Solutions to Elliptic and Parabolic Problems via Finite Difference Based Unsupervised Small Linear Convolutional Neural Networks

Adrian Celaya^{a,b}, Keegan Kirk^a, David Fuentes^b, Beatrice Riviere^a

^aRice University, Houston, 77005, TX, USA ^bThe University of Texas MD Anderson Cancer Center, Houston, 77003, TX, USA

Abstract

In recent years, there has been a growing interest in leveraging deep learning and neural networks to address scientific problems, particularly in solving partial differential equations (PDEs). However, current neural network-based PDE solvers often rely on extensive training data or labeled input-output pairs, making them prone to challenges in generalizing to out-of-distribution examples. To mitigate the generalization gap encountered by conventional neural network-based methods in estimating PDE solutions, we formulate a fully unsupervised approach, requiring no training data, to estimate finite difference solutions for PDEs directly via small convolutional neural networks. Our proposed algorithms demonstrate comparable accuracy to the true solution for several selected elliptic and parabolic problems compared to the finite difference method.

Keywords: Convolutional neural networks, unsupervised learning, partial differential equations, finite difference

1. Introduction

Partial differential equations (PDEs) of elliptic and parabolic type are ubiquitous in the mathematical modeling of many physical phenomena, and thus see wide application to many real-world problems. Classical numerical methods (finite difference methods, finite element methods, etc.) introduce a computational mesh over which matrix representations of differential operators are defined. This results in large linear systems whose efficient solution presents a number of computational challenges. Recently, neural networks

have gained popularity in solving PDEs, partially thanks to the universal approximation property of neural networks [1, 2]. Popular approaches include Physics Informed Neural Networks (PINNs) [3], the Deep Galerkin method [4], and the Deep Ritz method [5, 6]. These methods incorporate the underlying physics of the problem by introducing the differential operators of the PDE or an equivalent variational principle into the loss function. The common feature of these methods is the requirement of training data as these methods belong to the class of supervised learning methods. This requires sufficient training data, which is often generated using more classical numerical algorithms. An important drawback is the poor generalization in neural networks for supervised learning.

Our proposed method differs from the methods described above because it does not require training data. This helps mitigate the generalization problem of existing approaches. Our unsupervised learning method is an iterative method that computes successive approximations of the solution of the PDE given a source function and boundary data. The iterations are constructed to minimize a finite difference based loss function. To reduce the computational cost, the PocketNet paradigm is applied to the neural network architecture [7]. The method is presented for both steady-state and time-dependent problems. Finite difference methods have inspired the construction of various neural networks [8, 9], but again these methods need training data. The recent work [10] solves elliptic PDEs with constant diffusion and it shares the following similarities with our proposed numerical method: no training data is needed and the loss function is derived from the five-point stencil of the finite difference approximation. However, our loss function is different as it incorporates the Dirichlet boundary condition in a weighted fashion. We also define the method for elliptic PDEs with nonconstant diffusion coefficient and we extend it to time-dependent problems. Another key feature of our method is the use of small neural networks, which makes it computationally efficient.

An outline of the paper is as follows. Section 2 introduces the model problem, the network architecture and the loss function. Section 3 describes the algorithms for both steady-state and time-dependent problems. Results and discussion are presented in Section 4. Conclusions follow.

2. Material and Methods

We propose a fully unsupervised method for estimating finite differences solutions to partial differential equations via convolutional neural networks. In contrast to existing deep learning-based methods, our approach is fully unsupervised. In other words, our approach does not require training data and estimates the solution to a given PDE directly via the optimization process, also called the training process.

2.1. Elliptic Problems

Let $\Omega \subset \mathbb{R}^2$ be an open set. Given a function $f: \Omega \to \mathbb{R}$ and diffusion coefficients κ , we consider the 2D Poisson problem with solution $u: \overline{\Omega} \to \mathbb{R}$ such that

$$-\nabla \cdot (\kappa \nabla u) = f \text{ in } \Omega,$$

$$u = g \text{ on } \partial \Omega.$$
 (1)

We begin with the case when κ is constant and assume that $\kappa = 1$. For readibility, we assume that Ω is the square domain $(0, L)^2$, that is partitioned into a uniform grid of $N \times N$ squares with size h = L/N. The methodology presented in the paper can be easily extended to rectangular domains. We first recall the standard finite difference method based on the five-point stencil, applied to (1). Let \mathcal{V}^0 (resp. \mathcal{V}^0) be the set of indices (i, j) such that the point $(x_i, y_j) = (ih, jh)$ belongs to the interior (resp. boundary) of Ω .

$$\frac{4u_{i,j} - u_{i+1,j} - u_{i-1,j} - u_{i,j+1} - u_{i,j-1}}{h^2} = f(x_i, y_j), \quad \forall (i, j) \in \mathcal{V}^0, \quad (2)$$

$$u_{i,j} = g(x_i, y_j), \quad \forall (i, i) \in \mathcal{V}^{\partial}.$$
 (3)

The finite difference solution u_h is a vector with entries $u_{i,j}$ satisfying (2)-(3) and it is known that for smooth enough solutions u and small enough h, the value $u_{i,j}$ is a good approximation of $u(x_i, y_j)$. One can rewrite (2) in terms of a discrete convolutional operator \star and convolutional kernel K_{Δ} defined as:

$$K_{\Delta} = \frac{1}{h^2} \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}. \tag{4}$$

We use the shorthand notation $f_{i,j} = f(x_i, y_j)$. The equivalent form for (2) is

$$(K_{\Delta} \star u_h)_{i,j} = f_{i,j}, \quad \forall (i,j) \in \mathcal{V}^0.$$
 (5)

For a given kernel $K \in \mathbb{R}^{3\times 3}$, the convolution operator \star is defined by

$$(K \star u_h)_{i,j} = \sum_{p=-1}^{1} \sum_{q=-1}^{1} K_{p,q} u_{i+p,j+q}.$$
 (6)

With this convolution-based discretization, we can now reformulate the finite difference approximation of (1) as a convex optimization problem given by

$$\underset{u_h}{\operatorname{arg\,min}} \sum_{(i,j)\in\mathcal{V}^0} \left((K_{\Delta} \star u_h)_{i,j} - f_{i,j} \right)^2 + \sum_{(i,j)\in\mathcal{V}^\partial} \left((u_h)_{i,j} - g_{i,j} \right)^2. \tag{7}$$

We obtain an approximate solution \hat{u} to (7) by training a neural network $\mathcal{N}_{\theta}: \mathbb{R}^{N \times N} \to \mathbb{R}^{N \times N}$ with trainable parameters $\boldsymbol{\theta} \in \mathbb{R}^{M}$. We propose to train \mathcal{N}_{θ} using the unsupervised loss function $\mathcal{L}_{\alpha}: \mathbb{R}^{N \times N} \to \mathbb{R}$ given by

$$\mathcal{L}_{\alpha}(\hat{u}) = \alpha \sum_{(i,j)\in\mathcal{V}^{0}} \left((K_{\Delta} \star \hat{u})_{i,j} - f_{i,j} \right)^{2} + (1 - \alpha) \sum_{(i,j)\in\mathcal{V}^{\partial}} \left(\hat{u}_{i,j} - g_{i,j} \right)^{2}, \quad (8)$$

where $\alpha = h^2/4$ is a weighting term. The idea behind this loss function is to directly estimate the finite difference approximation to (1) without using any training data. Specifically, we use f as the input to the CNN and we perform \mathcal{M} updates of the parameters in the neural network \mathcal{N}_{θ} , using the loss function defined above. Algorithm 1 in Section 3 outlines this training procedure in more detail.

