

Réseau neuronal guidé par la physique pour modéliser la propagation troposphérique / Modelling the tropospheric propagation with a physics-guided neural network surrogate

Thomas Bonnafont², Abdelmalek Toumi¹, Emma Marcos¹, Ali Khenchaf¹

¹Lab-STICC, UMR CNRS 6285, ENSTA, Institut Polytechnique de Paris, Brest, {prenom.nom}@ensta.fr

²DGA, 60 Bd. du général Martial Valin, Paris 75509, France {thomas.bonnafont.enac}@gmail.com

Mots clés (en français et en anglais) : équation d'onde parabolique, U-Net, apprentissage machine, prédiction conforme ; parabolic wave equation, U-Net, machine learning, conformal prediction

Résumé/Abstract

La prédiction de la propagation des ondes électromagnétiques (EM) dans la troposphère est importante pour diverses applications, notamment la conception de systèmes de communication et l'optimisation du placement des antennes. Dans ce travail, nous proposons un réseau neuronal guidé par la physique pour la prédiction de la propagation EM. Pour cela, nous introduisons une technique de génération de données basée sur le modèle physique sous-jacent. De plus, à l'aide de méthodes de prédiction conforme, nous pouvons fournir à la fois le champ électromagnétique estimé et un intervalle de confiance pour la prédiction. Des simulations numériques dans la bande VHF démontrent l'efficacité et les avantages de notre approche.

Predicting the propagation of electromagnetic (EM) waves in the troposphere is important for various applications, including the design of communication systems and the optimization of antenna placement. In this work, we propose a physics-guided neural network for EM propagation prediction. To do this, we introduce a data generation technique based on the underlying physical model. Furthermore, using conformal prediction methods, we can provide both the estimated electromagnetic field and a confidence interval for the prediction. Numerical simulations in the VHF band demonstrate the effectiveness and advantages of our approach.

1 Introduction

Modelling the long-range propagation above irregular terrain is crucial for various applications, such as optimizing antenna placement or ensuring reliable maritime communications. In particular, fast numerical simulations are needed to account for various environments or to calibrate the systems. In this context, one can rely on the knife-edge model [1] but the latter lacks accuracy for very irregular terrains. On the other hand, the parabolic wave equation (PWE) [2,3] model is adequate for precise modeling of forward wave communication while accounting for the terrain and the refraction. In particular, fast and accurate numerical methods have been proposed to solve the PWE, such as split-step Fourier (SSF) [2,3,4] or split-step wavelet (SSW) [5,6]. Nonetheless, for very large scenarios or a large number of simulations, these methods are still time-consuming.

To overcome this problem, machine-learning based surrogates have been proposed. Indeed, after a training time, which can be large depending on the amount of data, they allow for near real-time prediction. In particular, for modelling electromagnetic propagation, physics-guided machine-learning methods such as Deep-Ray [7] have been proposed. Indeed, they used a well-defined, physics-informed, synthetic artificial database to train the model. For example, in [7,8] they proposed a surrogate model for a time-consuming ray-tracing method, while in [9] they proposed a proxy for an integral equation solver to model the propagation in rural areas. Nonetheless, the latter showed accurate results only if the terrain used to train the model was similar to the terrain for the future prediction. We could also cite [10] and [11] which proposed surrogate model for SSF either with terrain or refractive index as input.

Here, we propose a surrogate model for SSW with only the terrain as input. As in [9], we restrict ourselves to the 1D case, but we depart from [9,10] since we aim at proposing a general framework to construct the training data using physical insight from the model. Since the knowledge of the underlying physical model is used to construct the training data, as for Deep-Ray, our surrogate is said physic-guided. Indeed, contrary to physic-informed network, we do not modify the training loss. In addition, conformal prediction techniques [12,13] are used to quantify the uncertainty of the surrogate model predictions. Furthermore, a fine-tuning strategy is proposed to show that it is possible to adapt the proposed network at a low-cost to new data.

