

Meta-learning pseudo-differential operators with deep neural networks



Jordi Feliu-Fabà^a, Yuwei Fan^{b,*}, Lexing Ying^c

^a ICME, Stanford University, Stanford, CA 94305, United States of America

^b Department of Mathematics, Stanford University, Stanford, CA 94305, United States of America

^c Department of Mathematics and ICME, Stanford University, Stanford, CA 94305, United States of America

ARTICLE INFO

Article history:

Received 16 June 2019

Received in revised form 4 November 2019

Accepted 3 February 2020

Available online 6 February 2020

Keywords:

Deep neural networks

Convolutional neural networks

Nonstandard wavelet form

Meta-learning

Green's functions

Radiative transfer equation

ABSTRACT

This paper introduces a meta-learning approach for parameterized pseudo-differential operators with deep neural networks. With the help of the nonstandard wavelet form, the pseudo-differential operators can be approximated in a compressed form with a collection of vectors. The nonlinear map from the parameter to this collection of vectors and the wavelet transform are learned together from a small number of matrix-vector multiplications of the pseudo-differential operator. Numerical results for Green's functions of elliptic partial differential equations and the radiative transfer equations demonstrate the efficiency and accuracy of the proposed approach.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Many physical models for scientific and engineering applications can be written in a general form

$$L_\eta u(x) = f(x), \quad x \in \Omega \subset \mathbb{R}^d \quad (1.1)$$

for a domain Ω with appropriate boundary conditions, where L_η is often a partial differential or integral operator parameterized by a parameter function $\eta(x)$. Solving for $u(x)$ for a given $f(x)$ amounts to representing the inverse operator (sometimes also known as the Green's function) $G_\eta = L_\eta^{-1}$ either explicitly or implicitly via an efficient algorithm. Representing G_η , even if implicitly, can be computationally challenging, especially for multidimensional problems. The past few decades have witnessed steady progresses in developing efficient algorithms for this.

Problem statement This paper is concerned with a more ambitious task: representing the nonlinear map from η to G_η

$$\mathcal{M} : \eta \rightarrow G_\eta = L_\eta^{-1}, \quad (1.2)$$

when the operator G_η is a pseudo-differential operator (PDO) [60]. Although L_η and G_η can be linear operators, this map \mathcal{M} from η to G_η is highly nonlinear.

* Corresponding author.

E-mail addresses: jfeliu@stanford.edu (J. Feliu-Fabà), ywfan@stanford.edu (Y. Fan), lexing@stanford.edu (L. Ying).

Background In the recent years, deep learning has become the most versatile and effective tool in artificial intelligence and machine learning, witnessed by impressive achievements in computer vision [36,61,28], speech and natural language processing [29,56,51,57,11], drug discovery [43] or game playing [54,15,58]. Recent reviews on deep learning and its impacts on other fields can be found in for example [38,53]. At the center of deep learning, the model of deep neural networks (NNs) provides a flexible framework for approximating high-dimensional functions, while allowing for efficient training and good generalization properties in practice [42,47].

More recently, several groups have started applying NNs to partial differential equations (PDEs) and integral equations (IEs) arising from physical systems. In one direction, the NN model has been used to approximate solutions of high-dimensional PDEs [37,55,14,49,4,6,13,25,32,41]. In a somewhat orthogonal direction, the NNs have been utilized to approximate the high-dimensional parameter-to-solution of various PDEs and IEs [31,26,18,17,19,33,25,2,39,20].

Another topic from machine learning that is particularly relevant to this work is *meta-learning* or *learning-to-learn* [52, 3,27,21,59]. A meta-learning system learns to produce learning models for new tasks and scenarios from their metadata with zero or minimum amount of new data, by leveraging the common structure among different tasks. Due to the low requirements on new data points, meta-learning has gained a lot of attention in recent years in applications such as vision and reinforcement learning.

Main idea Following these recent advances in applying NNs to physical models, this paper takes a deep learning approach for representing the map in Eq. (1.2). The most straightforward solution would be to take a supervised learning approach, i.e., trying to learn the map $\mathcal{M} : \eta \rightarrow G_\eta$ from a large set of training data $\{(\eta_i, G_{\eta_i})\}_i$. However, since it is often difficult or even impossible to compute and store G_{η_i} due to the enormous discretization size, this straightforward supervised learning approach is not practical for Eq. (1.2).

Without explicit access to G_η , we take a meta-learning approach, i.e., learning to produce, for each new η , an NN approximation to G_η . To do this, we are faced with two key difficulties.

- How should we represent the output G_η for an arbitrary input η ?
- How should we represent the training data?

To address the first question, G_η should be represented in a compressed form. For pseudo-differential operators, several compressed representations exist, including hierarchical matrices [22–24], discrete symbol calculus [10], etc. In this paper, we choose to represent G_η with the nonstandard wavelet form introduced in [5]. The main advantage of the nonstandard wavelet form is that the nonzero entries of this compressed representation are simply organized into a small number of vectors. More precisely,

$$G_\eta \approx WS[C_\eta]W^T,$$

where W is a redundant form of a wavelet transform, C_η stands for the collection of vectors that contain the nonzero entries of the compressed form, and S is a certain operator that generates a sparse matrix from the vector collection C_η . Compared to [5], a key difference is that the current approach allows for W to be fine-tuned for the map \mathcal{M} .

To address the second question, instead of explicitly representing G_{η_i} , the training data consists of samples of the form

$$(\eta_i, \{f_{ij}, u_{ij}\}),$$

where $u_{ij} = G_{\eta_i} f_{ij}$. For a fixed η_i , such data can be obtained by solving the equation $L_{\eta_i} u_{ij} = f_{ij}$ for each f_{ij} , possibly with a fast algorithm.

Putting these two pieces together, the meta-learning approach of this paper learns two following key objects from the training data of form $\{(\eta_i, \{f_{ij}, u_{ij}\}_j)\}_i$:

- a map from η to the vector collection C_η ,
- the \mathcal{M} -dependent wavelet transform W .

Once trained, for a given test input η the architecture calculates C_η and returns a linear NN that implements $WS[C_\eta]W^T \approx G_\eta$.

Organization The rest of this paper is organized as follows. Section 2 briefly reviews the nonstandard wavelet form, used for representing G_η . In Section 3, the NN architecture of the meta-learning approach is discussed in detail. Section 4 applies the proposed NN to the Green's function of elliptic PDEs, in both the Schrödinger form and the divergence form. The application to the radiative transfer equation is presented in Section 5.

2. Nonstandard wavelet form

This section summarizes the nonstandard wavelet form proposed in [5]. To make things concrete, compactly supported orthonormal Daubechies wavelets [8] are used as the basis functions as an example.

2.1. Wavelet transform

In the one-dimensional multiresolution analysis, one starts by defining a scaling function $\varphi(x)$ that generates, through dyadic translations and dilations, a family of functions

$$\varphi_k^{(\ell)}(x) = 2^{\ell/2} \varphi(2^\ell x - k), \quad \ell = 0, 1, 2, \dots, \quad k \in \mathbb{Z}. \tag{2.1}$$

For each scale ℓ , the functions $\{\varphi_k^{(\ell)}\}$ form a Ritz basis for a space V_ℓ , which satisfies a nested relationship $V_\ell \subset V_{\ell+1}$. This nested property of $\{V_\ell\}$ implies the following *dilation relation* of the scaling function

$$\varphi(x) = \sqrt{2} \sum_{i \in \mathbb{Z}} h_i \varphi(2x - i). \tag{2.2}$$

For the Daubechies' wavelets [8], the scaling function $\varphi(x)$ has a compact support $[0, 2p - 1]$ for a given positive integer p and therefore the coefficients $\{h_i\}$ are only nonzero for $i = 0, \dots, 2p - 1$. The scaling function also satisfies the orthonormal condition

$$\int_{\mathbb{R}} \varphi(x - a) \varphi(x - b) dx = \delta_{a,b}, \quad \forall a, b \in \mathbb{Z}, \tag{2.3}$$

which leads to an orthonormal condition for the coefficients $\{h_i\}$

$$\sum_{i \in \mathbb{Z}} h_i^2 = 1, \quad \sum_{i \in \mathbb{Z}} h_i h_{i+2m} = 0, \quad m \in \mathbb{Z} \setminus \{0\}. \tag{2.4}$$

Given the scaling function $\varphi(x)$, another important component of the multiresolution analysis is the wavelet function $\psi(x)$, defined by

$$\psi(x) = \sqrt{2} \sum_{i \in \mathbb{Z}} g_i \varphi(2x - i), \tag{2.5}$$

where $g_i = (-1)^{1-i} h_{1-i}$ for $i \in \mathbb{Z}$. A simple calculation shows that the support of $\psi(x)$ is $[-p + 1, p]$ and $\{g_i\}$ is nonzero only for $i = -2p + 2, \dots, 1$, based on the support of the φ and the nonzero entries pattern of $\{h_i\}$. The *Daubechies wavelets* are then defined as

$$\psi_k^{(\ell)}(x) = 2^{\ell/2} \psi(2^\ell x - k), \quad \ell = 0, 1, 2, \dots, \quad k \in \mathbb{Z}. \tag{2.6}$$

For a function $v(x) \in L^2(\mathbb{R})$, its scaling and wavelet coefficients $d_k^{(\ell)}$ and $v_k^{(\ell)}$ are defined as the inner product with the scaling functions and the wavelets

$$v_k^{(\ell)} := \int v(x) \varphi_k^{(\ell)}(x) dx, \quad d_k^{(\ell)} := \int v(x) \psi_k^{(\ell)}(x) dx. \tag{2.7}$$

Using the recursive relationships of the scaling function Eq. (2.2) and the wavelet function Eq. (2.5), one obtains a recursive relationship of the scaling and wavelet coefficients

$$v_k^{(\ell)} = \sum_{i \in \mathbb{Z}} h_i v_{2k+i}^{(\ell+1)}, \quad d_k^{(\ell)} = \sum_{i \in \mathbb{Z}} g_i v_{2k+i}^{(\ell+1)}. \tag{2.8}$$

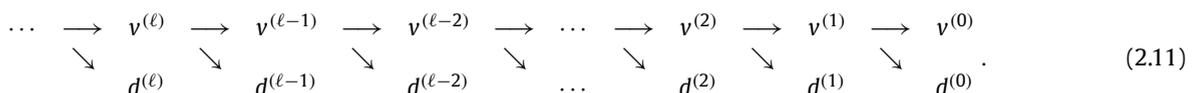
By defining $v^{(\ell)} = \left(v_k^{(\ell)} \right)_{k \in \mathbb{Z}}$ and $d^{(\ell)} = \left(d_k^{(\ell)} \right)_{k \in \mathbb{Z}}$, Eq. (2.8) can be written in a matrix form

$$v^{(\ell)} = \left(W_s^{(\ell)} \right)^T v^{(\ell+1)}, \quad d^{(\ell)} = \left(W_w^{(\ell)} \right)^T v^{(\ell+1)}, \tag{2.9}$$

where the operators $W_s^{(\ell)}$ and $W_w^{(\ell)} : \ell^2(\mathbb{Z}) \rightarrow \ell^2(\mathbb{Z})$ are banded with a bandwidth $2p$ due to the support of $\{h_i\}$ and $\{g_i\}$. By introducing the orthogonal operator $W^{(\ell)} = \begin{pmatrix} W_w^{(\ell)} & W_s^{(\ell)} \end{pmatrix}$, Eq. (2.9) can be rewritten as

$$\begin{pmatrix} d^{(\ell)} \\ v^{(\ell)} \end{pmatrix} = \left(W^{(\ell)} \right)^T v^{(\ell+1)}, \quad v^{(\ell+1)} = W^{(\ell)} \begin{pmatrix} d^{(\ell)} \\ v^{(\ell)} \end{pmatrix}. \tag{2.10}$$