In the case where κ is piecewise constant, $\kappa = (\kappa_{i,j})$, we use the following finite difference discretization

$$-\frac{1}{h^2} \left(\kappa_{i+\frac{1}{2},j} (u_{i+1,j} - u_{i,j}) - \kappa_{i-\frac{1}{2},j} (u_{i,j} - u_{i-1,j}) \right) - \frac{1}{h^2} \left(\kappa_{i,j+\frac{1}{2}} (u_{i,j+1} - u_{i,j}) - \kappa_{i,j-\frac{1}{2}} (u_{i,j} - u_{i,j-1}) \right) = f_{i,j}, \quad (9)$$

where the interface values $\kappa_{i+\frac{1}{2},j}$, $\kappa_{i-\frac{1}{2},j}$, $\kappa_{i,j+\frac{1}{2}}$, and $\kappa_{i,j-\frac{1}{2}}$ are the harmonic averages of $\kappa_{i,j}$ given by

$$\kappa_{i+\frac{1}{2},j} = \frac{2}{\frac{1}{\kappa_{i+1,j}} + \frac{1}{\kappa_{i,j}}} \qquad \kappa_{i-\frac{1}{2},j} = \frac{2}{\frac{1}{\kappa_{i,j}} + \frac{1}{\kappa_{i-1,j}}} \\
\kappa_{i,j+\frac{1}{2}} = \frac{2}{\frac{1}{\kappa_{i,j+1}} + \frac{1}{\kappa_{i,j}}} \qquad \kappa_{i,j-\frac{1}{2}} = \frac{2}{\frac{1}{\kappa_{i,j}} + \frac{1}{\kappa_{i,j-1}}}.$$
(10)

The interface values necessitate the imposition of a half-grid (also called dual grid) onto our original grid. As a result, we modify our neural network prediction via a dilation operator $D: \mathbb{R}^{N \times N} \to \mathbb{R}^{(2N-1) \times (2N-1)}$ such that

$$D(\hat{u})_{i,j} = \begin{cases} \hat{u}_{1+\frac{i-1}{2},1+\frac{j-1}{2}}, & \text{if } i \text{ and } j \text{ are odd} \\ 0, & \text{otherwise.} \end{cases}$$
 (11)

Additionally, we map κ from the original grid to the dual grid by applying the following steps:

Step 1: apply the dilation operator to obtain $\kappa^* \in \mathbb{R}^{(2N-1)\times(2N-1)}$:

$$\kappa^* = D(\kappa). \tag{12}$$

Step 2: update the values of κ^* on the half-grid:

$$\kappa_{i,j}^* = \begin{cases}
2\left(\frac{1}{\kappa_{i-1,j}^*} + \frac{1}{\kappa_{i+1,j}^*}\right)^{-1}, & \text{if } i \text{ odd and } j \text{ even} \\
2\left(\frac{1}{\kappa_{i,j-1}^*} + \frac{1}{\kappa_{i,j+1}^*}\right)^{-1}, & \text{if } i \text{ even and } j \text{ odd}
\end{cases}$$
(13)

Step 3: set to zero the values of κ^* that would correspond to the values on the original grid:

$$\tilde{\kappa} = \left(\mathbf{1}_{(2N-1)\times(2N-1)} - D(\mathbf{1}_{N\times N})\right) \odot \kappa^*. \tag{14}$$

In the equation above, $\mathbf{1}_{N\times N}$ is an $N\times N$ matrix with all entries equal to one, and \odot is the Hadamard or pointwise product.

We now introduce the following convolution kernels

which will allow us to rewrite (9) as follows:

$$-\frac{1}{h^2} \left((T_{\uparrow} \star D(u)) \odot (P_{\uparrow} \star \tilde{\kappa}) - (T_{\downarrow} \star D(u)) \odot (P_{\downarrow} \star \tilde{\kappa}) \right. \\ + \left. (T_{\to} \star D(u)) \odot (P_{\to} \star \tilde{\kappa}) - (T_{\leftarrow} \star D(u)) \odot (P_{\leftarrow} \star \tilde{\kappa}) \right)_{i,j} = f_{i,j}, \forall i, j \in \mathcal{V}^0,$$

$$(16)$$

which is then used in the first term of the loss function (8) to train our neural network. The same training process outlined in Algorithm 1 applies here. Note that in (16), the convolution operator \star uses a stride of two because of the dual grid and the use of larger convolution kernels.

2.2. Parabolic Problems

We generalize the method described in the previous section to timedependent problems. We propose a modified loss function that incorporates the backward Euler method to address the time derivative. Consider the following parabolic problem on a square domain Ω and for the time interval $[t_0, T]$.

$$\frac{\partial u}{\partial t} - \Delta u = f \quad \text{in} \quad \Omega \times (t_0, T],$$

$$u = g \quad \text{on} \quad \partial \Omega \times (t_0, T],$$

$$u(\cdot, t_0) = u^0 \quad \text{in} \quad \Omega \times \{t_0\}.$$
(17)

Let $\tau > 0$ be the time step size. At each time step, $t^n = n\tau$, the finite difference solution denoted by u_h^n solves the following optimization problem

$$\underset{u_{h}^{n}}{\operatorname{arg\,min}} \sum_{(i,j)\in\mathcal{V}^{0}} \left(u_{i,j}^{n} - u_{i,j}^{n-1} - \tau \left(f_{i,j}^{n} - (K_{\Delta} \star u_{h}^{n})_{i,j} \right) \right)^{2} + \sum_{(i,j)\in\mathcal{V}^{\partial}} \left(u_{i,j}^{n} - g_{i,j}^{n} \right)^{2}.$$
(18)

Above we use the short-hand notation $f_{i,j}^n = f(x_i, y_j, t^n)$ and $g_{i,j}^n = g(x_i, y_j, t^n)$. Here, the unknowns $u_{i,j}^n$ are the approximations of the solution u evaluated at the point (x_i, y_j) and time t^n . We then use the following loss function to train our CNN at each time step

$$\mathcal{L}_{\alpha,\tau}(\hat{u}^n, \hat{u}^{n-1}, t^n) = \alpha \sum_{(i,j) \in \mathcal{V}^0} \left(\hat{u}_{i,j}^n - \hat{u}_{i,j}^{n-1} - \tau \left(f_{i,j}^n - (K_\Delta \star \hat{u}^n)_{i,j} \right) \right)^2 + (1 - \alpha) \sum_{(i,j) \in \mathcal{V}^\partial} \left(\hat{u}_{i,j}^n - g_{i,j}^n \right)^2.$$
(19)

Again, we will choose the weight $\alpha = h^2/4$. The key difference between (8) and (19) is that we minimize at each time step. The input to the CNN at each time step is the previous solution \hat{u}^{n-1} and the output is the solution at the time step t^n . Algorithm 2 in Section 3 outlines this training procedure in more detail.