The remainder of this article is organized as follows. Section 2 is devoted to the introduction of the neural network model, and its training. In Section 3, numerical experiments are performed to assess the quality of the prediction in different scenario. Section 4 concludes the paper and gives perspectives for future work.

2 The neural network split-step wavelet surrogate

2.1 The parabolic wave equation model and split-step wavelet

Let us begin by recalling the underlying physical model, i.e., the PWE [2], and briefly introduce the SSW solver [5,6] we are using here.

In the context of tropospheric long-range propagation, the PWE is a widely used model. Indeed, it corresponds to an approximation of the Helmholtz equation where only the forward propagation is accounted for. This is also widely used in acoustics and optics. Here, we consider the wide-angle version of the PWE [2] which can be written as

$$\frac{\partial u}{\partial x} = -jk_0 \left(\sqrt{\frac{1}{k_0^2} \frac{\partial^2}{\partial z^2} + 1} - 1 \right) u - jk_0(n-1)u, \quad (1)$$

with k_0 the free-space wave number, n the refractive index and u the reduced field [2,3], i.e., the slowly varying envelope of the field. The advantage of this equation is that along the propagation direction, i.e., x here, it corresponds to an ordinary differential equation that can be efficiently solved.

In particular, an efficient way to solve Equation (1) is to use split-step methods [2,3,4,5,6]. The core idea of these iterative methods is to split the propagation into two parts, first, a step of propagation in free-space where the first term of the right-hand side (RHS) is accounted for, and then a step accounting for the refraction, i.e. the second term in the RHS, is accounted for [2,3]. In particular, SSF [2] uses the fact that the second derivative over z operator is diagonal in the Fourier domain. This process is then repeated iteratively.

Here we choose to use split-step wavelet (SSW) since it has been shown to have a lower time and spatial complexity than SSF. The idea is to iteratively go back and forth between the wavelet and spatial domains to solve Equation (1) efficiently [5,6]. One step of propagation from x to $x + \Delta x$ can be summarized as

$$u_{x+\Delta x} = A R L W^{-1} P C W u_x, \quad (2)$$

where W corresponds to the fast-wavelet transform, W^{-1} to its inverse, C to a hard-threshold compression, P to the wavelet sparse free-space propagator, and L and R to the phase-screen and terrain operators, respectively. The first one (L) allows to account for the refractive index efficiently at each step in the spatial domain. Numerous methods exist to account for the terrain with the PWE [2]. In this work we use the staircase model which is convenient since it only corresponds to translations. Finally, A corresponds to an apodization operator to avoid spurious reflection at the top of the domain. It shall be noted that the threshold for C and the values used to render the free-space propagation operator sparse are calculated to respect a desired accuracy at the maximum range [14].

This step is then repeated until we reach the maximum range. In the following we denote by $\Phi_{SSW}(I)$ the SSW prediction for a set of inputs I .

2.2 The proposed neural network and its training

The core idea of our work is to develop a surrogate model of $\Phi_{SSW}(I)$ denoted by Φ_θ using a neural network. In this notation θ corresponds to the trainable parameters of the model, i.e. weights and biases, with a dimensionality $|\theta|$. In addition, we restrict ourselves to solely terrain inputs, i.e., the refractive index is assumed to be constant, and to the 1D case where from a terrain t_x , a 1D vector, we predict the field in dB at the altitude of the antenna. Indeed, this is a simpler case that allows to validate the approach while also keeping a database of low memory size. Note that in this case it is natural to assume n constant since with the PWE model it does not vary along the propagation direction and thus will not add any information.

2.2.1 The network architecture

Let us first introduce the architecture of the network. Following [9,10], we use an adapted U-Net architecture that is pictured in Figure 1. Indeed, the latter has shown promising results in the prediction of electromagnetic fields due to its multilevel ability to represent the input data. Here, as in [9], a 1D version is used since our inputs are 1D vectors.

