





# BiLO: Bilevel Local Operator Learning for PDE Inverse Problems

Ray Zirui Zhang <sup>a,\*</sup>, Christopher E. Miles <sup>a</sup>, Xiaohui Xie <sup>b</sup>,  
John S. Lowengrub <sup>a,c,\*</sup>

<sup>a</sup> Department of Mathematics, University of California, Irvine, CA, USA

<sup>b</sup> Department of Computer Science, University of California, Irvine, CA, USA

<sup>c</sup> Department of Biomedical Engineering, University of California, Irvine, CA, USA

## ARTICLE INFO

### Keywords:

Bilevel optimization  
PDE inverse problems  
Neural operators  
Scientific machine learning

## ABSTRACT

We propose a new neural network based method for solving inverse problems for partial differential equations (PDEs) by formulating the PDE inverse problem as a bilevel optimization problem. At the upper level, we minimize the data loss with respect to the PDE parameters. At the lower level, we train a neural network to locally approximate the PDE solution operator in the neighborhood of a given set of PDE parameters, which enables an accurate approximation of the descent direction for the upper level optimization problem. The lower level loss function includes the least-square penalty of both the residual and its derivative with respect to the PDE parameters. We apply gradient descent simultaneously on both the upper and lower level optimization problems, leading to an effective and fast algorithm. The method, which we refer to as BiLO (Bilevel Local Operator learning), is also able to efficiently infer unknown functions in the PDEs through the introduction of an auxiliary variable. We provide a theoretical analysis that justifies our approach. Through extensive experiments over multiple PDE systems, we demonstrate that our method enforces strong PDE constraints, is robust to sparse and noisy data, and eliminates the need to balance the residual and the data loss, which is inherent to the soft PDE constraints in many existing methods.

## 1. Introduction

A fundamental task across various scientific and engineering fields is to infer the unknown parameters of a partial differential equation (PDE) from observed data. Applications include seismic imaging [1–3], electrical impedance tomography [4,5], personalized medicine [6–9], and climate modeling [10]. PDE inverse problems are commonly addressed within the frameworks of PDE-constrained optimization (PDECO) [11] or Bayesian inference [12]. In the PDE constrained optimization framework, the objective is to minimize the difference between the observed data and the PDE solution, and the PDE is enforced as a constraint using adjoint or deep learning methods. In the Bayesian inference framework, the inverse problem is formulated as a statistical inference problem, where the goal is to estimate the posterior distribution of the parameters given the data. This usually requires sampling parameter space and solving the forward PDE multiple times.

This is the first paper in a two-part series. Here in Part I, we develop a constrained optimization framework for solving PDE inverse problems using deep learning. In Part II, we extend this approach to Bayesian inference frameworks [13].

\* Corresponding authors.

E-mail addresses: [zirui.zhang@uci.edu](mailto:zirui.zhang@uci.edu) (R.Z. Zhang), [jlowengr@uci.edu](mailto:jlowengr@uci.edu) (J.S. Lowengrub).

<https://doi.org/10.1016/j.jcp.2026.114679>

Received 27 July 2025; Received in revised form 6 January 2026; Accepted 9 January 2026

Available online 11 January 2026

0021-9991/© 2026 Elsevier Inc. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

### 1.1. Related work

The **Adjoint Method** is widely used for computing the gradients of the objective function with respect to the PDE parameters using numerical PDE solvers in the PDE-constrained optimization framework. This method provides accurate gradients and strong PDE constraints. However, the adjoint method requires explicitly deriving the adjoint equation and solving both forward and adjoint equations at each iteration, which is complex and computationally expensive, especially for nonlinear or high-dimensional problems [11,14].

**Physics-Informed Neural Networks (PINNs)** have emerged as novel methods for solving inverse problems in a PDE constrained optimization framework [7,7,15–19,19–22]. PINNs represent emergent PDE solutions using neural networks and embed both the data and the PDE into the loss function through a mesh-free approach. By minimizing the total loss, PINNs effectively solve the PDE, fit the data, and infer the parameters simultaneously, showcasing integration of mathematical models with data-driven learning processes. A related approach, **Optimizing a Discrete Loss (ODIL)**, utilizes conventional numerical discretizations of the PDEs and the loss is minimized over the parameters and the PDE solutions at the grid points rather than the weights of a neural network [23,24]. However, in these methods, the PDE is enforced as a soft constraint, which requires balancing the residual and the data loss, and can lead to a trade-off between fitting the data and solving the PDE accurately.

**Neural Operators (NOs)** aim to approximate the PDE solution operator (parameter-to-solution map) and can serve as surrogate models for the forward PDE solvers [25–27]. Once these surrogates are established, they can be integrated into a Bayesian inference framework or other optimization algorithms to solve inverse problems, leveraging the speed of evaluating a neural network [28–31]. Some examples of operator learning frameworks include the Fourier Neural Operator (FNO) [32–34], Deep Operator Network (DeepONet) [35,36], In-Context Operator (ICON) [37], among others, e.g. [5,38]. However, when used to solve inverse problems, neural operators can encounter challenges when the ground truth is out of the distribution of the training dataset.

There are many other methods for PDE inverse problems using deep learning; see [39–41] for more comprehensive reviews.