2.3. Network Architecture

We select as our architecture the widely used U-Net [11]. The U-Net was initially designed for image segmentation tasks but is successful in several scientific machine learning tasks [12, 13, 14]. The architecture consists of convolution block operators (blocks), a downsampling (or encoding) path, and an upsampling (or decoding) path. A channel-wise concatenation operation links each layer in the downsampling and upsampling paths. Each block consists of two 2D convolution layers. Unlike most implementations of the U-Net architecture [11, 15, 16], we omit normalization (i.e., batch or instance normalization) and do not apply a non-linear activation function like ReLU and use the identity function instead (see Section 4). Additionally, we use the PocketNet approach proposed in [7] and leave the number of feature maps (channels) at each resolution within our architecture constant, namely

32. Figure 1 sketches the U-Net architecture for different network depths. We define the depth, d, of each network to be the number of downsampling operations present in the network. The size of the feature tensor is halved from depth d to depth d+1.

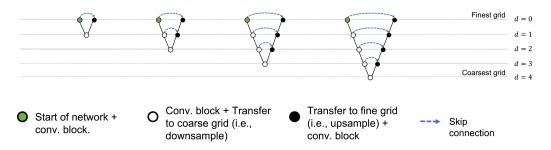


Figure 1: Sketch of U-Net architecture for different network depths: $0 \le d \le 4$.

2.4. Data

2.4.1. Elliptic Problem Data

We test our proposed approach for elliptic problems on four different test cases. Unless otherwise specified, the domain is the unit square.

The Bubble Function. In the first test case, (1) is solved with data chosen such that the exact solution is the bubble function

$$u(x,y) = x(x-1)y(y-1). (20)$$

This function is symmetric, and we apply homogeneous zero Dirichlet boundary conditions. Note that for this function, one can show that the five-point stencil scheme given by (2) applied to the bubble function gives exactly $-\Delta u$ evaluated at the grid point (x_i, y_j) , which implies that the finite difference approximation is exact.

The "Peak" Function. In the second case, we want to test the effectiveness of our strategy when the solution has high gradients and is non-symmetric. Hence, we select data such that the exact solution is the "peak" function

$$u(x,y) = 0.0005 (x(x-1)y(y-1))^{2} e^{10x^{2}+10y}.$$
 (21)

The Exponential-Trigonometric Function. We test our approach on a function with non-homogeneous Dirichlet boundary conditions in the third case. We select data such that the exact solution is the exponential-trigonometric function given by

$$u(x,y) = e^{-x^2 - y^2} \sin(3\pi x) \sin(3\pi y) + x.$$
 (22)

Non-Constant and Discontinuous Diffusion. To test how our approach fairs on problems with low regularity, we consider an example with a non-constant, discontinuous diffusion coefficient taken from [17]. The problem is posed on the square domain $\Omega = (-1,1)^2$ which we divide into four subdomains Ω_i corresponding to the quadrants of the Cartesian plane. On each subdomain, κ is constant and takes the values $\kappa_1 = \kappa_3 = 5$ and $\kappa_2 = \kappa_4 = 1$. The exact solution takes the form

$$r^{\beta}(a_i \sin(\beta \theta) + b_i \cos(\beta \theta)), \tag{23}$$

where (r, θ) are the polar coordinates of a given point in Ω , a_i , b_i are constants that depend on the subdomains (see [17] for exact values). This discontinuity in the diffusion coefficient introduces a singularity in the solution at the origin, namely the function belongs to the Sobolev space $H^{1+\beta}(\Omega)$, with $\beta \approx 0.5354$.

2.4.2. Parabolic Problem Data

We test our proposed approach for parabolic problems on three different test cases.

The Trigonometric Function. In the first two cases, (17) is solved with data such that the exact solution is given by the trigonometric function

$$u(x, y, t) = \cos(t)\sin(n\pi x)\sin(n\pi y). \tag{24}$$

In the first case, we select n = 1, which gives a symmetric function with a single peak. In the second case, we set n = 4, which results in 16 peaks and troughs in our domain.

The Gaussian Function. In the third case, we test our approach on a Gaussian function centered on the point $(\frac{1}{2}, \frac{1}{2})$

$$u(x, y, t) = \cos(t)e^{-50((2x-1)^2 + (2y-1)^2)}.$$
 (25)

2.5. Training and Testing Protocols

We use the Adam optimizer [18] and 5×5 convolutional kernels in each layer. The initial learning rate is 0.001 for the steady-state elliptic problems and 0.0001 for parabolic problems. During training, we use L2 regularization with a penalty of 10^{-7} and the norm of the network's gradient is clipped so that it is no greater than 10^{-2} . Our models are implemented in Python using TensorFlow (v2.12.0) and trained on an NVIDIA Quadro RTX 6000 GPU [19]. All network weights are initialized using the default TensorFlow initializers. All other hyperparameters are left at their default values. The code for our network architecture is available at https://github.com/aecelaya/pde-nets.

To assess the accuracy of our predicted solutions, we use the following norms of the error between the exact solution u and its approximate \hat{u} for steady-state problems:

$$||u - \hat{u}||_{2,h} = h \sqrt{\sum_{(i,j) \in \mathcal{V}^0 \cup \mathcal{V}^\partial} (u(x_i, y_j) - \hat{u}_{i,j})^2}$$

$$||u - \hat{u}||_{\infty} = \max_{(i,j) \in \mathcal{V}^0 \cup \mathcal{V}^\partial} |u(x_i, y_j) - \hat{u}_{i,j}|.$$
(26)

$$|u - \hat{u}||_{\infty} = \max_{(i,j) \in \mathcal{V}^0 \cup \mathcal{V}^0} |u(x_i, y_j) - \hat{u}_{i,j}|.$$
 (27)

For time-dependent problems, similar errors are computed at the final time T.

3. Calculation

We use Algorithm 1 to approximate a solution \hat{u} to (1). In this algorithm, we start by initializing the neural network \mathcal{N}_{θ} with trainable parameters θ . In this case, we use the U-Net architecture described in Section 2.3. We set the maximum number of optimization steps \mathcal{M} as the stopping criterion for the algorithm. We then begin the minimization (or training) process by generating a prediction from our neural network with the source term f as the input. Given the current prediction, we compute the loss value and update the network weights via backpropagation. If the loss value from the current prediction is less than the previous best value, we update our best loss and save the current prediction. Algorithm 1 fully outlines this procedure.

