

# Neural Network-Driven Domain Decomposition for Efficient Solutions to the Helmholtz Equation

Victorita Dolean<sup>[0000-0002-5885-1903]</sup>,  
Daria Hrebenshchykova<sup>[0009-0006-0190-7052]</sup>,  
Stéphane Lanteri<sup>[0000-0001-6394-0487]</sup>,  
Victor Michel-Dansac<sup>[0000-0002-3859-8517]</sup>

## 1 Introduction

Accurate simulation of wave propagation is crucial in many scientific and engineering applications, where the frequency-domain Helmholtz equation plays a central role. A particular challenge arises when solving this equation in free space, since the correct mathematical formulation requires the Sommerfeld radiation condition at infinity. This condition can be approximated by an absorbing boundary, such as a Robin boundary condition or perfectly matched layers (PMLs). Handling such open-domain problems accurately is essential to avoid reflections and ensure physically consistent wavefields.

We focus on solving a well-posed time-independent problem, in the form

$$\mathcal{D}[u](\mathbf{x}) = g(\mathbf{x}), \quad \mathbf{x} \in \Omega \subset \mathbb{R}^d, \quad (1)$$

where  $\mathcal{D}$  is some differential operator,  $u$  is the unknown solution and  $\mathbf{x}$  is the space variable, with  $\Omega$  being the space domain. Note that (1) does not contain a boundary operator; this is intentional as boundary condition will be handled by PMLs. To solve (1), we use Physics-Informed Neural Networks (PINNs) proposed

---

Victorita Dolean

Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands, e-mail: v.dolean.maini@tue.nl

Daria Hrebenshchykova

Atlantis project-team, Inria Research Center at Université Côte d'Azur, France, e-mail: daria.hrebenshchykova@inria.fr

Stéphane Lanteri

Atlantis project-team, Inria Research Center at Université Côte d'Azur, France, e-mail: stephane.lanteri@inria.fr

Victor Michel-Dansac

MACARON project-team, Université de Strasbourg, CNRS, Inria, IRMA, France, e-mail: victor.michel-dansac@inria.fr

in [6]. PINNs train a Neural Network (NN) to solve a differential equation by directly approximating the solution, i.e., defining  $v(\cdot, \boldsymbol{\theta}) \approx u(\cdot)$ , where  $\boldsymbol{\theta}$  represents a set of  $P$  trainable parameters of the neural network. The main advantages of PINNs are their mesh-free nature and the fact that we end up with a continuous functional approximation of the solution, defined over the entire domain, rather than a discretised solution on a fixed mesh as in conventional numerical methods. Based on [6], the following loss function is minimized to train the PINN

$$L(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^{N_I} \underbrace{(\mathcal{D}[v_{\boldsymbol{\theta}}](\mathbf{x}_i) - g(\mathbf{x}_i))^2}_{\text{PDE residual}}, \quad (2)$$

where  $\{\mathbf{x}_i\}_{i=1}^{N_I}$  is a set of collocation points sampled in the domain, and where we have defined  $v_{\boldsymbol{\theta}} = \mathbf{x} \mapsto v(\mathbf{x}, \boldsymbol{\theta})$ . As mentioned above, we do not include a boundary loss term here, as the PML formulation ensures wave absorption at the boundaries, removing the necessity for explicit boundary constraints as in standard PINNs.

We move from the standard PINNs formulation to its domain-decomposed extension, namely the Finite Basis PINNs (FBPINNs) introduced in [4], which decomposes the domain into overlapping subdomains. In this paper, we combine them with PML to simulate free-space wave propagation. The goal of this work is thus to assess the performance of the combination of FBPINNs and PMLs for solving the Helmholtz equation.

## 2 Mathematical Model

In this work, we consider the frequency-domain Helmholtz equation in free space, given by

$$-\Delta u(\mathbf{x}) - k^2 u(\mathbf{x}) = g(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (3)$$


for  $\mathbf{x} = (x, y)^{\top} \in \Omega = \mathbb{R}^d$  with  $d = 2$ . Here  $k = 2\pi/\lambda$  is the constant (scalar) wavenumber defined by wavelength  $\lambda = c/f$ , where  $c$  the wave speed and  $f$  the frequency. The unknown function is  $u : \mathbb{R}^2 \rightarrow \mathbb{C}$  and  $g : \mathbb{R}^2 \rightarrow \mathbb{C}$  is a source function. This problem naturally falls into the general framework (1). Physically valid solutions satisfy the Sommerfeld radiation condition at infinity

$$\lim_{|\mathbf{x}| \rightarrow \infty} |\mathbf{x}|^{\frac{d-1}{2}} \left( \frac{\partial u}{\partial |\mathbf{x}|} - iku \right) = 0.$$

### 2.1 Perfectly Matched Layer

To approximate the Sommerfeld radiation condition, the problem is reformulated using a PML transformation introduced by Bérenger [1] in 1994. The PML offers a

more robust approximation than e.g. Robin boundary conditions by creating a finite, artificial absorbing layer around the simulation area.



