

# Multilevel domain decomposition-based architectures for physics-informed neural networks.

V. Dolean and A. Heinlein and S. Mishra and B. Moseley

Research Report No. 2023-26  
June 2023

Seminar für Angewandte Mathematik  
Eidgenössische Technische Hochschule  
CH-8092 Zürich  
Switzerland

---

1                   **MULTILEVEL DOMAIN DECOMPOSITION-BASED**  
2                   **ARCHITECTURES FOR PHYSICS-INFORMED NEURAL**  
3                   **NETWORKS**

4   VICTORITA DOLEAN\*, ALEXANDER HEINLEIN†, SIDDHARTHA MISHRA‡, AND BEN  
5   MOSELEY§

6   **Abstract.** Physics-informed neural networks (PINNs) are a popular and powerful approach for  
7   solving problems involving differential equations, yet they often struggle to solve problems with high  
8   frequency and/or multi-scale solutions. Finite basis physics-informed neural networks (FBPINNs)  
9   improve the performance of PINNs in this regime by combining them with an overlapping domain  
10   decomposition approach. In this paper, the FBPINN approach is extended by adding multiple levels  
11   of domain decompositions to their solution ansatz, inspired by classical multilevel Schwarz domain  
12   decomposition methods (DDMs). Furthermore, analogous to typical tests for classical DDMs, strong  
13   and weak scaling studies designed for measuring how the accuracy of PINNs and FBPINNs behaves  
14   with respect to computational effort and solution complexity are carried out. Our numerical results  
15   show that the proposed multilevel FBPINNs consistently and significantly outperform PINNs across  
16   a range of problems with high frequency and multi-scale solutions. Furthermore, as expected in  
17   classical DDMs, we show that multilevel FBPINNs improve the scalability of FBPINNs to large  
18   numbers of subdomains by aiding global communication between subdomains.

19   **Key words.** Physics-informed neural networks, overlapping domain decomposition methods,  
20   multilevel methods, multi-scale modeling, spectral bias, forward modeling, differential equations

21   **1. Introduction.** Scientific machine learning (SciML) [3, 52, 12, 2, 35] is an  
22   emerging and rapidly growing field of research. The central goal of SciML is to provide  
23   accurate, efficient, and robust tools for carrying out scientific research by tightly  
24   combining scientific understanding with machine learning (ML). The field has provided  
25   many such tools which have enhanced traditional approaches, from accelerating  
26   simulation algorithms to discovering new scientific phenomena.

27   One popular SciML approach are physics-informed neural networks (PINNs)  
28   [25, 40]. PINNs solve forward and inverse problems related to differential equations  
29   by using a neural network to directly approximate the solution to the differential  
30   equation. They are trained by using a loss function which minimizes the residual of  
31   the differential equation over a set of collocation points. The initial concepts behind  
32   PINNs were introduced in [25], and later re-implemented and extended in [40]. One  
33   of the advantages of PINNs over traditional methods for solving differential equations  
34   such as finite difference (FD) and finite element methods (FEM) is that they provide  
35   a mesh-free approach, paving the way for the application of problems with complex  
36   geometry or in very high spatial dimensions; cf. [33]. Furthermore, they can easily be  
37   extended to solve inverse problems by incorporating observational data.

38   Since their invention, PINNs have been employed across a wide range of domains  
39   [12, 22]. For example, they have been used to solve forward and inverse problems  
40   in geophysics [36], fluid dynamics [21, 6, 41], and optics [10]. Many extensions of  
41   PINNs have also been proposed. For example, PINNs have been extended to carry  
42   out uncertainty quantification [54], learn fast surrogate models [49, 55], and carry out

---

\*University of Strathclyde, UK.

†Delft Institute of Applied Mathematics, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands. [a.heinlein@tudelft.nl](mailto:a.heinlein@tudelft.nl)

‡Seminar for Applied Mathematics (SAM) D-MATH, and ETH AI Center, ETH Zürich, Switzerland. [smishra@ethz.ch](mailto:smishra@ethz.ch)

§Seminar for Applied Mathematics (SAM) D-MATH, and ETH AI Center, ETH Zürich, Switzerland. [benjamin.moseley@ai.ethz.ch](mailto:benjamin.moseley@ai.ethz.ch)

43 equation discovery [11].

44 However, PINNs suffer from a number of limitations. One is that, compared to  
 45 traditional methods, their convergence properties are poorly understood, although  
 46 some work have started to explore this [34, 42, 51]. Another limitation is that, com-  
 47 pared to traditional methods, the computational cost of training PINNs is relatively  
 48 high, especially when they are only used for forward modeling [22]. Finally, a major  
 49 limitation of PINNs is that they often struggle to solve problems with high frequency  
 50 and/or multi-scale solutions [37, 50]. Typically, as higher frequencies and multi-scale  
 51 features are added to the solution, the accuracy of PINNs usually rapidly reduces and  
 52 their computational cost rapidly increases in a super-linear fashion [37].

53 There are multiple reasons for this behavior. One is the spectral bias of neural  
 54 networks, which is the well-studied property that neural networks tend to learn high  
 55 frequencies much slower than low frequencies [53, 39, 4, 9]. Another is that, as higher  
 56 frequencies and more multi-scale features are added, more collocation points and a  
 57 larger neural network with significantly more free parameters are typically required  
 58 to accurately approximate the solution. This creates a significantly more complex  
 59 optimization problem when training the PINN.

60 Recently, [37] proposed finite basis physics-informed neural networks (FBPINNs),  
 61 which aim to improve the performance of PINNs in this regime by using an overlapping  
 62 domain decomposition (DD) approach. In particular, instead of using a single neural  
 63 network to approximate the solution to the differential equation, many smaller neural  
 64 networks were placed in overlapping subdomains and summed together to represent  
 65 the solution. On the one hand, FBPINNs can be seen as a domain decomposition-  
 66 based network architecture for PINNs. On the other hand, by taking this “divide and  
 67 conquer” approach, the global PINN optimization problem is transformed into many  
 68 smaller local optimization problems, which are coupled implicitly due to the overlap  
 69 of the subdomains and their globally defined loss function. The results in [37] show  
 70 that this significantly improves the accuracy and reduces the training cost of PINNs  
 71 when solving differential equations with high frequency and multi-scale solutions.

72 In this work, we significantly extend FBPINNs by incorporating *multilevel model-*  
 73 *ing* into their design. In particular, instead of using a single domain decomposition in  
 74 their solution ansatz, we add multiple levels of overlapping domain decompositions.  
 75 This idea is inspired by classical DDMs, where coarse levels are required for numeri-  
 76 cal scalability when using large numbers of subdomains. Furthermore, to assess the  
 77 performance of multilevel FBPINNs, we define strong and weak scaling tests for mea-  
 78 suring how the accuracy of PINNs and FBPINNs scale with computational effort and  
 79 solution complexity, analogous to the strong and weak scaling tests commonly used  
 80 in classical DDMs.

81 Given these extensions, the performance of PINNs, (one-level) FBPINNs, and  
 82 multilevel FBPINNs across a range of high frequency and multi-scale problems is  
 83 investigated. Across all these tasks, we find that multilevel FBPINNs significantly  
 84 outperform both PINNs and FBPINNs in terms of accuracy and computational cost.  
 85 As expected in classical DDMs, we show that multilevel FBPINNs improve the scal-  
 86 ability of FBPINNs when a large number of subdomains are used by aiding global  
 87 communication between subdomains.

88 The remainder of this work is structured as follows. In [subsection 1.1](#), we discuss  
 89 related work on combining ML, PINNs, and DD, and in [subsections 1.2](#) and [1.3](#), we  
 90 give a brief overview of neural networks and PINNs. Then, we define FBPINNs and  
 91 extend them to multilevel FBPINNs in [section 2](#). Our strong and weak scaling tests  
 92 and corresponding numerical results on the performance of PINNs, FBPINNs, and

93 multilevel FBPINNs across a range of high frequency and multi-scale problems are  
 94 discussed in [section 3](#). Finally, in [section 4](#), we discuss the implications and limitations  
 95 of our work and further research directions.

96 **1.1. Related work.** In general, the idea of combing ML with classical DDMs is  
 97 not new; for early works on using ML to predict the geometrical location of constraints  
 98 in adaptive finite element tearing and interconnecting (FETI) and balancing domain  
 99 decomposition by constraints (BDDC) methods; see [\[17\]](#). An overview of the first  
 100 attempts on combining DD and ML can be found in [\[18\]](#).

101 For specifically combining PINNs with DD, some of the first methods in this area  
 102 were the deep domain decomposition method (D3M) [\[27\]](#), the deep-learning-based  
 103 domain decomposition method (DeepDDM) [\[29, 28\]](#), and its two-level variant [\[32\]](#),  
 104 which use PINNs to solve local problems and overlapping Schwarz steps to iteratively  
 105 connect them based on Lions’ parallel Schwarz algorithm [\[30\]](#). At the same time,  
 106 a series of other extensions, like cPINNs and XPINNs [\[20\]](#) were proposed, which  
 107 similarly divide the domain and use PINNs to solve each local problem; here, typically  
 108 a nonoverlapping domain decomposition is used. The advantage of all these methods  
 109 is their high potential for parallelism, but the downside is the increasing complexity of  
 110 the local loss functions as additional terms are required to enforce coupling between  
 111 subdomains.

112 In contrast, FBPINNs do not require local loss functions nor any additional loss  
 113 terms since they use a globally-defined solution ansatz and loss function [\[37\]](#). Gated-  
 114 PINNs introduced in [\[45\]](#) are perhaps most similar to FBPINNs, where several local  
 115 networks, called experts, are used for training and the domain decomposition itself is  
 116 learned for better efficiency. The idea of learnable domains was also recently exploited  
 117 in XPINNs to improve their performance [\[19\]](#).

118 **1.2. Neural networks.** We first provide a basic definition of a neural network.  
 119 For the purpose of this work, we simply consider a neural network to be a mathemat-  
 120 ical function with some learnable parameters. More precisely, the network is defined  
 121 as  $u(\mathbf{x}, \boldsymbol{\theta}) : \mathbb{R}^{d_x} \times \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}^{d_u}$ , where  $\mathbf{x}$  are some inputs to the network,  $\boldsymbol{\theta}$  are a set  
 122 of learnable parameters, and  $d_x$ ,  $d_\theta$ , and  $d_u$  are the dimensionality of the network’s  
 123 inputs, parameters, and outputs. In a traditional supervised learning setting, learn-  
 124 ing typically consists of fitting the network function to some training data containing  
 125 example inputs and outputs, by minimizing a loss function with respect to  $\boldsymbol{\theta}$  which  
 126 penalizes the difference between the network’s outputs and the training data.

127 The exact form of the network function is determined by the neural network’s  
 128 architecture. In this work, we solely use feedforward fully connected networks (FCNs)  
 129 [\[16\]](#). In this case, the network function is given by

$$130 \quad (1.1) \quad u(\mathbf{x}, \boldsymbol{\theta}) = f_n \circ \dots \circ f_i \circ \dots \circ f_1(\mathbf{x}, \boldsymbol{\theta})$$

131 where now  $\mathbf{x} \in \mathbb{R}^{d_0}$  is the input to the FCN,  $u \in \mathbb{R}^{d_n}$  is the output of the FCN,  $n$  is the  
 132 number of layers (depth) of the FCN, and  $f_i(\mathbf{x}, \boldsymbol{\theta}) = \sigma_i(W_i \mathbf{x} + \mathbf{b}_i)$  where  $\boldsymbol{\theta}_i = (W_i, \mathbf{b}_i)$ ,  
 133  $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$  are known as weight matrices,  $\mathbf{b}_i \in \mathbb{R}^{d_i}$  are known as bias vectors,  $\sigma_i$   
 134 are element-wise activation functions commonly chosen as rectified linear unit (ReLU),  
 135 hyperbolic tangent, or identity functions, and  $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_i, \dots, \boldsymbol{\theta}_n)$  are the set of  
 136 learnable parameters of the network. Note that only the nonlinear activation functions  
 137  $\sigma_i$  facilitate nonlinearity of the network function.

138 **1.3. Physics-informed neural networks.** Physics-informed neural networks  
 139 (PINNs) [\[25, 40\]](#) use neural networks to solve problems related to differential equa-

140 tions. In particular, PINNs focus on solving boundary value problems of the form

$$141 \quad (1.2) \quad \begin{aligned} \mathcal{N}[u](\mathbf{x}) &= f(\mathbf{x}), & \mathbf{x} \in \Omega \subset \mathbb{R}^d, \\ \mathcal{B}_k[u](\mathbf{x}) &= g_k(\mathbf{x}), & \mathbf{x} \in \Gamma_k \subset \partial\Omega \end{aligned}$$

142 where  $\mathcal{N}[u](\mathbf{x})$  is some differential operator,  $u(\mathbf{x})$  is the solution, and  $\mathcal{B}_k(\cdot)$  are a  
 143 set of boundary conditions (BCs) which ensure uniqueness of the solution. For the  
 144 sake of simplicity, we consider BCs in a broad sense; we do not explicitly distinguish  
 145 between initial and boundary conditions, and the  $\mathbf{x}$  variable can include time. Equa-  
 146 tion (1.2) can describe many different differential equation problems, including linear  
 147 and non-linear problems, time-dependent and time-independent problems, and those  
 148 with irregular, higher-order, and cyclic boundary conditions.