The procedure for computing the wavelet and scaling coefficients can be illustrated in the following diagram



The discussion until now is concerned with the wavelets on \mathbb{R} . It is straightforward to extend it to the functions defined on a finite domain with periodic boundary condition. If the function $v(x)$ is periodic on a finite domain, for instance, $[0, 1]$, then the only modification is that all the shifts and scaling in the x variable are done modulus the integer. When working with periodic functions, the procedure in Eq. (2.11) usually stops at a coarse level $L_0 = O(\log_2(p))$ before the wavelet and scaling functions start to overlap itself.

$$\begin{array}{ccccccccccc} \dots & \longrightarrow & v^{(\ell+1)} & \longrightarrow & v^{(\ell)} & \longrightarrow & v^{(\ell-1)} & \longrightarrow & \dots & \longrightarrow & v^{(L_0)} \\ & & \searrow \\ & & d^{(\ell+1)} & & d^{(\ell)} & & d^{(\ell-1)} & & \dots & & d^{(L_0)} \end{array} \quad (2.12)$$

2.2. Nonstandard wavelet form for integral operator

Let A be an integral operator with kernel $a(x, y)$, applied to periodic functions defined on $[0, 1]$, i.e.,

$$u = Av, \quad \text{equivalently} \quad u(x) = \int a(x, y)v(y) dy. \quad (2.13)$$

Denote by $A^{(L)} = (A_{k_1, k_2}^{(L)}) \in \mathbb{R}^{2^L \times 2^L}$ the Galerkin projection of A to the space V_L , for a sufficiently deep level L , i.e.

$$A_{k_1, k_2}^{(L)} = \int \int \varphi_{k_1}^{(L)}(x)a(x, y)\varphi_{k_2}^{(L)}(y) dx dy.$$

The nonstandard form described in [5] is a remarkably efficient way to compress the matrix $A^{(L)}$.

The main step for the nonstandard form is to treat $A^{(L)}$ as an image and use the 2D multiresolution analysis

$$\begin{aligned} D_{1, k_1, k_2}^{(\ell)} &:= \iint \psi_{k_1}^{(\ell)}(x)a(x, y)\psi_{k_2}^{(\ell)}(y) dx dy, & D_{2, k_1, k_2}^{(\ell)} &:= \iint \psi_{k_1}^{(\ell)}(x)a(x, y)\varphi_{k_2}^{(\ell)}(y) dx dy, \\ D_{3, k_1, k_2}^{(\ell)} &:= \iint \varphi_{k_1}^{(\ell)}(x)a(x, y)\psi_{k_2}^{(\ell)}(y) dx dy, & A_{k_1, k_2}^{(\ell)} &:= \iint \varphi_{k_1}^{(\ell)}(x)a(x, y)\varphi_{k_2}^{(\ell)}(y) dx dy, \end{aligned} \quad (2.14)$$

for $\ell = L_0, \dots, L - 1$, and $k_1, k_2 = 0, \dots, 2^\ell - 1$. For convenience, these coefficients are organized into the matrix form as

$$A^{(\ell)} = (A_{k_1, k_2}^{(\ell)})_{k_1, k_2=0, \dots, 2^\ell-1}, \quad D_j^{(\ell)} = (D_{j, k_1, k_2}^{(\ell)})_{k_1, k_2=0, \dots, 2^\ell-1}, \quad j = 1, 2, 3. \quad (2.15)$$

In this setting, a similar recursive relation to Eq. (2.10) can be obtained

$$\begin{pmatrix} D_1^{(\ell)} & D_2^{(\ell)} \\ D_3^{(\ell)} & A^{(\ell)} \end{pmatrix} = (W^{(\ell)})^T A^{(\ell+1)} W^{(\ell)}, \quad \ell = L_0, \dots, L - 1. \quad (2.16)$$

If A is a Calderon-Zygmund operator, the entries of the matrices $D_j^{(\ell)}$ with $j = 1, 2, 3$ decay rapidly away from the diagonal. For a prescribed relative accuracy ϵ , each matrix $D_j^{(\ell)}$ can be approximated by a band matrix by truncating at a band of width $O(\log(1/\epsilon))$. Since the bandwidth is independent of the specific choices of ℓ, j , or the mesh size $N = 2^L$, the nonstandard form of A stores only $O(N)$ nonzero entries. The readers are referred to [5] for more details. With a slight abuse of notation, the matrices $D_j^{(\ell)}$ are assumed to be pre-truncated in what follows.

One can assemble all the matrices $D_j^{(\ell)}$ and $A^{(L_0)}$ together, by defining the matrix $S^{(L)}$ in a recursive way as

$$S^{(L_0)} = \begin{pmatrix} D_1^{(L_0)} & D_2^{(L_0)} \\ D_3^{(L_0)} & A^{(L_0)} \end{pmatrix}, \quad S^{(\ell+1)} = \begin{pmatrix} D_1^{(\ell+1)} & D_2^{(\ell+1)} & 0 \\ D_3^{(\ell+1)} & 0 & 0 \\ 0 & 0 & S^{(\ell)} \end{pmatrix}, \quad \ell = L_0, \dots, L - 1. \quad (2.17)$$

The matrix $S := S^{(L)}$ is the nonstandard form of the matrix $A = A^{(L)}$ satisfying

$$A = WSW^T. \quad (2.18)$$

Here W is the extended wavelet transform matrix, defined in the recursive form as

$$T^{(L_0)} = W^{(L_0)}, \quad T^{(\ell+1)} = \begin{pmatrix} W^{(\ell+1)} & W_s^{(\ell+1)}T^{(L_0)} \end{pmatrix}, \quad \ell = L_0, \dots, L - 1, \quad W := T^{(L)}. \quad (2.19)$$

Fig. 1 illustrates the matrices W and S along with the formulation Eq. (2.18).

To clarify the notations, we denote by $W_s^{(\ell)}$ and $W_w^{(\ell)}$ the transform matrices defined in Eq. (2.9) for the scaling and wavelet parts on level ℓ , respectively. $W^{(\ell)} = (W_w^{(\ell)} \quad W_s^{(\ell)})$ is the wavelet transform matrix at the level ℓ .

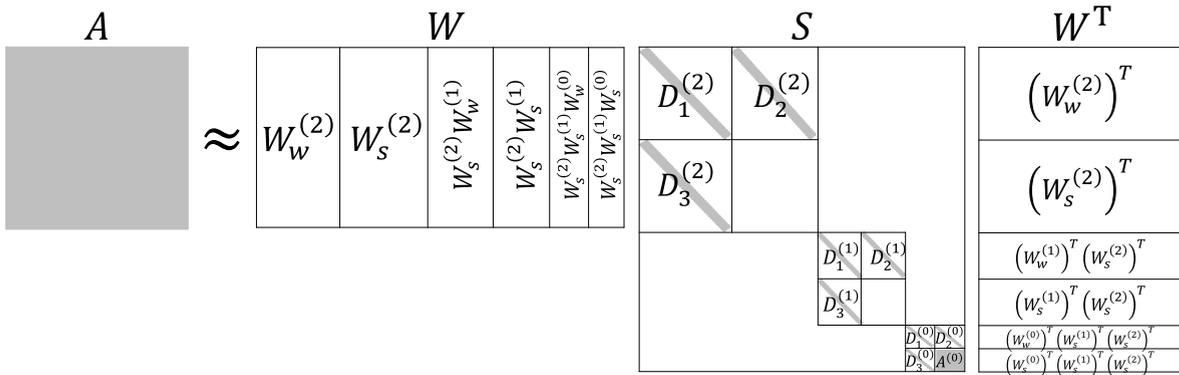


Fig. 1. Illustration of $A = WSW^T$ with $L_0 = 0$ and $L = 3$. The sparsity pattern of S is marked in gray.

2.3. Matrix-vector multiplication in the nonstandard form

With a Galerkin discretization of Eq. (2.13) at level L , the matrix-vector multiplication takes the form

$$u^{(L)} = A^{(L)} v^{(L)}. \tag{2.20}$$

The nonstandard form allows for accelerating the evaluation of Eq. (2.20). Using the nonstandard form $A^{(L)} = WSW^T$ obtained above, the matrix-vector multiplication

$$u^{(L)} = WSW^T v^{(L)} \tag{2.21}$$

can be split into four steps:

1. $A^{(L)} \rightarrow S$: generate the nonstandard form S from the matrix $A^{(L)}$ or the kernel $a(x, y)$;
2. $v^{(L)} \rightarrow \hat{v} := W^T v^{(L)}$: apply (forward) wavelet transform on $v^{(L)}$ to get \hat{v} ;
3. $\hat{u} := S\hat{v}$: evaluate the matrix-vector multiplication in the nonstandard form;
4. $\hat{u} \rightarrow u^{(L)} := W\hat{u}$: apply inverse wavelet transform on \hat{u} to obtain $u^{(L)}$.

The first step is computed using Eq. (2.16) if the matrix $A^{(L)}$ is given. The second step follows Eq. (2.10). The third step can be written as

$$\begin{pmatrix} w^{(\ell)} \\ s^{(\ell)} \end{pmatrix} = \begin{pmatrix} D_1^{(\ell)} & D_2^{(\ell)} \\ D_3^{(\ell)} & D_4^{(\ell)} \end{pmatrix} \begin{pmatrix} d^{(\ell)} \\ v^{(\ell)} \end{pmatrix}, \tag{2.22}$$

where $D_4^{(\ell)} = 0$ for $\ell = L_0 + 1, \dots, L - 1$ and $D_4^{(L_0)} = A^{(L_0)}$. The fourth step is essentially an inverse wavelet transform, implemented as

$$u^{(L_0)} = 0, \quad u^{(\ell+1)} = W^{(\ell)} \begin{pmatrix} w^{(\ell)} \\ s^{(\ell)} + u^{(\ell)} \end{pmatrix}, \quad \ell = L_0, \dots, L - 1. \tag{2.23}$$

A step-by-step description of these four steps are summarized in Algorithm 1.

2.4. The multidimensional case

The matrix-vector multiplication in the nonstandard form can be easily extended to the multidimensional case with the help of multidimensional orthogonal wavelets (see [44] for more details). For instance, in the two-dimensional setting, one defines at each scale ℓ three different types of wavelets of the form

$$\psi_{1,k}^{(\ell)}(x, y) = \varphi_{k_1}^{(\ell)}(x)\psi_{k_2}^{(\ell)}(y), \quad \psi_{2,k}^{(\ell)}(x, y) = \psi_{k_1}^{(\ell)}(x)\varphi_{k_2}^{(\ell)}(y), \quad \psi_{3,k}^{(\ell)}(x, y) = \psi_{k_1}^{(\ell)}(x)\psi_{k_2}^{(\ell)}(y), \tag{2.24}$$

with $k = (k_1, k_2) \in \mathbb{Z}^2$. Using these three types of wavelets, the transform matrix at each scale ℓ used in Eq. (2.10) is redefined to be $W^{(\ell)} = \begin{pmatrix} W_{w,1}^{(\ell)} & W_{w,2}^{(\ell)} & W_{w,3}^{(\ell)} & W_s^{(\ell)} \end{pmatrix}$. The 2D analog of Eq. (2.10) contains three types of wavelet coefficients

$$\begin{pmatrix} d_1^{(\ell)} \\ d_2^{(\ell)} \\ d_3^{(\ell)} \\ v^{(\ell)} \end{pmatrix} = \left(W^{(\ell)} \right)^T v^{(\ell+1)}, \quad v^{(\ell+1)} = W^{(\ell)} \begin{pmatrix} d_1^{(\ell)} \\ d_2^{(\ell)} \\ d_3^{(\ell)} \\ v^{(\ell)} \end{pmatrix}. \tag{2.25}$$

Algorithm 1 Compute $u = Av$ using the nonstandard form of the wavelet transform.