### Main contributions

In this paper, we focus on solving PDE inverse problems in the PDE-constrained optimization framework using deep learning methods. The contributions are as follows:

- We formulate the PDE inverse problem as a bilevel optimization problem, where the upper level problem minimizes the data loss with respect to the PDE parameters, and the lower level problem involves training a neural network to approximate the PDE solution operator locally at given PDE parameters, enabling direct and accurate computation of the descent direction for the upper level optimization problem.
- At the lower level problem, we introduce the “residual-gradient” loss, which is the least-square penalty of the derivative of the residual with respect to the PDE parameters. We show that this loss term compels the neural network to approximate the PDE solution for a small neighborhood of the PDE parameters, thus a “local operator”.
- Extensive experiments over multiple PDE systems demonstrate that our novel formulation is both more accurate and more robust than other existing methods. It exhibits stronger PDE fidelity, robustness to sparse and noisy data, and eliminates the need to balance the residual and the data loss, a common issue in PDE-based soft constraints.
- We solve the bilevel optimization problem using gradient descent simultaneously on both the upper and lower level optimization problems, leading to an effective and fast algorithm.
- We extend our method to infer unknown functions that are also parameterized by neural networks through an auxiliary variable. This bypasses the need to learn a high-dimensional local operator.
- We rigorously analyze the difference between the exact gradient of the upper-level loss and the approximate gradient that results from inexact minimization of the lower level problem. We establish an error bound for the difference between the gradients that provides a theoretical foundation for our approach.

Our approach combines elements of PINNs, operator learning, and the adjoint method. Our method is related to PINNs in that both use neural networks to represent the solution to the PDE, both use automatic differentiation to compute the PDE residual, and both aim to solve the PDE and infer the parameters simultaneously. However, in PINNs, the PDE-constraint is enforced as a regularization term (or soft constraint), leading to a trade-off between fitting the data and solving the PDE accurately. Compared with operator learning, which solves the PDE for a wide range of parameters and requires a large amount of synthetic data for training, our method only learns the operator local to the PDE parameters at each step of the optimization process and does not require a synthetic dataset for training. Similar to the adjoint method, we aim to approximate the descent direction for the PDE parameters with respect to the data loss, but we do not require deriving and solving the adjoint equation.

The outline of this paper, Part I of our study on solving PDE inverse problems using deep learning methods, is as follows. In [Section 2](#), we present and analyze the BiLO method and compare the formulation with other approaches (PINNs, NOs). In [Section 3](#), we apply the BiLO method to a collection of PDE inverse problems (elliptic, parabolic, hyperbolic) and compare the results to PINNs and NO. In [Section 4](#), we summarize our results. In the Appendices, we present details of the numerical analysis, numerical implementations, computational cost, sensitivity analyses (to hyperparameters), and additional numerical results.

## 2. Method

### 2.1. Bilevel local operator learning (BiLO) for PDE inverse problems

In this section, we present BiLO for solving PDE-constrained optimization problems where we aim to infer the PDE parameters from observed data. Let  $u : \Omega \rightarrow \mathbb{R}$  be a function defined over a domain  $\Omega \subset \mathbb{R}^d$ , and  $\hat{u}$  be the observed data, which might be noisy. For time-dependent problems, we treat time  $t$  as a special component of  $\mathbf{x}$ , and  $\Omega$  includes the temporal domain. We consider the following PDE-constrained optimization problem:

$$\begin{aligned} \min_{\theta} \quad & \|u - \hat{u}\|_2^2 \\ \text{s.t.} \quad & \mathcal{F}(D^k u(\mathbf{x}), \dots, Du(\mathbf{x}), u(\mathbf{x}), \theta) = \mathbf{0} \end{aligned} \quad (1)$$

where  $D^k$  is the  $k$ -th order derivative operator,  $\theta$  represents the PDE parameters, and  $\mathcal{F}$  denotes equality constraints that include the PDE, the boundary and initial conditions, and interface conditions, if needed, such as in elliptic interface problems.

For notational simplicity, we define the residual function of the PDE constraint to be

$$r(\mathbf{x}, \theta) := \mathcal{F}(D^k u(\mathbf{x}, \theta), \dots, Du(\mathbf{x}, \theta), u(\mathbf{x}, \theta), \theta) \quad (2)$$

The dependence of  $r$  on  $u$  is implicit. The local operator  $u$  is characterized by the following two conditions:

- **Condition 1:**  $r(\mathbf{x}, \theta) = \mathbf{0}$ ,
- **Condition 2:**  $\nabla_{\theta} r(\mathbf{x}, \theta) = \mathbf{0}$ .

Condition 1 means that  $u$  should solve the PDE at  $\theta$ . Condition 2 suggests that small variation of  $\theta$  should lead to small variation of the residual. If the conditions are satisfied, then the derivative of the data loss with respect to  $\theta$  will approximate the descent direction, and we can find the optimal  $\theta$  by minimizing the data loss with respect to  $\theta$  using a gradient descent algorithm.