149 To solve (1.2), PINNs use a neural network to directly approximate the solution,  
 150 i.e.,  $u(\mathbf{x}, \boldsymbol{\theta}) \approx u(\mathbf{x})$ . Note, for simplicity throughout this work, we use the same nota-  
 151 tion for the true solution and the neural network. It is important to note that PINNs  
 152 provide a functional approximation to the solution, and not a discretized solution  
 153 similar to that provided by traditional methods such as finite difference methods, and  
 154 as such PINNs are a mesh-free approach for solving differential equations. Following  
 155 the approach proposed by [40], the following loss function is minimized to train the  
 156 PINN,

$$157 \quad (1.3) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{\lambda_I}{N_I} \sum_{i=1}^{N_I} \underbrace{\left( \mathcal{N}[u](\mathbf{x}_i, \boldsymbol{\theta}) - f(\mathbf{x}_i) \right)^2}_{\text{PDE residual}} + \sum_{k=1}^{N_k} \frac{\lambda_B^k}{N_B^k} \sum_{i=1}^{N_B^k} \underbrace{\left( \mathcal{B}_k[u](\mathbf{x}_i^k, \boldsymbol{\theta}) - g_k(\mathbf{x}_i^k) \right)^2}_{\text{BC residual}}.$$

158 where  $\{\mathbf{x}_i\}_{i=1}^{N_I}$  is a set of collocation points sampled in the interior of the domain,  
 159  $\{\mathbf{x}_j^k\}_{j=1}^{N_B^k}$  is a set of points sampled along each boundary condition, and  $\lambda_I$  and  $\lambda_B^k$   
 160 are well-chosen scalar weights that ensure the terms in the loss function are well balanced.  
 161 Intuitively, one can see that by minimizing the PDE residual, the method tries to  
 162 ensure that the solution learned by the network obeys the underlying PDE, and by  
 163 minimizing the BC residual, the method tries to ensure that the learned solution is  
 164 unique by matching it to the BCs. Importantly, a sufficient number of collocation  
 165 and boundary points must be chosen such that the PINN is able to learn a consistent  
 166 solution across the domain.

167 Iterative schemes are typically used to optimize this loss function. Usually, vari-  
 168 ants of the gradient descent (GD) method, such as the Adam optimizer [24], or quasi-  
 169 Newton methods, such as the limited-memory Broyden–Fletcher–Goldfarb–Shanno  
 170 (L-BFGS) algorithm [31] are employed. These methods require the computation of  
 171 the gradient of the loss function with respect to the network parameters, which can  
 172 be computed easily and efficiently using automatic differentiation [23] provided in mod-  
 173 ern deep learning libraries [1, 38, 5]. Note that gradients of the network output with  
 174 respect to its inputs are also typically required to evaluate the PDE residual in the loss  
 175 function, and can similarly be obtained and further differentiated through to update  
 176 the network’s parameters using automatic differentiation.

177 **1.3.1. Hard constrained PINNs.** A downside of training PINNs with the  
 178 loss function given by (1.3) is that the BCs are *softly* enforced. This means the  
 179 learned solution may deviate from the BCs because the BC term may not be fully  
 180 minimized. Furthermore, it can be challenging to balance the different objectives  
 181 of the PDE and BC terms in the loss function, which can lead to poor convergence  
 182 and solution accuracy [51, 46]. An alternative approach, as originally proposed by




FIG. 1. *Scaling high frequency problems to low frequency problems using domain decomposition. FBPINNs decompose the domain into many subdomains, and use neural networks within each subdomain to learn the local solution. The input coordinates to each network are normalized to the range  $[-1, 1]$  over their individual subdomains. When solving problems with high frequency solutions, this effectively scales each local problem from a high frequency problem to a lower frequency problem, and helps reduce the network’s spectral bias.*

183 [25], is to enforce BCs in a *hard* fashion by using the neural network as part of a  
 184 solution ansatz. More precisely, the solution to the differential equation is instead  
 185 approximated by  $[Cu](\mathbf{x}, \boldsymbol{\theta}) \approx u(\mathbf{x})$  where  $C$  is an appropriately selected constraining  
 186 operator which analytically enforces the BCs [37, 26].

187 To give a simple example, suppose we want to enforce  $u(x = 0) = 0$  when solving a  
 188 one-dimensional ordinary differential equation (ODE). The constraining operator and  
 189 solution ansatz could be chosen as  $[Cu](x, \boldsymbol{\theta}) = \tanh(x)u(x, \boldsymbol{\theta}) \approx u(\mathbf{x})$ . The rationale  
 190 behind this is that the function  $\tanh(x)$  is zero at 0, forcing the BC to always be  
 191 obeyed, but non-zero away from 0, allowing the network to learn the solution away  
 192 from the BC.

193 In this approach, the BCs are always satisfied and therefore the BC term in the  
 194 loss function (1.3) can be removed, meaning that the PINN can be trained using the  
 195 simpler unconstrained loss function,

$$196 \quad (1.4) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N (n[Cu](\mathbf{x}_i, \boldsymbol{\theta}) - f(\mathbf{x}_i))^2.$$

197 where  $\{\mathbf{x}_i\}_{i=1}^N$  is a set of collocation points sampled in the interior of the domain.  
 198 Note that, in general, there is no unique way of choosing the constraining operator,  
 199 and the definition of a suitable constraining operator for complex geometries and/or  
 200 complex BCs may be difficult or sometimes even impossible, i.e., this strategy is  
 201 problem dependent; in this case, one may resort to the soft enforcement of boundary  
 202 conditions (1.3) instead.

203 **2. Methods.** In this section, we define FBPINNs (subsection 2.1) and extend  
 204 them to multilevel FBPINNs (subsection 2.2). We also discuss the similarities and  
 205 differences of FBPINNs and multilevel FBPINNs to classical DDMs (subsection 2.2.2).

206 **2.1. Finite basis physics-informed neural networks.** As discussed in sec-  
 207 tion 1, a major challenge when training PINNs is that, when higher frequencies  
 208 and multi-scale features are added to the solution, the accuracy of PINNs usually  
 209 rapidly reduces and their computational cost rapidly increases in a super-linear fash-  
 210 ion [37, 50].

211 In the FBPINN approach [37], instead of using a single neural network to represent  
 212 the solution, many smaller neural networks are confined in overlapping subdomains  
 213 and summed together to represent the solution. By taking this “divide and conquer”




FIG. 2. Plot of a square domain  $\Omega$  decomposed into four overlapping subdomains, using a uniform rectangular decomposition.

214 approach, the global PINN optimization problem is transformed into many smaller  
 215 coupled local optimization problems.

216 Furthermore, FBPINNs ensure that the inputs to each subdomain network are  
 217 normalized over their individual subdomain. When solving problems with high fre-  
 218 quency solutions, this effectively scales each local problem from a high frequency  
 219 problem to a lower frequency problem, and helps limit the effect of spectral bias;  
 220 Figure 1 explains this effect further.

**2.1.1. Mathematical definition.** We now provide a mathematical definition of FBPINNs. First, the global solution domain  $\Omega$  is decomposed into  $J$  overlapping subdomains  $\{\Omega_j\}_{j=1}^J$ ; cf. Figure 2. Then, for each subdomain  $\Omega_j$ , a space of network functions is defined,

$$\mathcal{V}_j = \{v_j(\mathbf{x}, \boldsymbol{\theta}_j) \mid \mathbf{x} \in \Omega_j, \boldsymbol{\theta}_j \in \Theta_j\},$$

221 where  $v_j(\mathbf{x}, \boldsymbol{\theta}_j)$  is a neural network placed in each subdomain and  $\Theta_j = \mathbb{R}^{K_j}$  is the  
 222 linear space of all possible network parameters. Here,  $K_j$  is the number of local  
 223 network parameters which is determined by the network architecture.

Next, each subdomain network is confined to its subdomain by multiplying each network with a window function  $\omega_j(\mathbf{x})$ , where  $\text{supp}(\omega_j) \subset \Omega_j$ . Note the neural network functions used in  $\mathcal{V}_j$  generally can have global support, and the window functions are used to restrict them to their individual subdomains. Furthermore, we impose that the window functions form a partition of unity, i.e.,

$$\sum_{j=1}^J \omega_j \equiv 1 \quad \text{on } \Omega.$$

Given the space of network functions and the window functions, we define a global space decomposition given by  $\mathcal{V}$  as

$$\mathcal{V} = \sum_{j=1}^J \omega_j \mathcal{V}_j.$$

224 This space decomposition allows for decomposing any given function  $u \in \mathcal{V}$  as follows

225 (2.1) 
$$u = \sum_{j=1}^J \omega_j v_j \quad \text{or} \quad u(\mathbf{x}, \boldsymbol{\theta}) = \sum_{j=1}^J \omega_j v_j(\mathbf{x}, \boldsymbol{\theta}_j),$$



226 respectively.

227 FBPINNs solve the boundary value problem (2.1) by using equation (2.1) to  
 228 approximate the solution, and we refer to (2.1) as the FBPINN solution. From a  
 229 PINN perspective, the FBPINN solution can simply be thought of as a specific type of  
 230 neural network architecture for the PINN which sums together many locally-confined  
 231 networks to generate the output solution.

The same scheme for training PINNs is used to train the FBPINN. More specifically, the FBPINN solution (2.1) is substituted into the PINN loss function (1.3) and the same iterative optimization scheme is used to learn the parameters  $\{\boldsymbol{\theta}_j\}_{j=1}^J$  of each subdomain network. FBPINNs can also be trained with hard BCs by using the same constraining operator approach described in subsection 1.3.1. In particular, substituting the FBPINN solution (2.1) into the hard-constrained loss function (1.4) yields the loss function

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N (n[c \sum_{j=1}^J \omega_j v_j](\mathbf{x}_i, \boldsymbol{\theta}_j) - f(\mathbf{x}_i))^2.$$

232 Note, naively computing the FBPINN solution (2.1) can be very expensive as it  
 233 requires a summation over all subdomain networks at each collocation point. However,  
 234 this cost can be significantly reduced by noting that only subdomains which contain  
 235 the respective collocation point need to be included; more details on our software  
 236 implementation are provided in Appendix A.

237 **2.2. Multilevel FBPINNs.** Multilevel FBPINNs extend FBPINNs by adding  
 238 multiple levels of domain decompositions to their solution ansatz. They are inspired by  
 239 classical multilevel DD methods, where coarse levels are generally required for numerical  
 240 scalability when using large numbers of subdomains, and multilevel approaches  
 241 may significantly improve performance; see, for instance, [47, 15]. Our hypothesis  
 242 is that multilevel modeling similarly improves the performance of FBPINNs. The  
 243 generalization of FBPINNs to two levels was briefly discussed in [14] and we fully  
 244 introduce the concept here.

A multilevel FBPINN is defined as follows. First, we define  $L$  levels of domain decompositions, where each level,  $l$ , defines an overlapping domain decomposition of  $\Omega$  with  $J^{(l)}$  subdomains, i.e.,

$$D^{(l)} = \left\{ \Omega_j^{(l)} \right\}_{j=1}^{J^{(l)}},$$

245 for  $j = 1, \dots, J^{(l)}$ . Without loss of generality, let  $J^{(1)} = 1$ , that is, on the first level,  
 246 we only have one subdomain  $\Omega_j^{(1)} = \Omega$ . Moreover, we let  $J^{(1)} < J^{(2)} < \dots < J^{(L)}$ ,  
 247 meaning that the number of subdomains increases from one to the next level.