Require: $A^{(L)} = A$, $v^{(L)} = v$, $0 \leq L_0 < L$;
Ensure: $u = u^{(L)}$;

▷ **Step 1:** generate nonstandard form of $A^{(L)}$
1: **for** ℓ from $L-1$ to L_0 by -1 **do**
2: $\begin{pmatrix} D_1^{(\ell)} & D_2^{(\ell)} \\ D_3^{(\ell)} & A^{(\ell)} \end{pmatrix} = (W^{(\ell)})^T A^{(\ell+1)} W^{(\ell)}$
3: **end for**

▷ **Step 2:** forward wavelet transform on $v^{(L)}$
4: **for** ℓ from $L-1$ to L_0 by -1 **do**
5: $\begin{pmatrix} d^{(\ell)} \\ v^{(\ell)} \end{pmatrix} = (W^{(\ell)})^T v^{(\ell+1)}$
6: **end for**

▷ **Step 3:** matrix-vector multiplication
7: $D_4^{(\ell)} = 0$ for $\ell = L_0 + 1, \dots, L-1$ and $D_4^{(L_0)} = A^{(L_0)}$
8: **for** ℓ from L_0 to $L-1$ **do**
9: $\begin{pmatrix} w^{(\ell)} \\ s^{(\ell)} \end{pmatrix} = \begin{pmatrix} D_1^{(\ell)} & D_2^{(\ell)} \\ D_3^{(\ell)} & D_4^{(\ell)} \end{pmatrix} \begin{pmatrix} d^{(\ell)} \\ v^{(\ell)} \end{pmatrix}$
10: **end for**

▷ **Step 4:** Inverse wavelet transform
11: $u^{(L_0)} = 0$
12: **for** ℓ from L_0 to $L-1$ **do**
13: $u^{(\ell+1)} = W^{(\ell)} \begin{pmatrix} w^{(\ell)} \\ s^{(\ell)} + u^{(\ell)} \end{pmatrix}$
14: **end for**
15: **return** $u = u^{(L)}$

Similarly, the recursive relation Eq. (2.16) can be extended as well

$$\begin{pmatrix} D_1^{(\ell)} & D_2^{(\ell)} & D_3^{(\ell)} & D_4^{(\ell)} \\ D_5^{(\ell)} & D_6^{(\ell)} & D_7^{(\ell)} & D_8^{(\ell)} \\ D_9^{(\ell)} & D_{10}^{(\ell)} & D_{11}^{(\ell)} & D_{12}^{(\ell)} \\ D_{13}^{(\ell)} & D_{14}^{(\ell)} & D_{15}^{(\ell)} & A^{(\ell)} \end{pmatrix} = (W^{(\ell)})^T A^{(\ell+1)} W^{(\ell)}, \quad \ell = L_0, \dots, L-1, \quad (2.26)$$

where $D_j^{(\ell)}$, $j = 1, \dots, 15$ are all sparse matrices with only $O(4^\ell)$ non-negligible entries in each. The matrix-vector multiplication follows the steps of Algorithm 1, with these necessary changes.

3. Meta-learning approach

The plan is to apply the nonstandard form to the operator G_η in Eq. (1.2)

$$u = G_\eta f, \quad u(x) = \int g_\eta(x, y) f(y) dy. \quad (3.1)$$

With a slight abuse of notations, the same letters are used to denote the discretizations. The discrete version of Eq. (3.1) takes the form

$$u = G_\eta f, \quad u, f, \eta \in \mathbb{R}^N \text{ and } G_\eta \in \mathbb{R}^{N \times N}, \quad (3.2)$$

with $N = 2^L$. The main goal of this paper is to construct a neural network to learn the map $\eta \rightarrow G_\eta$.

Following Eq. (2.18) and applying the wavelet transform to the matrix G_η leads to

$$G_\eta \approx W S_\eta W^T, \quad (3.3)$$

where W is the extended wavelet transform matrix, independent of the parameter η . Since each block of matrix S_η is a band matrix, the nonzero entries of each block can be represented by a set of vectors. Let us define $C_\eta^{(\ell)}$, of size $2^\ell \times n_c$, to be the collection of these vectors of S_η at level ℓ , with n_c dependent on the bandwidth and ℓ . By introducing the collection of vectors $C_\eta := \{C_\eta^{(\ell)}\}_{\ell=L_0, \dots, L-1}$, S_η is uniquely determined by C_η , i.e.,

$$S_\eta \equiv \mathcal{S}[C_\eta]$$

for a fixed embedding operator \mathcal{S} determined by the sparsity pattern of S_η .

Given a set of training samples of the form

$$(\eta_i, \{f_{ij}, u_{ij}\}), \quad (3.4)$$

where $u_{ij} = G_{\eta_i} f_{ij}$ can be obtained by solving $L_{\eta_i} u_{ij} = f_{ij}$ with right hand side f_{ij} , the meta-learning approach first learns both the map $\eta \rightarrow C_\eta$ and the wavelet transform matrix W . Once they are ready, given any new η , G_η can be approximated by evaluating the map $\eta \rightarrow C_\eta$ and representing (3.3) in an NN form.

3.1. Neural network architecture

Using the factorization of G_η in Eq. (3.3), one can factorize $u_{ij} = G_{\eta_i} f_{ij}$ as

$$u_{ij} \approx W S_{\eta_i} W^T f_{ij}. \quad (3.5)$$

Algorithm 2 Neural network architecture for meta-learning $u = G_\eta f$.

Require: $f^{(L)} = f, \eta^{(L)} = \eta, \alpha_1, \alpha_2 \in \mathbb{N}$;
Ensure: u ;
 ▷ **Module 1:** Learn the map $\eta \rightarrow S_\eta$
 1: **for** ℓ from $L - 1$ to L_0 by -1 **do**
 2: $C_\eta^{(\ell)} = \text{ConvNet}[\ell, \alpha_1, n_{\text{cnn}}](\eta)$
 3: **end for**
 4: Generate $D_j^{(\ell)}, j = 1, 2, 3, \ell = L_0, \dots, L - 1$ and $A^{(L_0)}$ from $C_\eta^{(\ell)}, \ell = L_0, \dots, L - 1$
 ▷ **Module 2:** forward wavelet transform on $v^{(L)}$
 5: $v^{(L)} = f^{(L)}$
 6: **for** ℓ from $L - 1$ to L_0 by -1 **do**
 7: $(d^{(\ell)}, v^{(\ell)}) = \text{FWT}[\alpha_2](v^{(\ell+1)})$
 8: **end for**
 ▷ **Module 3:** matrix-vector multiplication in the nonstandard form
 9: $D_4^{(\ell)} = 0$ for $\ell = L_0 + 1, \dots, L - 1$ and $D_4^{(L_0)} = A^{(L_0)}$
 10: **for** ℓ from L_0 to $L - 1$ **do**
 11:
$$\begin{pmatrix} w^{(\ell)} \\ s^{(\ell)} \end{pmatrix} = \begin{pmatrix} D_1^{(\ell)} & D_2^{(\ell)} \\ D_3^{(\ell)} & D_4^{(\ell)} \end{pmatrix} \begin{pmatrix} d^{(\ell)} \\ v^{(\ell)} \end{pmatrix}$$

 12: **end for**
 ▷ **Module 4:** Inverse wavelet transform
 13: $u^{(L_0)} = 0$
 14: **for** ℓ from L_0 to $L - 1$ **do**
 15: $u^{(\ell+1)} = \text{IWT}[\alpha_2](\{w^{(\ell)}, s^{(\ell)} + u^{(\ell)}\})$
 16: **end for**
 17: Average over the channel direction of $u^{(L)}$ to give u ;
 18: **return** u

Similar to the matrix-vector multiplication in Section 2.3 of the nonstandard form, we propose a neural network for meta-learning Eq. (3.5) with four modules:

1. $\eta \rightarrow S_\eta$: a module learns the map $\eta \rightarrow C_\eta$ and then generates the banded sparse matrix S_η from C_η (denoted as $S_\eta = \mathcal{S}[C_\eta]$);
2. $v^{(L)} \rightarrow \hat{v} := W^T v^{(L)}$: a module applies the forward wavelet transform to $v^{(L)}$ to generate \hat{v} ;
3. $\hat{u} := S_\eta \hat{v}$: a module evaluates the matrix-vector multiplication in the nonstandard form;
4. $\hat{u} \rightarrow u^{(L)} := W \hat{u}$: a module applies the inverse wavelet transform on \hat{u} to generate $u^{(L)}$.

Instead of computing C_η from the full operator G_η as described in Section 2.3, the first module forms C_η directly from the parameter η using a deep NN. This module can be split into two steps: (1) carrying out the map $\eta \rightarrow C_\eta^{(\ell)}$ for each scale ℓ ; (2) constructing the nonstandard form S_η from $C_\eta := \{C_\eta^{(\ell)}\}_{\ell=L_0, \dots, L-1}$. The NN architecture for the map $\eta \rightarrow C_\eta^{(\ell)}$ is often problem-dependent. For many applications, including the ones to be considered in Sections 4 and 5, the problem is often translation-invariant, i.e., for any translation operator T ,

$$u = G_\eta f \quad \text{implies} \quad (Tu) = G_{(T\eta)}(Tf). \tag{3.6}$$

For such problems, a convolutional NN is often used for its efficiency and robustness.

The second and fourth modules perform the forward and inverse wavelet transforms (as in Section 2.3), respectively, for a specific wavelet basis. The selection of an effective wavelet basis is often problem-dependent. The capability of learning a problem-dependent wavelet transform from data is essential for the accuracy of the NN architecture.

Combining these four modules results in the architecture summarized in Algorithm 2. An illustration is given in Fig. 2. Below we describe details of the layers and parameters used in this architecture.

