

ERDC/CRREL MP-21-4

Cold Regions Research and
Engineering Laboratory



**US Army Corps
of Engineers®**
Engineer Research and
Development Center



A Physics-Informed Neural Network for Sound Propagation in the Atmospheric Boundary Layer

Chris L Pettit and D. Keith Wilson

June 2021

The U.S. Army Engineer Research and Development Center (ERDC) solves the nation's toughest engineering and environmental challenges. ERDC develops innovative solutions in civil and military engineering, geospatial sciences, water resources, and environmental sciences for the Army, the Department of Defense, civilian agencies, and our nation's public good. Find out more at www.erdclibrary.on.worldcat.org/discovery.

To search for other technical reports published by ERDC, visit the ERDC online library at <https://erdclibrary.on.worldcat.org/discovery>.

A Physics-Informed Neural Network for Sound Propagation in the Atmospheric Boundary Layer

Chris L. Pettit

*Department of Aerospace Engineering
U.S. Naval Academy
Annapolis, MD 21402*

D. Keith Wilson

*Cold Regions Research and Engineering Laboratory
U.S. Army Engineer Research and Development Center
72 Lyme Road
Hanover, NH 03755*

Final report

Approved for public release; distribution is unlimited.

Prepared for U.S. Army Corps of Engineers
Washington, DC 201314

Under Program Element SAB201, Project Number AB2, Task Number 01

Preface

This study was conducted for the U.S. Army Corps of Engineers, Engineering Research and Development Center (USACE-ERDC), under Geospatial Research Engineering basic research program funds, Program Element Number SAB201, Project AB2, and Task Number 01. The technical monitor was Mr. Ritchie L. Rodebaugh, Technical Director for Geospatial Research and Engineering.

The work was performed by the Signature Physics Branch (Dr. M. Andrew Niccolai, Chief) of the Research and Engineering Division (Dr. George W. Calfas, Chief), U.S. Army Engineer Research and Development Center, Cold Regions Research Engineering Laboratory (ERDC-CRREL). At the time of publication, the Deputy Director for ERDC-CRREL was Mr. David Ringelberg and the Director was Dr. Joseph Corriveau.

This paper was originally published in *Proceedings of Meetings on Acoustics (POMA)* by the Acoustical Society of America on 14 December 2020.

The Commander of ERDC was COL Teresa A. Schlosser and the Director was Dr. David W. Pittman.

DISCLAIMER: The contents of this report are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval of the use of such commercial products. All product names and trademarks cited are the property of their respective owners. The findings of this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

DESTROY THIS REPORT WHEN NO LONGER NEEDED. DO NOT RETURN IT TO THE ORIGINATOR.

A Physics-informed Neural Network for Sound Propagation in the Atmospheric Boundary Layer

ABSTRACT

We describe what we believe is the first effort to develop a physics-informed neural network (PINN) to predict sound propagation through the atmospheric boundary layer. PINN is a recent innovation in the application of deep learning to simulate physics. The motivation is to combine the strengths of data-driven models and physics models, thereby producing a regularized surrogate model using less data than a purely data-driven model. In a PINN, the data-driven loss function is augmented with penalty terms for deviations from the underlying physics, e.g., a governing equation or a boundary condition. Training data are obtained from Crank-Nicholson solutions of the parabolic equation with homogeneous ground impedance and Monin-Obukhov similarity theory for the effective sound speed in the moving atmosphere. Training data are random samples from an ensemble of solutions for combinations of parameters governing the impedance and the effective sound speed. PINN output is processed to produce realizations of transmission loss that look much like the Crank-Nicholson solutions. We describe the framework for implementing PINN for outdoor sound, and we outline practical matters related to network architecture, the size of the training set, the physics-informed loss function, and challenge of managing the spatial complexity of the complex pressure.

1. INTRODUCTION

A. BACKGROUND

Complex natural and built-up propagation environments are a focus in near-ground signal propagation modeling owing to their simulation challenges and operational importance. Outdoor sound propagation in the atmospheric surface layer (ASL) is a good example,¹ but the physics-based and data-driven perspectives considered herein apply to other forms of signal propagation, e.g., optical and RF, as well. In the outdoor acoustics case, uncertainty in wave propagation predictions depends on many factors. Here we consider the source frequency, the range of variations in the mean wind and temperature profile parameters of the ASL, and the ranges of variations in acoustic impedance parameters of the ground. The ASL parameters govern how the atmosphere refracts sound energy. The impedance parameters govern how the ground reflects and absorbs acoustic energy.

Probabilistic machine learning (ML) and statistical learning (SL) methods¹⁻⁴ are helping to establish viable and testable ways of accounting for the many sources of uncertainty. Pettit and Wilson² used proper orthogonal decomposition (POD) for compact representations of the process's variability from an ensemble of transmission loss (TL) realizations, and cluster-weighted models of the joint probability density function of each POD coefficient and the governing parameters. A Crank-Nicolson parabolic equation (CNPE) solver was combined with Monin-Obukhov similarity theory (MOST) to model sound propagation in the atmospheric surface layer. Full-field sensitivities, i.e., sensitivities at every point in the computational grid, were computed by analytically differentiating the conditional expected value function of each POD coefficient and projecting the sensitivities onto the POD basis. Statistics of the full-field sensitivities were estimated by sampling the surrogate model throughout the parameter space.

Hart, et al.³ expanded this line of work using the same outdoor sound propagation model as the benchmark for testing the predictive skill of several engineering, machine learning, and statistical learning models. The ML and SL models were trained on samples from the CNPE-MOST model computed for various combinations of atmospheric refraction and ground impedance conditions, i.e., at various points in the parameter space. The trained models were then tested for their predictive skill at a much denser set of points in the parameter space, and their results were compared with predictions common engineering models of sound propagation, e.g., Harmonoise and Nord2000. All of the ML and SL models achieved much higher skill scores than the engineering models, and the differences between the skill scores of the ML and SL models were insignificant relative to the margin over the engineering models.