We use Algorithm 2 to approximate a solution \hat{u}^{N_T} to (17), where N_T is an integer such that $T = N_T \tau$. Like with Algorithm 1, we randomly initialize a U-Net and set a maximum number of optimization steps as a stopping

```
Algorithm 1 Unsupervised CNN Training For (1)
     Input Right-hand side f and boundary condition g
     Output Approximate solution to (7), \hat{u}_*
 1: Randomly initialize the network parameters \theta
 2: Set maximum iterations \mathcal{M}
 3: k \leftarrow 0
                                                                        ▶ Iteration counter
 4: \ell_* \leftarrow +\infty

    Store smallest loss value

 5: while k < \mathcal{M} do
         \hat{u} \leftarrow \mathcal{N}_{\theta}(f)
                                                                6:
         \ell \leftarrow \mathcal{L}_{\alpha}(\hat{u})
                                                                            7:
         Update \theta using backpropagation on \ell
 8:
 9:
         if \ell < \ell_* then
                                                                    ▷ Save best prediction
             \hat{u}_* \leftarrow \hat{u}
10:
             \ell_* \leftarrow \ell
                                                            ▶ Update smallest loss value
11:
         k \leftarrow k + 1
12:
```

criterion. However, in this case, the maximum number of steps applies to each time step. We then set the initial condition as the previous solution. We apply nearly the same steps for each time step as with Algorithm 1. Namely, we produce a candidate prediction from our neural network with the previous solution as the input. We compute the loss value and update the network weights via backpropogation on the loss. If the loss value is lower than the previously observed lowest loss value, then we update our lowest loss value and save the current prediction. At the end of this process, we set the previous solution to the best current solution. This process continues for every time step. Algorithm 2 fully outlines this procedure. Note that in this algorithm, we do not reinitialize our neural network weights at each time step. Instead, we use the previous network configuration as the initial state for the next time step.

4. Results and Discussion

Using the methods described in Section 2 and Algorithm 1, we approximate the finite difference solution to the elliptic problems defined in Section 2.4.1 with constant diffusion coefficients (i.e., $\kappa=1$). Table 1 shows the accuracy of our unsupervised predictions for a varying grid size and number of optimization steps. We fix the depth of our U-Net architecture to three in

Algorithm 2 Unsupervised CNN Training For (17)

16:

```
Input Number of time steps N_T, final time T, initial condition u^0,
     right hand side f, and boundary condition g
     Output Approximate solution to (18) at final time \hat{u}^{N_T}
 1: Randomly initialize the network parameters \theta
 2: Set maximum iterations \mathcal{M}
 3: \tau \leftarrow T/N_T
                                                                  ▶ Initialize time step size
 4: n \leftarrow 1

→ Time step counter

 5: while n \leq N_T do
         k \leftarrow 0
                                                                          ▶ Iteration counter
 6:
         \ell_* \leftarrow +\infty

    Store smallest loss value

 7:
         while k < \mathcal{M} do
 8:
             w \leftarrow \mathcal{N}_{\theta}(\hat{u}^{n-1})
 9:
                                                                  \ell \leftarrow \mathcal{L}_{\alpha,\tau}(w, \hat{u}^{n-1}, n\tau)

    Compute loss

10:
             Update \theta using backpropagation on \ell
11:
             if \ell < \ell_* then
12:
                  \hat{u}^n \leftarrow w
                                                                      > Save best prediction
13:
                  \ell_* \leftarrow \ell
                                                              ▶ Update smallest loss value
14:
15:
             k \leftarrow k + 1
         n \leftarrow n+1
```

this case. For reference, Table 3 shows the accuracy of the finite difference method for varying grid sizes. For the peak and exponential-trigonometric functions, our approach provides a good approximation to the finite difference solution. We do not, however, see this for the bubble function example. In that case, the solution of the finite difference method is exact. Under our selected settings (i.e., depth and optimization steps), our method stops at an error of approximately 10^{-6} in the $||\cdot||_{2,h}$ norm for the bubble function example. The cause of this discrepancy between our method applied to the bubble function and the finite difference solution may be that current neural network optimizers like Adam are stochastic, gradient-based, and do not have optimal step sizes. We hypothesize that our neural network is getting stuck in a neighborhood around the global minimum because of these previously listed limitations of optimizers like Adam. Hence, using secondorder, Hessian-based optimizers or better step-size computations may allow for faster convergence to more accurate solutions for our neural networkbased approach.

Figure 2 shows the true solution, predicted solution, and absolute difference for the bubble, peak, and exponential-trigonometric cases on a 128×128 grid. This figure shows that for every case, our unsupervised algorithm produces visually indistinguishable solutions from the true solution in each case.

We also want to study the network depth's effect on our predictions' accuracy. Table 2 shows the accuracy of our unsupervised neural network predictions for varying grid sizes and U-Net depths. We set the number of optimization steps to 4,000. Here, we see that the depth of the U-Net architectures does not appear to have a significant effect on the accuracy of our predictions. The only notable exceptions are for the bubble and exponential-trigonometric functions on a 128×128 grid with a depth equal to two. This discrepancy may be caused by the network not capturing sufficiently rich features on the finer grid. The fact that we see a decrease in the errors between our predictions and the true solution to those comparable to the finite difference method as we increase the depth to three or greater supports this hypothesis. Note that for coarser grids (i.e., fewer grid points), we do not test higher network depths since the size of the features (the output of each layer) at the coarser grids in such networks would be smaller than the size of the convolutional kernels in the network.

The results of Algorithm 1 for the non-constant diffusion problem defined in Section 2.4.1 are shown in Table 4. Here, we see more significant errors than with the constant diffusion problems. Figure 3 shows the true solution, predicted solution, and absolute difference for the non-constant diffusion problem on a 128×128 grid. This figure shows that the errors are mostly concentrated around the discontinuities in κ . We do not compare this to the finite difference method because the exact solution exhibits a singularity at the origin. Indeed, the gradient of the exact solution blows up at the point (0,0).

Because the finite difference solution solves the convex minimization problem (7), the finite difference approximation's accuracy limits our approach's accuracy. This explains the relatively poor performance in the case of discontinuous diffusion. It is well known that, due to the singularity at the origin, the exact solution to the discontinuous diffusion problem belongs to the Sobolev space $H^{1+\beta}(\Omega)$, where $\beta \approx 0.5354$ [17]. Hence, the finite difference method, which requires more regularity, performs poorly.

Using the methods described in Section 2 and Algorithm 2, we approximate the finite difference solution to the parabolic problems defined in Section

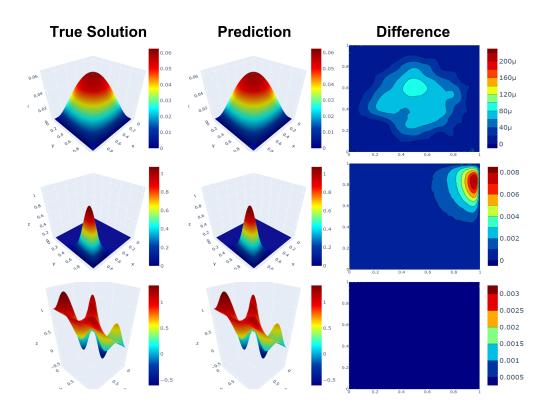


Figure 2: (Top) Bubble function. (Middle) "Peak" function. (Bottom) Exponential trigonometric function. From left to right, contour plots of true solution, predicted solution, and absolute difference. All predictions and solutions are on a 128×128 grid. Note that $\mu=10^{-6}$ in the colorbar for the bubble function difference.

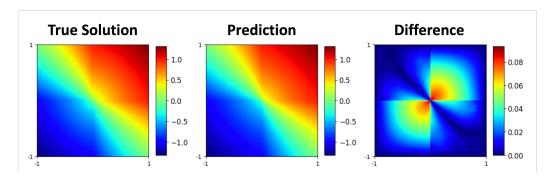


Figure 3: From left to right, contour plots of true solution, predicted solution, and absolute difference for the non-constant diffusion problem.