Implementation details The input, output, and intermediate data of the NN architecture are all represented with 2-tensors. For a tensor of size $N \times \alpha$, we refer to N as the spatial dimensions and α as the channel dimension. The main tool is the *convolutional layer*. Given an input tensor ξ of size $N \times \alpha$, the convolutional layers outputs a tensor ζ of size $N' \times \alpha'$ obtained via

$$\zeta_{i,c'} = \phi \left(\sum_{j=is}^{is+w-1} \sum_{c=0}^{\alpha-1} W_{j;c',c} \xi_{j,c} + b_{c'} \right), \quad i = 0, \dots, N' - 1, \quad c' = 0, \dots, \alpha' - 1, \tag{3.7}$$

where w is the *window size*, s is the *stride* and ϕ is the *activation function*, usually chosen to be a linear function, a rectified-linear unit (ReLU) function, or a sigmoid function. We denote this convolutional layer as

$$\zeta = \text{Conv1d}[\alpha', w, s, \phi](\xi). \tag{3.8}$$

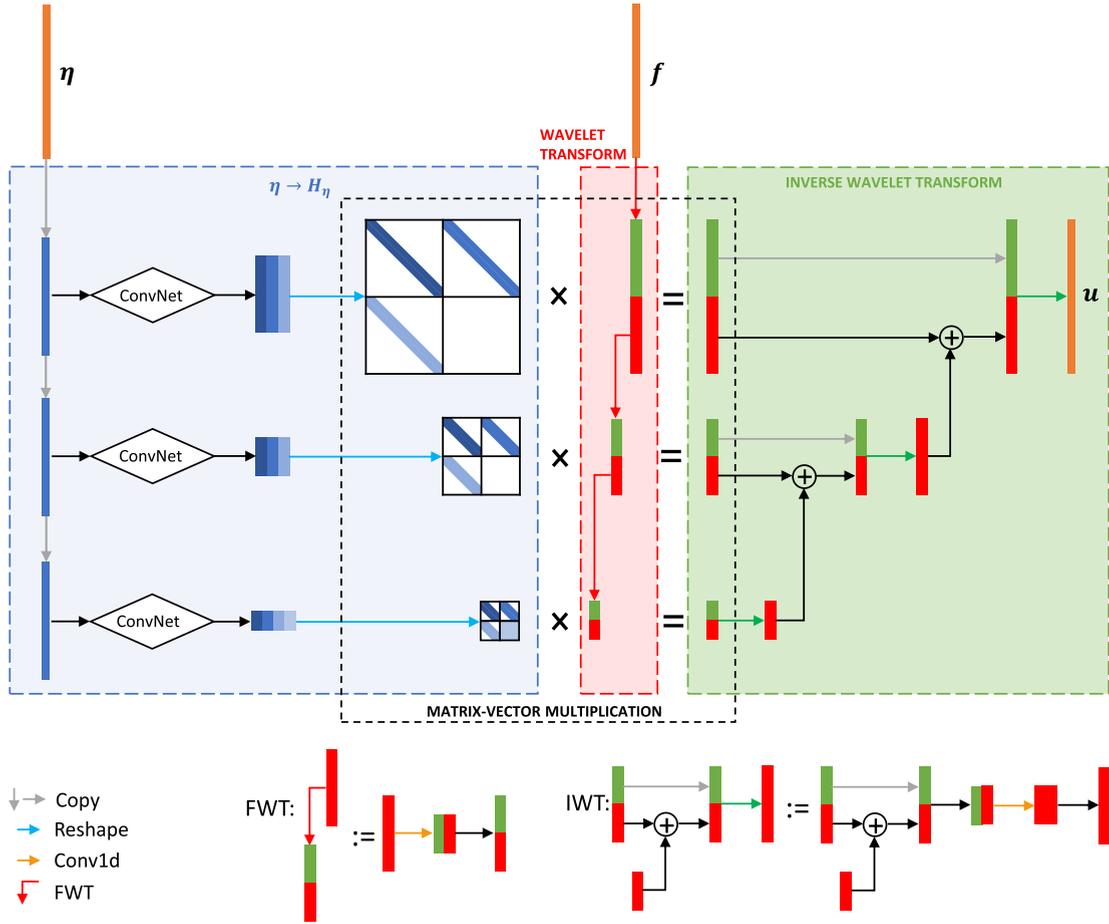


Fig. 2. Illustration of the neural network architecture from Algorithm 2, with input vectors η and f and output vector u . Each of the four modules in Algorithm 2 is represented by blocks with dashed contours. From left to right, the map $\eta \rightarrow S_\eta$ is the blue block, the forward wavelet transform applied to f is the red block, the sparse matrix-vector multiplication with S_η is the transparent block, and the inverse wavelet transform is the green block. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

The basic building blocks and layers used in Algorithm 2 are listed below.

- $\eta \rightarrow C_\eta^{(\ell)} : C_\eta^{(\ell)} = \text{ConvNet}[\ell, \alpha, n_{\text{cnn}}](\eta)$. As discussed above, it is often a convolutional NN if the system Eq. (3.1) is translation invariant. Since the spatial size of η is greater than that of $C_\eta^{(\ell)}$, $\text{ConvNet}[\ell, \alpha, n_{\text{cnn}}]$ consists of n_{cnn} convolutional layers and several downsampling or pooling layers.
- Forward wavelet transform at level ℓ : $(d^{(\ell)}, v^{(\ell)}) = \text{FWT}[\alpha](v^{(\ell+1)})$. This is the NN representation of the first equation in Eq. (2.10). It is implemented as $f^{(\ell)} = \text{Conv1d}[2\alpha, 2p, 2, \text{id}](v^{(\ell)})$, where the first α and the last α channels of $f^{(\ell)}$ are assigned to $d^{(\ell)}$ and $v^{(\ell)}$, respectively.
- Inverse wavelet transform at level ℓ : $u^{(\ell+1)} = \text{IWT}[\alpha](w^{(\ell)}, s^{(\ell)})$. This is the NN representation of the second equation in Eq. (2.10). The expression $[w^{(\ell)}, s^{(\ell)}]$ stands for concatenating the 2-tensors $w^{(\ell)}$ and $s^{(\ell)}$ of size $2^{(\ell)} \times \alpha$ to a 2-tensor of size $2^{(\ell)} \times 2\alpha$ along the channel dimension. This layer first applies the inverse transform, implemented by $\text{Conv1d}[2\alpha, p, 1, \text{id}]$, and then reshapes the output of size $2^\ell \times 2\alpha$ to a 2-tensor of size $2^{\ell+1} \times \alpha$ by a column-first ordering.

The generation of $D_j^{(\ell)}$ and $A^{(L_0)}$ from $C_\eta^{(\ell)}$ in Line 4 and the matrix-vector multiplication in Line 11 of Algorithm 2 require some discussion. Fig. 3 illustrates two approaches for evaluating the matrix-vector multiplication of a band matrix whose nonzero entries are stored in a set of vectors. The left figure corresponds to the case in Algorithm 2, while the right one is used in the actual implementation. To avoid the copying and shifting of $d^{(\ell)}$ and $v^{(\ell)}$, it is convenient to set $\alpha_2 = \alpha_1 = \alpha$. Though there are slightly more NN parameters in this case, this implementation change allows for a more flexible NN that can learn faster.



Fig. 3. Two approaches to evaluate the multiplication of a band matrix (stored with vectors) with a vector. The left figure corresponds to the case in Algorithm 2 and the right one is used in the implementation. The symbol \odot stands for element-wise multiplication.

3.2. The multidimensional case

Let us focus on the 2D case. The input, the output and the intermediate data are all 3-tensors of size $N_1 \times N_2 \times \alpha$, where $N = (N_1, N_2)$ is the spatial dimension and α is the channel dimension. The convolutional layer takes the form

$$\zeta_{i,c'} = \phi \left(\sum_{j_1=i_1s}^{i_1s+w-1} \sum_{j_2=i_2s}^{i_2s+w-1} \sum_{c=0}^{\alpha-1} W_{j_1c',c} \xi_{j_1,c} + b_{c'} \right), \quad i_1 = 0, \dots, N_1' - 1, i_2 = 0, \dots, N_2' - 1, c' = 0, \dots, \alpha' - 1, \quad (3.9)$$

where the input ξ is of size $N_1 \times N_2 \times \alpha$ and the output ζ is of size $N_1' \times N_2' \times \alpha'$. Here, the same stride s and window size w are used in both dimensions. We denote this convolutional layer as

$$\zeta = \text{Conv2d}[\alpha, w, s, \phi](\xi). \quad (3.10)$$

Algorithm 2 can be easily extended to the 2D case, following the same way that Algorithm 1 was extended in Section 2.4. Since there are three different types of wavelets in 2D, the layers in Algorithm 2 are redefined as follows:

- $\eta \rightarrow C_\eta^{(\ell)}$ module: $C_\eta^{(\ell)} = \text{ConvNet}[\ell, \alpha, n_{\text{cnn}}](\eta)$. This module is often a two-dimensional convolutional NN with several downsampling or pooling layers.
- Wavelet transform at level ℓ : $(d_1^{(\ell)}, d_2^{(\ell)}, d_3^{(\ell)}, v^{(\ell)}) = \text{FWT}[\alpha](v^{(\ell+1)})$. This is implemented using $f^{(\ell)} = \text{Conv2d}[4\alpha, 2p, 2, \text{id}](v^{(\ell)})$. The first, second, third and last α channels of $f^{(\ell)}$ are assigned to $d_1^{(\ell)}, d_2^{(\ell)}, d_3^{(\ell)}$ and $v^{(\ell)}$, respectively.
- Inverse wavelet transform at level ℓ : $u^{(\ell+1)} = \text{IWT}[\alpha]([w_1^{(\ell)}, w_2^{(\ell)}, w_3^{(\ell)}, s^{(\ell)}])$. This is implemented by first computing $\text{Conv2d}[4\alpha, p, 1, \text{id}]([d_1^{(\ell)}, d_2^{(\ell)}, d_3^{(\ell)}, v^{(\ell)} + u^{(\ell)}])$, and then reshaping the output of size $2^\ell \times 2^\ell \times 4\alpha$ to a 3-tensor of size $2^{\ell+1} \times 2^{\ell+1} \times \alpha$. The reshape operation is performed as follows: (1) reshape the output to a 5-tensor of size $2^\ell \times 2^\ell \times 2 \times 2 \times \alpha$ by splitting the last dimension; (2) permute the second and third dimensions to obtain a 5-tensor of size $2^\ell \times 2 \times 2^\ell \times 2 \times \alpha$; (3) group the first and second dimensions, and the third and fourth dimensions, respectively, to obtain the resulting 3-tensor of size $2^{\ell+1} \times 2^{\ell+1} \times \alpha$.

4. Elliptic partial differential equations

This section applies the meta-learning approach described in Section 3 to the Green's functions of elliptic PDEs, both in the Schrödinger form and in the divergence form.

4.1. Schrödinger form

Consider the equation

$$-\Delta u(x) + \eta(x)u(x) = f(x), \quad x \in \Omega = [0, 1]^d, \quad (4.1)$$

with a periodic boundary condition, where $\eta(x) > 0$ is the potential and $f(x)$ is the source term. Following the notations of Section 1,

$$L_\eta = -\Delta + \eta(x), \quad G_\eta = L_\eta^{-1}, \quad u = G_\eta f. \quad (4.2)$$

Since the problem Eq. (4.1) is translation-invariant due to the periodic boundary condition, the map $\eta \rightarrow C_\eta^{(\ell)}$ can be represented with a convolutional NN. In what follows, we first derive the explicit dependence of $C_\eta^{(\ell)}$ on η using a linear perturbative analysis and then report some numerical studies.