Physics-informed or physics-guided neural networks (PINN)⁵ are a more recent advance in machine learning applied to computational mechanics. The motivation in PINN is to combine the strengths of data-driven deep learning (DL) models and physics-driven models to obtain a qualitative improvement in the accuracy of physical process simulations, while requiring many fewer training data than a purely data-driven model, e.g., a standard neural network (NN) model. Bianco, et al.⁴ briefly described deep learning and surveyed recent applications in acoustics. Goodfellow, et al.,⁶ described deep learning as an approach to "... gathering knowledge from experience ..." in a way that "... avoids the need for human operators to formally specify all the knowledge that the computer needs ..." to solve a problem. Instead, a DL model learns to mimic trends in data by processing the inputs to the model through multiple iterations, i.e., layers, of a simple nonlinear map. The resulting architecture may be able to mimic nonlinearities of arbitrary order. In practice, each layer may learn to resolve particular features in the data; i.e., the latent low-order aspects of the data may emerge via the network's architecture as the DL model converges. PINN models described by Raissi, et al.⁵ had a simple architecture called a fully-connected feedforward neural network (FFNN) or multilayer perceptron.⁶ We use the same architecture herein; see Section 2 for details.

Central to supervised learning of any NN model, whether deep or not, is a loss function, which quantifies how well a NN in a given stage of training predicts the training data. PINN augments a data-fit loss

function with penalty terms to quantify deviations from essential physical constraints, such as satisfying a conservation law or a boundary condition. Enforcing these physics constraints during training is akin to placing an implicit but well-informed prior distribution on the range of parameters the model can reach during training; i.e., the physics-based constraints can regularize the training by guiding the model toward valid output using fewer training data than would otherwise be required. Recent results demonstrate the efficacy of this approach for UQ of heat and mass transfer phenomena involving nonlinear conservation laws and fundamental physical constraints on the predicted results.^{5,7} Outdoor sound propagation involves similar classes of models and phenomena, e.g., reflection and transmission of sound waves and shock waves by porous ground surfaces. These similarities suggest that PINN has the potential to improve performance and practicality of machine learning models as surrogates to rapidly predict sound propagation.

A recent literature search by the authors found no published studies using PINN to simulate solutions of the Helmholtz equation (or a parabolic approximation thereof) having significant refraction and interference effects, as commonly observed in outdoor and underwater acoustics. Muingiole and Wilson⁸ trained a two-layer FFNN to predict point-to-point transmission loss in outdoor sound using simulation data from a non-dimensionalized narrow-angle CNPE-MOST model with locally reacting ground impedance. The NN was trained for various combinations of source and receiver heights and separation distances, but the NN was not used to simulate the whole acoustic field between the source and receiver. The closest PINN application the authors found was in the inference of optical impedance in a metamaterial based on scattering predictions from a high-fidelity finite element model.⁹ It should be noted that the PINN model was not used to predict scattered optical field; rather, the Helmholtz equation served as the physics-informed loss function, and the optical impedance of the meta-material was a training parameter that the optimizer was free to adjust as the loss function was reduced to better satisfy the Helmholtz equation.

The current study appears to be the first attempt to use PINN models for outdoor acoustics modeling. The literature features some PINN applications to textbook acoustics problems; e.g., Perdikaris and Wang¹⁰ described PINN models of the Helmholtz equation in a square domain and the wave equations in a 1-D domain with constant wave speed and simple boundary conditions, but nothing as complicated as outdoor sound with spatially-varying sound speed and various ground impedance conditions. Moreover, the PINNs were trained to represent only combinations of parameter values, not a wide range.

B. FOCUS OF THIS STUDY

PINN was tested for modeling outdoor sound propagating through a refracting atmosphere given training realizations of acoustic pressure fields from an equitably sampled subset of atmospheric and ground conditions. The two main goals for PINN training for this problem were to (1) predict the acoustic pressure fields given the values of some model parameters, and (2) infer, as part of training the network to reduce the physics-informed loss function, the values of underlying model parameters not provided in the network's input, such as sensible heat flux or soil porosity.

After this project was under way, the authors learned of a different but relevant study by Lee and Dowling,¹¹ who sought to train data-driven FFNN to solve the Helmholtz equation in typical acoustics applications, including free-space propagation, Lloyd's mirror, and ideal waveguides. Lee and Dowling did not use PINN or attempt to train a single NN across a variety of propagation conditions, but conclusions from their study help to shed some light on observations from the PINN model work described herein.




Figure 1: Sound propagation effects in the atmospheric boundary layer.

Table 1: Parameters in the Monin-Obukhov similarity model for wind speed and temperature profiles and the relaxation model for ground acoustic impedance.