Next, we define spaces of network functions for each level,

$$\mathcal{V}_j^{(l)} = \left\{ v_j^{(l)}(\mathbf{x}, \boldsymbol{\theta}_j^{(l)}) \mid \mathbf{x} \in \Omega_j^{(l)}, \boldsymbol{\theta}_j^{(l)} \in \Theta_j^{(l)} \right\}, \quad j = 1, \dots, J^{(l)}, \quad l = 1, \dots, L,$$

as well as a partition of unity for each level using window functions,  $\omega_j^{(l)}$ , with

$$\text{supp}(\omega_j^{(l)}) \subset \Omega_j^{(l)} \quad \text{and} \quad \sum_{j=1}^{J^{(l)}} \omega_j^{(l)} \equiv 1 \quad \text{on } \Omega \quad \forall l.$$



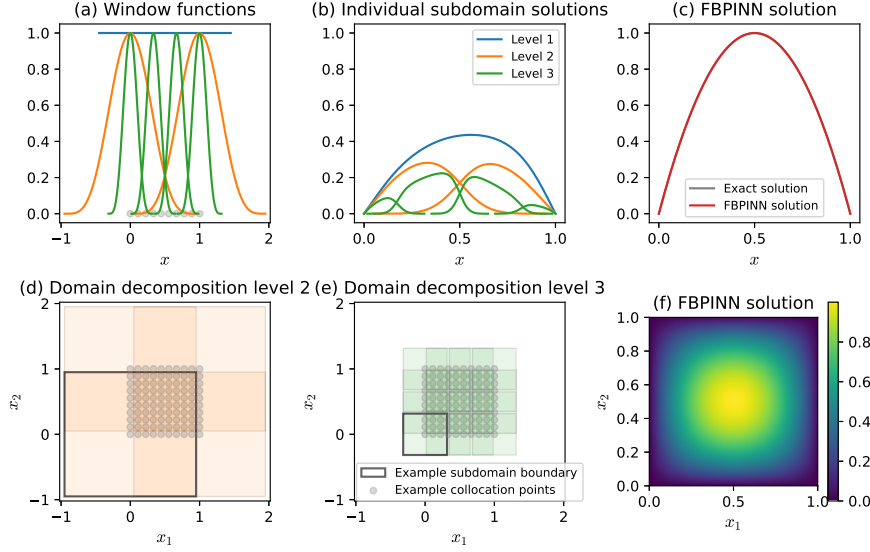


FIG. 3. *Example of a multilevel FBPINN solving Laplace's equation in one and two dimensions. For the 1D problem, the multilevel FBPINN uses  $L = 3$  levels, where each level has 1, 2 and 4 subdomains respectively. The window functions,  $\tilde{\omega}_j^{(l)}(x)$ , used for each level are shown in (a), the individual solutions learned by each subdomain network are shown in (b), and the multilevel FBPINN solution is shown in (c). For the 2D problem, the multilevel FBPINN uses  $L = 3$  levels, where each level has  $1 \times 1$ ,  $2 \times 2$  and  $4 \times 4$  subdomains respectively, using a uniform rectangular domain decomposition. The domain decompositions for level 2 and level 3 are plotted in (d) and (e), and the multilevel FBPINN solution is shown in (f). Note the subdomain boundaries and window functions extend past the problem domain (in this case,  $[0, 1]^d$ ). Example collocation points used to train the multilevel FBPINN are plotted in (a), (d) and (e).*

We can then define a global space decomposition,

$$\mathcal{V} = \frac{1}{L} \sum_{l=1}^L \sum_{j=1}^{J^{(l)}} \omega_j^{(l)} \mathcal{V}_j^{(l)},$$

248 and use this space decomposition to decompose any given function  $u \in \mathcal{V}$  as follows,

$$249 \quad (2.2) \quad u = \frac{1}{L} \sum_{l=1}^L \sum_{j=1}^{J^{(l)}} \omega_j^{(l)} v_j^{(l)} \quad \text{or} \quad u(\mathbf{x}, \boldsymbol{\theta}) = \frac{1}{L} \sum_{l=1}^L \sum_{j=1}^{J^{(l)}} \omega_j^{(l)} v_j^{(l)}(\mathbf{x}, \boldsymbol{\theta}_j^{(l)}).$$

250 We refer to (2.2) as the multilevel FBPINN solution. Note, the original FBPINN  
 251 solution described in subsection 2.1 can be obtained by simply setting  $L = 1$ ; we refer  
 252 to these as one-level FBPINNs going forward.

253 Analogously, we can train multilevel FBPINNs by using the same training scheme  
 254 as PINNs and inserting (2.2) into the PINN loss function. When using the hard-  
 255 constrained PINN loss function (1.4), this yields the corresponding multilevel FBPINN  
 256 loss function

$$257 \quad (2.3) \quad \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N (n [c \frac{1}{L} \sum_{l=1}^L \sum_{j=1}^{J^{(l)}} \omega_j^{(l)} v_j^{(l)}](\mathbf{x}_i, \boldsymbol{\theta}_j^{(l)}) - f(\mathbf{x}_i))^2.$$

**2.2.1. Example of a multilevel FBPINN.** We now show a simple example of a multilevel FBPINN to aid understanding. In particular, we use a multilevel FBPINN to solve the Laplacian boundary value problem,

$$\begin{aligned} -\Delta u &= f \quad \text{in } \Omega = [0, 1]^d, \\ u &= 0 \quad \text{on } \partial\Omega. \end{aligned}$$

258 First we consider the 1D case ( $d = 1$ ), and set  $f = 8$ . Then the exact solution is  
 259 given by  $u(x) = 4x(1 - x)$ .

260 We create an  $L = 3$  level FBPINN to solve this problem, with  $J^{(1)} = 1$ ,  $J^{(2)} = 2$ ,  
 261 and  $J^{(3)} = 4$ . Each level uses a uniform domain decomposition given by

262 (2.4) 
$$\Omega_j^{(l)} = \begin{cases} [0.5 - \delta/2, 0.5 + \delta/2] & l = 1, \\ \left[ \frac{(j-1) - \delta/2}{J^{(l)} - 1}, \frac{(j-1) + \delta/2}{J^{(l)} - 1} \right] & l > 1, \end{cases}$$

263

where  $\delta$  is defined as the *overlap ratio* and is fixed at a value of  $\delta = 1.9$ . Note that an overlap ratio of less than 1 means that the subdomains are no longer overlapping. The subdomain window functions form a partition of unity for each level and are given by

$$\omega_j^{(l)} = \frac{\hat{\omega}_j^{(l)}}{\sum_{j=1}^{J^{(l)}} \hat{\omega}_j^{(l)}} \quad \text{where} \quad \hat{\omega}_j^{(l)}(x) = \begin{cases} 1 & l = 1 \\ [1 + \cos(\pi(x - \mu_j^{(l)})/\sigma_j^{(l)})]^2 & l > 1, \end{cases}$$

264 where  $\mu_j^{(l)} = (j - 1)/(J^{(l)} - 1)$  and  $\sigma_j^{(l)} = (\delta/2)/(J^{(l)} - 1)$  represent the center and  
 265 half-width of each subdomain respectively. Note that FBPINNs are not restricted to  
 266 these particular window functions or partition of unities and any other choice could  
 267 be used instead. The window functions for each level are plotted in [Figure 3 \(a\)](#). A  
 268 FCN (1.1) with 1 hidden layer and 16 hidden units is placed in each subdomain, and  
 269 the  $x$  inputs to each subdomain network are normalized to the range  $[-1, 1]$  over their  
 270 individual subdomains.

271 The multilevel FBPINN is trained using the hard-constrained loss function (2.3)  
 272 with a constraining operator given by  $[Cu](x, \theta) = \tanh(x/\sigma) \tanh((1 - x)/\sigma)u(x, \theta)$   
 273 and  $\sigma = 0.2$ . The loss function is minimized using the Adam optimizer with a learning  
 274 rate of  $1 \times 10^{-3}$  and  $N = 80$  uniformly-spaced collocation points across the domain.

275 The resulting multilevel FBPINN solution is shown in [Figure 3 \(c\)](#), and the in-  
 276 dividual subdomain network solutions (with the constraining operator and window  
 277 function applied) are shown in [Figure 3 \(b\)](#). In this case, we find the FBPINN closely  
 278 matches the exact solution.

279 Next we consider the 2D case ( $d = 2$ ), and set  $f(x_1, x_2) = 32(x_1(1 - x_1) + x_2(1 -$   
 280  $x_2))$ . Then the exact solution is given by  $u(x_1, x_2) = 16(x_1(1 - x_1)x_2(1 - x_2))$ .

281 In this case we create a  $L = 3$  level FBPINN to solve this problem, using a  
 282 uniform rectangular domain decomposition for each level with  $J^{(1)} = 1 \times 1 = 1$ ,  
 283  $J^{(2)} = 2 \times 2 = 4$ , and  $J^{(3)} = 4 \times 4 = 16$ , as shown in [Figure 3 \(d\)](#) and (e). The size  
 284 of each subdomain along each dimension is defined similar as in (2.4) using, again, an  
 285 overlap ratio of  $\delta = 1.9$ . The subdomain window functions are given by

(2.5) 
$$\omega_j^{(l)} = \frac{\hat{\omega}_j^{(l)}}{\sum_{j=1}^{J^{(l)}} \hat{\omega}_j^{(l)}}, \quad \text{where} \quad \hat{\omega}_j^{(l)}(\mathbf{x}) = \begin{cases} 1 & l = 1 \\ \prod_i^d [1 + \cos(\pi(x_i - \mu_{ij}^{(l)})/\sigma_{ij}^{(l)})]^2 & l > 1, \end{cases}$$

287 where  $\mu_{ij}^{(l)}$  and  $\sigma_{ij}^{(l)}$  represent the center and half-width of each subdomain along each  
 288 dimension, respectively. An FCN (1.1) with 1 hidden layer and 16 hidden units is

289 placed in each subdomain, and the  $\mathbf{x}$  inputs to each subdomain network are normalized  
 290 to the range  $[-1,1]$  along each dimension over their individual subdomains.

Similar to above, the multilevel FBPINN is trained using the hard-constrained loss function (2.3), using a constraining operator given by

$$[\mathcal{C}u](\mathbf{x}, \boldsymbol{\theta}) = \tanh(x_1/\sigma) \tanh((1-x_1)/\sigma) \tanh(x_2/\sigma) \tanh((1-x_2)/\sigma)u(\mathbf{x}, \boldsymbol{\theta}),$$

291 with  $\sigma = 0.2$ . The loss function is minimized using the Adam optimizer with a learning  
 292 rate of  $1 \times 10^{-3}$  and  $N = 80 \times 80 = 6,400$  uniformly-spaced collocation points across  
 293 the domain.

294 The resulting multilevel FBPINN solution is shown in Figure 3 (f). Similar to the  
 295 1D case, we find the multilevel FBPINN solution closely matches the exact solution.

296 **2.2.2. Multilevel FBPINNs versus classical multilevel DDMs.** Whilst  
 297 multilevel FBPINNs are inspired by classical multilevel DDMs, a number of differences  
 298 and similarities exist between these approaches. We believe it is insightful to briefly  
 299 discuss these below.

300 Most classical DDMs can be described in terms of the abstract Schwarz frame-  
 301 work [44, 47]. Similar to FBPINNs, this framework is based on a decomposition of a  
 302 global function space  $V$  into local spaces  $\{V_j\}_{j=1}^J$  defined on overlapping subdomains  
 303  $\Omega_j$ , where

$$304 \quad (2.6) \quad V = \sum_{j=1}^J R_j^\top V_j.$$

305 Here,  $R_j^\top : V_j \rightarrow V$  is an interpolation respectively prolongation operator from the  
 306 local into the global space. These notions can be defined in a similar fashion at  
 307 the continuous and discrete level. For the sake of simplicity, we suppose here a  
 308 variational discretisation of the PDE to solve. The space decomposition (2.6) allows  
 309 for decomposing any given discrete function  $u \in V$  as

$$310 \quad (2.7) \quad u = \sum_{j=1}^J R_j^\top v_j, \quad v_j \in V_j;$$

311 due to the overlap, this decomposition is generally not unique. Schwarz DDMs are  
 312 then based on solving local overlapping problems corresponding to the local spaces  
 313  $\{V_j\}_{j=1}^J$  and merging them via the prolongation operators  $R_j^\top$ .

314 Classical one-level Schwarz methods based on this framework are typically not  
 315 scalable to large numbers of subdomains. In particular, since information is only  
 316 transported via the overlap, their rate of convergence will deteriorate when increasing  
 317 the number of subdomains [47]. In order to fix this, multilevel methods add coarser  
 318 problems to the Schwarz framework to facilitate the global transfer of information; in  
 319 particular, the coarsest level typically corresponds to a global problem.

320 We note that:

- 321 • In classical Schwarz methods, the global discretization space  $V$  is often fixed  
 322 first, and then, the local spaces  $\{V_j\}_{j=1}^J$  are constructed. In FBPINNs, we do  
 323 the opposite; we define a local space of neural network functions on an overlap-  
 324 ping domain decomposition  $\{\Omega_j\}_{j=1}^J$  and construct the global discretization  
 325 space from them.