In more detail, it consists of two parts. First, the encoder, or the descending part, which allows a multilevel projection of the inputs into a latent space. The latter is comprised of 5 stages of convolution, batch normalization, ReLU activation and max pooling. In particular, as in [7,8], we use atrous, or dilated, convolutions in order to obtain information at a larger scale with the same number of parameters. Since the information at each stage is of different levels, the dilation factor is increased throughout the network. Using a grid search method, we found that a kernel of size 5 with a dilation of 12 for the first two stages and 24 for the rest is optimal in our case. Other methods, such as Bayesian optimization, could have been used for this task and will be tested in the future work.

Then, the decoder, or ascending stage, that uses the information of the latent space to reconstruct a vector, which here would represent the field u in dB. Since it is a symmetric to the decoder, it also comprises 5 stages using the same operations except that transpose convolution is summed with direct link from each level to ascend from stage to stage. At the end, we output the field prediction, μ_θ , and an associated confidence score denoted by σ_θ . More details on this will be given in Section 2.2.3.

It should be noted that this architecture resembles SSW, where the descending part could be seen as the direct wavelet transform (also a multilevel representation) and the ascending part as the propagation and inverse transform. Now, one question remains: *how can we optimize the parameters of the network to obtain a good surrogate, i.e. close to SSW but capable of generalizing?* For this purpose, we need two major ingredients: a well-suited loss function and a good sampling of the underlying manifold of the terrain (or good training data).

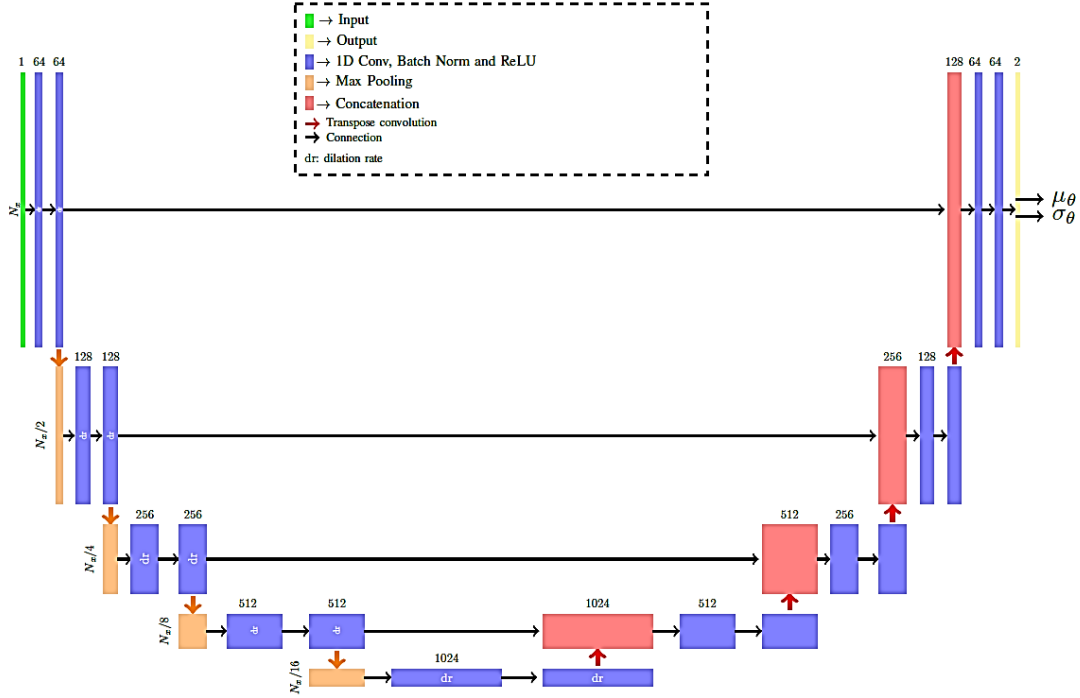


Figure 1: Architecture of the proposed network with two outputs to perform conformal predictions.