Parameter	Symbol	Units	Values
Porosity	Ω	-	0.525
Flow resistivity	σ	$\text{kPa} \cdot \text{s} \cdot \text{m}^{-2}$	$[150, 600] \times 10^3$
Surface heat flux	Q_s	$\text{W} \cdot \text{m}^{-2}$	0
Friction velocity	u_*	$\text{m} \cdot \text{s}^{-1}$	0.3
Wind direction	α	deg	$[0, 180]$
Roughness length	z_0	m	0.05

2. PHYSICS-INFORMED NEURAL NETWORK MODEL OF OUTDOOR SOUND PROPAGATION

A. OUTDOOR SOUND PROPAGATION MODEL

To train and test a PINN for outdoor sound propagation, we use results from a numerical model of propagation in a refracting atmospheric surface layer (ASL), as depicted in Fig. 1. The same model and conditions were employed by the authors in a previous study on global sensitivity analysis.² The model depends on the six parameters listed in Table 1. Two-dimensional sound propagating from a single frequency point source 1.0 m above the ground is propagated with a wide-angle parabolic equation (PE) approximation,¹² which is solved with the Crank-Nicolson scheme. The domain extends 1000 m downrange of the source and 400 m above the ground. Turbulence is not represented and the ground is assumed to be flat. The ground's acoustic impedance is represented by a relaxation model in terms of the flow resistivity, σ , and porosity, Ω , of the material.¹³ The mean wind and temperature profiles in the ASL are represented through Monin-Obukhov similarity theory¹⁴ in terms of the friction velocity, u_* , surface heat flux, Q_s , wind direction, α , and roughness length, z_0 . The primary outputs for a given set of parameters are the real and imaginary components of the complex pressure, which are then combined to compute the transmission loss (TL) in dB.

Results are shown herein for two source frequencies: 50 Hz and 150 Hz. PINNs were trained separately with ensembles for each frequency value. No attempt was made to train a single PINN to accommodate different frequencies.

B. ENSEMBLE GENERATION AND PROCESSING

As in Pettit and Wilson,² a reduced version of the sound propagation problem with two random parameters was studied to simplify the development, testing, and interpretation of PINN models. Only the wind direction, α , and flow resistivity, σ , were allowed to vary, as described in Table 1. The parameter ranges in Table 1 were not meant to cover every condition that might be encountered during sound propagation outdoors; for example, the static flow resistivity was limited to moderately absorptive surfaces such as soil. Our intent is to examine the ability of PINN to function as a reliable reduced-order model throughout a sufficiently wide range of conditions that might be encountered in practice.

Latin hypercube sampling (LHS) was used to select specific samples of the parameter vector, (σ, α) , and proper orthogonal decomposition (POD) was used to compute a mean-square optimal representation of the ensemble of response fields produced by LHS. For the sake of brevity, we provide only a minimal description of the LHS and POD steps; Pettit and Wilson² should be consulted for more details.

A total of $N = 500$ latin hypercube samples were selected from the parameter ranges listed in Table 1. For each latin hypercube sample vector, the mean wind profile, mean temperature profile, and acoustic impedance of the ground were computed. The PE model was solved to predict the associated sound pressure field for a single-frequency point source at $(x = 0 \text{ m}, z = 1 \text{ m})$. Each realization had $N_g = N_x \times N_z$ grid points, where N_x is the number of range points and N_z is the number of height points. For subsequent analysis, the grid points were mapped into an $N_g \times 1$ column array, referred to in the POD literature as a snapshot.

Let $S(x) = [s_1(x)|s_2(x)|\dots|s_N(x)]$ be the $N_g \times N$ collection of snapshots arranged as column vectors. For $x \in \mathbb{R}^2$ or higher dimensions, the POD computation is facilitated by rearranging each multidimensional snapshot as a column vector to fit this structure. POD represents the energy of the snapshots in terms of a mean-square optimal set of orthonormal vectors or modes, $\Phi^{N_g \times N} = [\varphi_1(x)|\varphi_2(x)|\dots|\varphi_N(x)]$, which are computed as linear combinations of the snapshots as described by the following matrix form:

$$\Phi = SV \tag{1}$$

and $\Phi^T \Phi = I^{N \times N}$, where $V^{N \times N} = [v_1|v_2|\dots|v_N]$ must be computed to maximize the energy of the projection of the snapshots onto the vectors in Φ . V must satisfy the eigenvalue problem

$$S^T S V = V \Lambda, \tag{2}$$

where each $\lambda_n \geq 0$ because $S^T S$ is symmetric and non-negative definite. The normalization condition $\Phi^T \Phi = I$ and Eqn. (1) require that the eigenvectors V be normalized such that $v_n = \widehat{v}_n / \sqrt{\lambda_n}$ and $\widehat{V}^T \widehat{V} = I$. The eigenvalue problem in Eqn. (2) is not costly to solve because $S^T S$ is an $N \times N$ matrix and $N \ll N_g$ here, as in most applications.

The importance of each mode in a snapshot is reflected in its corresponding generalized Fourier coefficient, i.e.,

$$a_q^{(n)} = \varphi_q^T s_n \quad q = 1, 2, \dots, N \quad \text{or} \quad A = \Phi^T S. \tag{3}$$

Individual snapshots can be recovered by inverting the transform, i.e., $s_n = \Phi a^{(n)}$. The surrogate model is made even more efficient by truncating the modal approximation so that only the high energy modes are retained, i.e., $q = 1, 2, \dots, N_r$, where $N_r < N$. Often no more than the strongest 10% to 20% of the modes are needed to retain over 99% of the process's energy.

C. PINN MODEL DEVELOPMENT AND TRAINING

We begin by summarizing the acoustical quantities and equations essential to understanding how we trained PINN models for outdoor sound. The complex pressure field due to a harmonic source of circular

frequency ω is $p_c(x, z, \omega) = p_r(x, z, \omega) + ip_i(x, z, \omega)$. The complex carrier (or complex amplitude) is $A(x, z) = A_r(x, z) + iA_i(x, z) = p_c(x, z, \omega)e^{-ik_0x}$, where $k_0 = \omega/c_0$ and c_0 is the reference sound speed. The complex carrier satisfies the parabolic equation (PE) for planar propagation,

$$\frac{\partial^2 A}{\partial x^2} + 2ik_0 \frac{\partial A}{\partial x} + \frac{\partial^2 A}{\partial z^2} + ((k(z))^2 - k_0^2) A = 0. \quad (4)$$

where $k(z) = \omega/c_{\text{eff}}(z)$ and $c_{\text{eff}}(z)$ is the effective sound speed profile above the ground according to MOST. The transmission loss (TL) is

$$\text{TL}(x, z) = -20 \log_{10} \frac{\bar{p}(x, z)}{\bar{p}_s}, \quad (5)$$

where \bar{p}_r and \bar{p}_s are the RMS pressure measured at the receiver and at 1 m from the source, respectively.