- 326 • In classical Schwarz methods, the local functions  $v_j \in V_j$  are generally not de-  
 327 fined on the global domain  $\Omega$  outside the overlapping subdomain  $\Omega_j$ ; the pro-  
 328 longation operators  $R_j^\top$  extend the local functions to  $\Omega$  such that  $\text{supp}(R_j^\top v_j) \subset$   
 329  $\Omega_j, \forall v_j \in V_j$ . On the other hand, in FBPINNs, the local neural network func-  
 330 tions  $v_j$  generally have global support, and the window functions  $\omega_j$  are used  
 331 to confine them to their subdomains. This difference stems from the fact  
 332 that the local neural networks are not based on a spatial discretization but a  
 333 function approximation; cf. [subsection 1.3](#). Nonetheless, both the prolonga-  
 334 tion operators and the window functions ensure locality; cf. (2.1) and (2.7).  
 335 Note that the prolongation operators in the restricted additive Schwarz (RAS)  
 336 method [8] also include a partition of unity, such that they are very close to  
 337 the window functions in FBPINNs.
- 338 • A key difference is how the boundary value problem is solved. Whereas in  
 339 domain decomposition methods, local subdomain problems are explicitly de-  
 340 fined and solved in a global iteration, in FBPINNs, the global loss function  
 341 is minimized. Moreover, classical DDMs can exploit properties of the sys-  
 342 tem to be solved. For instance, if the PDE is linear elliptic, convergence  
 343 guarantees for classical DDMs can be derived; cf. [47, 15]. In FBPINNs, we  
 344 always have to solve a non-convex optimization problem (1.3) or (1.4), which  
 345 makes the derivation of convergence bounds difficult. Note that there are  
 346 also nonlinear overlapping domain decomposition methods, for instance, ad-  
 347 ditive Schwarz preconditioned inexact Newton (ASPIN) and additive Schwarz  
 348 preconditioned exact Newton (ASPIN) methods [7].

349 **3. Numerical results.** In this section, we assess the performance of multilevel  
 350 FBPINNs. In particular, we investigate the accuracy and computational cost of using  
 351 multilevel FBPINNs to solve various differential equations, and compare them to  
 352 PINNs and one-level FBPINNs.

353 First, in [subsection 3.1](#), we introduce the problems studied. Then, in [subsec-](#)  
 354 [tion 3.2](#) we introduce a notion of strong and weak scaling, inspired by classical DDMs,  
 355 for assessing how the accuracy of FBPINNs and PINNs scales with computational ef-  
 356 fort and solution complexity. In [subsection 3.3](#), we list the common implementation  
 357 details used across all experiments. Finally, in [subsection 3.4](#) we present our numerical  
 358 results.

359 **3.1. Problems studied.** The following problems are used to assess the perfor-  
 360 mance of multilevel FBPINNs;

361 **3.1.1. Homogeneous Laplacian problem in two dimensions.** First, we  
 362 consider the 2D homogeneous Laplacian problem already presented above, namely

$$363 \quad (3.1) \quad \begin{aligned} -\Delta u &= f && \text{in } \Omega = [0, 1]^2, \\ u &= 0 && \text{on } \partial\Omega, \end{aligned}$$

where

$$f(x_1, x_2) = 32(x_1(1 - x_1) + x_2(1 - x_2)).$$

In this case, the exact solution is given by

$$u(x_1, x_2) = 16(x_1(1 - x_1)x_2(1 - x_2)).$$

364 This problem is used to carry out simple ablation tests of the multilevel FBPINN.  
 365 In particular, we assess how varying the number of levels and subdomains as well as the

366 overlap ratio and size of the subdomain networks (architecture) affects the multilevel  
367 FBPINN performance.

368 **3.1.2. Multi-scale Laplacian problem in two dimensions.** Next, we con-  
369 sider a multi-scale variant of the Laplacian problem (3.1) above by using the source  
370 term

$$371 \quad (3.2) \quad f(x_1, x_2) = \frac{2}{n} \sum_{i=1}^n (\omega_i \pi)^2 \sin(\omega_i \pi x_1) \sin(\omega_i \pi x_2).$$

Then, the exact solution is given by

$$u(x_1, x_2) = \frac{1}{n} \sum_{i=1}^n \sin(\omega_i \pi x_1) \sin(\omega_i \pi x_2).$$

372 In this case, multi-scale frequencies are contained in the solution, and the values  
373 of  $n$  and  $\omega_i$  allow us to choose the number of components and the frequency of each  
374 component. We use this problem to assess how the performance of the multilevel  
375 FBPINN scales when more multi-scale components are added to the solution.

376 **3.1.3. Helmholtz problem in two dimensions.** Finally, we study the 2D  
377 Helmholtz problem

$$378 \quad (3.3) \quad \begin{aligned} \Delta u - k^2 u &= f \quad \text{in } \Omega = [0, 1]^2, \\ u &= 0 \quad \text{on } \partial\Omega, \\ f(\mathbf{x}) &= e^{-\frac{1}{2}(\|\mathbf{x} - 0.5\|/\sigma)^2}, \end{aligned}$$

379 with a constant (scalar) wave number,  $k$ . Here, homogeneous Dirichlet boundary  
380 conditions and a Gaussian point source with a scalar width,  $\sigma$ , placed in the center  
381 of the domain are used. Note that, for this problem, the exact solution is not known,  
382 and instead, we compare our models to the solution obtained from FD modeling, as  
383 described in [Appendix B](#).

384 In this case, the solution contains complex patterns of standing waves where the  
385 dominant frequency of the solution depends on the wave number,  $k$ . We use this prob-  
386 lem to test the multilevel FBPINN on a more realistic problem. We first carry out  
387 some simple ablation tests by assessing how varying the number of levels, subdomains,  
388 overlap ratio and size of the subdomain networks affects the multilevel FBPINN per-  
389 formance. Then, we assess how the performance of the multilevel FBPINN scales  
390 when the value of  $k$  is increased.

391 **3.2. Definition of strong and weak scaling.** For both the multi-scale Lapla-  
392 cian and Helmholtz problems, we carry out strong and weak scaling tests. These assess  
393 how the accuracy of the multilevel FBPINN scales with computational effort and so-  
394 lution complexity and are inspired by the strong and weak scaling tests commonly  
395 used in classical DD. They are defined in the following way;

- 396 • *Strong scaling:* We fix the complexity of the problem and increase the model  
397 capacity. For optimal scaling, we expect the convergence rate and/or accuracy  
398 to improve at the same rate as the increase of model capacity.
- 399 • *Weak scaling:* We increase the complexity of the problem and the model  
400 capacity at the same rate. For optimal scaling, we expect the convergence  
401 rate and/or accuracy to stay approximately constant.




FIG. 4. Hierarchy of levels used in the multilevel FBPINN. For all the multilevel FBPINNs tested we use an exponential level structure. This means that the number of subdomains in each level is given by  $2^{d(l-1)}$ , where  $l$  is the level number and  $d$  is the dimensionality of the domain. Our hypothesis is that this helps the multilevel FBPINN model solutions with frequency components that span multiple orders of magnitude.

402 For all our tests, increasing the model capacity means increasing the number of  
 403 levels, number of subdomains, size of the subdomain networks, and/or the number  
 404 of collocation points used. The exact factors varied and their rates of increase are  
 405 detailed in the relevant results sections below. Note all of the multilevel FBPINNs  
 406 tested have been trained on a single GPU, and hence we only show strong and weak  
 407 scaling tests with respect to model capacity and not hardware parallelization.

408 **3.3. Common implementation details.** Some of the implementation details  
 409 of the multilevel FBPINNs, one-level FBPINNs, and PINNs tested are the same across  
 410 all tests. These details are presented here; some are only changed for ablation studies,  
 411 in which case they are described in the relevant results section below.

412 *Level structure.* Firstly, all multilevel FBPINNs use an exponentially increasing  
 413 number of subdomains per level. In particular, we choose  $J^{(l)} = 2^{d(l-1)}$  for  $l = 1, \dots, L$ .  
 414 This level structure is shown in Figure 4. This constraint is chosen so that the  
 415 multilevel FBPINN is able to contain an exponentially large number of subdomains  
 416 with a relatively small number of levels; our hypothesis is that this helps the multilevel  
 417 FBPINN model solutions with frequency components that span multiple orders of  
 418 magnitude.

419 *Domain decomposition.* All FBPINNs tested use a uniform rectangular domain  
 420 decomposition for each level, with all multilevel FBPINNs having  $2^{l-1}$  subdomains  
 421 along each dimension. The size of each subdomain along each dimension is defined  
 422 similar to (2.4), i.e., all 2D domain decompositions look similar to those shown in  
 423 Figure 3 (d) and (e). Furthermore, all FBPINNs use the same subdomain window  
 424 functions, given by (2.5).

425 *Network architecture.* All the FBPINNs tested use FCNs with identical architec-  
 426 tures as their subdomain networks. The PINNs tested also use FCNs as their network  
 427 architecture. For all the FBPINNs tested, the  $\mathbf{x}$  inputs to each subdomain network  
 428 are normalized to the range  $[-1,1]$  along each dimension over their individual subdo-  
 429 mains. For the PINNs tested, the  $\mathbf{x}$  inputs are normalized to the range  $[-1,1]$  along  
 430 each dimension over the global domain.

431 *Loss function and optimization.* All FBPINNs and PINNs tested use the hard-  
 432 constrained variants of their loss functions. All tests use the Adam optimizer with a




FIG. 5. Ablation tests using the homogeneous Laplacian problem. The convergence curve of a baseline multilevel FBPINN is plotted when changing the number of levels (top right), overlap ratio (bottom left), and number of hidden units for each subdomain network (bottom right). The baseline model has  $L = 3$  levels, an overlap ratio of  $\delta = 1.9$ , and 16 hidden units for each subdomain network. The exact solution is shown (top left). Convergence curves of two other benchmarks are shown; a PINN (bottom right), and one-level FBPINNs with varying numbers of subdomains (top right). The lists which label each model in the top right plot contain the number of subdomains along each dimension for each level in the model.

433 learning rate of  $1 \times 10^{-3}$ . For fairness, the same constraining operator is used across all  
 434 models tested on a given problem. Furthermore, exactly the same collocation points  
 435 are used for training whenever multiple models are compared on a given problem. This  
 436 is similarly the case for all testing points used after training. All models are evaluated  
 437 using the normalized L1 test loss, given by  $\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{M} \sum_i^M \|u(\mathbf{x}_i, \boldsymbol{\theta}) - u(\mathbf{x}_i)\|/\sigma$ ,  
 438 where  $M$  is the number of test points and  $\sigma$  is the standard deviation of the set of  
 439 true solutions  $\{u(\mathbf{x}_i)\}_i^M$ .

440 *Software and hardware implementation.* All FBPINNs and PINNs tested are im-  
 441 plemented using a common training framework written in JAX [5]. Further details  
 442 on our software implementation are given in Appendix A. All models are trained on  
 443 a single NVIDIA RTX 3090 GPU.

444 **3.4. Results.** Here, we will discuss the results for the model problems described  
 445 in subsection 3.1.

446 **3.4.1. Homogeneous Laplacian problem in two dimensions.** First, we  
 447 carry out simple ablation tests of the multilevel FBPINN using the 2D homogeneous  
 448 Laplacian problem described in subsection 3.4.1.

449 To carry out our ablation tests, we first train a baseline multilevel FBPINN  
 450 to solve this problem, using  $L = 3$  levels, an overlap ratio along each dimension  
 451 of  $\delta = 1.9$ , and FCNs with 1 hidden layer and 16 hidden units for each subdo-



452 main network. The multilevel FBPINN is trained using the constraining operator  
 453  $[Cu](\mathbf{x}, \boldsymbol{\theta}) = \tanh(x_1/\sigma) \tanh((1 - x_1)/\sigma) \tanh(x_2/\sigma) \tanh((1 - x_2)/\sigma)u(\mathbf{x}, \boldsymbol{\theta})$  with  
 454  $\sigma = 0.2$ . Here,  $N = 80 \times 80 = 6,400$  uniformly-spaced collocation points and  
 455  $M = 350 \times 350$  uniformly-spaced test points across the global domain are used to  
 456 train and test the multilevel FBPINN, respectively.

457 Given this baseline model, we then vary different hyperparameters over a range  
 458 of values and measure the change in performance. This is carried out for the number  
 459 of levels ranging from  $L = 2$  to 5, the overlap ratio ranging from 1.1 to 2.7, and the  
 460 number of hidden units in the subdomain network ranging from 2 to 32. Our results  
 461 are shown in Figure 5. We observe that the accuracy of the multilevel FBPINN  
 462 does not depend significantly on the number of levels, likely because in this case the  
 463 solution is very simple. However, its accuracy increases as the overlap ratio increases,  
 464 likely because there is more communication between the subdomain networks, which is  
 465 similar to what is expected in classical DDMs. Furthermore, its accuracy increases as  
 466 the number of free parameters of the subdomain networks increases. This is expected  
 467 as the capacity of the model increases. Thus, the multilevel FBPINN has similar  
 468 characteristics to classical DDMs for this problem.