To further clarify these conditions, we view  $\mathcal{F}$  as a differential operator acting on the state  $u$  and parameters  $\theta$ ; the PDE is then expressed by the operator equation  $\mathcal{F}[u, \theta] = \mathbf{0}$ . We defer details of underlying function space or regularity assumptions to the appendix when theoretical results are presented. The ‘‘global solution operator’’, or the parameter-to-solution map,  $\theta \rightarrow u(\cdot, \theta)$ , satisfies  $\mathcal{F}[u(\cdot, \theta), \theta] = 0$  for all  $\theta$ . Finding the global solution operator is generally challenging, and is unnecessary if we only want to find the descent direction to update the current  $\theta$  in an optimization algorithm. If we consider a local perturbation  $\theta \rightarrow \theta + \delta\theta$ , a formal Taylor expansion yields

$$\mathcal{F}[u(\mathbf{x}, \theta + \delta\theta), \theta + \delta\theta] = \mathcal{F}[u(\mathbf{x}, \theta), \theta] + \left( \frac{\delta\mathcal{F}}{\delta u} \frac{\partial u}{\partial \theta} + \frac{\partial \mathcal{F}}{\partial \theta} \right) \delta\theta + O(\delta\theta^2), \quad (3)$$

where  $\delta\mathcal{F}/\delta u$  denotes the Fréchet derivative of  $\mathcal{F}$  with respect to  $u$ . Then condition 1 corresponds to setting the zero-th order term in the expansion to zero:

$$r(\mathbf{x}, \theta) = \mathcal{F}[u(\mathbf{x}, \theta), \theta] = 0,$$

while condition 2 corresponds to setting the first-order term to zero:

$$\nabla_{\theta} r(\mathbf{x}, \theta) = \frac{d}{d\theta} \mathcal{F}[u(\mathbf{x}, \theta), \theta] = \frac{\delta\mathcal{F}}{\delta u} \frac{\partial u}{\partial \theta} + \frac{\partial \mathcal{F}}{\partial \theta} = 0.$$

Thus, a local operator solves the PDE at  $\theta$  and is a second-order accurate approximation in  $\delta\theta$  to the PDE under small perturbations of  $\theta$ .

We approximate the operator locally at  $\theta$  using a neural network  $u(\mathbf{x}, \theta; W)$ , where  $W$  denotes the weights of the network. The objective function (1) leads to the data loss:

$$\mathcal{L}_{\text{data}}(\theta, W) = \frac{1}{|\mathcal{T}_{\text{data}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{data}}} |u(\mathbf{x}, \theta; W) - \hat{u}(\mathbf{x})|^2, \quad (4)$$

where  $\mathcal{T}_{\text{data}}$  denotes the set of measurement locations for observed data. The residual loss is given by

$$\mathcal{L}_{\text{res}}(\theta, W) := \frac{1}{|\mathcal{T}_{\text{res}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{res}}} |r(\mathbf{x}, \theta; W)|^2. \quad (5)$$

where  $\mathcal{T}_{\text{res}}$  is the set of collocation points for evaluating the residual. We introduce the following ‘‘residual-gradient loss’’

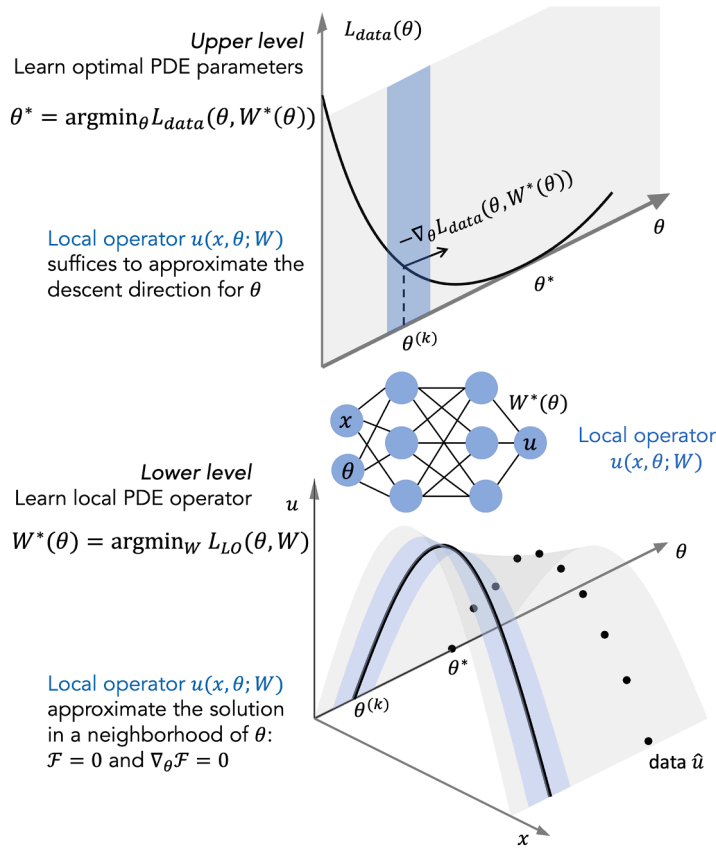
$$\mathcal{L}_{\text{rgrad}}(\theta, W) = \frac{1}{|\mathcal{T}_{\text{res}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{res}}} |\nabla_{\theta} r(\mathbf{x}, \theta)|^2. \quad (6)$$

We define the ‘‘local operator loss’’ as the sum of the residual loss and the residual-gradient loss with weight  $w_{\text{rgrad}}$ :

$$\mathcal{L}_{\text{LO}}(\theta, W) = \mathcal{L}_{\text{res}}(\theta, W) + w_{\text{rgrad}} \mathcal{L}_{\text{rgrad}}(\theta, W) \quad (7)$$