Mathematical analysis When η is close to a fixed homogeneous background $\eta_0 > 0$, it is convenient to write

$$L_\eta = L_0 - E_\eta, \quad L_0 = -\Delta + \eta_0, \quad E_\eta = \text{diag}(-\eta + \eta_0). \quad (4.3)$$

Let $G_0 = L_0^{-1}$ be the Green's function of L_0 with the periodic boundary condition. Using the Neumann series for the resolvent $(I - G_0 E_\eta)^{-1}$ with $|\eta(x) - \eta_0|$ sufficiently small, one can write the Green's function G_η as a perturbative expansion

$$G_\eta = (L_0 - E_\eta)^{-1} = G_0 + G_0 E_\eta G_0 + G_0 E_\eta G_0 E_\eta G_0 + \dots \quad (4.4)$$

For sufficiently small $|\eta(x) - \eta_0|$, the operator G_η can be approximated by its linear part as

$$G_\eta \approx G_0 + G_0 E_\eta G_0. \quad (4.5)$$

Let $g_0(x)$ and $g_\eta(x)$ be the kernel of G_0 and G_η , respectively. Since G_0 is the Green's function of $-\Delta + \eta_0$ with the periodic boundary condition, the kernel g_0 is translation-invariant, i.e., $g_0(x, y) = g_0(x - y)$. The wavelet-wavelet coefficients of G_η at level ℓ take the form

$$\begin{aligned} D_{1,k_1,k_2}^{(\ell)} &= \iint \psi_{k_1}^{(\ell)}(x) g_\eta(x, y) \psi_{k_2}^{(\ell)}(y) dx dy \\ &\approx \iint \psi_{k_1}^{(\ell)}(x) g_0(x - y) \psi_{k_2}^{(\ell)}(y) dx dy + \iiint \psi_{k_1}^{(\ell)}(x) g_0(x - z) (\eta(z) - \eta_0) g_0(z - y) \psi_{k_2}^{(\ell)}(y) dx dy dz \\ &= \iint \psi_{k_1}^{(\ell)}(x) g_0(x - y) \psi_{k_2}^{(\ell)}(y) dx dy + \int \tilde{\psi}_{k_1}^{(\ell)}(z) \tilde{\psi}_{k_2}^{(\ell)}(z) (\eta(z) - \eta_0) dz, \end{aligned} \quad (4.6)$$

where $\tilde{\psi}_k^{(\ell)}(z) := (g_0 * \psi_k^{(\ell)})(z)$. For a fixed diagonal of $D_1^{(\ell)}$ with $k_2 = k_1 + c$ for a constant c , Eq. (4.6) states that the map from η to $D_{1,k_1,k_1+c}^{(\ell)}$ for all possible k_1 is simply a convolution with an addition of a term independent of η , which can be simply represented by the Conv1d layer in Eq. (3.7). It is straightforward to extend the conclusion to $D_{j,k_1,k_2}^{(\ell)}$, $j = 2, 3$ and $A_{k_1,k_2}^{(L)}$.

When $|\eta(x) - \eta_0|$ is not small, one can account for the nonlinearities neglected in the perturbative analysis by using multiple convolutional layers and making use of nonlinear activation functions. In other words, it is natural to approximate the map $\eta \rightarrow C_\eta^{(\ell)}$ using a convolutional NN with enough layers and an appropriate window size [40,30,46].

Moreover, since the matrix G_η is a symmetric matrix, $D_1^{(\ell)}$ and $A^{(L_0)}$ are symmetric and $(D_2^{(\ell)})^\top = D_3^{(\ell)}$. In the implementation, the symmetry is enforced by generating $D_3^{(\ell)}$ from $D_2^{(\ell)}$, and replacing $D_1^{(\ell)}$ (or $A^{(L_0)}$) by $\frac{1}{2}(D_1^{(\ell)} + (D_1^{(\ell)})^\top)$ (or $\frac{1}{2}(A^{(L_0)} + (A^{(L_0)})^\top)$, respectively. Since the Schrödinger form considered in this section includes the periodic boundary condition, the convolutional layers are all implemented with periodic padding.

Numerical results The NN discussed above is implemented in Keras [7] (running on top of TensorFlow [1]). The parameters of the NN are initialized randomly from the normal distribution. The loss function is set to be the mean squared error

$$\frac{1}{N_{\text{samples}}} \|u - u^{\text{NN}}\|_{\ell^2}, \quad (4.7)$$

where the exact solution, obtained by solving Eq. (4.2), is denoted as u and the NN prediction as u^{NN} . N_{samples} denotes the number of samples. The NN is trained until convergence using the Nadam optimizer [12] with the learning rates equal to 10^{-3} for the 1D case and 10^{-4} for the 2D case. The batch size is set to be one percent of the number of training samples. The support of the scaling function φ is chosen to be $2p = 6$. The number of levels $L - L_0$ in the wavelet transform is 6 for the 1D case and 4 for the 2D case.

The data set contains 5,000 different η and for each η Eq. (4.2) is solved with 20 randomly generated f using the central difference scheme. Therefore, the number of training samples corresponds to the number of different $\{f_{ij}\}$, rather than different $\{\eta_i\}$. Half of the generated data is used for training data, while the other half is reserved for testing. The accuracy of the NN is measured by the relative error in the ℓ^2 norm

$$\epsilon = \frac{\|u - u^{\text{NN}}\|_{\ell^2}}{\|u\|_{\ell^2}}. \quad (4.8)$$

The training error ϵ_{train} and test error ϵ_{test} are calculated by averaging the relative error over all training and test samples, respectively. The number of parameters in the NN is denoted by N_{params} . The operator error ϵ_{op} is calculated by averaging the relative 2-norm error of the matrix

$$\frac{\|G_\eta - G_\eta^{\text{NN}}\|_{\ell^2 \rightarrow \ell^2}}{\|G_\eta\|_{\ell^2 \rightarrow \ell^2}} \quad (4.9)$$

over samples of the exact inverse operator G_η and its NN approximation G_η^{NN} .

For the 1D case, the domain $\Omega = [0, 1]$ is discretized by a uniform Cartesian grid with 320 points. The positive potential $\eta(x)$ is generated by (1) sampling independently from $\mathcal{N}(0, 1)$ on a uniform grid with 40 points, (2) interpolating to the

Table 1
Relative error in approximating the solution to the Schrödinger form in the 1D case.

α	n_{cnn}	N_{params}	ϵ_{train}	ϵ_{test}	ϵ_{op}
5	5	30201	4.43e-3	4.74e-3	2.49e-3
5	7	38061	4.83e-3	5.18e-3	4.28e-3
7	5	58717	4.09e-3	4.35e-3	3.28e-3
7	7	74089	4.11e-3	4.42e-3	2.18e-3

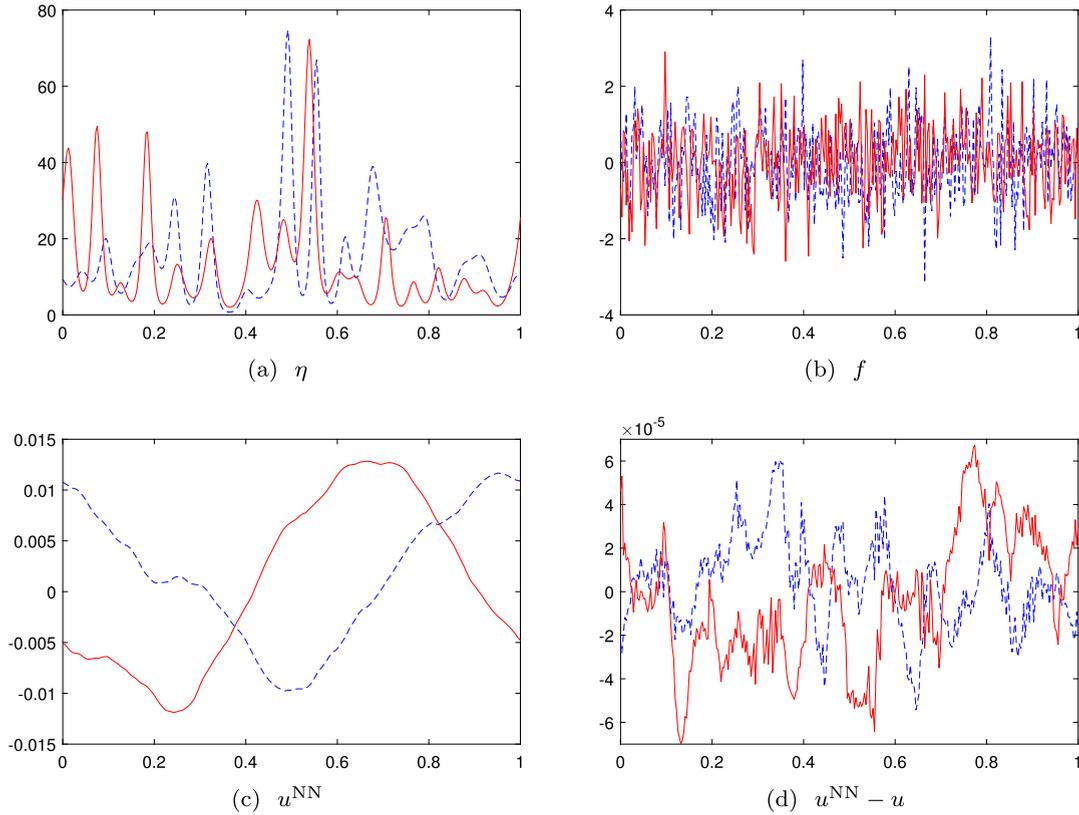


Fig. 4. Two samples (one in red line and the other in blue dashed line) from the test set for the potentials η , the source terms f , the predictions u^{NN} with $\alpha = 5$ and $n_{\text{cnn}} = 5$, and their corresponding error for the Schrödinger form in the 1D case.

Table 2
Relative error in approximating the solution to the Schrödinger form in the 2D case.

α	n_{cnn}	N_{params}	ϵ_{train}	ϵ_{test}	ϵ_{op}
11	5	930447	2.21e-2	2.18e-2	4.17e-3
15	5	1226071	2.12e-2	2.10e-2	2.04e-3

320-point grid via a Fourier interpolation, and (3) point-wise exponentiating followed by a factor of 10 scaling. The source term $f(x)$ is generated by sampling independently from $\mathcal{N}(0, 1)$. The results for different values of α (channel number) and n_{cnn} (layer number) are reported in Table 1. The best approximation of the operator, obtained with $\alpha = 5$ and $n_{\text{cnn}} = 5$, results in a test error of 4.7×10^{-3} and an operator error of 2.5×10^{-3} with only 3×10^4 parameters. The operator error reported in Table 1 has been averaged among 100 different samples of G_η . Two random samples from the test data are illustrated in Fig. 4 along with the NN prediction. A representative sample of the inverse operator G_η and its NN approximation are displayed in Fig. 5.