Problem	\mathcal{M}	N =	= 32	N =	= 64	N = 128		
1 Toblem	JV 1	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u-\hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u-\hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	
	500	4.7521e-05	5.2953e-04	4.8521e-05	8.2744e-04	4.2407e-04	1.6607e-03	
	1,000	2.5330e-05	2.4341e-04	2.1550e-05	3.4875 e-04	1.3244e-04	5.3431e-04	
Bubble	2,000	1.4457e-05	1.4313e-04	9.1955e-06	1.4581e-04	4.0090e-05	3.0874e-04	
	4,000	3.2266e-06	3.0962e-05	3.1159e-06	5.4674 e - 05	7.6852e-06	8.5074 e-05	
	8,000	2.2661e-06	2.9351e-05	3.7016e-06	4.9231e-05	3.7105e-06	4.4149e-05	
	500	2.0066e-02	1.2398e-01	1.0361e-02	1.2776e-01	3.2611e-02	2.8655e-01	
	1,000	2.0603e-02	1.1937e-01	5.3352e-03	3.2162 e-02	5.0469e-03	2.6498e-02	
Peak	2,000	2.0594e-02	1.1934e-01	5.5620e-03	3.2176e-02	1.3510e-03	8.0407e-03	
	4,000	2.1275e-02	1.1932e-01	5.6542e-03	3.2192e-02	1.3851e-03	8.0365 e-03	
	8,000	2.0616e-02	1.1933e-01	5.5843e-03	3.2196e-02	1.4042e-03	8.0850 e-03	
	500	2.8224e-03	1.8145e-02	1.0084e-03	1.5877e-02	2.6316e-03	3.0236e-02	
	1,000	2.4582e-03	7.8004e-03	6.3506e-04	4.8903e-03	4.0585e-04	1.6215 e-02	
Exp-Trig	2,000	2.3809e-03	7.8190e-03	5.9138e-04	1.8863e-03	1.8116e-04	3.6817e-03	
	4,000	2.4604e-03	7.8169e-03	5.9671e-04	2.4966e-03	1.5110e-04	1.3024 e-03	
	8,000	2.3742e-03	7.8151e-03	5.8227e-04	1.8845 e - 03	1.4478e-04	5.7032e-04	

Table 1: Accuracy of unsupervised predictions for a varying grid size N, number of optimization steps \mathcal{M} , and the depth of the U-Net set to three.

Problem	d	# Params.	N = 32		N =	= 64	N = 128		
1 Toblem	a	# 1 aranis.	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u-\hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	
	2	283,713	3.2705 e06	3.8412e-05	2.8414e-06	4.1518e-05	5.4252e-04	1.9439e-03	
Bubble	3	412,225	3.2266 e - 06	3.0962 e-05	3.1159e-06	5.4674 e - 05	7.6852e-06	8.5074 e-05	
Бирріе	4	541,889	-	-	2.6481e-06	4.1081e-05	3.2062e-06	6.0881 e- 05	
	5	669,249	-	-	-	-	2.8731e-06	9.5620 e- 05	
	2	283,713	2.1274e-02	1.1932e-01	5.6116e-03	3.2174e-02	1.4157e-03	7.9525e-03	
Peak	3	412,225	2.1275e-02	1.1932e-01	5.6542e-03	3.2192 e-02	1.3851e-03	8.0365 e-03	
геак	4	541,889	-	-	5.6673e-03	3.2192e-02	1.4087e-03	8.0848 e-03	
	5	669,249	-	-	-	-	1.4042e-03	8.0918e-03	
	2	283,713	2.4689e-03	7.8191e-03	5.9526e-04	1.8814e-03	2.8053e- 02	7.5356e-02	
Erro Toia	3	412,225	2.4604 e - 03	7.8169e-03	5.9671e-04	2.4966e-03	1.5110e-04	1.3024 e-03	
Exp-Trig	4	541,889	-	-	5.9509e-04	1.8855e-03	1.5166e-04	1.8679 e - 03	
	5	669,249	-	-	-	-	1.4944e-04	1.2124 e-03	

Table 2: Accuracy of unsupervised predictions for varying grid sizes N and U-Net depths d. We set the number of optimization steps to 4,000.

2.4.2. The time step is $\tau = 0.1$. Except for the first time step, we set the number of optimization steps per time step to 250. Because we start with randomly initialized weights in the first time step, we set the number of optimization steps to 1,000 in the first time step. Table 6 shows the accuracy of our unsupervised predictions for varying grid sizes. For reference, Table

Problem -	N =	~=	N =	~ -	N = 128		
	$ u-u_h _{2,h}$	$ u-u_h _{\infty}$	$ u-u_h _{2,h}$	$ u-u_h _{\infty}$	$ u-u_h _{2,h}$	$ u-u_h _{\infty}$	
Bubble	1.3582e-16	2.8449e-16	2.4524e-16	6.8695 e-16	8.1134e-16	1.9776e-15	
Peak	1.7597e-02	1.1151e-01	5.3441e-03	3.1282 e- 02	1.3988e-03	7.9696e-03	
Exp-Trig	2.2986e-03	7.3266e-03	5.7262e-04	1.8251e-03	1.4303e-04	4.5588e-04	

Table 3: Finite difference errors for selected elliptic problems for comparison.

Problem	\mathcal{M}	N = 32		N = 64		N = 128	
		$\frac{1}{ u-\hat{u} _{2,h}}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$
	500	4.6597e-02	1.8380 e - 01	8.6890e-02	3.7069e-01	5.5651e-01	2.1105e+00
Non. Const.	1,000	3.9795e-02	1.6035 e-01	4.9640e-02	1.3376e-01	2.0766e-01	6.4809 e-01
	2,000	3.9718e-02	1.5799e-01	3.1458e-02	1.2455 e-01	8.7473e-02	2.8797e-01
Diff.	4,000	3.9713e-02	1.5794 e-01	3.1224e-02	1.2258e-01	4.9868e-02	1.1281e-01
	8,000	3.9710 e- 02	$1.5800 \mathrm{e}\text{-}01$	3.1045e-02	1.2221 e-01	3.1002e-02	9.3761 e- 02

Table 4: Accuracy of unsupervised predictions on the non-constant diffusion problem for a varying grid sizes N, number of optimization steps \mathcal{M} , and the depth of the U-Net set to three.

Problem	d	# Params.	N = 32		N =	= 64	N = 128	
	u		$\frac{1}{ u-\hat{u} _{2,h}}$	$ u - \hat{u} _{\infty}$	$ u-\hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$
	2	283,713	3.9709e-02	1.5798e-01	3.1294e-02	1.2265 e-01	3.3859e-02	1.4377e-01
Non. Const.	3	412,225	3.9713e-02	1.5794 e-01	3.1224e-02	1.2258 e-01	4.9868e-02	1.1281e-01
Diff.	4	541,889	-	-	3.1110e-02	1.2282 e-01	3.8937e-02	1.0907e-01
	5	669,249	-	-	-	-	3.6566e-02	1.0511e-01

Table 5: Accuracy of unsupervised predictions on non-constant diffusion problem for varying grid sizes N and U-Net depths d. We set the number of optimization steps to 4,000.