Finally, we propose to solve the following bilevel optimization problem:

$$\begin{cases} \theta^* = \arg \min_{\theta} \mathcal{L}_{\text{data}}(\theta, W^*(\theta)) \\ W^*(\theta) = \arg \min_W \mathcal{L}_{\text{LO}}(\theta, W) \end{cases} \quad (8)$$



**Fig. 1.** A schematic of BiLO. Top: The full PDE operator  $u(x, \theta)$  (gray) solves the PDE for all  $\theta$ , while the local operator (blue) approximates the solution in a small neighborhood of  $\theta$ . The local operator is sufficient for approximating the descent direction of the data loss. The figure uses the model boundary value problem  $-\theta u_{xx} = \sin(\pi x)$  with Dirichlet boundary condition. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In the upper level problem, we find the optimal PDE parameters  $\theta$  by minimizing the data loss with respect to  $\theta$ . In the lower level problem, we train a network to approximate the local operator  $u(x, \theta; W)$  by minimizing the local operator loss with respect to the weights of the neural network.

Boundary and initial conditions can be incorporated as additional loss terms, evaluated at respective domain boundaries. In some cases, these conditions can be enforced exactly by transforming the network output or specialized network architecture [42–44]. For example, on the domain  $\Omega = [0, 1]$ , one can impose the Dirichlet boundary condition  $u(0) = u(1) = 0$  by multiplying the output of the neural network by  $x(1 - x)$ . For simplicity of discussion, we focus on the residual loss and the data loss, and assume that the boundary conditions are enforced exactly.

Fig. 1 illustrates the idea of the BiLO framework. We consider functions of the form  $u(x, \theta)$ , where  $x \in \Omega$  and  $\theta \in \Theta$ , where  $\Theta$  is an admissible set of PDE parameters. We call such a function the “PDE solution operator” (hereafter referred to as the “operator”), if it solves the PDE for all  $\theta$ , that is, the map  $\theta \mapsto u(\cdot, \theta)$  is the parameter-to-solution map. Such an operator exists if the PDE solution is unique and continuous with respect to the parameters  $\theta$ . An example of such an operator  $u(x, \theta)$  is shown as a gray surface in Fig. 1, which solves the PDE  $-\theta u_{xx} = \sin(\pi x)$  with Dirichlet boundary condition for all  $\theta > 0$ . If such an operator is available, we can solve the optimization problem easily by minimizing the objective function using a gradient descent algorithm. However, finding the full operator  $u(x, \theta)$  is challenging and unnecessary. Since we are only interested in the descent direction to update  $\theta$ , a local approximation of the solution operator suffices (blue surface in Fig. 1), that is, the operator should approximate the PDE solution for a small neighborhood of a particular value of  $\theta$ .

### 2.2. Pre-training and fine-tuning

In this work, we assume access to an initial guess of the PDE parameters,  $\theta_0$ , as required by most gradient-based methods. In BiLO, the lower-level problem must be solved to compute the descent direction for the upper-level optimization. To this end, we introduce a pre-training phase in which we fix  $\theta = \theta_0$  and train the neural network to approximate the local solution operator at  $\theta_0$ . Since  $\theta$  is

fixed, this stage is not a bilevel optimization problem, as only the lower-level problem is solved. After pre-training, we solve the full bilevel optimization problem to infer the PDE parameters  $\theta$ , and we refer to this as the fine-tuning phase.

When available, a numerical solution of the PDE at  $\theta_0$ , denoted by  $u_0(\mathbf{x})$  and computed using a method such as finite difference (FDM) or finite element (FEM), can be used to accelerate pre-training. We define a pre-training data loss  $\mathcal{L}_{u_0}$  as the mean squared error between the numerical solution  $u_0$  and the neural network output at  $\theta_0$ :

$$\mathcal{L}_{u_0}(W) = \frac{1}{|\mathcal{T}_{\text{res}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{res}}} |u(\mathbf{x}, \theta_0; W) - u_0(\mathbf{x})|^2. \quad (9)$$

The pre-training objective is then:

$$\min_W \mathcal{L}_{\text{LO}}(\theta_0, W) + \mathcal{L}_{u_0}(W). \quad (10)$$

The use of  $\mathcal{L}_{u_0}$  is optional, but can significantly speed up the pre-training process. This is computationally inexpensive, as we only need one numerical solution. This strategy was been shown to be effective in [7], and is conceptually related to curriculum learning [45], where the network first learns to approximate a simpler solution.

### 2.3. Network architecture

The network architecture follows the standard multilayer perceptron (MLP) design. For the scalar parameter case, the embedding layer maps the inputs  $\mathbf{x} \in \mathbb{R}^d$  and the unknown PDE parameters  $\theta \in \mathbb{R}^m$  to a high-dimensional vector  $\mathbf{h}^{(0)} \in \mathbb{R}^p$ , using an affine transformation followed by a non-linear activation function  $\sigma$  (tanh in this work):

$$\mathbf{h}^{(0)} = \sigma(W_{\mathbf{x}}^{(0)} \mathbf{x} + W_{\theta}^{(0)} \theta + \mathbf{b}), \quad (11)$$