The ground is assumed to be flat with $z = 0$. The characteristic impedance of the ground was computed as

$$Z_s = \frac{p(z=0)}{\frac{1}{i\omega\rho_s} \frac{\partial p}{\partial z} |_{z=0}}. \quad (6)$$

Two groups of neural networks were tested. All networks had the same basic, fully connected architecture using two to four hidden layers, with hyperbolic tangent activation functions at each neuron except those in the output layer.

In the first group, which we refer to as the PINN group, the input layer had four neurons to take in the flow resistivity, σ , and wind direction, α , of the realization, and a randomly sampled (x, z) grid location. The output layer in the PINN group had either two or three neurons. The first two neurons represented the predicted p_{real} and p_{imag} for the input combination of $(\sigma, \alpha; x, z)$. When the third output neuron was included, it represented the network's effectively independent prediction of the transmission loss. Of course, TL in reality is a redundant quantity when the complex pressure is known, but we wanted to test whether a PINN could predict the two quantities by learning their separate functional dependencies on the inputs, i.e., as if the relationship between p_c and TL were unknown to the analyst.

In the second group, which we refer to as the PINN-POD group, the focus shifted from predicting the complex pressure at a particular grid point to predicting the POD coefficients of the complex pressure for the realization as a whole. In this group, the only two input neurons were for σ and α , and the output layer was expanded to have one neuron for the generalized Fourier coefficient of each retained POD mode. Synthesis of the complex pressure and TL at each grid point became, in effect, a post-processing step.

The PINN-POD approach was pursued because the purely PINN models were less successful than desired. As some results shown below will illustrate, the root cause of difficulties with the PINN models is the fine-scale spatial features of the complex pressure field and how this field varies between realizations. After initiating this study, the authors learned of a recent study by Lee and Dowling.¹¹ They described similar phase- and BC-related difficulties when using a deep learning model to simulate the magnitude and phase fields in a simple environment having uniform sound speed, i.e., in conditions more benign to model than those attempted here. Neural networks are most effective when there are low-dimensional features to lock onto in the training data. Ideally, these features are present in the training data, such as in facial recognition applications; however, pre-processing may be needed to transform the training data in a way that elicits latent low-dimensional features. With this in mind, we chose PINN-POD to try to circumvent the spatial complexity issue. The POD process naturally embeds intricate spatial features in the POD modes, thereby significantly lowering the effective dimensionality of the problem.

We refer the reader to Raissi, et al.⁵ for the general formulation of a PINN model for a nonlinear PDE. The training algorithm followed in this study was based on the algorithm Raissi, et al. employed to simulate Schrodinger's equation and the 2-D incompressible Navier-Stokes equations. For the sake of comparison,

we will call their algorithm the basic PINN algorithm. Two significant changes were required to adapt the basic PINN algorithm to this application. One is described here; the other is described near the end of this section. The first significant change was to have a random ensemble of distinct 2-D realizations vice a time-stamped sequence 2-D realizations of the solution process, as in their Navier-Stokes application. Instead of having the time value as in input to the PINN model, we had the physical parameters specific to the training realization, e.g., the flow resistivity of the soil and the wind direction.

For the sake of brevity, we jump directly to the description of the training details and loss function components specific to the current application. $N = 500$ realizations from the CNPE-MOST model were computed for a given input frequency. In most cases, $N_t = 100$ realizations were selected randomly from the N -member ensemble for PINN or PINN-POD training, and the same training realizations were used in all cases for the sake of comparing the resulting predictions.

Loss functions must in general be computed at specific points in the physical domain and the parameter space. Data-driven loss terms considered in different stages of this study include the mean-square error (MSE) of the following quantities:

- Real and imaginary parts of $p_c(x, z)$ at collocation points
- Real and imaginary parts of $A_x(x, z)$ at collocation points
- $TL(x, z)$ at collocation points
- Predicted coefficients of the retained POD modes for the real and imaginary parts of the complex pressure

Physics-driven loss terms considered in different stages of the study included

- MSE of $A(x, z)$ satisfying the parabolic equation at a set of collocation points
- Z_s , as computed from the complex pressure at a set of ground collocation points, and then averaged for comparison with the parametric value. Z_s was computed based on Eqn. (47) of West, et al.,¹² with $\partial p/\partial z$ computed as a finite difference of the complex pressure field near the ground. Inferring a consistent (i.e., low-variance) estimate of Z_s from the ensemble of complex pressures predicted along the ground was a significant challenge that remains unresolved. We deem this difficulty to be a direct consequence of the underlying difficulties in accurately predicting the spatial variations of the phase. Both the PINN and PINN-POD approaches resulted in similar difficulties when predicting Z_s . Of the two approaches, PINN-POD was better owing to the inherent regularization of the POD, but still not accurate or consistent enough to be deemed successful.