7 shows the accuracy of the finite difference method with backward Euler for varying grid sizes. Here, we see that Algorithm 2 achieves comparable accuracy to the finite difference method with backward Euler. However, we do not see the same convergence as with the finite difference approach for the trigonometric functions. The small number of optimization steps can explain this lack of convergence and, again, the weakness of first-order optimizers like Adam. Figures 4 through 6 show the true solution, predicted solution, and absolute difference for the first ten time steps of each problem. These figures show that our unsupervised approach produces visually accurate solutions. Finally, Table 8 displays the errors for several choices of activation functions at time t=0.5. The errors are similar for ReLU, Tanh, Swish [20], and identity activation functions. These results indicate that linear activation functions (i.e., the identity) are sufficient for learning

solutions to time-dependent problems.

Problem	t	N = 32		N = 64		N = 128	
1 Toblem	ι	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$
	0.5	7.6825 e - 04	1.5388e-03	1.0661e-03	2.1703e-03	1.1562e-03	2.4109e-03
Trig. $(n=1)$	1.0	5.2059 e-04	1.0459 e-03	7.0987e-04	1.4189e-03	7.7541e-04	1.5364 e-03
111g. $(n-1)$	2.5	6.1710 e-04	1.2354 e-03	8.6825e-04	1.7483e-03	9.4440e-04	1.9382 e-03
	5.0	1.6629 e-04	3.3075 e-04	2.5692e-04	5.2649 e-04	3.0442e-04	6.2141e-04
	0.5	6.0011e-03	1.1977e-02	1.3953e-03	2.9254e-03	3.2371e-04	9.0033e-04
Trig. $(n=4)$	1.0	3.6984 e - 03	7.3841e-03	8.5877e-04	1.8117e-03	1.9062e-04	5.3841e-04
111g. $(n = 4)$	2.5	5.4592 e-03	1.0897e-02	1.3062e-03	3.1186e-03	8.4891e-04	2.8031e-03
	5.0	1.9352 e-03	3.8744e-03	4.8639e-04	1.3043 e-03	9.8322e-04	2.5854 e - 03
	0.5	3.0932e-03	3.8826e-02	7.1997e-04	1.0341e-02	1.8902e-04	2.5298e-03
Gaussian	1.0	1.9069e-03	2.3921e-02	4.4389e-04	6.3892 e-03	1.3421e-04	1.4701 e-03
Gaussian	2.5	2.8197e-03	3.5385 e-02	6.5362e-04	9.5078e-03	2.7103e-04	2.3168e-03
	5.0	1.0004 e-03	1.2579 e-02	2.8684e-04	2.6686e-03	3.1119e-04	9.0614 e-04

Table 6: Accuracy of unsupervised predictions for varying grid sizes and time steps for parabolic problems. The depth of the network is set to three, the number of optimization iterations at each time step to 250, and the time step to 0.1.

Problem	t	N = 32		N =	= 64	N = 128	
1 Toblem	ι	$\frac{1}{ u-\hat{u} _{2,h}}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$
	0.5	7.6693e-04	1.5299e-03	1.0560e-03	2.1106e-03	1.1255e-03	2.2507e-03
This (n 1)	1.0	5.2033e-04	1.0380 e-03	7.0797e-04	1.4151e-03	7.5314e-04	1.5060 e-03
Trig. $(n=1)$	2.5	6.1623 e-04	1.2293 e-03	8.6455e-04	1.7280 e - 03	9.2434e-04	1.8484e-03
	5.0	1.5534 e-04	3.0988e-04	2.3116e-04	4.6204 e - 04	2.4942e-04	4.9877e-04
	0.5	5.9961e-03	1.1961e-02	1.3890e-03	2.7763e-03	2.8791e-04	5.7573e-04
Trig. $(n-4)$	1.0	3.7021e-03	7.3852e-03	8.5647e-04	1.7119e-03	1.7638e-04	$3.5271 \mathrm{e}\text{-}04$
Trig. $(n=4)$	2.5	5.4540 e-03	1.0880 e-02	1.2656e-03	2.5296e-03	2.6449e-04	5.2889 e-04
	5.0	2.8215 e-03	5.6286 e - 03	6.5597e-04	1.3111e-03	1.3834e-04	2.7664 e-04
	0.5	3.0933e-03	3.8814e-02	7.1809e-04	1.0412e-02	1.8424e-04	2.4368e-03
Gaussian	1.0	1.9068e-03	2.3920 e-02	4.4286e-04	6.4085 e-03	1.1490e-04	1.4920 e-03
Gaussian	2.5	2.8194e-03	3.5387e-02	6.5422e-04	9.5075 e-03	1.6586e-04	2.2389e-03
	5.0	1.4619 e-03	1.8355 e-02	3.3909e-04	4.9403 e-03	8.4882e-05	1.1716 e-03

Table 7: Finite differences with backward Euler errors on selected parabolic problems for comparison. Here, the time step is 0.1.

Algorithm 1 may benefit from transfer learning (i.e., reusing weights from previous problems) with a supervised counterpart that is trained on a large

Problem	Activation	N = 32		N =	= 64	N = 128		
1 Toblem	Activation	$\frac{1}{ u-\hat{u} _{2,h}}$	$ u - \hat{u} _{\infty}$	$ u - \hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	$ u-\hat{u} _{2,h}$	$ u - \hat{u} _{\infty}$	
	ReLU	7.6775e-04	1.5420 e - 03	1.0746e-03	2.1598e-03	1.2400e-03	2.5946e-03	
Trig. $(n=1)$	Tanh	7.6743e-04	$1.5391 \mathrm{e}\text{-}03$	1.0590e-03	2.1269 e-03	1.1478e-03	2.3299e-03	
111g. $(n = 1)$	Swish	7.6905e-04	1.5538 e-03	1.0623e-03	2.0923e-03	1.1637e-03	2.4307e-03	
	Identity	7.6825 e-04	1.5388e-03	1.0661e-03	2.1703e- 03	1.1562e-03	2.4109e-03	
	ReLU	5.9967e-03	1.1968e-02	1.5607e-03	4.1957e-03	9.2973e-04	2.7800e-03	
Trig. $(n=4)$	Tanh	5.9970 e-03	1.1968e-02	1.3906e-03	2.8266e-03	3.1624e-04	8.8251 e-04	
111g. $(n - 4)$	Swish	6.0061 e-03	1.1990 e-02	1.6013e-03	3.9862 e-03	4.7380e-04	1.4040 e - 03	
	Identity	6.0011e-03	1.1977e-02	1.3953e-03	2.9254 e-03	3.2371e-04	9.0033e-04	
	ReLU	3.0932e-03	3.8828e-02	7.3665e-04	1.0873e-02	3.8312e-04	3.3270e-03	
Gaussian	Tanh	3.0930 e-03	3.8808e-02	7.1797e-04	1.0419 e - 02	1.9917e-04	2.4221e-03	
Gaussian	Swish	3.0933e-03	3.8830 e-02	7.4640e-04	1.0099e-02	6.0038e-04	4.4619 e - 03	
	Identity	3.0932 e-03	3.8826 e - 02	7.1997e-04	$1.0341 \mathrm{e}\text{-}02$	1.8902e-04	2.5298e-03	

Table 8: Accuracy of unsupervised predictions for varying grid sizes and activation functions at time 0.5. The depth of the network is set to three, the number of optimization steps to 250, and the time step is 0.1.

labeled dataset. We randomly initialize the neural network weights in Algorithms 1 and 2 (first time step only). This random initialization may result in predictions with high errors at the beginning of training. Without optimal step sizes in our optimizer, these poor solutions may result in slow convergence towards an optimal solution. Using transfer learning with a pre-trained network that is trained in a supervised setting (i.e., with labeled training data), the initial predictions from our algorithms may be already close to optimal, resulting in faster convergence. One could view the use of these pre-trained networks as a sort of preconditioner for our neural network-based approach. Indeed, we see the benefits of transfer learning with Algorithm 2 after the first time step. Instead of reinitializing the weights for the subsequent time steps, we resuse the weights from the previous solution. That allows us to use a small number of optimization steps (i.e., $\mathcal{M}=250$) at each time for parabolic problems.