where  $W_{\mathbf{x}}^{(0)} \in \mathbb{R}^{p \times d}$  is the embedding matrix for  $\mathbf{x}$ ,  $W_{\theta}^{(0)} \in \mathbb{R}^{p \times m}$  is the embedding matrix for  $\theta$ , and  $\mathbf{b} \in \mathbb{R}^p$  is the bias vector. This transformation serves as the standard input embedding, mapping both coordinates and PDE parameters into a shared latent space and enables the network  $u(\mathbf{x}, \theta; W)$  to model the nonlinear effect of  $\theta$  on  $u$  at each point  $\mathbf{x}$ . The embedding vector  $\mathbf{h}^{(0)}$  is then passed through a series of fully connected layers with activation functions:  $\mathbf{h}^{(l)} = \sigma(W^{(l)} \mathbf{h}^{(l-1)} + b^{(l)})$  for  $l = 1, \dots, L-1$ , where  $L$  is the total number of layers,  $W^{(l)} \in \mathbb{R}^{p \times p}$  and  $b^{(l)} \in \mathbb{R}^p$  are the weights and bias of the  $l$ -th layer. The final output is given by  $\mathbf{y} = W^{(L)} \mathbf{h}^{(L-1)} + b^{(L)}$ , where  $W^{(L)} \in \mathbb{R}^{1 \times p}$  and  $b^{(L)} \in \mathbb{R}$ . Collectively, the network is denoted as  $\mathcal{N}(\mathbf{x}, \theta; W)$ , where  $W = \{W^{(l)}, b^{(l)}\}_{l=0}^L$  denotes all the trainable weights of the neural network. In some cases, a final transformation is applied to the output of the neural network  $u(\mathbf{x}; W) = \tau(\mathcal{N}(\mathbf{x}, \theta; W), \mathbf{x})$ , to enforce boundary conditions [42–44]. For example, if the PDE is defined on a unit interval  $[0, 1]$  with Dirichlet boundary conditions  $u(0) = u(1) = 0$ , the BiLO solution can be represented as  $u(x, \theta; W) = \mathcal{N}(x, \theta; W)(1-x)x$ . This MLP architecture is presented as a basic example, and more advanced architectures can also be incorporated, including residual connections [46], random Fourier features [47], and other modified MLP designs [48].

### 2.4. Inferring an unknown function

We can extend our method to learn an unknown function  $f(\mathbf{x})$  in the PDE, such as a variable diffusion coefficient in the Poisson equation or an initial condition in the heat equation. In these cases, the following PDE constrained optimization problem is solved:

$$\begin{aligned} \min_f \quad & \|u - \hat{u}\|^2 + w_{\text{reg}} \|\nabla f\|^2 \\ \text{s.t.} \quad & \mathcal{F}(D^k u(\mathbf{x}), \dots, Du(\mathbf{x}), u(\mathbf{x}), f(\mathbf{x})) = \mathbf{0} \end{aligned} \quad (12)$$

where the PDE depends on an unknown function  $f$ . Given that these problems are ill-posed, regularization of the unknown function is often necessary. A typical choice is the L2-norm of the gradient of the unknown function, which penalizes non-smooth functions. The choice of an appropriate regularization form is important and problem-dependent. This paper assumes such choices are predetermined and are not aspects of the method under direct consideration.

Suppose  $f$  is parameterized by a neural network  $f(\mathbf{x}; V)$  with weights  $V$ . A straightforward extension from the scalar parameter case is to learn the local operator of the form  $u(\mathbf{x}, V)$ . However, this would be computationally expensive, as the weights  $V$  can be very high dimensional. We introduce an auxiliary variable  $z = f(\mathbf{x})$ , and find a local operator  $u(\mathbf{x}, z)$  such that  $u(\mathbf{x}, f(\mathbf{x}))$  solves the PDE locally at  $f$ . We define the following function  $a$ , which is the augmented residual function with an auxiliary variable  $z$ :

$$a(\mathbf{x}, z) := \mathcal{F}(D^k u(\mathbf{x}, z), \dots, Du(\mathbf{x}, z), u(\mathbf{x}, z), z) \quad (13)$$

where  $D = \nabla_{\mathbf{x}} + \nabla_{\mathbf{x}} z \nabla_z$  is the total derivative operator. If  $u$  is a local solution operator at  $f$ , then we should have: (1)  $a(\mathbf{x}, f(\mathbf{x})) = 0$  so that the function  $u(\mathbf{x}, f(\mathbf{x}))$  have zero residual, and (2)  $\nabla_z a(\mathbf{x}, f(\mathbf{x})) = 0$  so that small variations of  $f$  should lead to small variations in the residual, as in the scalar parameter case in Eq. (6). These two conditions translate into the corresponding residual loss and residual-gradient loss, similar to Eqs. (5) and (6). The residual loss is given by

$$\mathcal{L}_{\text{res}}(W, V) := \frac{1}{|\mathcal{T}_{\text{res}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{res}}} |a(\mathbf{x}, f(\mathbf{x}; V); W)|^2. \quad (14)$$

and the residual-gradient loss is

$$\mathcal{L}_{\text{rgrad}}(W, V) = \frac{1}{|\mathcal{T}_{\text{res}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{res}}} |\nabla_{\mathbf{z}} a(\mathbf{x}, f(\mathbf{x}; V); W)|^2 \quad (15)$$