**Fig. 1** Schematic illustration of the computational domains for PINNs (left) and FBPINNs (right) with PML. The blue rectangle denotes the physical domain  $\Omega$  where we want to find the solution. The red line indicates the outer boundary of the computational domain, including the PML region. The PML is thus located between the red and blue boundaries and ensures wave absorption. In the FBPINNs case (right), the physical domain together with the PML region is further divided into four overlapping subdomains  $\{\Omega_j\}_{j \in \{1, \dots, 4\}}$ .

We are looking for the solution in the region  $\Omega = [-L, L]^2$ , considering a symmetrical rectangular domain without loss of generality, since equivalent results can be obtained for non-symmetric rectangular domains. While the method can be adapted to more general domain geometries by modifying the functions defined below, such extensions are not considered in this work for the sake of conciseness. To apply the PML, the computational domain  $\Omega = [-L - L_{\text{PML}}, L + L_{\text{PML}}]^2$ , where  $L_{\text{PML}}$  represents the width of our absorbing layer. What do you think about the correction The computational domain is illustrated in Figure 1 (left), where the blue square corresponds to the physical region and the red layer represents the PML. At this stage, we consider a single computational domain. The implementation of a PML uses a coordinate stretching technique, thus, we define the coordinate stretching functions

$$D_x : x \mapsto \frac{1}{1 - i \frac{\sigma(x)}{k}} \quad \text{and} \quad D_y : y \mapsto \frac{1}{1 - i \frac{\sigma(y)}{k}}.$$

where

$$\sigma(x) = \sigma_0 k \left( \frac{\max(|x| - L, 0)}{L_{\text{PML}}} \right)^2$$

The tuning constant  $\sigma_0$  and the PML thickness are specified parameters in the numerical experiments. The function  $\sigma$  is defined following [2, 7], but other choices are also possible. Then, thanks to the absorbing layer, we no longer impose a boundary condition explicitly, since it is incorporated into the equation itself. As a result, the differential operator with PML  $\mathcal{D}_{\text{PML}}$ , associated with the Helmholtz equation, defined as

$$\mathcal{D}_{\text{PML}} = -D_x \frac{\partial}{\partial x} \left( D_x \frac{\partial}{\partial x} \right) - D_y \frac{\partial}{\partial y} \left( D_y \frac{\partial}{\partial y} \right) - k^2. \tag{4}$$

## 2.2 FBPINNs

This section defines the mathematical formulation of FBPINNs according to [4]. First, we split the global domain  $\Omega$  into  $J$  overlapping subdomains  $\{\Omega_j\}_{j \in \{1, \dots, J\}}$ . To achieve smoother and more accurate solutions, these subdomains are constructed with an overlap  $\eta$ , which will be specified in the numerical experiments. This overlap is introduced to avoid dealing explicitly with continuity constraints at subdomain interfaces. The approximate solution to our problem (1) is then defined as

$$\tilde{u}(\mathbf{x}, \boldsymbol{\theta}) = \sum_{j=1}^J \varphi_j(\mathbf{x}) v_j(\mathbf{x}; \boldsymbol{\theta}_j), \quad (5)$$

where  $\tilde{u}$  is the weighted combination of the local neural networks  $v_j$ , each trained on its corresponding subdomain  $\Omega_j$ . The window function  $\varphi_j$  restricts the influence of  $v_j$  to its local subdomain while ensuring a smooth transition between neighboring regions.

Each window function must be smooth, differentiable, and form a partition of unity, i.e.  $\sum_{j=1}^J \varphi_j(\mathbf{x}) = 1$ . Like in [3], we use a cosine-based window function

$$\varphi_j(\mathbf{x}) = \frac{\hat{\varphi}_j(\mathbf{x})}{\sum_{m=1}^J \hat{\varphi}_m(\mathbf{x})}, \quad \text{where} \quad \hat{\varphi}_j(\mathbf{x}) = \prod_{n=1}^2 \left[ 1 + \cos\left(\pi \frac{x_n - \mu_{j,n}}{\psi_{j,n}}\right) \right]^2.$$

Here,  $\mu_{j,n}$  denotes the center of the  $j$ -th subdomain along the  $n$ -th spatial axis, and  $\psi_{j,n} = \eta \Delta_n / 2$  defines its half-width, with  $\Delta_n$  the subdomain size along dimension  $n$ .