For the 2D case, the domain $\Omega = [0, 1]^2$ is discretized with a 80×80 uniform Cartesian mesh. The potential $\eta(x)$ is generated by (1) sampling independently from $\mathcal{N}(0, 1)$ on a uniform mesh with 10×10 points, (2) then interpolating to 80×80 points via a Fourier interpolation, and (3) point-wise exponentiating followed by appropriate scaling. The source term is sampled point-wisely from a standard Gaussian distribution. When trained with $\alpha = 11$ and $n_{\text{cnn}} = 5$, the NN achieves a test error of 2.2×10^{-2} and an operator error of 4.2×10^{-3} with only 9.3×10^5 parameters, as reported in Table 2. The operator error ϵ_{op} estimate is computed by averaging the error among 10 distinct samples of the inverse

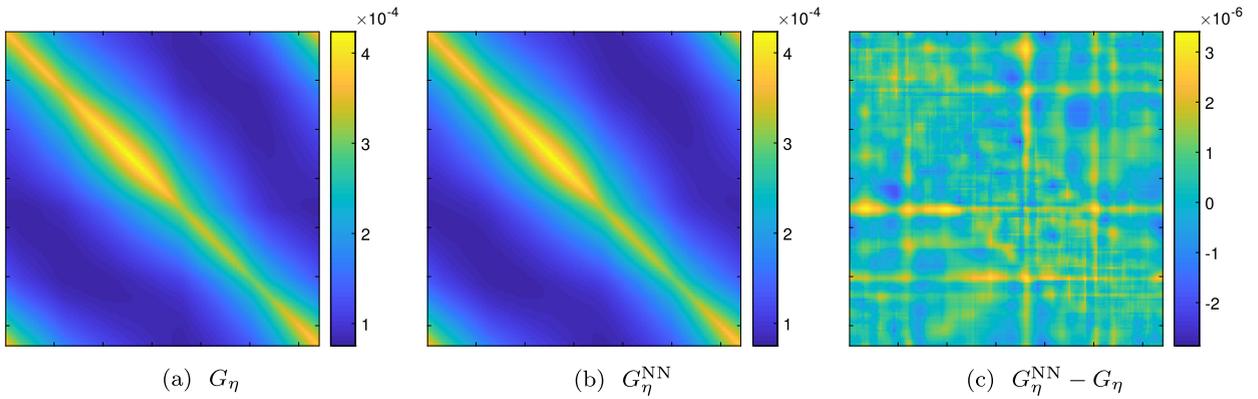


Fig. 5. Operator approximation with $\alpha = 5$ and $n_{\text{cnn}} = 5$ for the Schrödinger form in the 1D case.

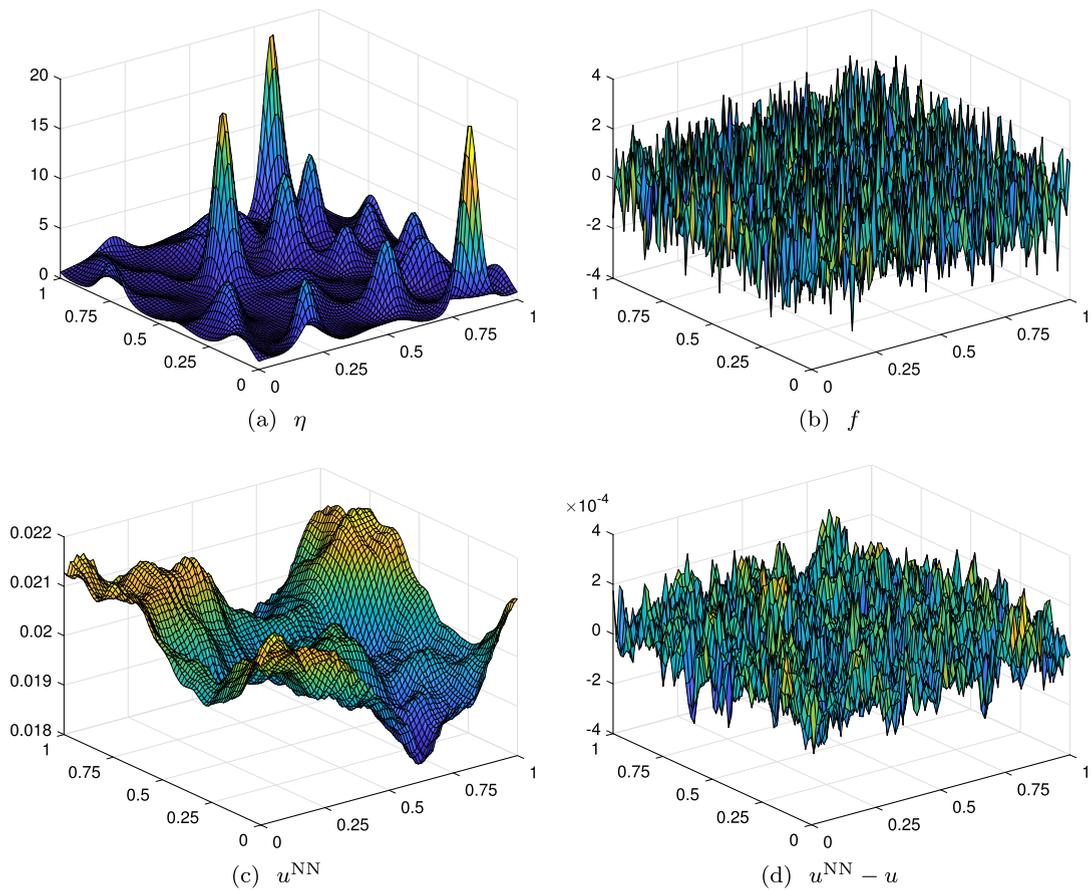


Fig. 6. A sample from the test set for the potential η , the source term f , the prediction u^{NN} with $\alpha = 11$ and $n_{\text{cnn}} = 5$, and its corresponding error for Schrödinger form in the 2D case.

operator G_η . The values of η and f of a representative sample are displayed in Fig. 6, along with the NN prediction and the error.

4.2. Divergence form

The same NN architecture is applied to the Green's functions of the divergence form

Table 3
Relative error in approximating the solution of divergence form in the 1D case.

α	n_{cnn}	N_{params}	ϵ_{train}	ϵ_{test}
7	5	58717	7.27e-3	7.76e-3
7	7	74089	7.46e-3	8.29e-3
9	5	96625	6.05e-3	6.88e-3
9	7	122005	6.83e-3	8.21e-3

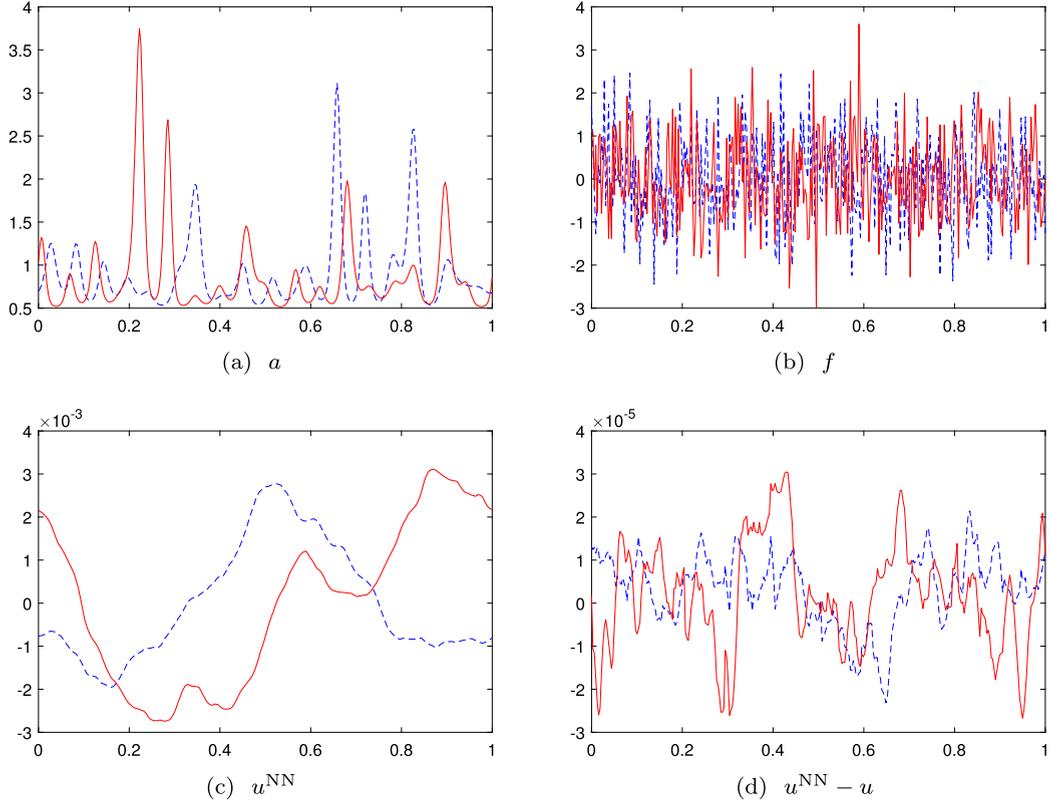


Fig. 7. Two samples (one in red line and the other in blue dashed line) from the test set for the distinct coefficients η , the source terms f , the predictions u^{NN} with $\alpha=9$ and $n_{\text{cnn}}=5$, and their corresponding error for the divergence form in the 1D case.

$$\begin{aligned}
 &-\nabla \cdot (\eta(x)\nabla u(x)) = f(x), \quad x \in [0, 1]^d, \\
 &\int_{[0, 1]^d} u(x) \, dx = 0,
 \end{aligned} \tag{4.10}$$

with $\eta(x) \geq \eta_0 > 0$ along with the periodic boundary condition. Following the notations of Section 1,

$$L_\eta = -\nabla \cdot \text{diag}(\eta)\nabla, \quad G_\eta = L_\eta^{-1} \text{ in the constraint of } \int_{[0, 1]^d} u(x) \, dx = 0. \tag{4.11}$$

When $\eta(x)$ is close to a fixed $\eta_0 > 0$, the operator can be decomposed as

$$L_0 = -\eta_0 \Delta, \quad E_\eta = \nabla \cdot \text{diag}(\eta(x) - \eta_0)\nabla, \quad L_\eta = L_0 - E_\eta. \tag{4.12}$$

Since the operator E_η is linearly dependent on η , it is easy to check that the discussion for the Schrödinger form holds for the divergence form case as well.

Numerical results The parameter field $\eta(x)$ is generated in a way similar to the potential of the Schrödinger form, with the difference that the scaling factor is set to 1/5 and an additive term of 0.5 is applied point-wise to avoid the ill-conditioning of G_η . The numerical results for different choices of α (channel number) and n_{cnn} (layer number) are summarized in Table 3.

For example, a test error of 6.9×10^{-3} is achieved at $\alpha = 9$ and $n_{\text{cnn}} = 5$ with 9.6×10^4 parameters. Two random samples from the test data are illustrated in Fig. 7.

5. Radiative transfer equation with isotropic scattering

The radiative transfer equation (RTE) is a fundamental model for describing particle propagation, with applications in many fields, such as neutron transport in reactor physics [48], light transport in atmospheric radiative transfer [45], heat transfer [35], and optical imaging [34]. The steady-state RTE in the homogeneous scattering regime is

$$\begin{aligned} \nu \cdot \nabla_x \varphi(x, \nu) + (\eta(x) + \eta_a(x))\varphi(x, \nu) &= \eta(x)u(x) + f(x), \quad \text{in } \Omega \times \mathbb{S}^{d-1}, \quad \Omega \subset \mathbb{R}^d, \\ \varphi(x, \nu) &= 0, \quad \text{on } \{(x, \nu) \in \partial\Omega \times \mathbb{S}^{d-1} : n(x) \cdot \nu < 0\}, \\ u(x) &= \frac{1}{4\pi} \int_{\mathbb{S}^{d-1}} \varphi(x, \nu) \, d\nu, \end{aligned} \quad (5.1)$$

where $\varphi(x, \nu)$ denotes the photon flux that depends on both space x and angle ν , $f(x)$ is the light source, $\eta(x)$ is the scattering coefficient, and $\eta_a(x)$ is the physical absorption coefficient. In many applications, it is reasonable to assume $\eta_a(x)$ to be constant. Below, we focus on the most challenging case $\eta_a(x) \equiv 0$.