2.2.2 Training the network: the data

The main idea behind the construction of our training dataset is to use the a priori knowledge we have on the model, i.e. physics-guided. Indeed, we are not physics-informed since the physical insight is only introduced in the training dataset. Here, we know that the terrain operator corresponds to modelling the obstacles as a series of rectangles, i.e. a staircase model. Therefore, a general training dataset can be constructed from rectangular obstacles. Nonetheless, for sharp terrain gradients, diffraction occurs; therefore, we also introduce triangular obstacles. In addition, this allows to have training data with obstacles constructed from a single stair, the rectangles, and from multiple ones, the triangles.

Furthermore, an orthogonal Latin hypercube sampling (LHS) [15] method can be used to generate the samples, by enforcing the widths and sizes of the different obstacles so that numerous obstacles are placed randomly in the computation domain. Indeed, this approach has the advantage to uniformly cover the underlying manifold with the least amount of data. It shall be noted that, compared to [9,10], this provides a very general framework and does not rely on producing data that resembles real terrains.

In particular, here, since we focus on rural terrain and maritime scenarios, the dataset is comprised of a total of 4000 samples of terrain with 2 to 5 randomly placed obstacles of the two different shapes, with a maximum height of 65m. Indeed, 1000 samples for each number of obstacles are created so that our dataset is balanced. The case of 1 obstacle is omitted since it is included in the other cases. For each of these terrains, we computed the associated field in dB at the altitude of the antenna, here 70m above the ground, using SSW. In Figure 2, we plotted one example of the sample terrain and the associated field in dB. It shall be noted that for a validation purpose, we also construct a validation dataset that comprises 6 randomly placed obstacles with the same shapes.

2.2.3 Training the network for conformal prediction: the loss

The second key ingredient to obtain a good surrogate model is the training loss. In addition, here we want a good prediction, but we also want to have a quantification on the certitude of this prediction. This is why, we perform conformal predictions. Indeed, as you can see in Figure 1, we do not only have one output, the field prediction, but two. The first

corresponds to the field prediction and the second to a confidence score over this prediction. This corresponds to making uncertainty quantification over the *aleatoric uncertainty* of our surrogate model.




Figure 2: a sample of the training dataset. (a) a terrain with four obstacles, and (b) the associated prediction at the altitude of the antenna.

Indeed, under the heteroscedasticity assumption [9,16], the output of the network given the data can be considered Gaussian such that

$$Y|t_x = N(\mu_\theta, \sigma_\theta), \quad (3)$$

with Y the predicted output. In this case, we can define $\sigma_\theta = V[Y|I]$, where V denotes the variance and I the input. From this, we can derive the following empirical risk to minimize, or loss function

$$\min_{\theta} E \left(\alpha \left(\frac{1}{N} \sum_n \frac{(\mu_\theta - Y_{true}[n])^2}{\sigma_\theta^2} + \log \sigma_\theta^2 \right) + \beta \left(\frac{1}{N} \sum_n |\mu_\theta - Y_{true}[n]| \right) \right), \quad (4)$$

where E is the expectation over the dataset, Y_{true} the true output and α and β two weights between 0 and 1. Indeed, the norm ℓ_1 is used to regularize the prediction. Indeed, the ℓ_2 norm is good for point-to-point accuracy, but its results can lack smoothness since it is prone to take outliers more into account. The value of both parameters has been obtained using Bayesian optimization [17] and set to $\alpha = 0.57$ and $\beta = 0.69$. Also, to avoid manipulating values of different scales, we prefer to output $\log \sigma_\theta^2$ instead of the variance. It shall be noted that this regularization is not common, in particular for conformal prediction, but it can be shown using a Laplace prior on the mean error that it is justified.