The FBPINNs loss function is based on (2), but now combined with PML, we no longer include the boundary residual term, since the PML formulation ensures wave absorption at the boundaries. Consequently, the computational domain becomes the extended region  $\Omega_{\text{PML}}$ , and the standard differential operator  $\mathcal{D}$  is replaced by the PML operator  $\mathcal{D}_{\text{PML}}$  defined in (4). The general structure of the FBPINNs combined with PML is shown in Figure 1 (right), where the global domain, including the PML, is partitioned into four overlapping subdomains. Given the ansatz defined by (5), the loss function accounting for both FBPINNs and PML takes the following form

$$L(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N (\mathcal{D}_{\text{PML}}[\tilde{u}_{\boldsymbol{\theta}}](\mathbf{x}_i) - g(\mathbf{x}_i))^2, \quad (6)$$

where  $\tilde{u}_{\boldsymbol{\theta}} = \mathbf{x} \mapsto \tilde{u}(\mathbf{x}, \boldsymbol{\theta})$ .

## 3 Numerical Results

In this section we present numerical results for 2D Helmholtz problem with a Gaussian point source  $g(x, y) = e^{-k(x^2+y^2)}$  in the domain  $\Omega = [-5, 5]^2$ . In our tests,

we set the wave speed  $c = 1$  and vary the frequency  $f$  such that the wavenumber  $k = 2\pi/\lambda$  does not coincide with the eigenfrequencies of the domain, ensuring the existence and uniqueness of the solution. We vary the PML width based on the wavelength  $\lambda = 2\pi/k$ , specifically  $0.25\lambda, 0.5\lambda$  and  $\lambda$ . To ensure that subdomains overlap, the overlap ratio must satisfy  $\eta > 1$ . If it is smaller than 1, the subdomains remain disjoint. In our tests we use an overlap  $\eta = 1.5$ . To help the network find the correct solution for the imaginary part of the problem, we start training with an incorrect source term for the imaginary part that contains small oscillations with frequency  $k$ . After a few epochs, we replace the source term with the correct one, which is zero in our case.

Some PINNs and FBPINNs settings and hyperparameters are fixed across all experiments. Specifically, we use a learning rate of 0.001, a standard multilayer perceptron, 5000 collocation points and *sine* activation function. We divide the domain into four subdomains when using FBPINNs.

We compare the Adam + L-BFGS optimization technique with the Energy Natural Gradient Descent (ENGD) method. Adam and L-BFGS are state-of-the-art optimizers for training PINNs, commonly used in the literature due to their reliability and convergence properties. ENGD, introduced in [5], arises by projecting the optimal function-space descent direction onto the tangent space of the neural-network manifold, yielding a preconditioned gradient in parameter space, that best approximates this optimal direction. It modifies the standard gradient update as follows:

$$\theta_{t+1} = \theta_t - G^\dagger(\theta_t) \nabla_\theta L(\theta_t),$$

where  $G^\dagger$  is the Moore-Penrose pseudoinverse of the matrix  $G$  with components

$$(G(\theta))_{ij} = \int_{\Omega} \partial_{\theta_i}(\mathcal{D}_{\text{PML}}[\tilde{u}_\theta])(\mathbf{x}) \partial_{\theta_j}(\mathcal{D}_{\text{PML}}[\tilde{u}_\theta])(\mathbf{x}) d\mathbf{x}$$

for all  $(i, j) \in \{1, \dots, P\}$ . For experiments with ENGD and Adam optimizer in combination with L-BFGS, where L-BFGS is initiated at a training ratio of 0.7, we use the architecture 3 hidden layers with 32 neurons each for PINNs and 3 hidden layers with 16 neurons in every layer for each NN of FBPINNs.

To evaluate the effect of domain decomposition, we compare the performance of standard PINNs and FBPINNs for different wavenumbers and PML thicknesses. In Tables 1 and 2, we report the relative  $L^2$  error, defined by

$$\varepsilon_{L^2} = \frac{\|u_{\text{approx}} - u_{\text{ref}}\|_2}{\|u_{\text{ref}}\|_2},$$

where  $u_{\text{approx}}$  and  $u_{\text{ref}}$  represent the approximate and reference solutions, respectively. We compute the errors for the real and the imaginary parts separately.