The various loss functions were summed to form a multiobjective loss function. The second of the two significant changes to the basic PINN algorithm was to add weights to the combined loss function, which can be written as

$$\text{MSE} = w_{\text{data}}\text{MSE}_{\text{data}} + w_{\text{physics}}\text{MSE}_{\text{physics}} + w_{\text{BC}}\text{MSE}_{\text{BC}} + w_{\text{IC}}\text{MSE}_{\text{IC}}. \quad (7)$$

In the basic algorithm of Raissi, et al., the applications were all nicely scaled such that the various loss terms had similar orders of magnitude, and thus were all similarly influential in guiding the training process iterations. In the outdoor sound application having realistic physical parameter values, we found that such convenient scaling was not possible. We attempted to compensate by adding the pre-set weights to the loss terms when summing them. This approach was partially successful but suffered from the same fundamental difficulty inherent to solving any multiobjective optimization problem: adjusting the weights of the objective function allows the optimizer to explore a Pareto frontier of possible solutions. Of course, what gets emphasized by the optimizer comes at the expense of the other losses. Perdikaris and Wang¹⁰ described

a approach to dynamically adjusting the loss function weights as the network is trained, but this approach was not implemented in the current study owing to time constraints as well as the difficulties noted above in inferring Z_s from the PINN and PINN-POD predictions.

D. PINN MODEL PROGRAMMING

The basic PINN algorithm of Raissi, et al. was implemented in Python using the TensorFlow (v1.x) deep learning library. Their Python code is available on GitHub at <https://github.com/maziarraissi/PINNs>. Our first implementation of PINN, i.e., without POD, was an extension of their code, and so used Python and TensorFlow. Our early difficulties in obtaining satisfactory PINN predictions were compounded by the challenges of PINN model development in the paradigm enforced by TensorFlow 1.x. Accordingly, we reprogrammed our PINN algorithm in PyTorch, which supports a more flexible and easier-to-debug approach to model development.

Various test cases showed that the PyTorch PINN implementation produced essentially the same results as the TensorFlow PINN implementation. This outcome corroborated the suspicion that the PINN approach to outdoor sound prediction and uncertainty quantification has some inherent challenges related to the physics and numerics of the problem. It was at this point that the PINN-POD approach was implemented. POD models were produced in MATLAB using well-tested code from previous studies (e.g., Pettit and Wilson²), exported to a MAT-file, and read into our Python code using a standard SciPy library function for reading MAT files.

A simple library of functions was written in Python (NumPy), TensorFlow, and PyTorch to support the various acoustics calculations required to implement MOST and the ground impedance model. The TensorFlow and PyTorch implementations were essential to pursuing parameter inference in parallel with PINN training, as described in Section 3.2.

3. RESULTS

Results are shown for two ensembles of CNPE-MOST realizations, one with a 50 Hz point source and one with a 150 Hz point source, as described in Section 2.1. In both cases, the PINN-POD results were based on retaining N_r POD modes to represent at least 99% of the variance in the N -member ensemble. In the authors' experience, POD effectively reduces the dimensionality of the outdoor sound realizations such that $N_r < N/10$, and often much less.

A. PREDICTIVE ACCURACY AND FREQUENCY DEPENDENCE

Testing showed that at least three hidden layers were needed to qualitatively match the visual characteristics of the TL contours across a wide range of test cases. Figure 2 shows TL contours from a 50 Hz source for a randomly selected $\{\sigma, \alpha\}$ test pair, i.e., a pair not represented in the training data. The upper plots show the CNPE-MOST prediction, which is shown in both columns for easy comparison with the PINN-POD predictions in the lower plots. Both PINN-POD models had three hidden layers, but one had 32 neurons in each layer and the other had 64 neurons in each layer. The 32-neuron plot might be deemed an adequate depending on the goals of the analysis. Doubling the number of neurons in the hidden layers smoothed the TL contours throughout the domain and slightly improved the downrange accuracy for $z < 50$ m. Both PINN models overestimated the TL downrange at higher elevations.

Figure 3 shows a different realization from the 50 Hz ensemble. The PINN accurately captured the general shape and locations of the TL contours, including the ducting effect downrange near the ground. As in Fig. 2, accuracy somewhat decreases downrange at higher elevations.

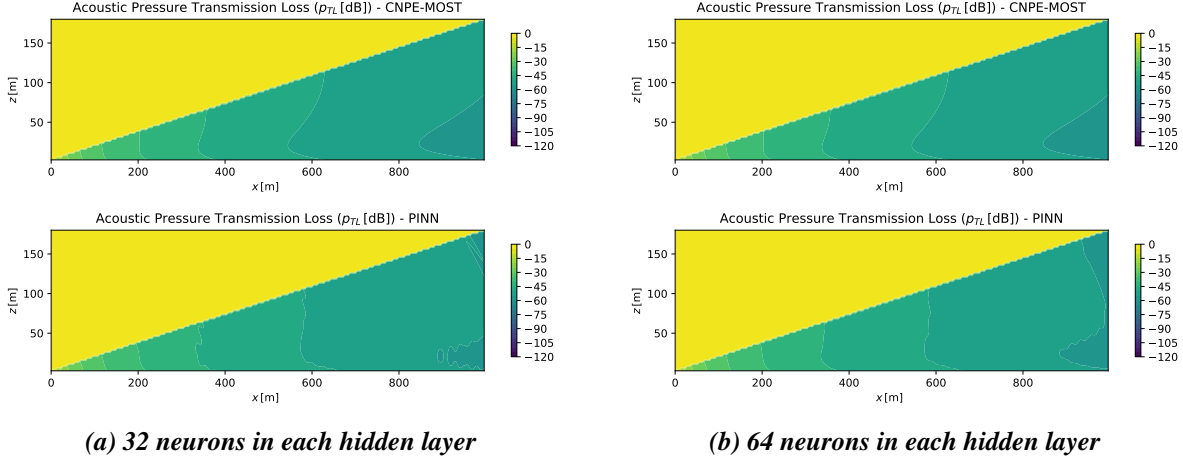


Figure 2: Transmission loss predicted for a 50 Hz point source at $(x, z) = (0 \text{ m}, 1.0 \text{ m})$. The upper figures, which are exactly the same, are the ground truth results from the CNPE-MOST model for wind direction $\alpha = 82.0^\circ$ and flow resistivity $\sigma = 1.78 \times 10^5 \text{ kPa s m}^{-2}$. The other physics parameters were set at the baseline values described in the text. The lower figures show PINN-POD predictions from networks having three hidden layers. Each NN had two input neurons and 10 output neurons, one for each POD mode. (a) 32 neurons each hidden layer. (b) 64 neurons in each hidden layer.