The numerical solution to the RTE has been extensively studied using the Monte Carlo methods and various discretization schemes for the differential-integral formulation Eq. (5.1) of RTE. However, these approaches often suffer from the high-dimensionality and non-smoothness of the photon-flux $\varphi(x, \nu)$. The recent numerical work in [9,16,50] follows the integral formulation by eliminating $\varphi(x, \nu)$ from the equation and keeping only $u(x)$ as unknown:

$$(I - K_\eta \eta) u = K_\eta f, \quad (5.2)$$

where the operator K_η is defined as

$$K_\eta f = \int_{y \in \Omega} k_\eta(x, y) f(y) \, dy, \quad k_\eta(x, y) = \frac{\exp\left(-|x - y| \int_0^1 \eta(x - s(x - y)) \, ds\right)}{4\pi |x - y|^{d-1}}. \quad (5.3)$$

The parameterized Green's function operator for the steady-state RTE is then

$$G_\eta = (I - K_\eta \eta)^{-1} K_\eta. \quad (5.4)$$

Since K_η is a dense operator, forming G_η following Eq. (5.4) is often computationally expensive. Instead, the meta-learning approach developed above allows for approximating the map from η to G_η directly.

Section 4 argues that the map $\eta \rightarrow C_\eta^{(\ell)}$ for the translation invariant operator can be represented by a convolutional NN. A key observation for the current setting is that the integral equation Eq. (5.2) can be extended to the whole domain by padding f and η with zero. As a result, the map from η to G_η can be represented by a convolutional NN *with zero padding*.

Numerical results The first test is concerned with the one-dimensional slab geometry, where the parameter η varies only in the x_1 direction (i.e., constant in the x_2 and x_3 directions). For this geometry, the integral equation Eq. (5.2) reduces to

$$(I - K_\eta^{(1)} \eta) u(x) = K_\eta^{(1)} f(x), \quad G_\eta = (I - K_\eta^{(1)} \eta)^{-1} K_\eta^{(1)}, \quad (5.5)$$

where x stands for only x_1 and the operator $K_\eta^{(1)}$ is defined as

$$\begin{aligned} K_\eta^{(1)} f(x) &= \int_{y \in \Omega} k_\eta^{(1)}(x, y) f(y) \, dy, \\ k_\eta^{(1)}(x, y) &= \frac{1}{2} \text{Ei} \left(-|x - y| \int_0^1 \eta(x - s(x - y)) \, ds \right), \quad \text{Ei}(x) = \int_{-\infty}^x \frac{e^t}{t} \, dt, \end{aligned} \quad (5.6)$$

with the domain $\Omega = [0, 1]$. In the implementation, $[-x_0, 1 + x_0]$ is discretized by a uniform Cartesian mesh with $N = 320$ points, where $x_0 > 0$ is selected such that there are 300 points in Ω . The scattering coefficient η is generated in the same way for $\eta(x)$ in Section 4 followed by appropriate rescaling. The source term $f(x)$, positive due to physical considerations, is generated by sampling independently from $\mathcal{U}(0, 1)$ instead of $\mathcal{N}(0, 1)$ and interpolated via Fourier interpolation. The values of η and f outside of Ω are set to be 0. The results for different values of α (channel number) and n_{cnn} (layer number) are summarized in Table 4. A test error of 2.9×10^{-3} is achieved with as few as 3.4×10^4 parameters with $\alpha = n_{\text{cnn}} = 5$. Two representative examples from the test set are shown in Fig. 8.

Table 4
Relative error in approximating the solution to the 1D RTE.

α	K	N_{params}	ϵ_{train}	ϵ_{test}
5	5	34131	2.48e-3	2.93e-3
5	7	41991	2.46e-3	3.01e-3
7	5	66403	1.92e-3	2.45e-3
7	7	81775	2.05e-3	2.36e-3

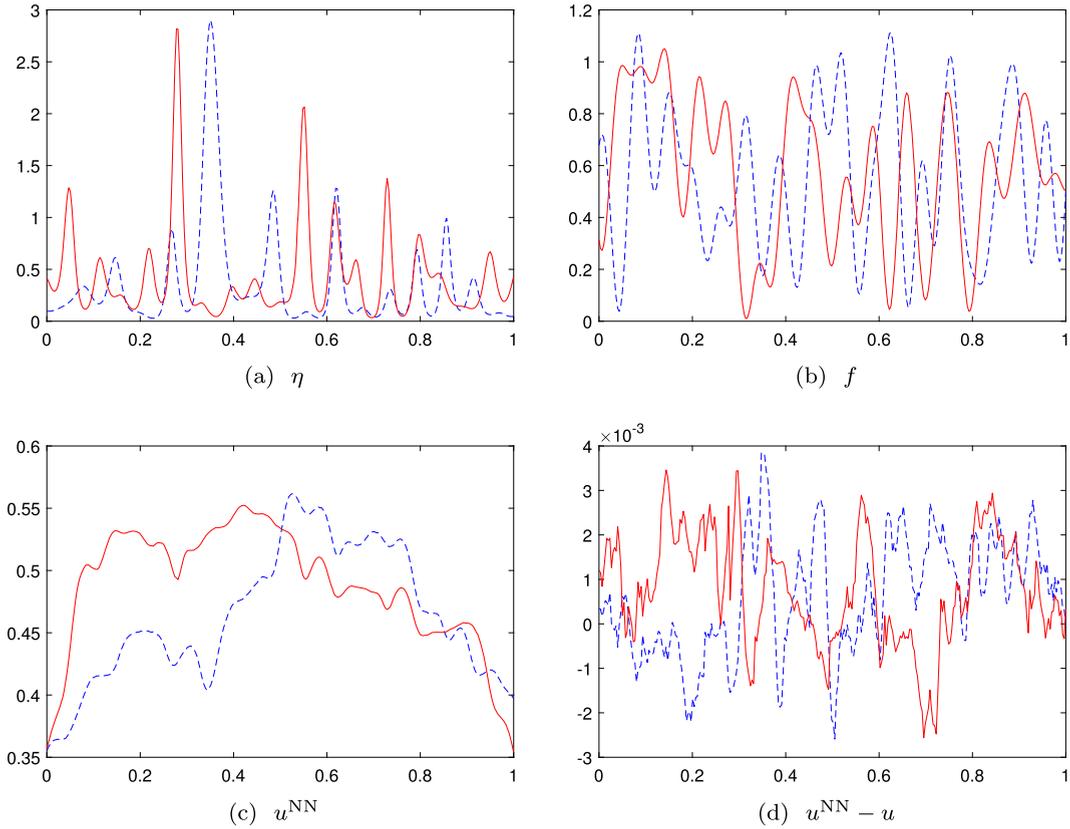


Fig. 8. Two samples (one in red line and the other in blue dashed line) from the test set with the scattering coefficients η , the source terms f , the predictions u^{NN} with $\alpha = 5$ and $n_{\text{cnn}} = 5$, and their corresponding error in the 1D RTE.

Table 5
Relative error in approximating the solution to the 2D RTE.

α	K	N_{params}	ϵ_{train}	ϵ_{test}
11	5	1287903	4.39e-3	4.39e-3
15	5	1663831	3.55e-3	3.55e-3

The second test is concerned with the 2D RTE. The domain $\Omega = [-x_0, 1 + x_0]^2$ is discretized with a uniform Cartesian grid with 80×80 points, where x_0 is chosen such that there are 70×70 points in Ω . The scattering coefficient is generated following the same way of η in Section 4 for the 2D case, followed by an appropriate rescaling. The source term $f(x)$ is generated by sampling independently from $\mathcal{U}(0, 1)$ instead of $\mathcal{N}(0, 1)$. The values of η and f outside of Ω are set to be 0. Results reported in Table 5 show that by setting $\alpha = 11$ and $n_{\text{cnn}} = 5$, the NN can achieve a test error of 4.4×10^{-3} with as few as 1.3×10^6 parameters. A representative sample from the test set is illustrated in Fig. 9.

6. Conclusions

This paper presented a meta-learning approach for learning the map from the equation parameter η to the pseudo-differential solution operator G_η . Motivated by the nonstandard wavelet form [5], the pseudo-differential operator is compressed to a collection of vectors. The nonlinear map from the parameter to this collection of vectors and the wavelet

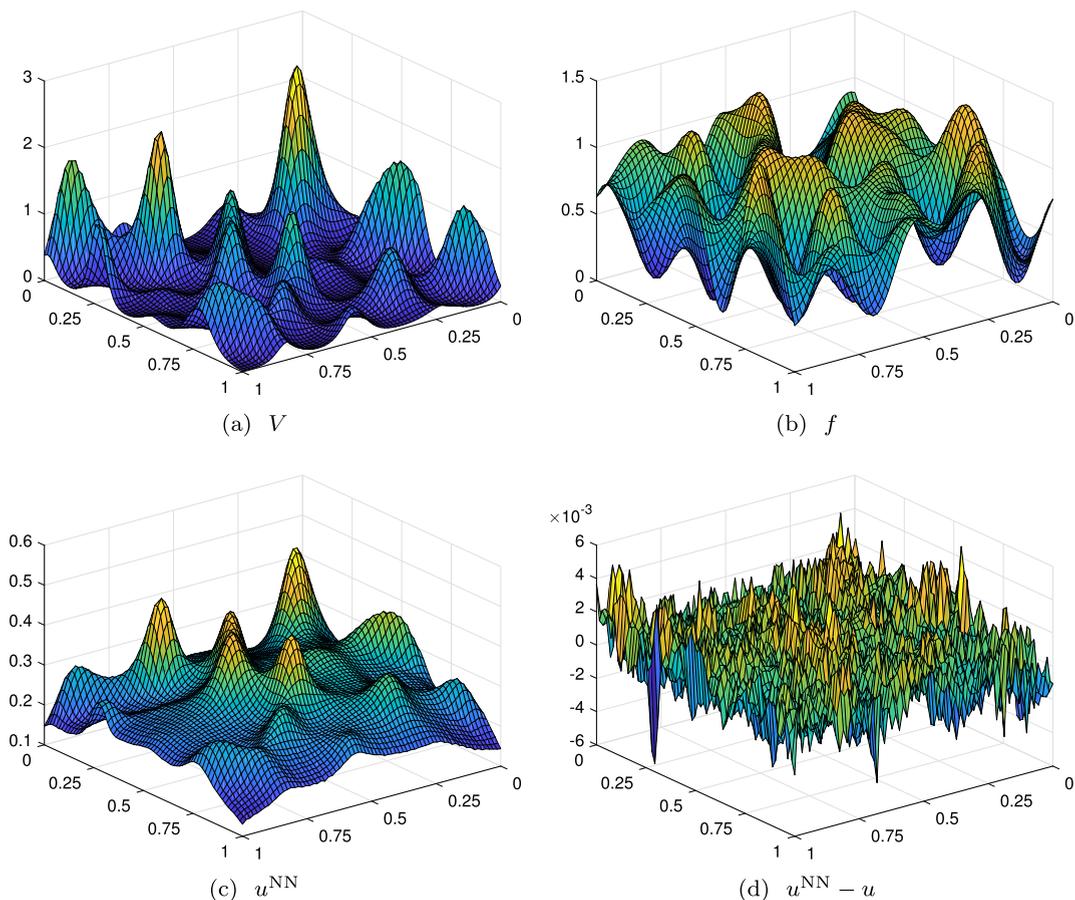


Fig. 9. A sample from the test set for the scattering coefficient η , the source term f , the prediction u^{NN} with $\alpha = 11$ and $n_{\text{cm}} = 5$, and the corresponding error for the 2D RTE.

transform are learned hand-in-hand in the meta-learning approach. Numerical studies are carried out for the Green's functions of elliptic PDEs as well as the radiative transfer equation.