It is important to note that our methods for estimating the finite difference solutions do not explicitly construct the finite difference matrix. Instead, we use our neural networks to map the source term f to the approximate finite difference solution. In this sense, we are learning the inverse mapping of the finite difference matrix. From a numerical point of view, this provides the advantage of not having to construct or store a finite difference matrix. Instead, we implement the required stencil(s) for our problem in the loss

function and apply them appropriately. Only having to implement the stencils and not construct a finite difference matrix is also an advantage from an implementation point-of-view.

The justification for the lack of non-linear activation functions in our network is that the relationship between the source term f and the finite differences solution u_h is given by a system of linear equations of the form $Au_h = f$. In other words, we know a priori that our goal is to learn a linear relationship between f and our estimated solution. In many machine learning applications, the exact nature of the input and output relationship is unknown and assumed to be highly non-linear. Hence, including non-linear activation functions results in a network that learns a non-linear relationship between inputs and outputs. The use of the identity function as an activation puts us outside of the scope of the universal approximation theorem [21, 22, 23]. However, in our case, we do not need our neural network to be a universal approximator. The role of the network is to learn a linear relationship for a single example.

We utilize the PocketNet approach proposed by [7] in our proposed algorithms. This approach takes advantage of the similarity between the U-Net architecture and geometric multigrid methods to drastically reduce the number of parameters, while maintaining the same accuracy as conventional CNNs for medical imaging and scientific machine learning tasks [7, 24, 25]. Additionally, we replace transposed convolution with bilinear upsampling. We find that these changes save time and memory and yield the same accuracy that we see using conventional CNNs (i.e., doubling the number of channels at every depth). These results indicate that smaller neural networks (in terms of parameters) can achieve high accuracy for scientific machine learning tasks.

We see in Tables 2 and 5 that the depth of the U-Net architecture does not generally have a significant effect on our results. This indicates that, regardless of depth, the U-Net architecture is sufficiently expressive to learn a mapping from the right hand side f to an approximation of the finite difference solution u_h . However, non-U-shaped architectures like the HRNet may also produce similar or improved results [26]. Additionally, the use of residual or dense connections within the convolutions of our architecture may also be beneficial [27, 28]. Such block designs have been shown to speed up convergence to lower loss values for neural networks [29]. Finally, modifying our existing architecture by adding deep supervision could also speed up convergence to lower loss values in fewer iterations [30, 31]. The use of the weighting parameter α in (8) and (19) is necessary to enforce the given boundary conditions. We arrived at our proposed values of α via a grid search over a range of possible values. However, finding optimal values of weighting parameters like α is an open question [32]. Future work will explore this topic further.

5. Conclusions

The results presented above show the effectiveness of our proposed unsupervised approaches for estimating the finite difference solution to elliptic and parabolic problems. Unlike current neural network-based methods for estimating the solutions of PDEs, our algorithms do not rely on training data. Instead, they leverage the machinery of neural networks to estimate the finite difference solution to a given problem directly. As a result, our approach does not suffer from the generalization gap seen in PINNs or other neural network-based approaches, which rely on training data.

Our approach could benefit finite difference solvers by producing better initial guesses and/or acting as a preconditioner. With a few iterations, Algorithm 1 can produce good initial guesses for iterative solvers, thereby reducing the number of iterations required to solve the linear system resulting from (2) or (9). This same idea can also apply to time-dependent problems, but with initial guesses being produced at each time step. Additionally, because we employ identity for activation functions, our neural networks are linear. Hence, it may be possible to represent a pretrained architecture as a matrix. This resulting matrix could then be used as a preconditioner for finite difference solvers. This use of preconditioners and further testing on other kinds of PDEs, like convection-diffusion and nonlinear problems, will be the object of future work.

6. Acknowledgements

The Department of Defense supports Adrian Celaya through the National Defense Science & Engineering Graduate Fellowship Program. Keegan Kirk is supported by the Natural Sciences and Engineering Research Council of Canada through the Postdoctoral Fellowship Program (PDF-568008). David Fuentes is partially supported by R21CA249373. Beatrice Riviere is partially supported by NSF-DMS2111459. This research was partially supported by the Tumor Measurement Initiative through the MD Anderson Strategic Research Initiative Development (STRIDE), NSF-2111147, and NSF-2111459.

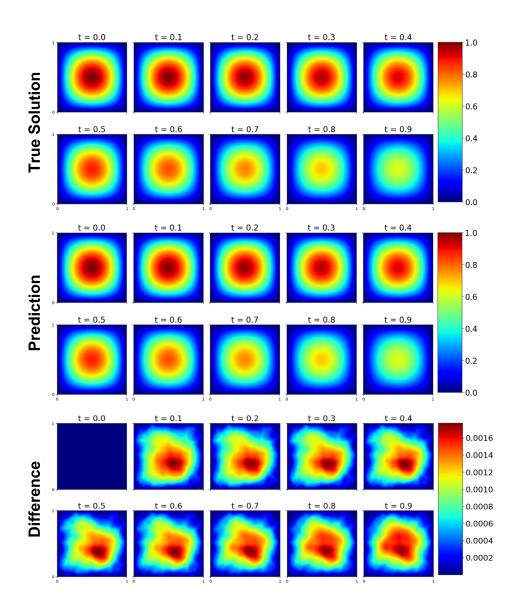


Figure 4: True solution (top), predicted solution (middle) and absolute difference (bottom) for the first ten time steps of the trigonometric problem presented in Section 2.4 with n = 1.

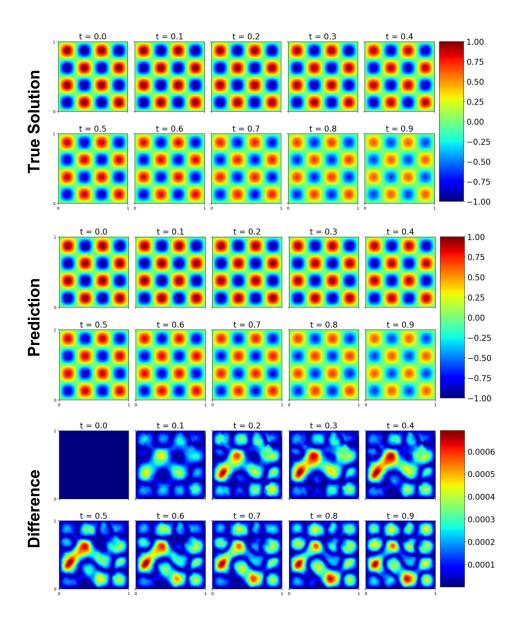


Figure 5: True solution (top), predicted solution (middle) and absolute difference (bottom) for the first ten time steps of the trigonometric problem presented in Section 2.4 with n = 4.