From Table 1 we can see that both methods perform similarly at low frequencies ( $k = 0.57$ ) for both methods of optimizations. However, as the wavenumber increases, the standard PINNs quickly deteriorate, showing errors close to unity already for  $k = 4.51$ , while the FBPINNs maintain a slightly better accuracy. We can

**Table 1** Relative  $L_2$  error between real part of approximate and reference solutions for different  $L_{\text{PML}}$ , comparing PINNs and FBPINNs performance using (a) Adam + L-BFGS and (b) ENGD.

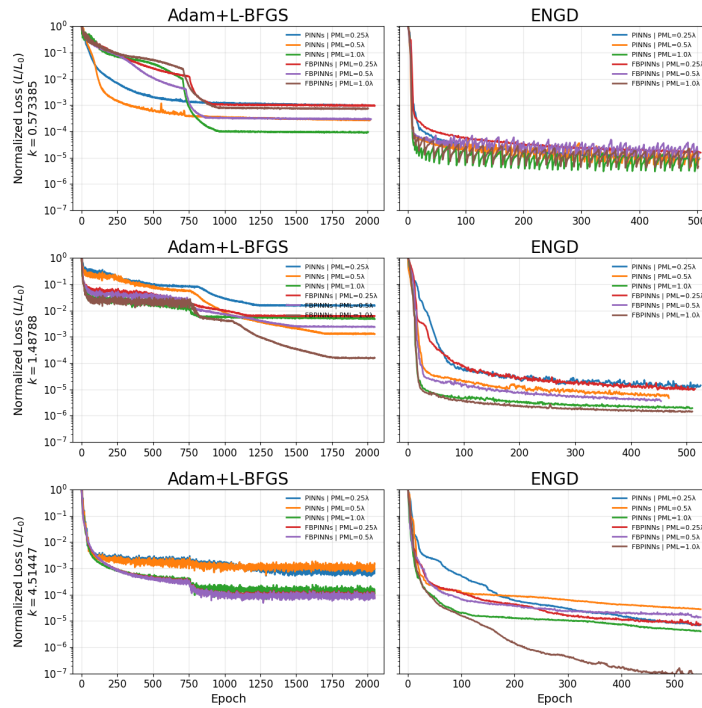
(a) Adam + L-BFGS optimizer					(b) ENGD optimizer				
Method	$k$	$0.25\lambda$	$0.5\lambda$	$1.0\lambda$	Method	$k$	$0.25\lambda$	$0.5\lambda$	$1.0\lambda$
PINNs	0.57	5.6e-2	1.6e-2	4.0e-3	PINNs	0.57	1.7e-1	1.0e-2	5.4e-3
	1.59	4.0e-1	2.9e-2	2.2e-2		1.59	6.7e-1	1.8e-2	4.9e-3
	4.51	9.1e-1	9.7e-1	9.7e-1		4.51	6.7e-1	3.0e-1	1.4e-2
FBPINNs	0.57	4.1e-2	1.5e-2	7.3e-3	FBPINNs	0.57	7.3e-1	6.1e-3	3.3e-3
	1.59	3.0e-1	1.6e-2	1.1e-2		1.59	8.7e-1	1.5e-2	4.1e-3
	4.51	9.3e-1	9.3e-1	9.5e-1		4.51	6.4e-1	1.7e-1	1.1e-2

**Table 2** Relative  $L_2$  error between imaginary part of approximate and reference solutions for different  $L_{\text{PML}}$ , comparing PINNs and FBPINNs performance using (a) Adam + L-BFGS and (b) ENGD.

(a) Adam + L-BFGS optimizer					(b) ENGD optimizer				
Method	$k$	$0.25\lambda$	$0.5\lambda$	$1.0\lambda$	Method	$k$	$0.25\lambda$	$0.5\lambda$	$1.0\lambda$
PINNs	0.57	1.0e-1	2.0e-2	6.2e-3	PINNs	0.57	2.7e-1	1.4e-2	1.8e-3
	1.59	6.4e-1	2.7e-1	8.8e-1		1.59	8.5e-1	2.4e-2	1.4e-3
	4.51	1.0e+0	1.0e+0	1.0e+0		4.51	9.0e-1	4.7e-1	8.2e-2
FBPINNs	0.57	5.6e-1	2.7e-2	1.8e-3	FBPINNs	0.57	5.9e-1	1.3e-2	1.3e-3
	1.59	5.1e-1	5.9e-2	3.3e-2		1.59	6.1e-1	2.0e-2	1.3e-3
	4.51	1.0e+0	1.0e+0	1.0e+0		4.51	9.0e-1	3.7e-1	1.6e-2