The data loss is similar to the parameter inference case in Eq. (4) and depends on both  $V$  and  $W$ . We also need the regularization loss, evaluated on  $\mathcal{T}_{\text{reg}}$ :

$$\mathcal{L}_{\text{reg}}(V) = \frac{1}{|\mathcal{T}_{\text{reg}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{reg}}} |\nabla_{\mathbf{x}} f(\mathbf{x}; V)|^2. \quad (16)$$

Finally, we solve the following bilevel optimization problem:

$$\begin{cases} V^* = \arg \min_V \mathcal{L}_{\text{data}}(W^*(V), V) + w_{\text{reg}} \mathcal{L}_{\text{reg}}(V) \\ W^*(V) = \arg \min_W \mathcal{L}_{\text{LO}}(W, V) \end{cases} \quad (17)$$

where  $\mathcal{L}_{\text{LO}} = \mathcal{L}_{\text{res}} + w_{\text{rgrad}} \mathcal{L}_{\text{rgrad}}$ . At the upper level, we minimize the data loss and the regularization loss with respect to the weights  $V$  of the unknown function, and at the lower level, we minimize the local operator loss with respect to the weights  $W$  of the local operator. The pre-training stage is similar to the parameter inference case. Given an initial guess of the unknown function  $f_0$ , and its corresponding numerical solution  $u_0$ , we can train the network  $f_V$  to approximate  $f_0$  by minimizing the MSE between  $f_V$  and  $f_0$ , and train the network  $u_W$  to be the local operator at  $f_0$  by minimizing the local operator loss and the MSE between  $u_W$  and  $u_0$ .

## 2.5. Algorithm and theoretical analysis

*Algorithm.* Solving bilevel optimization problems remains an active area of research [49–55]. In our case, the upper level problem is usually non-convex, and the lower level problem has a challenging loss landscape [45,56]. However, the lower level problem does not need to be solved to optimality at each iteration because the primary goal is to approximate the descent direction for the upper level problem. We propose to apply gradient descent to the upper and lower level optimization problems simultaneously. In Algorithm 1, we describe the optimization algorithm for inferring scalar parameters in the BiLO framework. The algorithm for inferring unknown functions is similar. We write the algorithm as simple gradient descent for notational simplicity, while in practice we use ADAM [57] or its variants. We can have two different learning rates for the two groups of variables  $W$  and  $\theta$ , denoted as  $\alpha_W$  and  $\alpha_\theta$ , respectively.

---

### Algorithm 1 Bi-level local operator for inferring scalar PDE parameters.

---

- 1: **Input:** Collections of collocation points  $\mathcal{T}_{\text{res}}$  and  $\mathcal{T}_{\text{data}}$ , initial guess of the PDE parameters  $\theta_0$
- 2: **Pre-train:** Solve the lower level problem at fixed  $\theta_0$ :

$$\min_W \mathcal{L}_{\text{LO}}(\theta_0, W) \quad (18)$$

- 3: **Fine-Tune:** Simultaneous gradient descent at the upper and lower levels in system (8).

$$\begin{cases} \theta^{(k+1)} = \theta^{(k)} - \alpha_\theta \nabla_\theta \mathcal{L}_{\text{data}}(\theta^{(k)}, W^{(k)}) \\ W^{(k+1)} = W^{(k)} - \alpha_W \nabla_W \mathcal{L}_{\text{LO}}(\theta^{(k)}, W^{(k)}) \end{cases} \quad (19)$$


---

*Theoretical analysis.* We next provide a theoretical characterization of our bilevel optimization method by analyzing the difference between the exact and approximate gradients of the upper-level loss. The approximate gradient arises from inexact minimization of the lower level problem. The exact gradient, or hypergradient, accounts for the total dependence of the system on the hyperparameter  $\theta$ , including the sensitivity of the ideal weights  $W^*(\theta)$ . In contrast, our simultaneous training algorithm uses an approximate gradient, which efficiently computes only the partial derivative with respect to  $\theta$  at the current weights  $W$ . Our analysis establishes two key results:

- **Consistency:** We demonstrate that under ideal conditions (i.e., the lower-level problem is solved exactly), the approximate gradient is identical to the true gradient. A precise statement of the corresponding theorem (Theorem 1) is given in Appendix A.1 along with its proof.
- **Approximation Error:** More practically, we establish an error bound. Theorem 2 (stated precisely and proved in Appendix A.2) guarantees that when the lower-level problem is solved to a tolerance  $\epsilon$ , the error between the approximate and true gradients is also bounded by  $\epsilon$ , assuming the PDE is well-behaved.

These theorems provide a solid theoretical foundation for our approach. Furthermore, our numerical experiments demonstrate the method's effectiveness under even less restrictive conditions than required by the theory.

## 2.6. Difference between BiLO, PINN, and NO

We next clarify the differences between BiLO, PINNs, and neural operators.

**Neural Operators** can serve as surrogate models for PDE solution operators, and can be used in algorithms that require solving the forward PDE multiple times, such as Bayesian inference, derivative-free optimization [30,58], and gradient-based optimization algorithms [28,30,59]. However, if the objective is to estimate parameters from limited data, the considerable initial cost for data generation and network training might seem excessive. The accuracy of specific PDE solutions depends on the accuracy of the neural operator, which may decrease if the true PDE parameters fall outside the training data's distribution [60]. This issue can be mitigated by instance-wise fine-tuning using the residual loss [33,36], though it introduces an additional trade-off: fine-tuning for one parameter set could reduce the operator's overall accuracy for other parameters and an "anchor loss" is thus required to maintain generalization [33]. Thus, in the context of finding the best estimate of the parameters given the data in a PDE-constrained optimization framework, we mainly compare BiLO with PINNs.