469 We carry out two other benchmark tests. First, we train a PINN with 3 hid-  
 470 den layers and 64 hidden units, and second, we train four one-level FBPINNs with  
 471  $J^{(1)} = 2, 4, 8$ , and 16 subdomains along each dimension, respectively. All other rele-  
 472 vant hyperparameters are kept the same as the baseline model. These results are also  
 473 shown in Figure 5. In these tests, the PINN is able to solve the problem, although  
 474 its final accuracy is lower than the baseline multilevel FBPINN and its convergence  
 475 curve is more unstable. Furthermore, the accuracy of the one-level FBPINN reduces  
 476 as more subdomains are added. This is analogous to the expected behavior of one-  
 477 level classical DDMs, which is not scalable to large numbers of subdomains, and shows  
 478 that coarse levels are required for scalability. It is therefore likely that the additional  
 479 levels in FBPINNs serve the same purpose as in classical DDMs, i.e., they allow direct  
 480 transfer of global information.

481 **3.4.2. Multi-scale Laplacian problem in two dimensions.** Next, we eval-  
 482 uate the strong and weak scalability of the multilevel FBPINN using the multi-scale  
 483 Laplacian problem described in subsection 3.4.2.

484 *Strong scaling test.* First, we carry out a strong scaling test. Here, the problem  
 485 complexity is fixed and we assess how the performance of the multilevel FBPINN  
 486 changes as the capacity of the model is increased. In particular, we fix the problem  
 487 complexity by choosing  $n = 6$  with  $\omega_i = 2^i$  for  $i = 1, \dots, n$  in (3.2). Thus, the solution  
 488 contains 6 multi-scale components with exponentially increasing frequencies. This  
 489 represents a much more challenging problem than the homogeneous problem studied  
 490 above. The exact solution in this case is shown in Figure 6.

491 We increase the capacity of the multilevel FBPINN by increasing the number of  
 492 levels, testing from  $L = 2$  to 7. For each test,  $(5 \times 2^{L-1}) \times (5 \times 2^{L-1})$  uniformly-  
 493 spaced collocation points are used, i.e., the density of collocation points inside the  
 494 subdomains in the highest level of each model is kept constant. The rest of the  
 495 hyperparameters of the multilevel FBPINN are kept fixed across all tests. Namely,  
 496 we use an overlap ratio of  $\delta = 1.9$  and FCNs with 1 hidden layer and 16 hidden units  
 497 for each subdomain network. All models are trained using the constraining operator  
 498  $[Cu](\mathbf{x}, \boldsymbol{\theta}) = \tanh(x_1/\sigma) \tanh((1 - x_1)/\sigma) \tanh(x_2/\sigma) \tanh((1 - x_2)/\sigma)u(\mathbf{x}, \boldsymbol{\theta})$  with  
 499  $\sigma = 1/\omega_n$ .  $M = 350 \times 350$  uniformly-spaced test points are used to test all models.

500 The results of this study are shown in Figure 6. We find that the accuracy

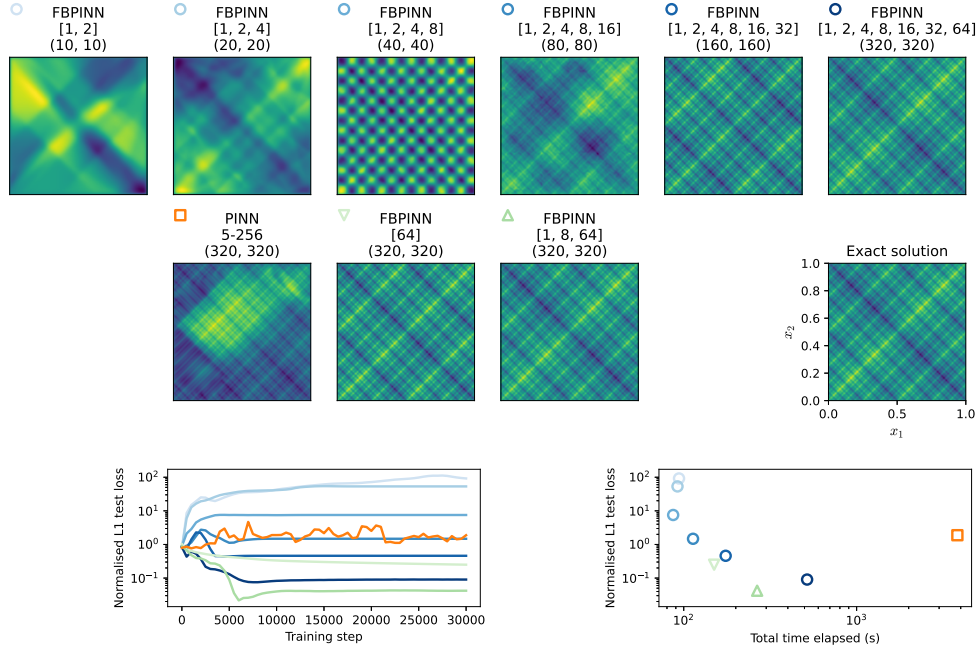


FIG. 6. Strong scaling test using the multi-scale Laplacian problem. In this test the problem complexity is fixed and the solution estimated using multilevel FBPINNs with increasing numbers of levels and collocation points are plotted (top row). The title of each plot describes the level structure (first line) and the number of collocation points along each dimension (second line). The color-coded convergence curves and training times for each model are shown (bottom row). The exact solution is shown (middle row). Plots of the solutions and convergence curves of a PINN, one-level FBPINN and three-level FBPINN benchmark are also shown (middle and bottom row).

501 of the multilevel FBPINN increases as the number of levels increases, where the  
 502  $L = 2, 3, 4$  and 5 models are unable to accurately model the solution, whilst the  
 503  $L = 6$  and 7 models are able to accurately model all of the frequency components. The  
 504 test shows that the multilevel FBPINN is able to solve a high frequency, multi-scale  
 505 problem, and exhibits strong scaling behavior somewhat similar to what is expected  
 506 by classical DDMs.

507 Three other benchmark tests are carried out for this problem. First, we train a  
 508 PINN with 5 hidden layers and 256 hidden units. Then, we train a one-level FBPINN  
 509 with  $J^{(1)} = 64$  subdomains along each dimension and a three-level FBPINN with  
 510  $J^{(1)} = 1$ ,  $J^{(2)} = 8$ , and  $J^{(3)} = 64$  subdomains along each dimension, respectively.  
 511 All other relevant hyperparameters are kept the same as the baseline model above.  
 512 These results are also shown in Figure 6. We find that the accuracy of the PINN is  
 513 poor, and it is only able to model some of the cycles in the solution. Furthermore  
 514 its convergence curve is very unstable, and its training time is an order of magnitude  
 515 larger than the  $L = 7$  level FBPINN tested. Its poor convergence is likely due to  
 516 spectral bias and the increasing complexity of the PINN's optimization problem, as  
 517 discussed in subsection 2.1 and [37]. This shows that the multilevel FBPINN strongly  
 518 outperforms the PINN for this problem. The one-level FBPINN is able to model the  
 519 solution, although its accuracy is less than the  $L = 7$  level FBPINN. Finally, for this  
 520 test the three-level FBPINN benchmark performs best, most accurately modeling the

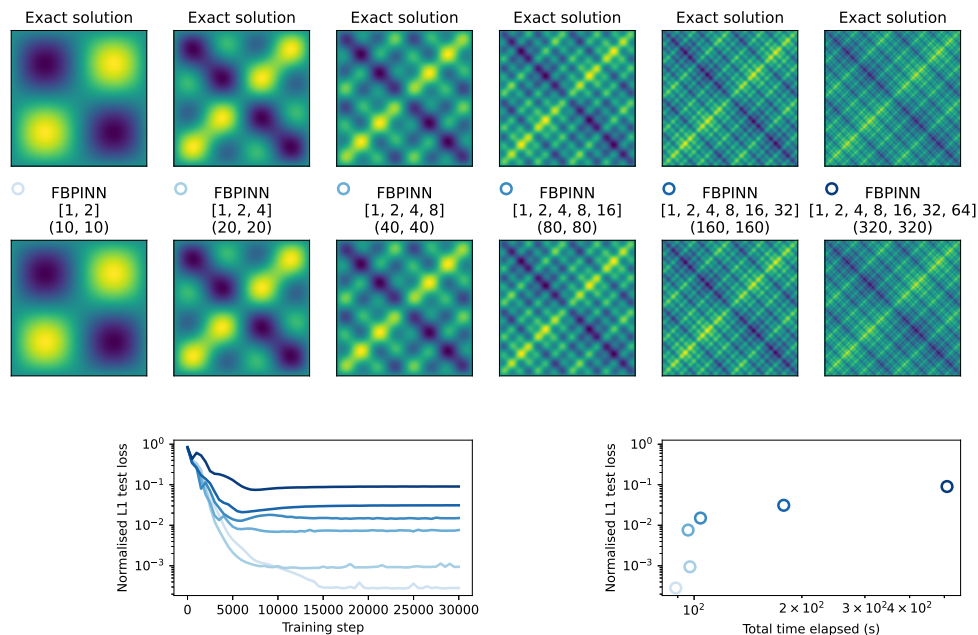


FIG. 7. *Weak scaling test using the multi-scale Laplacian problem. In this test the problem complexity is increased (in this case, the number of frequency components in the solution) (top row) and the solution estimated using multilevel FBPINNs with increasing numbers of levels and collocation points are plotted (middle row). The title of each plot describes the level structure (first line) and the number of collocation points along each dimension (second line). The color-coded convergence curves and training times for each model are shown (bottom row).*

521 solution. This suggests that a stronger coarsening ratio between the levels may be  
 522 beneficial; this is likely to depend on the problem.

523 *Weak scaling test.* Next, we carry out a weak scaling test. Here, both the problem  
 524 complexity and model capacity are scaled at the same rate, and we assess how the  
 525 performance of the multilevel FBPINN changes. We increase the model capacity in  
 526 exactly the same way as the strong scaling test above, i.e., the number of levels is  
 527 increased from  $L = 2$  to 7, where each test has  $(5 \times 2^{L-1}) \times (5 \times 2^{L-1})$  uniformly-  
 528 spaced collocation points. However, now the problem complexity is also scaled, such  
 529 that for each test  $n = L - 1$  and  $\omega_i = 2^i$  for  $i = 1, \dots, n$ . Note that the number of  
 530 subdomains per level and the frequency range of the solution both grow exponentially,  
 531 and the multilevel FBPINN is in alignment with the problem structure. All other  
 532 hyperparameters are fixed to the same values as the strong scaling test above.

533 The results of this test are shown in Figure 7. We find that the multilevel  
 534 FBPINNs are able to model all of the problems tested accurately, that is, model-  
 535 ing all of their frequency components. However, the normalized L1 accuracy of the  
 536 multilevel FBPINNs does reduce slightly as the problem complex increases. Thus in  
 537 this case the multilevel FBPINN exhibits near – but not perfect – weak scaling.

538 **3.4.3. Helmholtz problem in two dimensions.** Finally, we test the mul-  
 539 tiple-level FBPINN using the more complex Helmholtz problem described in [subsec-](#)  
 540 [tion 3.4.3](#). Again, we carry out ablation tests first and then carry out a weak scaling  
 541 study assessing how the performance of the multilevel FBPINN changes as the wave

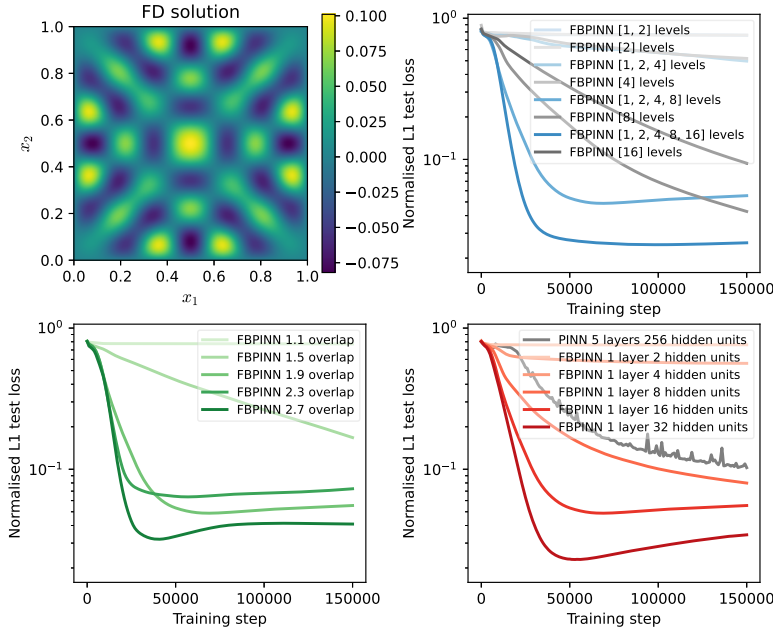


FIG. 8. Ablation tests using the Helmholtz problem. The convergence curve of a baseline multilevel FBPINN is plotted when changing the number of levels (top right), overlap ratio (bottom left), and number of hidden units for each subdomain network (bottom right). The baseline model has  $L = 4$  levels, an overlap ratio of  $\delta = 1.9$ , and 16 hidden units for each subdomain network. The solution obtained from FD modeling is shown (top left). Convergence curves of two other benchmarks are shown; a PINN (bottom right), and one-level FBPINNs with varying numbers of subdomains (top right). The lists which label each model in the top right plot contain the number of subdomains along each dimension for each level in the model.