Raising the source’s frequency should be expected to increase the challenge for any type of reduced-order model, including POD and NN, either on their own or combined as they were in this study. This is because raising the frequency proportionally decreases the distance between extrema in the complex pressure. The higher spatial density of information requires a corresponding increase in the number of POD modes. A source frequency of 150 Hz was found to require 30 modes (15 for p_{real} and 15 for p_{imag}) to meet the 99% energy (mean-square) threshold. Accordingly, in PINN-POD training, the number of neurons in the output layer was increased to 30 and the number of neurons in each hidden layer was increased to 150.

Figure 4 shows TL contours for a randomly chosen case not included in the PINN-POD training set. The PINN model in the left-hand column had three hidden layers, which was sufficient for 50 Hz source prediction, but evidently less so for the higher frequency. The PINN-POD model still captured the TL trends in the CNPE-MOST output, but downrange contours in the PINN-POD results are rather ragged and isolated artifacts are visible farther above the ground. The right-hand column shows the effects of increasing the NN depth to four hidden layers. Downrange contours are still less smooth than one might wish, but the artifacts at moderate elevations are significantly reduced. Needing to increase the depth of a DL model in image processing applications often is connected with increasing the spatial complexity of the features in the images. This is similar to the need here for more POD modes, with the additional modes being those that incorporate more localized features in the data.

In comparison, negligible improvements were observed in increasing from three to four hidden layers in the 50 Hz case. As a general rule in training DL models, adding neurons, either in a given hidden layer or by adding new layers, increases the number of unknowns to be adjusted during training. The increased model complexity may result in needing more data to achieve consistent convergence during optimization, and always results in longer training cycles. Adding physics-informed components to the loss function should help to regularize the model training process, but cannot be expected to eliminate the need for more data. Compare Fig. 4b with Fig. 5. The only difference between them is that the number of training realizations was doubled from 100 to 200 when training the PINN-POD model used to generate Fig. 5. The resulting TL contours are visibly smoother and obviously reproduce those in the CNPE-MOST results with much better

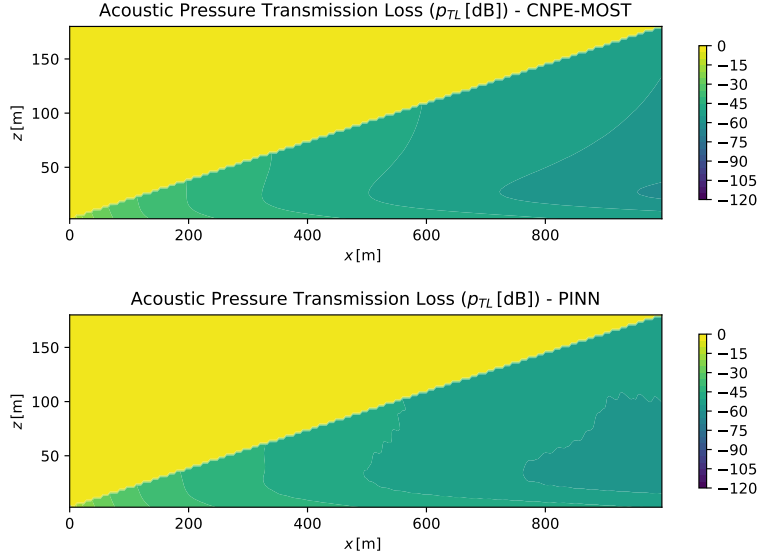


Figure 3: Transmission loss shown as in Figure 2b, but with inputs $\alpha = 32.9^\circ$ and flow resistivity $\sigma = 1.62 \times 10^5 \text{ kPa s m}^{-2}$

accuracy than when only 100 training cases were used.

B. INFERENCE OF PHYSICS PARAMETERS

Raissi, et al.⁵ showed in their Navier-Stokes application that the viscosity could be inferred in parallel as a simple by-product of training the network. Viscosity is an explicit parameter in the physics-informed loss function that represents how well the PINN predictions satisfy the Navier-Stokes equations. The optimizer can adjust it in parallel with adjusting the network’s weights and biases, and eventually converge to an accurate estimate of the viscosity value used when computing the training realizations.

As described above, PINN-POD was found to reliably mimic the spatial characteristics of the outdoor sound field across a wide range of physical parameters governing the wind and temperature profiles, and the impedance of the ground. It was hoped that the PINN-POD training process would, like the Navier-Stokes PINN model just described, be similarly successful in inferring latent physics-model parameters, such as the friction velocity and the surface heat flux. Unfortunately, the optimizer consistently failed to converge on realistic values of these latent physical constants in the outdoor sound model. The results from this part of our study stand in direct contrast to the successful parameter inference results obtained by Raissi, et al. Having run Raissi’s code independently, the authors verified that the viscosity inference results were genuinely good, so it seems that the difficulties are inherent to the complexities of the current application.

The inference difficulties in the outdoor sound application seem to be due to three possible factors: (1) the phase prediction challenges already described, (2) the way in which the physics parameters appear in the outdoor sound model, buried within functions of functions, and (3) the relatively low sensitivity of the global features of the sound field to small variations in the parameters.