Within the **PINN** framework, the solution of the PDE is represented by a deep neural network  $u(\mathbf{x}; W)$  [15,16,61]. Notice that the PDE parameters  $\theta$  are not input to the neural network. Therefore, the data loss does not depend on the PDE parameters  $\theta$  directly, and we write the data loss as  $\mathcal{L}_{\text{data}}^{\text{PINN}}$ .

$$\mathcal{L}_{\text{data}}^{\text{PINN}}(W) = \sum_{\mathbf{x} \in \mathcal{T}_{\text{data}}} (u(\mathbf{x}; W) - \hat{u}(\mathbf{x}))^2$$

and enforce the PDE constraints by minimizing the residual loss.

$$\mathcal{L}_{\text{res}}^{\text{PINN}}(W, \theta) = \sum_{\mathbf{x} \in \mathcal{T}_{\text{res}}} \mathcal{F}(D^k u(\mathbf{x}; W), \dots, u(\mathbf{x}; W), \theta)^2.$$

Solving an inverse problem using PINN involves minimizing an unconstrained optimization problem, where the objective function is the weighted sum of the residual loss and the data loss

$$\min_{W, \theta} \mathcal{L}_{\text{res}}^{\text{PINN}}(W, \theta) + w_{\text{data}} \mathcal{L}_{\text{data}}^{\text{PINN}}(W) \quad (20)$$

where  $w_{\text{data}}$  is the weight of the data loss. For simplicity of discussion, we assume the weight of the residual loss is always 1. In PINN, the PDE is enforced as a soft constraint or as a regularization term for fitting the data. The relationship between the PDE parameter and the data loss is indirect. If we consider the gradient descent dynamics for training of the PINN, we have

$$\begin{cases} \theta^{(k+1)} = \theta^{(k)} - \alpha_{\theta} \nabla_{\theta} \mathcal{L}_{\text{res}}^{\text{PINN}}(W^{(k)}, \theta^{(k)}) \\ W^{(k+1)} = W^{(k)} - \alpha_W \nabla_W (\mathcal{L}_{\text{res}}^{\text{PINN}}(W^{(k)}, \theta^{(k)}) + w_{\text{data}} \mathcal{L}_{\text{data}}^{\text{PINN}}(W^{(k)})) \end{cases} \quad (21)$$


The descent directions of the PDE parameters do not directly depend on the data loss  $\mathcal{L}_{\text{data}}^{\text{PINN}}$ .

**Challenges for PINNs.** Solving PDE inverse problems using PINNs can encounter challenges stemming from the soft PDE constraint in Eq. (20), especially when the data is sparse and noisy, or when the PDE model does not fully explain the data [7]. The soft PDE constraint can result in a trade-off between fitting the data and solving the PDE accurately. In addition, since the PDE parameters are updated in the descent direction of the residual loss, they can be biased toward parameters corresponding to very smooth solutions. Nevertheless, it is important to recognize that PINNs can indeed be effective for PDE inverse problems, if the weights are chosen properly [18,19,21].

There are many techniques to improve the performance of PINNs, such as adaptive sampling and weighting of collocation points [61–64], new architectures [48,65–67], new optimization algorithms [45,68], new loss functions [69–71], adaptive weighting of loss terms [48,72–74]. However, these techniques do not fundamentally change the soft PDE-constraints in the PINN framework. In our work, we propose a different optimization problem that does not involve a trade-off between the residual loss and the data loss, and our method can be used in conjunction with many of these techniques to improve the performance. Therefore, in the following numerical experiments, we do not use any of these techniques, and we focus on comparing the two different optimization formulations (BiLO and the soft PDE-constraints).

We note that while our lower-level problem also involve a weighted sum of two loss terms (the residual loss and the residual-gradient loss), the residual-gradient term does not compete with the residual loss in the way that the data loss does in PINNs. The minimizer of the data term is an interpolant of the observations, whereas the minimizer of the residual term is a function that satisfies the PDE; these optima typically do not coincide, creating an inherent tension between fitting data and solving the PDE in PINNs. In contrast, a local operator is a function characterized by vanishing residual and residual-gradient. Thus, minimizing these two terms does not pull the neural network in conflicting directions. In the ideal case with sufficient network capacity, both terms can be sufficiently minimized simultaneously. Therefore, the parameter  $w_{\text{rgrad}}$  conditions the optimization landscape rather than altering the location of the minimizer. We illustrate its effect in the Appendix.

The challenge of balancing trade-offs also motivated the Bilevel PINN (BPN) method developed in [75], which applies a bilevel optimization framework to PDE inverse problems by representing the PDE solution with a neural network, using the residual loss for the lower-level problem, and approximating the upper-level hypergradient with Broyden's method. In contrast, our approach incorporates the PDE parameter as part of the network input, with the lower-level problem focused on approximating the local operator, allowing more direct computation of the upper-level descent direction. We compare BPN and BiLO in Appendix B.



