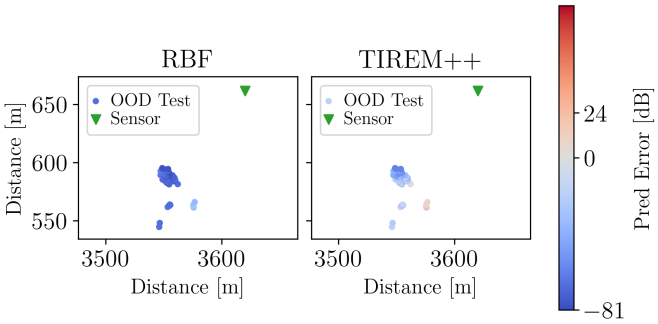

Fig. 8. RBF and TIREM++ prediction error for OOD Test samples for Sensor 15 in ANTW+Grid10. Positive error indicates higher predicted RSS. Without any training points nearby this sensor, TIREM++ gives better RSS predictions than RBF interpolation.

predictions are still skewed by an underinformed path loss model. We observe this phenomenon for a few other sensors, where TIREM++ is more accurate on high-RSS samples close to isolated stationary devices in OOD regions. This indicates that although the TIREM++ model is less accurate on average, it can be more accurate in cases where few training samples are available. This is exactly the result we would hope for from the TIREM++ approach.

However, TIREM++ does not always excel in circumstances that we might consider to be ideal. For each of the 132 configurations of stationary devices and grid separations, we found 26 scenarios where the stationary device was isolated from training samples. In these cases, we would expect TIREM++ to outperform other techniques, but we found it only outperformed other methods on 7 of the 26 cases. However, as shown in Fig. 8, these improvements can be significant.

A. Discussion

Over all of the 132 scenarios (44 sensors with 3 grid separations each), RBF is the most accurate model for 96/132 scenarios, or 72.7% of the time. The poor overall performance of the TIREM++ model is somewhat surprising, since this model uses detailed environmental features to make predictions, which far more information provided than for any of the other techniques. This may be due to model regularization or limited capacity which prevents the overfitting of the data, while also preventing high accuracy. If this is the case, an improved model may be able to achieve RBF-level accuracy while handling the few edge cases where TIREM++ is more accurate.

The significantly better performance of the RBF model can be attributed to a few key features. First, a combination of all training samples within a certain radius will be used to estimate RSS. RBF interpolation is based entirely on very local features, so the success of this method indicates that local information provided by training samples is extremely useful for predicting RSS values. This local information proves useful for prediction on samples that are near the OOD region boundaries.

TABLE II
USE CASE SCENARIOS OF RSS PREDICTION METHODS

Method	Application	Runtime [s] Train	Infer	Needs DSM	Data Needs
Linear	-	1.7	2.75	No	High
RBF	ID	9.8	65	No	High
CELF	ID/OOD	2.8	1.1	No	Medium
TIREM++	OOD, isolated BS	<800	3.8	Yes	Low

Second, the RBF model produces smooth and continuous interpolation for areas with no local information, as shown in Figure 2. As long as no RSS peaks occur in the unseen areas (such as when a transmitter is located there), the smoothness of the interpolation function allows the model to generalize well beyond the training data.

On the other hand, TIREM++ does not utilize the local information as well RBF, but it does provide additional context due to the dependence on TIREM predictions and geometric features. CELF also provides this context, though more limited due to the lack of environmental information. Consequently, CELF and TIREM++ models can be more accurate in scenarios where the stationary device is isolated from training samples.

As was shown in Fig. 5, high variance in training data may provide a greater challenge for RSS prediction. All of our datasets have a relatively limited collection scope, spanning at most several months, and in some cases as limited as 3 days. With a relatively short time for collection, the relationship between the RSS training data, the predicted RSS values, and the actual average channel conditions cannot be determined.

In Table II we present use case scenarios of the different RSS prediction methods, according to the results seen here. We provide rough runtimes for training and inference over all of SLC2. For TIREM++ we also include estimates of the time to calculate TIREM predictions and environmental features.

RBF is best used for In-Distribution (ID) predictions when training data is diverse and covers many locations, especially near transmitters. CELF is recommended when there is no DSM map of the environment, especially for OOD regions. TIREM++ is useful when training data is not diverse or does not cover base station locations, since the ML model does not depend on local information to learn the impact of environmental features.