This approach can be extended in several directions. First, this paper is only concerned with linear operators G_η . This work can be readily extended to nonlinear operators if a simple compressed representation (such as the collection of vectors used here) can be identified. Second, the ConvNet module for the map $\eta \rightarrow C_\eta^{(\ell)}$ can be replaced with the recently proposed multiscale NNs [19,17,18], which are more effective for certain global-scale convolutions.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The work of Y.F. and L.Y. is partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program. The work of J.F. is partially supported by Stanford Graduate Fellowship in Science & Engineering and by “la Caixa” Fellowship, sponsored by the “la Caixa” Banking Foundation of Spain under Fellowship LCF/BQ/AA16/11580045. The work of L.Y. is also partially supported by the National Science Foundation under award DMS-1818449. This work is also supported by the GCP Research Credits Program from Google and AWS Cloud Credits for Research program from Amazon.

References

- [1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., Tensorflow: a system for large-scale machine learning, in: OSDI, vol. 16, 2016, pp. 265–283.
- [2] M. Araya-Polo, J. Jennings, A. Adler, T. Dahlke, Deep-learning tomography, *Lead. Edge* 37 (1) (2018) 58–66.

- [3] Y. Bengio, S. Bengio, J. Cloutier, Learning a Synaptic Learning Rule, Université de Montréal, 1990, Département d'informatique et de recherche opérationnelle.
- [4] J. Berg, K. Nyström, A unified deep artificial neural network approach to partial differential equations in complex geometries, *Neurocomputing* 317 (2018) 28–41.
- [5] G. Beylkin, R. Coifman, V. Rokhlin, Fast wavelet transforms and numerical algorithms I, *Commun. Pure Appl. Math.* 44 (2) (1991) 141–183.
- [6] G. Carleo, M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* 355 (6325) (2017) 602–606.
- [7] F. Chollet, et al., Keras, <https://keras.io>, 2015.
- [8] I. Daubechies, Orthonormal bases of compactly supported wavelets, *Commun. Pure Appl. Math.* 41 (7) (1988) 909–996.
- [9] R. Dautray, J.-L. Lions, *Mathematical Analysis and Numerical Methods for Science and Technology: Volume 6 Evolution Problems II*, Springer-Verlag, Berlin Heidelberg, 2000.
- [10] L. Demanet, L. Ying, Discrete symbol calculus, *SIAM Rev.* 53 (1) (2011) 71–104.
- [11] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova Bert, Pre-training of deep bidirectional transformers for language understanding, arXiv preprint, arXiv: 1810.04805, 2018.
- [12] T. Dozat, Incorporating Nesterov Momentum Into Adam, 2016.
- [13] W. E, J. Han, A. Jentzen, Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations, *Commun. Math. Stat.* 5 (4) (2017) 349–380.
- [14] W. E, B. Yu, The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems, CoRR, arXiv:1710.00211 [abs], 2017.
- [15] R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, T. Green, C. Qin, A. Zidek, A. Nelson, A. Bridgland, H. Penedones, et al., De novo structure prediction with deeplearning based scoring, *Annu. Rev. Biochem.* 77 (2018) 363–382.
- [16] Y. Fan, J. An, L. Ying, Fast algorithms for integral formulations of steady-state radiative transfer equation, *J. Comput. Phys.* 380 (2019) 191–211.
- [17] Y. Fan, J. Feliu-Fabà, L. Lin, L. Ying, L. Zepeda-Núñez, A multiscale neural network based on hierarchical nested bases, *Res. Math. Sci.* 6 (21) (2019).
- [18] Y. Fan, L. Lin, L. Ying, L. Zepeda-Núñez, A multiscale neural network based on hierarchical matrices, *Multiscale Model. Simul.* 17 (4) (2019) 1189–1213.
- [19] Y. Fan, C. Orozco-Bohorquez, L. Ying, BCR-Net: a neural network based on the nonstandard wavelet form, *J. Comput. Phys.* 384 (2019) 1–15.
- [20] Y. Fan, L. Ying, Solving electrical impedance tomography with deep learning, *J. Comput. Phys.* 404 (2020) 109119.
- [21] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in: Proceedings of the 34th International Conference on Machine Learning, vol. 70, ICML'17, 2017, JMLR.org.
- [22] W. Hackbusch, A sparse matrix arithmetic based on \mathcal{H} -matrices. Part I: Introduction to \mathcal{H} -matrices, *Computing* 62 (2) (1999) 89–108.
- [23] W. Hackbusch, L. Grasedyck, S. Börm, An introduction to hierarchical matrices, *Math. Bohem.* 127 (229–241) (2002).
- [24] W. Hackbusch, B.N. Khoromskij, S. Sauter, On \mathcal{H}^2 -matrices, in: *In Lectures on Applied Mathematics*, Springer, 2000, pp. 9–29.
- [25] J. Han, A. Jentzen, W. E, Solving high-dimensional partial differential equations using deep learning, *Proc. Natl. Acad. Sci.* 115 (34) (2018) 8505–8510.
- [26] J. Han, L. Zhang, R. Car, et al., Deep potential: a general representation of a many-body potential energy surface, arXiv preprint, arXiv:1707.01478, 2017.
- [27] B. Hariharan, R. Girshick, Low-shot visual recognition by shrinking and hallucinating features, in: Proceedings of the International Conference on Computer Vision, ICCV, 2017, pp. 3018–3027.
- [28] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2016, pp. 770–778.
- [29] G. Hinton, L. Deng, D. Yu, G.E. Dahl, A.R. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T.N. Sainath, B. Kingsbury, Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups, *IEEE Signal Process. Mag.* 29 (6) (2012) 82–97.
- [30] A. Jacot, F. Gabriel, C. Hongler, Neural tangent kernel: convergence and generalization in neural networks, in: Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18, USA, 2018, pp. 8580–8589, Curran Associates Inc.
- [31] Y. Khoo, J. Lu, L. Ying, Solving parametric PDE problems with artificial neural networks, arXiv preprint, arXiv:1707.03351, 2017.
- [32] Y. Khoo, J. Lu, L. Ying, Solving for high-dimensional committor functions using artificial neural networks, *Res. Math. Sci.* 6 (1) (2019) 1.
- [33] Y. Khoo, L. Ying, SwitchNet: a neural network model for forward and inverse scattering problems, *SIAM J. Sci. Comput.* 41 (5) (2019) A3182–A3201.
- [34] A.D. Klose, U. Netz, J. Beuthan, A.H. Hielscher, Optical tomography using the time-independent equation of radiative transfer—part 1: forward model, *J. Quant. Spectrosc. Radiat. Transf.* 72 (5) (2002) 691–713.
- [35] R. Koch, R. Becker, Evaluation of quadrature schemes for the discrete ordinates method, *J. Quant. Spectrosc. Radiat. Transf.* 84 (4) (2004) 423–435.
- [36] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: Proceedings of the 25th International Conference on Neural Information Processing Systems, vol. 1, NIPS'12, USA, 2012, pp. 1097–1105, Curran Associates Inc.
- [37] I.E. Lagaris, A. Likas, D.I. Fotiadis, Artificial neural networks for solving ordinary and partial differential equations, *IEEE Trans. Neural Netw.* 9 (5) (1998) 987–1000.
- [38] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (436) (2015).
- [39] Y. Li, J. Lu, A. Mao, Variational training of neural network approximations of solution maps for physical models, arXiv preprint, arXiv:1905.02789, 2019.
- [40] S. Liang, R. Srikant, Why deep neural networks for function approximation?, in: 6th International Conference on Learning Representations, ICLR 2018, 2018.
- [41] Z. Liu, Y. Yang, Q. Cai, Neural network as a function approximator and its application in solving differential equations, *Appl. Math. Mech.* 40 (2) (2019) 237–248.
- [42] R. Livni, S. Shalev-Shwartz, O. Shamir, On the computational efficiency of training neural networks, in: *Advances in Neural Information Processing Systems*, 2014, pp. 855–863.
- [43] J. Ma, R.P. Sheridan, A. Liaw, G.E. Dahl, V. Svetnik, Deep neural nets as a method for quantitative structure–activity relationships, *J. Chem. Inf. Model.* 55 (2) (2015) 263–274, PMID, 25635324.
- [44] S. Mallat, A wavelet tour of signal processing: the sparse way, in: *A Wavelet Tour of Signal Processing: the Sparse Way*, third edition, Academic Press, Boston, 2008.
- [45] A. Marshak, A. Davis, *3D Radiative Transfer in Cloudy Atmospheres*, Springer Science & Business Media, 2005.
- [46] I. Ohn, Y. Kim, Smooth function approximation by deep neural networks with general activation functions, *Entropy* 21 (7) (2019) 627.
- [47] T. Poggio, H. Mhaskar, L. Rosasco, B. Miranda, Q. Liao, Why and when can deep—but not shallow—networks avoid the curse of dimensionality: a review, *Int. J. Autom. Comput.* 14 (5) (2017) 503–519.
- [48] G.C. Pomraning, *The Equations of Radiation Hydrodynamics*, Courier Corporation, 1973.
- [49] M. Raissi, P. Perdikaris, G.E. Karniadakis, Physics informed deep learning (part I): data-driven solutions of nonlinear partial differential equations, CoRR, arXiv:1711.10561 [abs], 2017.
- [50] K. Ren, R. Zhang, Y. Zhong, A fast algorithm for radiative transport in isotropic media, *J. Comput. Phys.* 399 (2019) 108958.
- [51] R. Sarikaya, G.E. Hinton, A. Deoras, Application of deep belief networks for natural language understanding, *IEEE/ACM Trans. Audio Speech Lang. Process.* 22 (4) (2014) 778–784.
- [52] J. Schmidhuber, *Evolutionary Principles in Self-Referential Learning*, PhD thesis, Institut F. Informatik, Tech. Univ. Munich, 1987.
- [53] J. Schmidhuber, Deep learning in neural networks: an overview, *Neural Netw.* 61 (2015) 85–117.

- [54] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis, Mastering the game of go with deep neural networks and tree search, *Nature* 529 (2016) 484–503.
- [55] J. Sirignano, K. Spiliopoulos, DGM: a deep learning algorithm for solving partial differential equations, *J. Comput. Phys.* 375 (2018) 1339–1364.
- [56] R. Socher, Y. Bengio, C.D. Manning, Deep learning for NLP (without magic), in: *The 50th Annual Meeting of the Association for Computational Linguistics*, Tutorial Abstracts, vol. 5, 2012.
- [57] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008.
- [58] O. Vinyals, T. Ewalds, S. Bartunov, P. Georgiev, A.S. Vezhnevets, M. Yeo, A. Makhzani, H. Küttler, J. Agapiou, J. Schrittwieser, et al., StarCraft II: a new challenge for reinforcement learning, *arXiv preprint*, arXiv:1708.04782, 2017.
- [59] J.X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J.Z. Leibo, R. Munos, C. Blundell, D. Kumaran, M. Botvinick, Learning to reinforcement learn, *arXiv preprint*, arXiv:1611.05763, 2016.
- [60] M.W. Wong, *An Introduction to Pseudo-Differential Operators*, vol. 6, World Scientific Publishing Company, 2014.
- [61] M.D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: *Computer Vision - ECCV 2014–13 European Conference*, 2014, pp. 818–833.