542 number,  $k$ , increases.

543 *Ablation tests.* For our ablation tests, we fix the problem parameters to be  $k =$   
 544  $2^4\pi/1.6$  and  $\sigma = 0.8/2^4$  in (3.3). Then, similar to subsection 3.4.1, we train a baseline  
 545 multilevel FBPINN to solve this problem, using  $L = 4$  levels, an overlap ratio along  
 546 each dimension of  $\delta = 1.9$ , and FCNs with 1 hidden layer and 16 hidden units for  
 547 each subdomain network. The multilevel FBPINN is trained using the constraining  
 548 operator  $[Cu](\mathbf{x}, \boldsymbol{\theta}) = \tanh(x_1/\sigma) \tanh((1-x_1)/\sigma) \tanh(x_2/\sigma) \tanh((1-x_2)/\sigma)u(\mathbf{x}, \boldsymbol{\theta})$   
 549 with  $\sigma = 1/k$ . We use  $N = 160 \times 160 = 25,600$  uniformly-spaced collocation points  
 550 and  $M = 320 \times 320$  uniformly-spaced test points to train and test the multilevel  
 551 FBPINN, respectively.

552 Given this baseline model, we then vary different hyperparameters over a range of  
 553 values and measure the change in performance. This is carried out for the number  
 554 of levels ranging from  $L = 2$  to 5, the overlap ratio ranging from 1.1 to 2.7, and  
 555 the number of hidden units in the subdomain network ranging from 2 to 32. Our  
 556 results are shown in Figure 8. We obtain similar results to the ablation tests carried  
 557 out in subsection 3.4.1 for the homogeneous Laplace problem. Namely, that the  
 558 accuracy of the multilevel FBPINN improves as the overlap ratio or the number of free  
 559 parameters of the subdomain networks increases. Furthermore, its accuracy improves  
 560 as the number of levels increases, likely because the solution contains relatively high  
 561 frequencies and multiple subdomains are needed.

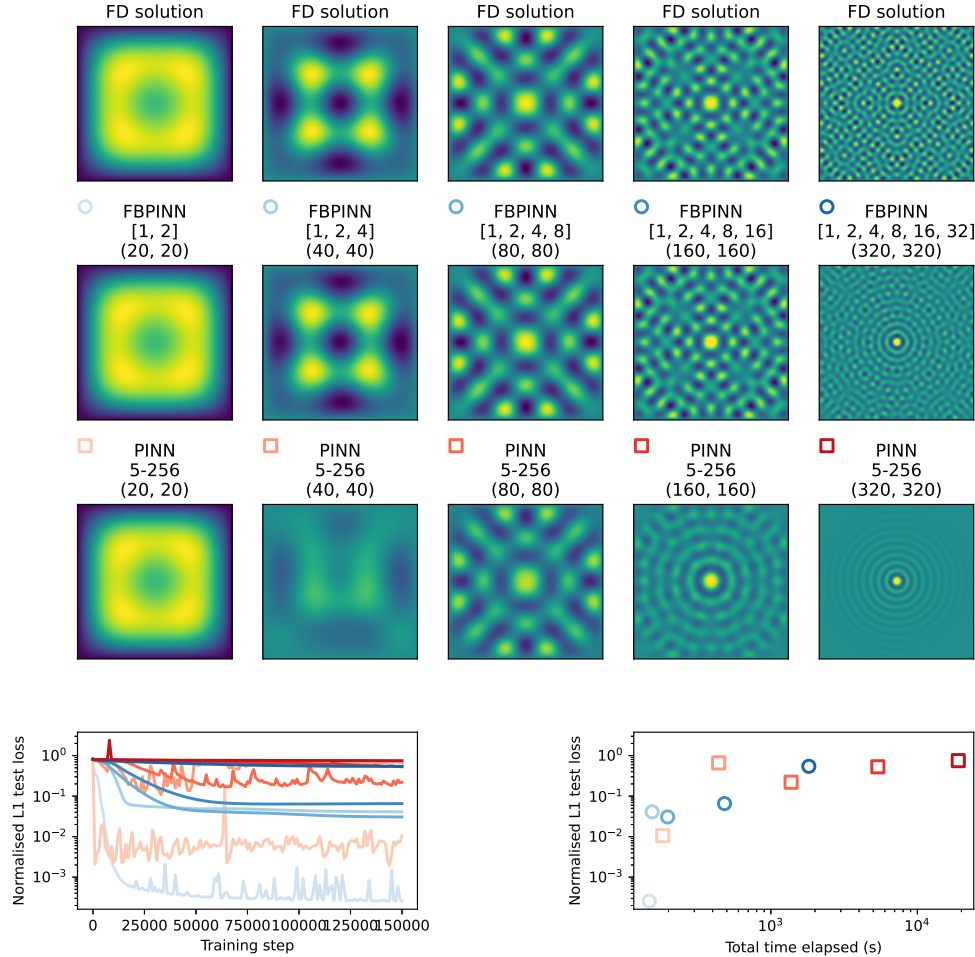


FIG. 9. Weak scaling test using the Helmholtz problem. In this test the problem complexity is increased (in this case, the wave number) (top row) and the solution estimated using multilevel FBPINNs with increasing numbers of levels and collocation points are plotted (second row). The title of each plot describes the level structure (first line) and the number of collocation points along each dimension (second line). The color-coded convergence curves and training times for each model are shown (bottom row). A PINN benchmark using a fixed network size and increasing numbers of collocation points is also shown (third and bottom row).

562 We carry out two other benchmark tests. First, we train a PINN with 5 hid-  
 563 den layers and 256 hidden units, and second, we train four one-level FBPINNs with  
 564  $J^{(1)} = 2, 4, 8,$  and 16 subdomains along each dimension, respectively. All other rel-  
 565 evant hyperparameters are kept the same as the baseline model. These results are  
 566 also shown in Figure 8. Here, the PINN converges poorly, which again highlights the  
 567 shortcomings of PINNs when solving more complex problems. Furthermore, the con-  
 568 vergence of all the one-level FBPINNs is much slower than the multilevel FBPINN,  
 569 and their final accuracy is worse. This again suggests that multiple levels are required  
 570 for scalability.



571 *Weak scaling test.* We carry out a weak scaling study, where both the problem  
 572 complexity and model capacity are scaled at the same rate. In a similar fashion to  
 573 the weak scaling test in [subsection 3.4.2](#), the capacity of the multilevel FBPINN is  
 574 increased by increasing the number of levels, testing from  $L = 2$  to 6. For each test,  
 575  $(10 \times 2^{L-1}) \times (10 \times 2^{L-1})$  uniformly-spaced collocation points are used. The problem  
 576 complexity for each test is increased by setting  $k = 2^L \pi / 1.6$  and  $\sigma = 0.8 / 2^L$  in [\(3.3\)](#).  
 577 All other hyperparameters are fixed to the same values as the baseline model used in  
 578 the ablation tests above.

579 The results of this test are shown in [Figure 9](#). We find that the multilevel FBPINN  
 580 is able to accurately model all the problems tested, except for the highest wave number  
 581 test. In this case, the multilevel FBPINN successfully models the dominant frequency  
 582 and overall concentricity of the solution but fails to model its more complex motifs.  
 583 In this case, we believe that the FBPINN is struggling to satisfy both the point source  
 584 and Dirichlet boundary conditions. Without the Dirichlet boundary condition, the  
 585 solution to [\(3.3\)](#) is that of a simple point source. For all tests, we notice that in the  
 586 first few training steps this is the solution learned by the multilevel FBPINN, which  
 587 is then updated to the correct solution after further training. Thus, it appears the  
 588 presence of the Dirichlet boundary condition leads to an optimization problem which  
 589 remains challenging. Further work is required to understand this behavior; one may  
 590 be able to address this problem by using scheduling strategies to incrementally train  
 591 the multilevel FBPINN, as proposed in [\[37\]](#).

592 Finally, we carry out the same weak scaling test but using a PINN instead of a  
 593 multilevel FBPINN. For each test, the PINN’s architecture is kept fixed at 5 layers  
 594 and 256 hidden units whilst the number of collocation points and problem complexity  
 595 is increased in the same way as the previous test. All other relevant hyperparameters  
 596 are kept the same. The result of this study is shown in [Figure 9](#). In this case, we find  
 597 that the PINN is unable to accurately model any of the solutions, and its training  
 598 time is an order of magnitude larger than the multilevel FBPINN. Thus, the multilevel  
 599 FBPINN still strongly outperforms the PINN for this problem.

600 **4. Discussion.** Across all the problems studied, we find that the multilevel  
 601 FBPINNs consistently outperform the one-level FBPINNs and PINNs tested. The  
 602 multilevel FBPINNs are more accurate than the one-level FBPINNs when a large  
 603 number of subdomains are used, suggesting that coarse levels are required for scala-  
 604 bility by improving the global communication. Furthermore, the multilevel FBPINNs  
 605 significantly outperform the PINNs across all problems tested.

606 We have only started to investigate multilevel FBPINNs in this work and there  
 607 are a range of ways they could be extended. One interesting direction would be to  
 608 investigate more complex domain decompositions and level hierarchies; in this work,  
 609 we restrict ourselves to uniform rectangular decompositions with an exponentially  
 610 increasing number of subdomains with respect to the levels. It is likely that irregular  
 611 domain decompositions would be useful for complex problem geometries, and domain  
 612 decompositions which are tailored to the structure of the solution are likely to help  
 613 where the solution has a large amount of variation. Taking this further, it may  
 614 be possible to learn the domain decomposition itself, for example, by learning the  
 615 parameters of the window functions. This would remove the need to know about the  
 616 solution structure beforehand, and be similar to, e.g., adaptive meshes in traditional  
 617 methods.

618 Furthermore, we only consider one type of window function and partition of unity  
 619 in this work, and it would be useful to assess the impact of different partitioning

620 schemes. We only use small and identical FCNs as subdomain networks, and it would  
 621 be interesting to understand the performance of other network architectures. For  
 622 example, it may be useful to use different size networks for different subdomains if part  
 623 of the solution is more complex and requires a higher capacity model than elsewhere.  
 624 For the Helmholtz problem, we tested sinusoidal activation functions similar to [43,  
 625 28], although we did not notice a significant improvement.

626 Another valuable direction would be to study the theoretical convergence prop-  
 627 erties of multilevel FBPINNs. A major limitation of PINNs compared to classical  
 628 DDMs is that their convergence properties are still poorly understood. In particu-  
 629 lar, whilst the multilevel FBPINN exhibits good scaling properties for the Laplacian  
 630 problems studied, it remains unclear why the optimization of the high wave number  
 631 Helmholtz problem is challenging; note that the convergence of classical DDMs for  
 632 high wave number Helmholtz problems is also not fully understood.

633 A limitation of the multilevel FBPINNs tested is that, despite them being over an  
 634 order of magnitude more efficient than the PINNs tested, their training times are still  
 635 likely to be slower than many traditional methods, such as numerical solvers for finite  
 636 difference or finite element systems. Fundamentally, this is because (FB)PINNs yield  
 637 a non-convex optimization problem, which is relatively expensive compared to the linear  
 638 solves which traditional methods typically rely on. FBPINNs could be extended  
 639 in various ways to reduce their training cost; one direction, as suggested in [37], is to  
 640 provide more inputs to the subdomain networks, such as BCs and PDE coefficients,  
 641 and train across a range of these inputs so that the multilevel FBPINN learns a fast  
 642 surrogate model which does not need to be retrained for each new solution. Another  
 643 option is to implement multi-GPU training; in [37] a parallel FBPINN training al-  
 644 gorithm with minimal communication between subdomains is suggested, which may  
 645 allow highly scalable training. We note that classical numerical solvers can also be  
 646 efficiently parallelized, for instance by using domain decomposition methods. Finally,  
 647 it remains important to test the performance of FBPINNs on 3D problems; only 2D  
 648 problems were studied here, and adding more dimensions is likely to significantly in-  
 649 crease the number of collocation points and subdomains required. These limitations  
 650 will be addressed in future work.

651 **Code availability.** All the code for reproducing the original FBPINN paper [37]  
 652 is available here: <https://github.com/benmoseley/FBPINNs>. All the code for training  
 653 multilevel FBPINNs and reproducing this work will be released on publication.