Finally, from the variance, we want to derive an associated confidence interval, which is defined as

$$IC_\delta(t_x) = [\mu_\theta(t_x) - q_\delta \sigma_\theta(t_x), \mu_\theta(t_x) + q_\delta \sigma_\theta(t_x)], \quad (5)$$

where the quantile q_δ is calculated to obtain the confidence interval IC for a certain risk $1 - \delta$. Contrary to [9,16], we do not use the quantiles for the usual Gaussian distribution, but instead use conformal prediction strategies [12,13] to estimate it so that it has a good coverage over the predictions. To do so, after the training process, we have to calibrate the model over a calibration set. Over this set, we calculate the following residuals

$$R_i = \frac{|\mu_\theta[i] - Y_{cal}[i]|}{\sigma_\theta^2[i]}, \quad (6)$$

with $Y_{cal}[i]$ the true output over the calibration set. Then, we concatenate all the residuals and compute the quantile as the $[(1 - \delta)(N_x N_{cal} + 1)]$ smallest value of R_i , where N_x is the size of the terrain vector and N_{cal} the number of data used for the calibration.

To conclude this part, to expedite the training we set the bias of the last layer to the mean of the true field in dB over the training data and to 1 for the variance. In addition, the risk minimization is performed through a usual ADAM [18] optimizer with a batch size of 8 for 80 epochs. In particular, 80% of the training samples are used for parameter optimization while 20% are used to test the network.

3 Numerical simulations

It shall be noted that all the tests are performed on a standard desktop computer and on CPU. Therefore, the computation time for the surrogate model would drastically reduce if a GPU was used.

3.1 On the validation dataset

First, let us validate the surrogate model over the test dataset. Since, these data are comprised of the same shape as the training ones, we expect to have good prediction. The calibration part is performed by dividing our validation set into two parts, one to compute the quantile and one to evaluate our approach. Thus, we also expect to have a good coverage in this case. In Figure 3 we plot a sample of an evaluation terrain as well as the associated prediction in orange, with its confidence interval, and the true field computed with SSW in blue. It shall be noted that for one sample the prediction is of 0.03s, which is near real time, but also that for the overall evaluation set the time does not increase since all is performed in parallel with GPU.

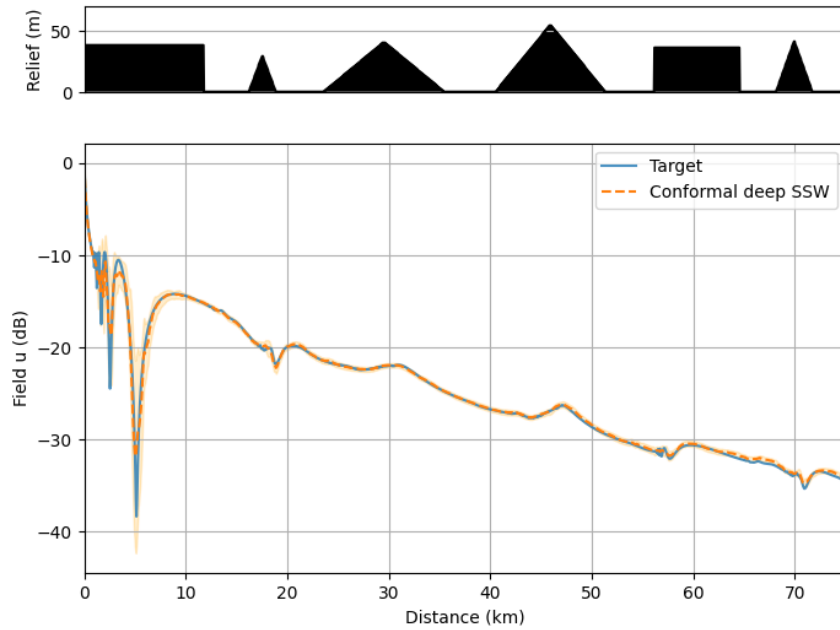


Figure 3: conformal prediction over one sample of the evaluation set. The prediction time was 0.03s for this sample.