VI. CONCLUSION

In conclusion, accurately predicting path loss in wireless channels is fundamental for dynamic network planning and coexistence. In this work, we explored various techniques for predicting RSS in wireless channels, focusing on OOD scenarios where test inputs are separated from training data.

Our evaluation included Linear and RBF interpolation methods, a loss field estimation technique (CELF), and a physics-informed machine learning model (TIREM++). Surprisingly, simple interpolation methods consistently outperformed more complex models across different datasets and with different amounts of distribution shift represented by the grid size. In particular, the RBF interpolation is most accurate, due to effectively leveraging local information from training data while providing a smoother interpolation surface compared to other models, providing accurate predictions in OOD regions.

While CELF and TIREM++ occasionally showed greater accuracy in scenarios with isolated stationary devices in OOD regions, their overall performance was inconsistent. CELF in particular is ill-suited for OOD scenarios with large distribution shifts since all measurements are uninformative of the OOD region. Nevertheless, CELF is consistently more accurate than TIREM++ for smaller distribution shifts. TIREM++ failed to consistently outperform RBF interpolation, likely due to the noisy input features and overall lack of smoothness in OOD regions.

Our findings underscore the importance of simplicity and local information in path loss prediction for wireless channels. RBF interpolation, with its ability to capture local variations and generalize well beyond training data, emerges as a robust and efficient method for RSS prediction. Future research could explore ways to enhance interpolation methods by integrating physics-based or ML insights, particularly in scenarios with limited training data. In addition, we plan to evaluate additional ML models used for RSS prediction.

REFERENCES

[1] S. Tadik, K. M. Graves, M. A. Varner, C. R. Anderson, D. Johnson, S. K. Kasera, N. Patwari, J. Van der Merwe, and G. D. Durgin, "Digital spectrum twins for enhanced spectrum sharing and other radio applications," *IEEE Journal of Radio Frequency Identification*, 2023.

[2] J. Wang, M. G. Weldegebriel, and N. Patwari, "Channel estimation via loss field: Accurate site-trained modeling for shadowing prediction," in *2024 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, IEEE, 2024. To appear.

[3] J.-H. Lee, O. G. Serbetci, D. P. Selvam, and A. F. Molisch, "Pmnet: Robust pathloss map prediction via supervised learning," *arXiv preprint arXiv:2211.10527*, 2022.

[4] R. Levie, Ç. Yapar, G. Kutyniok, and G. Caire, "Radiounet: Fast radio map estimation with convolutional neural networks," *IEEE Transactions on Wireless Communications*, vol. 20, no. 6, pp. 4001–4015, 2021.

[5] S. Tadik, M. A. Varner, F. Mitchell, and G. D. Durgin, "Augmented rf propagation modeling," *IEEE Journal of Radio Frequency Identification*, vol. 7, pp. 211–221, 2023.

[6] F. Mitchell, N. Patwari, S. K. Kasera, and A. Bhaskara, "Learning-based techniques for transmitter localization: A case study on model robustness," in *20th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, 2023.

[7] M. Arnold and M. Alloulah, "Benchmarking learnt radio localisation under distribution shift," *arXiv preprint arXiv:2210.01930*, 2022.

[8] F. Mitchell, P. Smith, A. Bhaskara, and S. K. Kasera, "Exploring adversarial attacks on learning-based localization," in *Proceedings of the 2023 ACM Workshop on Wireless Security and Machine Learning*, pp. 15–20, 2023.

[9] F. Mitchell, A. Baset, S. K. Kasera, and A. Bhaskara, "A Dataset of Outdoor RSS Measurements for Localization," *Zenodo https://doi.org/10.5281/zenodo.7259895*, Oct. 2022.

[10] M. Aernouts, R. Berkvens, K. Van Vlaenderen, and M. Weyn, "Sigfox and LoRaWAN Datasets for Fingerprint Localization in Large Urban and Rural Areas," *Zenodo https://doi.org/10.5281/zenodo.1193563*, Mar. 2018.

[11] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, *et al.*, "Scipy 1.0: fundamental algorithms for scientific computing in python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.

[12] D. Eppink and W. Kuebler, "Tirem/sem handbook.," tech. rep., ELECTROMAGNETIC COMPATIBILITY ANALYSIS CENTER ANNAPOLIS MD, 1994.