observe a similar situation with the imaginary part from Table 2, but by comparing the results, we can conclude that despite our assistance to the imaginary part in the form of preliminary training with another source term, it is still approximated worse than the real part. We assume that this effect appears because there is a local minimum where the imaginary part tends to vanish, which makes the optimization harder for this part. The difference between PINNs and FBPINNs is clearly visible in the convergence curves in Figure 2, where the FBPINNs show a faster decrease in loss and achieve a lower final value, especially for high frequencies. Nevertheless, when the weight of PML is too thin (for example  $L_{\text{PML}} = 0.25\lambda$ ), both approaches struggle to converge due to insufficient wave absorption at the boundaries.

We also compare the performance of Adam + L-BFGS to ENGD. At low frequencies, both optimizers provide similar results, although ENGD demonstrating faster stabilization of the loss that we can see from Figure 2. As the frequency increases, the difference becomes clearer: while the standard optimizer struggles to converge, ENGD succeeds in maintaining stable training and correctly reproducing the oscillatory solution. From the real and imaginary part approximations shown in Figure 3, we observed that ENGD captures the fine wave structures missed by Adam + L-BFGS, even in the imaginary part. Regarding the errors reported in Tables 1 and 2, the results are most evident when we move to the frequency of  $k = 4.5$ , since in this case, especially for the imaginary part, Adam + L-BFGS gives the worst error, while ENGD achieve an accurate approximation. The training speed of ENGD is about six times lower than that of Adam + L-BFGS for both PINNs and FBPINNs. This is mainly due to the additional computation of the energy matrix at each iteration,



**Fig. 2** Normalized physics loss during training for Adam + L-BFGS (left) and ENGD (right). Rows correspond to different  $k$  and curves compare PINNs vs FBPINNs across PML thicknesses.

which significantly increases the computational cost. The use of double precision arithmetic further contributes to the slowdown. In practice, switching to hardware that handles double precision more efficiently could reduce this gap, although the method would still remain slower overall.



**Fig. 3** Approximation and reference solutions of the real and imaginary parts for the Helmholtz problem with  $k = 4.51$  comparing ADAM+L-BFGS and ENGD optimization methods for PINNs and FBPINNs.

## 4 Conclusion

In this work, we showed that FBPINNs combined with PML boundaries can effectively handle wave absorption and provide more accurate solutions than standard PINNs. While the improvement remains moderate, the obtained results confirm that the approach is effective and has potential for further refinement. The use of natural gradient preconditioning also helped to stabilize training and achieve better convergence. In future work, we plan to refine the combination of FBPINNs and PML by introducing internal absorbing layers and to extend the framework toward multilevel FBPINNs as in [3] for more complex wave problems (higher frequencies, heterogeneous media, . . .).

**Acknowledgements** As part of the “France 2030” initiative, this work has benefited from a national grant managed by the French National Research Agency (Agence Nationale de la Recherche) attributed to the Exa-MA project of the NumPEX PEPR program, under the reference ANR-22-EXNU-0002.

**Competing Interests** The authors have no conflicts of interest to declare that are relevant to the content of this chapter.

## References

1. Berenger, J.P.: A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.* **114**(2), 185–200 (1994)
2. Bootland, N., Borzooei, S., Dolean, V., Tournier, P.H.: Numerical assessment of PML transmission conditions in a domain decomposition method for the Helmholtz equation, pp. 445–453. Springer Nature Switzerland (2024)
3. Dolean, V., Heinlein, A., Mishra, S., Moseley, B.: Multilevel domain decomposition-based architectures for physics-informed neural networks. *Comput. Methods Appl. Mech. Eng.* **429**, 117116 (2024)
4. Moseley, B., Markham, A., Nissen-Meyer, T.: Finite basis physics-informed neural networks (FBPINNs): a scalable domain decomposition approach for solving differential equations. *Adv. Comput. Math.* **49**(4) (2023)
5. Müller, J., Zeinhofer, M.: Achieving high accuracy with PINNs via energy natural gradient descent. In: Proceedings of the 40th International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 202, pp. 25471–25485. PMLR (2023)
6. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019)
7. Song, C., Alkhalifah, T., Bin Waheed, U.: A versatile framework to solve the Helmholtz equation using physics-informed neural networks. *Geophys. J. Int.* **228** (2021)