Both friction velocity and surface heat flux are hiding in the physics-informed loss terms, but far from direct visibility. Both parameters are latent in the effective wave number variation with height above the ground, and thus compose part of the parabolic equation loss function. Both can be deduced from the complex pressure predictions, but the derivations and resulting relations do not fit within the space constraints of this paper. Numerical tests indicated fairly low sensitivities of the mean-square physics-informed loss to these parameters, at least when the loss function weights were set to values that produced acceptable

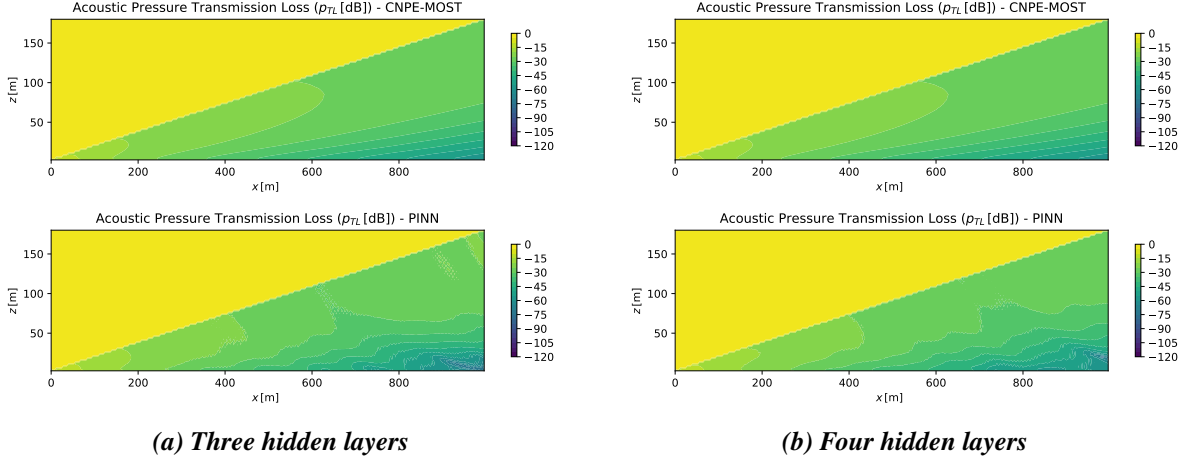


Figure 4: Transmission loss predicted for a 150 Hz point source at $(x, z) = (0 \text{ m}, 1.0 \text{ m})$. The upper figures, which are exactly the same, are the ground truth results from the CNPE-MOST model for wind direction $\alpha = -145^\circ$ and flow resistivity $\sigma = 4.59 \times 10^5 \text{ kPa s m}^{-2}$. The other physics parameters were set at the baseline values described in the text. The lower figures show PINN-POD predictions from networks having three or hidden layers. Each NN had two input neurons and 10 output neurons, one for each POD mode.

accuracy in the transmission loss field. Accordingly, the optimizer did not prioritize them during back-propagation and tended to reduce the overall loss by adjusting the network weights and biases to match $\text{TL}(x, z)$.

Significantly increasing the weights of the physics-informed loss terms forced the optimizer to target those terms more when reducing the overall loss function. In some cases, depending on the initial guesses of the friction velocity and porosity, this led to somewhat more realistic but still largely inaccurate estimates of these parameters. However, owing to the inherent trade-off in multi-objective optimization, emphasizing the physics-informed loss function terms led to less respect for the data-informed terms. Accordingly, the predicted POD coefficients and the resulting transmission loss were much less accurate.

C. NETWORK ARCHITECTURE EFFECTS

As this was an initial capability study, a trial-and-error approach was used to pick the number of network layers and the number of neurons in each layer. Results from the two source frequencies tested here indicated that three hidden layers were needed, with roughly twice as many neurons in each hidden layer as the number of POD coefficients, to accurately reproduce the POD coefficients in the test data. Increasing the depth to four or five layers did not significantly improve accuracy, and in some cases decreased it owing to the relatively small size of the training set. Adjusting the depth of the network and the size of each layer did not significantly improve the poor inference of the latent model parameters.

4. CONCLUSIONS

Two distinct types of PINN models were formulated for the outdoor sound application. The POD-based PINN models approximated the transmission loss fields fairly well because the spatial complexity of the sound field is optimally encoded in the POD modes. The result is analogous to using Fourier series to represent period functions. However, both PINN models failed to produce accurate estimates of the underlying complex pressure field or physics parameters latent in the physics-informed loss function. The primary

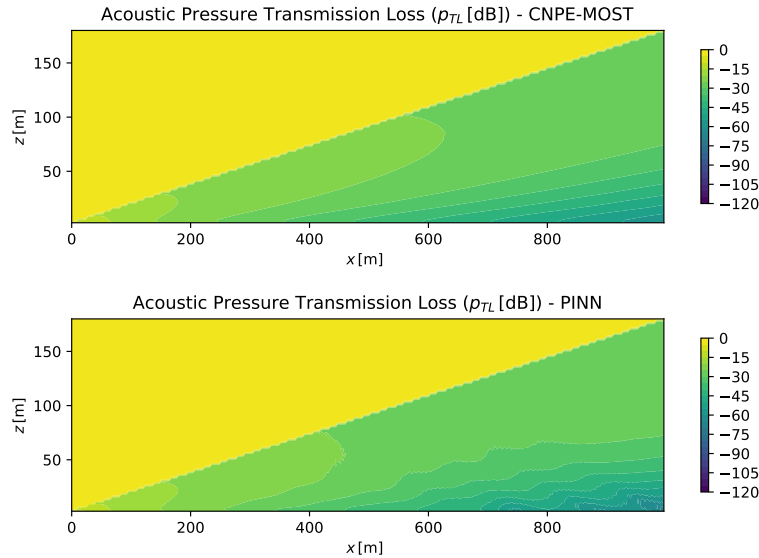


Figure 5: Transmission loss shown as in Figure 4b, but with the number of training realizations increased from 100 to 200.

challenge in using deep learning models, whether physics-informed or not, appears to be in managing the spatial complexity of the complex pressure.