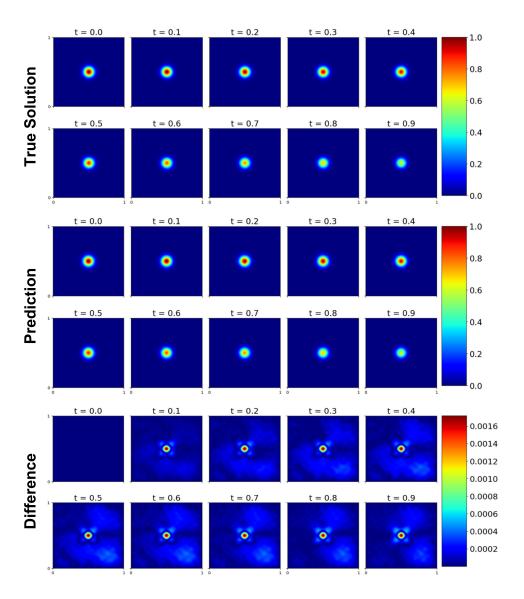


Figure 6: True solution (top), predicted solution (middle) and absolute difference (bottom) for the first ten time steps of the Gaussian problem presented in Section 2.4.

References

- [1] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," <u>Neural Networks</u>, vol. 2, pp. 359–366, 1989.
- [2] K.-I. Funahashi, "On the approximate realization of continuous mappings by neural networks," Neural Networks, vol. 2, pp. 183–192, 1989.
- [3] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," <u>Journal of Computational Physics</u>, vol. 378, pp. 686–707, 2019.
- [4] K. Sirignano, J.and Spiliopoulos, "DGM: A deep learning algorithm for solving partial differential equations," <u>Journal of Computational</u> Physics, vol. 375, pp. 1339–1364, 2018.
- [5] E. Wang and B. Yu, "The Deep Ritz method: A deep learning-based numerical algorithm for solving variational problems," <u>Communications</u> in Mathematical Sciences, vol. 6, pp. 1–12, 2018.
- [6] J. Müller and M. Zeinhofer, "Error estimate for the Deep Ritz method with boundary penalty," <u>Machine Learning Research</u>, vol. 145, pp. 1–20, 2022.
- [7] A. Celaya, J. A. Actor, R. Muthusivarajan, E. Gates, C. Chung, D. Schellingerhout, B. Riviere, and D. Fuentes, "Pocketnet: A smaller neural network for medical image analysis," <u>IEEE Transactions on Medical Imaging</u>, vol. 42, no. 4, pp. 1172–1184, 2022.
- [8] P.-H. Chiu, J. C. Wong, C. Ooi, M. H. Dao, and Y.-S. Ong, "Canpinn: A fast physics-informed neural network based on coupled-automatic-numerical differentiation method," <u>Computer Methods in Applied Mechanics and Engineering</u>, vol. 395, p. 114909, 2022.
- [9] K. L. Lim, R. Dutta, and M. Rotaru, "Phyics informed neural network using finite difference method," 2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2022.

- [10] X. Zhao, Z. Gong, Y. Zhang, W. Yao, and X. Chen, "Physics-informed convolutional neural networks for temperature field prediction of heat source without labeled data," <u>Engineering Applications of Artificial Intelligence</u>, vol. 117, p. 105516, 2023.
- [11] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in Medical Image Computing and Computer-Assisted Intervention-MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, pp. 234–241, Springer, 2015.
- [12] A. Celaya, B. Denel, Y. Sun, M. Araya-Polo, and A. Price, "Inversion of time-lapse surface gravity data for detection of 3-D CO2 plumes via deep learning," <u>IEEE Transactions on Geoscience and Remote Sensing</u>, vol. 61, pp. 1–11, 2023.
- [13] X. Yang, X. Chen, and M. M. Smith, "Deep learning inversion of gravity data for detection of CO₂ plumes in overlying aquifers," <u>Journal of Applied Geophysics</u>, vol. 196, p. 104507, 2022.
- [14] Y. Sun, B. Denel, N. Daril, L. Evano, P. Williamson, and M. Araya-Polo, "Deep learning joint inversion of seismic and electromagnetic data for salt reconstruction," <u>SEG Technical Program Expanded Abstracts</u> 2020, pp. 550–554, 2020.
- [15] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: learning dense volumetric segmentation from sparse annotation," in Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19, pp. 424–432, Springer, 2016.
- [16] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation," <u>Nature Methods 2020 18:2</u>, vol. 18, pp. 203–211, 12 2020.
- [17] B. Riviere and M. F. Wheeler, "A posteriori error estimates for a discontinuous Galerkin method applied to elliptic problems. log number: R74,"

- Computers & Mathematics with Applications, vol. 46, no. 1, pp. 141–163, 2003.
- [18] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [19] F. Chollet et al., "Keras." https://keras.io, 2015.
- [20] P. Ramachandran, B. Zoph, and Q. V. Le, "Searching for activation functions," arXiv preprint arXiv:1710.05941, 2017.
- [21] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," Neural networks, vol. 2, no. 5, pp. 359–366, 1989.
- [22] G. Cybenko, "Approximation by superpositions of a sigmoidal function," Mathematics of control, signals and systems, vol. 2, no. 4, pp. 303–314, 1989.
- [23] K.-I. Funahashi, "On the approximate realization of continuous mappings by neural networks," <u>Neural networks</u>, vol. 2, no. 3, pp. 183–192, 1989.
- [24] J. He and J. Xu, "MgNet: A unified framework of multigrid and convolutional neural network," <u>Science China Mathematics</u>, vol. 62, no. 7, pp. 1331–1354, 2019.
- [25] A. Celaya and M. Araya-Polo, "Joint inversion of time-lapse surface gravity and seismic data for monitoring of 3D CO _2 plumes via deep learning," arxiv.preprint.org/arxiv.2310.04430, 2023.
- [26] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in CVPR, 2019.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in <u>European conference on computer vision</u>, pp. 630–645, Springer, 2016.
- [28] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in <u>Proceedings of the IEEE conference on computer vision and pattern recognition</u>, pp. 4700–4708, 2017.

- [29] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, "Visualizing the loss landscape of neural nets," <u>Advances in neural information</u> processing systems, vol. 31, 2018.
- [30] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," in Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (G. Lebanon and S. V. N. Vishwanathan, eds.), vol. 38 of Proceedings of Machine Learning Research, (San Diego, California, USA), pp. 562–570, PMLR, 09–12 May 2015.
- [31] R. Li, X. Wang, G. Huang, W. Yang, K. Zhang, X. Gu, S. N. Tran, S. Garg, J. Alty, and Q. Bai, "A comprehensive review on deep supervision: Theories and applications," <u>arXiv preprint arXiv:2207.02376</u>, 2022.
- [32] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, "Scientific machine learning through physics—informed neural networks: Where we are and what's next," <u>Journal of Scientific Computing</u>, vol. 92, no. 3, p. 88, 2022.