654 **Appendix A. Software implementation.** All FBPINNs and PINNs are  
 655 implemented using a common training framework written using the JAX automatic  
 656 differentiation library [5]. When training FBPINNs, computing the FBPINN solu-  
 657 tion (either (2.1) or (2.2)) naively can be very expensive. This is because evaluating  
 658 the solution at each collocation point involves summing over all subdomain networks  
 659 and all levels. However, the cost of this summation can be significantly reduced by  
 660 exploiting that, because the output of all subdomain networks is zero outside of the  
 661 corresponding subdomains, only subdomains which contain each collocation point  
 662 contribute to the summation. Practically, this can be carried out by pre-computing a  
 663 mapping describing which subdomains contain each collocation point before training  
 664 and only evaluating the corresponding subdomain networks during training. Another  
 665 important efficiency gain in our software implementation is that the outputs of each  
 666 subdomain network are computed in parallel on the GPU by using JAX’s `vmap` func-  
 667 tionality. This is important as the FBPINNs tested use small subdomain networks  
 668 that if evaluated sequentially would not fully utilize the GPU’s parallelism.



669 **Appendix B. Finite difference solver for Helmholtz equation.** We use a  
 670 finite difference (FD) solver to compute a reference solution for the Helmholtz problem  
 671 studied in subsection 3.4.3. For all the problem variants studied, we discretize the  
 672 Laplacian operator in (3.3) using a 5-point stencil, and we discretize the solution using  
 673 a  $320 \times 320$  uniformly-spaced mesh over the problem domain. This turns (3.3) into a  
 674 set of linear equations, which are solved using the `scipy.sparse.linalg` [48] sparse  
 675 direct solver, that is, using UMFPACK [13].

676

## REFERENCES

- 677 [1] M. ABADI, A. AGARWAL, P. BARHAM, E. BREVDO, Z. CHEN, C. CITRO, G. S. CORRADO,  
 678 A. DAVIS, J. DEAN, M. DEVIN, S. GHEMAWAT, I. GOODFELLOW, A. HARP, G. IRVING,  
 679 M. ISARD, Y. JIA, R. JOZEFOWICZ, L. KAISER, M. KUDLUR, J. LEVENBERG, D. MANE,  
 680 R. MONGA, S. MOORE, D. MURRAY, C. OLAH, M. SCHUSTER, J. SHLENS, B. STEINER,  
 681 I. SUTSKEVER, K. TALWAR, P. TUCKER, V. VANHOUCKE, V. VASUDEVAN, F. VIEGAS,  
 682 O. VINYALS, P. WARDEN, M. WATTENBERG, M. WICKE, Y. YU, AND X. ZHENG, *Tensor-*  
 683 *Flow: Large-Scale Machine Learning on Heterogeneous Distributed Systems*, Mar. 2016,  
 684 <https://doi.org/10.48550/arXiv.1603.04467>.
- 685 [2] S. ARRIDGE, P. MAASS, O. ÖKTEM, AND C.-B. SCHÖNLIEB, *Solving inverse problems us-*  
 686 *ing data-driven models*, *Acta Numerica*, 28 (2019), pp. 1–174, [https://doi.org/10.1017/](https://doi.org/10.1017/s0962492919000059)  
 687 [s0962492919000059](https://doi.org/10.1017/s0962492919000059).
- 688 [3] N. BAKER, F. ALEXANDER, T. BREMER, A. HAGBERG, Y. KEVREKIDIS, H. NAJM, M. PARASHAR,  
 689 A. PATRA, J. SETHIAN, S. WILD, K. WILLCOX, AND S. LEE, *Workshop Report on Basic*  
 690 *Research Needs for Scientific Machine Learning: Core Technologies for Artificial*  
 691 *Intelligence*, tech. report, USDOE Office of Science (SC) (United States), feb 2019,  
 692 <https://doi.org/10.2172/1478744>, <http://www.osti.gov/servlets/purl/1478744/>.
- 693 [4] R. BASRI, D. JACOBS, Y. KASTEN, AND S. KRITCHMAN, *The convergence rate of neural networks*  
 694 *for learned functions of different frequencies*, in *Advances in Neural Information Processing*  
 695 *Systems*, vol. 32, Neural information processing systems foundation, jun 2019, [http://arxiv.](http://arxiv.org/abs/1906.00425)  
 696 [org/abs/1906.00425](http://arxiv.org/abs/1906.00425), <https://arxiv.org/abs/1906.00425>.
- 697 [5] J. BRADBURY, R. FROSTIG, P. HAWKINS, M. J. JOHNSON, C. LEARY, D. MACLAURIN, G. NEC-  
 698 ULA, A. PASZKE, J. VANDERPLAS, S. WANDERMAN-MILNE, AND Q. ZHANG, *JAX: compos-*  
 699 *able transformations of Python+NumPy programs*, 2018, <http://github.com/google/jax>.
- 700 [6] S. CAI, Z. WANG, F. FUEST, Y. J. JEON, C. GRAY, AND G. E. KARNIADAKIS, *Flow over an*  
 701 *espresso cup: Inferring 3-D velocity and pressure fields from tomographic background ori-*  
 702 *ented Schlieren via physics-informed neural networks*, *Journal of Fluid Mechanics*, 915  
 703 (2021), p. 102, [https://doi.org/10.1017/jfm.2021.](https://doi.org/10.1017/jfm.2021.135)  
 704 [135](https://doi.org/10.1017/jfm.2021.135), <https://arxiv.org/abs/2103.02807>.
- 705 [7] X.-C. CAI AND D. E. KEYES, *Nonlinearly preconditioned inexact Newton algorithms*, *SIAM*  
 706 *Journal on Scientific Computing*, 24 (2002), pp. 183–200, [https://doi.org/10.1137/](https://doi.org/10.1137/S106482750037620X)  
 707 [S106482750037620X](https://doi.org/10.1137/S106482750037620X).
- 708 [8] X.-C. CAI AND M. SARKIS, *A restricted additive Schwarz preconditioner for general sparse*  
 709 *linear systems*, *SIAM J. Sci. Comput.*, 21 (1999), pp. 792–797, [https://doi.org/10.1137/](https://doi.org/10.1137/S106482759732678X)  
 710 [S106482759732678X](https://doi.org/10.1137/S106482759732678X).
- 711 [9] Y. CAO, Z. FANG, Y. WU, D.-X. ZHOU, AND Q. GU, *Towards Understanding the Spectral Bias*  
 712 *of Deep Learning*, *IJCAI*, (2021), <http://arxiv.org/abs/1912.01198>, [https://arxiv.org/abs/](https://arxiv.org/abs/1912.01198)  
 713 [1912.01198](https://arxiv.org/abs/1912.01198).
- 714 [10] Y. CHEN, L. LU, G. E. KARNIADAKIS, AND L. DAL NEGRO, *Physics-informed neural net-*  
 715 *works for inverse problems in nano-optics and metamaterials*, *Optics Express*, 28 (2020),  
 716 p. 11618, <https://doi.org/10.1364/oe.384875>, <https://arxiv.org/abs/1912.01085>.
- 717 [11] Z. CHEN, Y. LIU, AND H. SUN, *Physics-informed learning of governing equations from*  
 718 *scarce data*, *Nature Communications*, 12 (2021), pp. 1–13, [https://doi.org/10.1038/](https://doi.org/10.1038/s41467-021-26434-1)  
 719 [s41467-021-26434-1](https://doi.org/10.1038/s41467-021-26434-1), <https://www.nature.com/articles/s41467-021-26434-1>, [https://arxiv.](https://arxiv.org/abs/2005.03448)  
 720 [org/abs/2005.03448](https://arxiv.org/abs/2005.03448).
- 721 [12] S. CUOMO, V. S. DI COLA, F. GIAMPAOLO, G. ROZZA, M. RAISSI, AND F. PICCIALLI, *Scien-*  
 722 *tific Machine Learning Through Physics-Informed Neural Networks: Where we are and*  
 723 *What's Next*, *Journal of Scientific Computing*, 92 (2022), pp. 1–62, [https://doi.org/10.](https://doi.org/10.1007/s10915-022-01939-z)  
 724 [1007/s10915-022-01939-z](https://doi.org/10.1007/s10915-022-01939-z), <https://link.springer.com/article/10.1007/s10915-022-01939-z>,  
 725 <https://arxiv.org/abs/2201.05624>.
- 726 [13] T. A. DAVIS, *Umfpack-an unsymmetric-pattern multifrontal method with a column pre-ordering*

- 727 *strategy*, ACM Trans. Math. Software, 30 (2004), pp. 196–199.
- 728 [14] V. DOLEAN, A. HEINLEIN, S. MISHRA, AND B. MOSELEY, *Finite basis physics-informed neural*  
 729 *networks as a schwarz domain decomposition method*, 2022, [https://doi.org/10.48550/](https://doi.org/10.48550/ARXIV.2211.05560)  
 730 [ARXIV.2211.05560](https://doi.org/10.48550/ARXIV.2211.05560).
- 731 [15] V. DOLEAN, P. JOLIVET, AND F. NATAF, *An introduction to domain decomposition methods*,  
 732 Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2015, [http://](http://dx.doi.org/10.1137/1.9781611974065.ch1)  
 733 [dx.doi.org/10.1137/1.9781611974065.ch1](http://dx.doi.org/10.1137/1.9781611974065.ch1). Algorithms, theory, and parallel implementation.
- 734 [16] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep Learning*, MIT Press, 2016.
- 735 [17] A. HEINLEIN, A. KLAWONN, M. LANSER, AND J. WEBER, *Machine learning in adaptive domain*  
 736 *decomposition methods - predicting the geometric location of constraints*, SIAM Journal  
 737 on Scientific Computing, 41 (2019), pp. A3887–A3912.
- 738 [18] A. HEINLEIN, A. KLAWONN, M. LANSER, AND J. WEBER, *Combining machine learning and*  
 739 *domain decomposition methods for the solution of partial differential equations – a review*,  
 740 GAMM-Mitteilungen, 44 (2021), p. e202100001.
- 741 [19] Z. HU, A. D. JAGTAP, G. E. KARNIADAKIS, AND K. KAWAGUCHI, *Augmented physics-informed*  
 742 *neural networks (apinns): A gating network-based soft domain decomposition methodology*,  
 743 2022, <https://arxiv.org/abs/2211.08939>.
- 744 [20] A. D. JAGTAP AND G. E. KARNIADAKIS, *Extended physics-informed neural networks (XPINNs):*  
 745 *A generalized space-time domain decomposition based deep learning framework for nonlin-*  
 746 *ear partial differential equations*, Communications in Computational Physics, 28 (2020),  
 747 pp. 2002–2041, <https://doi.org/10.4208/CICP.OA-2020-0164>.
- 748 [21] X. JIN, S. CAI, H. LI, AND G. E. KARNIADAKIS, *NSFnets (Navier-Stokes flow nets): Physics-*  
 749 *informed neural networks for the incompressible Navier-Stokes equations*, Journal of  
 750 Computational Physics, 426 (2021), p. 109951, <https://doi.org/10.1016/j.jcp.2020.109951>,  
 751 <https://arxiv.org/abs/2003.06496>.
- 752 [22] G. E. KARNIADAKIS, I. G. KEVREKIDIS, L. LU, P. PERDIKARIS, S. WANG, AND L. YANG, *Physics-*  
 753 *informed machine learning*, Nature Reviews Physics, (2021), pp. 1–19, [https://doi.org/10.](https://doi.org/10.1038/s42254-021-00314-5)  
 754 [1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5), [www.nature.com/natrephys](http://www.nature.com/natrephys).
- 755 [23] H. J. KELLEY, *Gradient Theory of Optimal Flight Paths*, ARS Journal, 30 (1960), pp. 947–  
 756 954, <https://doi.org/10.2514/8.5282>.
- 757 [24] D. P. KINGMA AND J. BA, *Adam: A Method for Stochastic Optimization*, Jan. 2017, [https:](https://doi.org/10.48550/arXiv.1412.6980)  
 758 [//doi.org/10.48550/arXiv.1412.6980](https://doi.org/10.48550/arXiv.1412.6980). arXiv:1412.6980 [cs].
- 759 [25] I. E. LAGARIS, A. LIKAS, AND D. I. FOTIADIS, *Artificial neural networks for solving ordinary and*  
 760 *partial differential equations*, IEEE Transactions on Neural Networks, 9 (1998), pp. 987–  
 761 1000, <https://doi.org/10.1109/72.712178>, <https://arxiv.org/abs/9705023>.
- 762 [26] C. LEAKE AND D. MORTARI, *Deep Theory of Functional Connections: A New Method*  
 763 *for Estimating the Solutions of Partial Differential Equations*, Machine Learning and  
 764 Knowledge Extraction, 2 (2020), pp. 37–55, <https://doi.org/10.3390/make2010004>, <https://europepmc.org/articles/PMC7259480>  
 765 [https://europepmc.org/article/pmc/pmc7259480](https://europepmc.org/articles/PMC7259480).
- 766 [27] K. LI, K. TANG, T. WU, AND Q. LIAO, *D3m: A deep domain decomposition method for partial*  
 767 *differential equations*, IEEE Access, 8 (2019), pp. 5283–5294.
- 768 [28] W. LI, Z. WANG, T. CUI, Y. X. NULL, AND X. XIANG, *Deep Domain Decomposition Methods:*  
 769 *Helmholtz Equation*, Advances in Applied Mathematics and Mechanics, 15 (2023), pp. 118–  
 770 138, <https://doi.org/10.4208/aamm.OA-2021-0305>.
- 771 [29] W. LI, X. XIANG, AND Y. XU, *Deep domain decomposition method: Elliptic problems*, in  
 772 Mathematical and Scientific Machine Learning, PMLR, 2020, pp. 269–286.
- 773 [30] P.-L. LIONS, *On the Schwarz alternating method. I*, in First International Symposium on  
 774 Domain Decomposition Methods for Partial Differential Equations (Paris, 1987), SIAM,  
 775 Philadelphia, PA, 1988, pp. 1–42.
- 776 [31] D. C. LIU AND J. NOCEDAL, *On the limited memory BFGS method for large scale opti-*  
 777 *mization*, Mathematical Programming, 45 (1989), pp. 503–528, [https://doi.org/10.1007/](https://doi.org/10.1007/BF01589116)  
 778 [BF01589116](https://doi.org/10.1007/BF01589116).
- 779 [32] V. MERCIER, S. GRATTON, AND P. BOUDIER, *A coarse space acceleration of deep-ddm*, 2021,  
 780 <https://doi.org/10.48550/ARXIV.2112.03732>.
- 781 [33] S. MISHRA AND R. MOLINARO, *Physics informed neural networks for simulating radiative trans-*  
 782 *fer*, J. Quant. Spectroscopy and Radiative Transfer, 270 (2021).
- 783 [34] S. MISHRA AND R. MOLINARO, *Estimates on the generalization error of physics-informed*  
 784 *neural networks for approximating PDEs*, IMA Journal of Numerical Analysis,  
 785 00 (2022), pp. 1–43, <https://doi.org/10.1093/imanum/drab093>, [https://academic.oup.](https://academic.oup.com/imanum/advance-article/doi/10.1093/imanum/drab093/6503953)  
 786 [com/imanum/advance-article/doi/10.1093/imanum/drab093/6503953](https://academic.oup.com/imanum/advance-article/doi/10.1093/imanum/drab093/6503953), [https://arxiv.org/](https://arxiv.org/abs/2006.16144)  
 787 [abs/2006.16144](https://arxiv.org/abs/2006.16144).
- 788 [35] B. MOSELEY, *Physics-informed machine learning: from concepts to real-world applications*,