From this Figure, we can see that the proposed surrogate works well and that using conformal prediction, we do not only have the field prediction but also have a confidence interval. In addition, on the overall evaluation dataset we obtain a mean error 0.3dB as well as a coverage of 95.5% for the confidence interval. Therefore, as expected we obtain a very good accuracy for the prediction as well as a good coverage.

3.2 Test on real terrain data in France

For this case, the surrogate model is used to predict the field in dB from Paris to Chartres, two French cities separated by 80 km. The terrain between the two cities is obtained from the elevation data of the IGN [19]. To validate the method, the prediction is compared to the one computed with SSW, which also allows us to compute the coverage of our confidence interval. In this case, we expect that the prediction error increases while the coverage decreases. As before, we plot the terrain, and the associated predictions, either with SSW or from the network, in Figure 4.

First, we can see that the prediction is quite accurate, with a mean squared error of 1.03 dB, even if the terrain is realistic. This shows the generalization capability of our physics-guided approach. Indeed, the overall variations are well accounted for while only the local extrema differ. In addition, even if the confidence interval coverage is not as good, around 85%, it shows a good coverage where the mean is not close to the SSW prediction. Finally, with a prediction time below 0.03 s, near real time, this proxy model can be used as an accurate and efficient first glance of the field when planning new antenna location placement or radio-communication link.

3.3 Fine tuning for maritime scenarios

For the last numerical test, we consider a highly different terrain. Indeed, we consider the propagation in a maritime environment where the waves must be accounted for. The latter are modeled as random Gaussian terrain. In this case, our pre-trained network is expected to fail. To correct it without needing to retrain a new network with Gaussian terrains as training inputs from scratch, we propose to use a transfer learning strategy. Indeed, we have seen that our surrogate model performs well on different input data, thus it just needs to be adapted to this new data, as is performed for image processing for example.




Figure 4: conformal prediction over a real terrain with the surrogate model. Prediction time is still 0.03s.

To do so, we construct a new dataset of 1300 samples of random Gaussian terrain and optimize only the parameters of the ascending stage. The training phase is then performed over only 20 epochs. After that, we test our surrogate model over 200 newly generated samples. The latter are split in two parts: one for the calibration and one for the evaluation. We expect the fine-tuned surrogate to perform better than the pre-trained surrogate both in terms of accuracy and coverage.

In Figure 5, we plot one sample of random Gaussian terrain as well as the predictions from SSW, the pre-trained network and the fine-tuned one, with their associated confidence intervals. It shall be noted that the prediction time is as before of 0.03 s.




Figure 5: Field (dB) predictions computed with SSW or the surrogate models.

This figure clearly highlights the improvement obtained through fine-tuning. Indeed, the overall error of the prediction error has drastically reduced, with the overall variations of the field being accounted for, while the confidence interval exhibits a good coverage. In addition, over the evaluation dataset we have a mean error of 4.2 dB and a coverage of 97% while for the pre-trained network the mean error was 18.1 dB. In conclusion, as in image processing, the fine-tuning strategy allows to improve the surrogate model for drastically different input at a low cost. It also paves the way to improving surrogate models with only a few experimental data.

4 Conclusion

In conclusion, this work demonstrates that a physics-guided neural network, trained on synthetic data from a PWE/SSW solver and coupled with a conformal prediction approach, can predict long-distance electromagnetic propagation in near real time while providing a reliable quantification of uncertainty, including over real terrain. In particular, the derived surrogate can be used in place of low accuracy real-time prediction method such as the knife-edge model.

Ongoing work focuses on extending the proposed framework to explicitly account for refraction and two-dimensional configurations, and on improving the calibration of conformal prediction in order to increase the coverage of confidence intervals. In addition, further validation tests in realistic scenario to see the limit of the model.