ACKNOWLEDGMENTS

This work was supported by the U.S. Army Engineer Research and Development Center, Geospatial Research Engineering basic research program. Permission to publish was granted by the Director, Cold Regions Research and Engineering Laboratory. Any opinions expressed in this paper are those of the authors, and are not to be construed as official positions of the funding agency or the Department of the Army unless so designated by other authorized documents.

REFERENCES

- ¹ D.K. Wilson, C.L. Pettit, and V.E. Ostashev. “Sound propagation in the atmospheric boundary layer”, *Acoustics Today*, **11(2)**, 44-52 (2015).
- ² C.L. Pettit and D.K. Wilson. “Full-field sensitivity analysis through dimension reduction and probabilistic surrogate models”, *Probabilistic Engineering Mechanics*, (**25**), 380-392 (2010).
- ³ C. Hart, N. Reznicek, D.K. Wilson, C.L. Pettit, and E. Nykaza. “Comparisons between physics-based, engineering, and statistical learning models for outdoor sound propagation”, *Journal of the Acoustical Society of America*, **139**, 2640-2655 (2016).
- ⁴ M. Bianco, P. Gerstoft, J. Traer, E. Ozanich, M. Roch, S. Gannot, and C. Deledalle. “Machine learning in acoustics: Theory and application”, *Journal of the Acoustical Society of America*, **146**, 3590-3628 (2019).

-
- ⁵ M. Raissi, P. Perdikaris, and G.E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”, *Journal of Computational Physics*, **378**, 686-707 (2019).
- ⁶ I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*, MIT Press, 2016.
- ⁷ Y. Yang and P. Perdikaris, “Adversarial uncertainty quantification in physics-informed neural networks”, *Journal of Computational Physics*, **394**, 136-152 (2019).
- ⁸ M. Mungiole and D.K. Wilson. *Prediction of outdoor sound transmission loss with an artificial neural network*, *Applied Acoustics*, **67**, 324-345 (2006).
- ⁹ Y. Chen, L. Lu, G.E. Karniadakis, and L. Dal Negro. “Physics-informed neural networks for inverse problems in nano-optics and metamaterials”, *Optics Express*, **28(8)**, 11618 (2020).
- ¹⁰ P. Perdikaris and S. Wang. “Overcoming gradient pathologies in constrained neural networks”, presented at UCLA Institute of Pure and Applied Mathematics, October 29, 2019.
- ¹¹ B. Lee and D. Dowling. “Comparing analytical and machine learning techniques for predicting acoustic fields in simple environments”, *Meeting of the Acoustical Society of America*, San Diego, December 2019.
- ¹² M. West, K. Gilbert, and R.A. Sack. “A tutorial on the parabolic equation (PE) model used for long range sound propagation in the atmosphere”, *Applied Acoustics*, **37**, 31-49, (1992).
- ¹³ D.K. Wilson. “Simple relaxational models for the acoustical properties of porous media”, *Applied Acoustics*, **50**, 171-188 (1997).
- ¹⁴ H.A. Panofsky and J.A. Dutton. *Atmospheric Turbulence: Models and Methods for Engineering Applications*, Wiley & Sons, New York, 1983.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) June 2021		2. REPORT TYPE Final		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE A Physics-Informed Neural Network for Sound Propagation in the Atmospheric Boundary Layer				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER SAB201	
Chris L Pettit and D. Keith Wilson				5d. PROJECT NUMBER AB2	
				5e. TASK NUMBER 01	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army Engineer Research & Development Center Cold Regions Research & Engineering Laboratory 72 Lyme Road, Hanover, NH 03755				8. PERFORMING ORGANIZATION REPORT NUMBER ERDC/CRREL MP-21-4	
US Naval Academy Dept of Aerospace Engineering Annapolis, MD 21402				10. SPONSOR/MONITOR'S ACRONYM(S) USACE	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Headquarters, US Army Corps of Engineers 441 G Street NW Washington, DC 20314-1000					
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES This paper was originally published in <i>Proceedings of Meetings on Acoustics (POMA)</i> by the Acoustical Society of America on 14 December 2020.					
14. ABSTRACT We describe what we believe is the first effort to develop a physics-informed neural network (PINN) to predict sound propagation through the atmospheric boundary layer. PINN is a recent innovation in the application of deep learning to simulate physics. The motivation is to combine the strengths of data-driven models and physics models, thereby producing a regularized surrogate model using less data than a purely data-driven model. In a PINN, the data-driven loss function is augmented with penalty terms for deviations from the underlying physics, e.g., a governing equation or a boundary condition. Training data are obtained from Crank-Nicholson solutions of the parabolic equation with homogeneous ground impedance and Monin-Obukhov similarity theory for the effective sound speed in the moving atmosphere. Training data are random samples from an ensemble of solutions for combinations of parameters governing the impedance and the effective sound speed. PINN output is processed to produce realizations of transmission loss that look much like the Crank-Nicholson solutions. We describe the framework for implementing PINN for outdoor sound, and we outline practical matters related to network architecture, the size of the training set, the physics-informed loss function, and challenge of managing the spatial complexity of the complex pressure.					
15. SUBJECT TERMS Sound—Propagation, Outdoor sounds, Boundary layer (Meteorology), Boundary layer (Meteorology)					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)