- 789 PhD thesis, University of Oxford, 2022, <https://doi.org/10.13039/501100000266>.
- 790 [36] B. MOSELEY, A. MARKHAM, AND T. NISSEN-MEYER, *Solving the wave equation with physics-*  
791 *informed deep learning*, arXiv, (2020), <http://arxiv.org/abs/2006.11894>, <https://arxiv.org/abs/2006.11894>.
- 792
- 793 [37] B. MOSELEY, A. MARKHAM, AND T. NISSEN-MEYER, *Finite Basis Physics-Informed Neural Net-*  
794 *works (FBPINNs): a scalable domain decomposition approach for solving differential equa-*  
795 *tions*, arXiv, (2021), <https://arxiv.org/abs/2107.07871v1><http://arxiv.org/abs/2107.07871>,  
796 <https://arxiv.org/abs/2107.07871>.
- 797 [38] A. PASZKE, S. GROSS, F. MASSA, A. LERER, J. BRADBURY, G. CHANAN, T. KILLEEN, Z. LIN,  
798 N. GIMELSHEIN, L. ANTIGA, A. DESMAISON, A. KOPF, E. YANG, Z. DEVITO, M. RAISON,  
799 A. TEJANI, S. CHILAMKURTHY, B. STEINER, L. FANG, J. BAI, AND S. CHINTALA, *PyTorch:*  
800 *An Imperative Style, High-Performance Deep Learning Library*, in *Advances in Neural In-*  
801 *formation Processing Systems*, vol. 32, Curran Associates, Inc., 2019, [https://proceedings.](https://proceedings.neurips.cc/paper/2019/hash/bdca288fee7f92f2bfa9f7012727740-Abstract.html)  
802 [neurips.cc/paper/2019/hash/bdca288fee7f92f2bfa9f7012727740-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/bdca288fee7f92f2bfa9f7012727740-Abstract.html) (accessed  
803 2023-04-05).
- 804 [39] N. RAHAMAN, A. BARATIN, D. ARPIT, F. DRAXLOR, M. LIN, F. A. HAMPRECHT, Y. BEN-  
805 GIO, AND A. COURVILLE, *On the spectral bias of neural networks*, in *36th Interna-*  
806 *tional Conference on Machine Learning, ICML 2019*, vol. 2019-June, International Ma-  
807 *chine Learning Society (IMLS)*, jun 2019, pp. 9230–9239, <http://arxiv.org/abs/1806.08734>,  
808 <https://arxiv.org/abs/1806.08734>.
- 809 [40] M. RAISSI, P. PERDIKARIS, AND G. E. KARNIADAKIS, *Physics-informed neural networks: A*  
810 *deep learning framework for solving forward and inverse problems involving nonlinear*  
811 *partial differential equations*, *Journal of Computational Physics*, 378 (2019), pp. 686–707,  
812 <https://doi.org/10.1016/j.jcp.2018.10.045>, <https://doi.org/10.1016/j.jcp.2018.10.045>.
- 813 [41] M. RAISSI, A. YAZDANI, AND G. E. KARNIADAKIS, *Hidden fluid mechanics: Learning velocity*  
814 *and pressure fields from flow visualizations*, *Science*, 367 (2020), pp. 1026–1030, <https://doi.org/10.1126/science.aaw4741>.
- 815
- 816 [42] Y. SHIN, J. DARBON, AND G. E. KARNIADAKIS, *On the Convergence of Physics Informed*  
817 *Neural Networks for Linear Second-Order Elliptic and Parabolic Type PDEs*, *Communi-*  
818 *cations in Computational Physics*, 28 (2020), pp. 2042–2074, [https://doi.org/10.4208/cicp.](https://doi.org/10.4208/cicp.oa-2020-0193)  
819 [oa-2020-0193](https://doi.org/10.4208/cicp.oa-2020-0193), <http://arxiv.org/abs/2004.01806>, <https://arxiv.org/abs/2004.01806>.
- 820 [43] V. SITZMANN, J. N. P. MARTEL, A. W. BERGMAN, D. B. LINDELL, AND G. WETZSTEIN, *Implicit*  
821 *Neural Representations with Periodic Activation Functions*, *NeurIPS 2020*, (2020), <http://arxiv.org/abs/2006.09661>,  
822 <https://arxiv.org/abs/2006.09661>, <https://arxiv.org/abs/2006.09661>.
- 823 [44] B. F. SMITH, P. E. BJØRSTAD, AND W. D. GROPP, *Domain decomposition*, Cambridge Uni-  
824 *versity Press, Cambridge*, 1996.
- 825 [45] P. STILLER, F. BETHKE, M. BÖHME, R. PAUSCH, S. TORGE, A. DEBUS, J. VORBERGER,  
826 M. BUSSMANN, AND N. HOFFMANN, *Large-Scale Neural Solvers for Partial Differential*  
827 *Equations*, in *Driving Scientific and Engineering Discoveries Through the Convergence of*  
828 *HPC, Big Data and AI*, J. Nichols, B. Verastegui, A. B. Maccabe, O. Hernandez, S. Parete-  
829 *Koon*, and T. Ahearn, eds., Cham, 2020, Springer International Publishing, pp. 20–34.
- 830 [46] L. SUN, H. GAO, S. PAN, AND J. X. WANG, *Surrogate modeling for fluid flows based on physics-*  
831 *constrained deep learning without simulation data*, *Computer Methods in Applied Me-*  
832 *chanics and Engineering*, 361 (2020), <https://doi.org/10.1016/j.cma.2019.112732>, [http://](http://arxiv.org/abs/1906.02382)  
833 [arxiv.org/abs/1906.02382](http://arxiv.org/abs/1906.02382)<http://dx.doi.org/10.1016/j.cma.2019.112732>, <https://arxiv.org/abs/1906.02382>.
- 834
- 835 [47] A. TOSELLI AND O. WIDLUND, *Domain decomposition methods—algorithms and theory*, vol. 34  
836 *of Springer Series in Computational Mathematics*, Springer-Verlag, Berlin, 2005, <https://doi.org/10.1007/b137868>.
- 837
- 838 [48] P. VIRTANEN, R. GOMMERS, T. E. OLIPHANT, M. HABERLAND, T. REDDY, D. COURNAPEAU,  
839 E. BUROVSKI, P. PETERSON, W. WECKESSER, J. BRIGHT, S. J. VAN DER WALT, M. BRETT,  
840 J. WILSON, K. J. MILLMAN, N. MAYOROV, A. R. J. NELSON, E. JONES, R. KERN, E. LAR-  
841 *son*, C. J. CAREY, İ. POLAT, Y. FENG, E. W. MOORE, J. VANDERPLAS, D. LAXALDE,  
842 J. PERKTOLD, R. CIMRMAN, I. HENRIKSEN, E. A. QUINTERO, C. R. HARRIS, A. M.  
843 *Archibald*, A. H. RIBEIRO, F. PEDREGOSA, P. VAN MULBREGT, AND SCI-PY 1.0 CONTRI-  
844 *butors*, *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature*  
845 *Methods*, 17 (2020), pp. 261–272, <https://doi.org/10.1038/s41592-019-0686-2>.
- 846 [49] S. WANG, H. WANG, AND P. PERDIKARIS, *Learning the solution operator of parametric par-*  
847 *tial differential equations with physics-informed DeepONets*, *Science Advances*, 7 (2021),  
848 pp. 8605–8634, <https://doi.org/10.1126/sciadv.abi8605>, [https://www.science.org/doi/full/](https://www.science.org/doi/full/10.1126/sciadv.abi8605)  
849 [10.1126/sciadv.abi8605](https://www.science.org/doi/full/10.1126/sciadv.abi8605), <https://arxiv.org/abs/2103.10974>.
- 850 [50] S. WANG, H. WANG, AND P. PERDIKARIS, *On the eigenvector bias of Fourier feature net-*

- 851            *works: From regression to solving multi-scale PDEs with physics-informed neural net-*  
852            *works*, *Computer Methods in Applied Mechanics and Engineering*, 384 (2021), p. 113938,  
853            <https://doi.org/10.1016/j.cma.2021.113938>, <https://arxiv.org/abs/2012.10047>.
- 854 [51] S. WANG, X. YU, AND P. PERDIKARIS, *When and why PINNs fail to train: A neural tangent*  
855            *kernel perspective*, *Journal of Computational Physics*, 449 (2022), p. 110768, <https://doi.org/10.1016/j.jcp.2021.110768>, <https://arxiv.org/abs/2007.14527>.
- 856 [52] J. WILLARD, X. JIA, S. XU, M. STEINBACH, AND V. KUMAR, *Integrating Scientific Knowl-*  
857            *edge with Machine Learning for Engineering and Environmental Systems*, *ACM Comput-*  
858            *ing Surveys*, 55 (2022), <https://doi.org/10.1145/3514228>, [https://dl.acm.org/doi/10.1145/](https://dl.acm.org/doi/10.1145/3514228)  
859            [3514228](https://dl.acm.org/doi/10.1145/3514228), <https://arxiv.org/abs/2003.04919>.
- 860 [53] Z. Q. J. XU, Y. ZHANG, T. LUO, Y. XIAO, AND Z. MA, *Frequency principle: Fourier analysis*  
861            *sheds light on deep neural networks*, *Communications in Computational Physics*, 28 (2020),  
862            pp. 1746–1767, <https://doi.org/10.4208/CICP.OA-2020-0085>, [https://arxiv.org/abs/1901.](https://arxiv.org/abs/1901.06523)  
863            [06523](https://arxiv.org/abs/1901.06523).
- 864 [54] L. YANG, X. MENG, AND G. E. KARNIADAKIS, *B-PINNs: Bayesian physics-informed neural*  
865            *networks for forward and inverse PDE problems with noisy data*, *Journal of Computational*  
866            *Physics*, 425 (2021), p. 109913, <https://doi.org/10.1016/j.jcp.2020.109913>.
- 867 [55] Y. ZHU, N. ZABARAS, P. S. KOUTSOURELAKIS, AND P. PERDIKARIS, *Physics-constrained*  
868            *deep learning for high-dimensional surrogate modeling and uncertainty quantifica-*  
869            *tion without labeled data*, *Journal of Computational Physics*, 394 (2019), pp. 56–81,  
870            <https://doi.org/10.1016/j.jcp.2019.05.024>, <http://arxiv.org/abs/1901.06314>[http://dx.doi.](http://dx.doi.org/10.1016/j.jcp.2019.05.024)  
871            [org/10.1016/j.jcp.2019.05.024](http://dx.doi.org/10.1016/j.jcp.2019.05.024), <https://arxiv.org/abs/1901.06314>.
- 872