Acknowledgment

Les auteurs souhaitent remercier la Région Bretagne et l'AID/DGA pour leur soutien aux travaux menés dans le cadre du projet MOPAM.

References

- [1] Deygout, J. (1966). Multiple knife-edge diffraction of microwaves. *IEEE Transactions on Antennas and Propagation*, 14(4), 480-489.
- [2] Levy, M. (2000). *Parabolic equation methods for electromagnetic wave propagation* (No. 45). IET.
- [3] Dockery, G. D. (2002). Modeling electromagnetic wave propagation in the troposphere using the parabolic equation. *IEEE Transactions on Antennas and Propagation*, 36(10), 1464-1470.
- [4] Dockery, D., & Kuttler, J. R. (1996). An improved impedance-boundary algorithm for Fourier split-step solutions of the parabolic wave equation. *IEEE Transactions on Antennas and Propagation*, 44(12), 1592-1599.
- [5] Zhou, H., Douvenot, R., & Chabory, A. (2020). Modeling the long-range wave propagation by a split-step wavelet method. *Journal of Computational Physics*, 402, 109042.
- [6] Bonnafont, T., Douvenot, R., & Chabory, A. (2021). A local split-step wavelet method for the long range propagation simulation in 2D. *Radio science*, 56(2), 1-11.
- [7] Bakirtzis, S., Chen, J., Qiu, K., Zhang, J., & Wassell, I. (2022). EM DeepRay: An expedient, generalizable, and realistic data-driven indoor propagation model. *IEEE Transactions on Antennas and Propagation*, 70(6), 4140-4154.
- [8] Bakirtzis, S., Qiu, K., Zhang, J., & Wassell, I. (2022, March). DeepRay: Deep learning meets ray-tracing. In *2022 16th European Conference on Antennas and Propagation (EuCAP)* (pp. 1-5). IEEE.
- [9] Brennan, C., & McGuinness, K. (2023, March). Site-specific deep learning path loss models based on the method of moments. In *2023 17th European Conference on Antennas and Propagation (EuCAP)* (pp. 1-5). IEEE.
- [10] Huang, S., Qin, H., Hou, W., Zhang, X., & Zhang, X. (2025). Generalizable physics-guided convolutional neural network for irregular terrain propagation. *IEEE Transactions on Antennas and Propagation*.
- [11] Wessinger, S. E., Smith, L. N., Gull, J., Gehman, J., Beever, Z., & Kammerer, A. J. (2025). A Deep Learning Framework for Two-Dimensional, Multi-Frequency Propagation Factor Estimation. *arXiv preprint arXiv:2505.15802*.
- [12] Angelopoulos, A. N., & Bates, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*.
- [13] Dheur, V., Fontana, M., Estievenart, Y., Desobry, N., & Taieb, S. B. (2025). A unified comparative study with generalized conformity scores for multi-output conformal regression. *arXiv preprint arXiv:2501.10533*.
- [14] Bonnafont, T., Douvenot, R., & Chabory, A. (2021). Determination of the thresholds in the split-step wavelet method to assess accuracy for long-range propagation. *URSI Radio Science Letters*, 3.
- [15] Stein, M. (1987). Large sample properties of simulations using Latin hypercube sampling. *Technometrics*, 29(2), 143-151.

- [16] Ethier, J. (2025). Heteroscedastic Neural Networks for Path Loss Prediction with Link-Specific Uncertainty. *arXiv preprint arXiv:2511.23243*.
- [17] Singh, G. S., & Acerbi, L. (2023). PyBADS: Fast and robust black-box optimization in Python. *arXiv preprint arXiv:2306.15576*.
- [18] Kingma, D. P. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [19] Elevation lines data of the "Institut nationale de l'informations Géographique et Forestière" (IGN)."
<https://www.geoportail.gouv.fr/>. Accessed: 01-12-2025.