**Fig. 2.** Evaluating the local operator  $u(x, D_0 + \delta D, \rho_0 + \delta \rho)$  at  $t = 1$  after pretraining at  $D_0$  and  $\rho_0$  provides an approximation to the corresponding FDM solutions.

### 3. Numerical experiments

We evaluate the effectiveness of our method on a diverse set of PDE inverse problems, encompassing elliptic, parabolic, and hyperbolic systems. Our test cases include challenging scenarios such as nonlinear dynamics, singular forcing terms, and inverse problems in Glioblastoma using patient data. In our experiments, we denote the neural network solution (obtained via BiLO, PINN, or NO) as  $u_{\text{NN}}$ , and the numerical solution computed with the estimated parameters using the Finite Difference Method (FDM) as  $u_{\text{FDM}}$ , which is computed to high accuracy and serves as the exact solution to measure the accuracy of the neural network solution. To add synthetic noise, we consider Gaussian noise with mean 0 and standard deviation  $\sigma$ .

For each numerical experiment, we solve the optimization problem 5 times with different random seeds, which affect both the initialization of the neural network and the noise (if applicable). Although each realization of the noise may yield a different globally optimal PDE parameter  $\theta^*$ , the average of the estimated parameters across multiple runs should still be close to the ground truth parameter  $\theta_{\text{GT}}$ . Therefore, we report the mean and standard deviation of the error between the estimated quantities and the ground truth. We empirically determined  $w_{\text{grad}} = 0.1$  and  $\alpha_w = \alpha_\theta = 0.001$  to be effective across our numerical experiments.

#### 3.1. Fisher-KPP equation

We aim to infer the unknown parameters  $D$  and  $\rho$  in the following nonlinear reaction-diffusion equation (Fisher-KPP equation) as in [76]:

$$\begin{cases} u_t(x, t) = 0.01 D u_{xx}(x, t) + \rho u(1 - u) \\ u(x, 0) = \frac{1}{2} \sin(\pi x)^2, \\ u(0, t) = u(1, t) = 0. \end{cases} \quad (22)$$

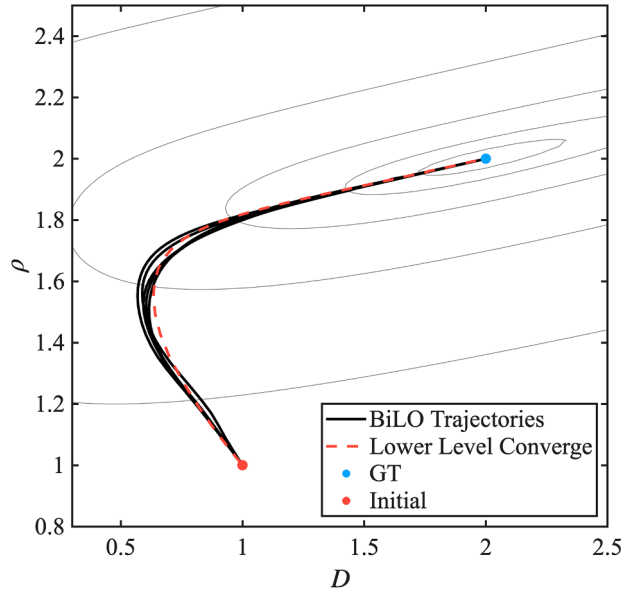
The initial guesses of the PDE parameters are  $D_0 = 1$  and  $\rho_0 = 1$ , and the ground truth parameters are  $D_{\text{GT}} = 2$  and  $\rho_{\text{GT}} = 2$ . Thus the relative error of the initial guesses are 50%. This equation has been used to model various biological phenomena, such as the growth of tumors [77,78] or the spreading of misfolded proteins [8,79,80].

##### 3.1.1. Visualizing BiLO

In Fig. 2, we visualize the local operator  $u(x, D, \rho; W)$  after pre-training with  $D_0 = 1$  and  $\rho_0 = 1$ . We consider the variation  $(\delta D, \delta \rho) = (0.5, 0)$  and  $(0, 0.2)$  and evaluate the neural network at  $u(x, D_0 + \delta D, \rho_0 + \delta \rho)$ . The FDM solutions of the PDE corresponding to the neighboring parameters are also shown. The neural network approximates the solution corresponding to the neighboring parameters well, and the neural network is able to learn the local operator of the PDE. Note that this figure is intended to show that the neural network varies consistently with the PDE constraint. In practice, we do not evaluate the network at nearby parameter values. Rather, Conditions 1 and 2 provide the correct gradient direction for updating the parameters.

We show the trajectories of the parameters  $D$  and  $\rho$  during the fine-tuning process in Fig. 3 without noise. The contours are the data loss in log scale using the FDM solution for each parameter pair  $(D, \rho)$ . They serve as a reference loss landscape for visualization. The dashed red line shows the trajectory obtained by solving the lower level problem to a small tolerance ( $10^{-4}$ ) before each update of the PDE parameters at the upper level. This trajectory can be considered as the gradient descent path on the reference loss landscape. Each BiLO trajectory (black line) corresponds to a different random initialization of the neural network and is obtained by the





**Fig. 3.** Trajectories of the parameters  $D$  and  $\rho$  during fine-tuning. The contours show the data-misfit loss in log scale computed using a finite-difference solution and serve as a reference loss landscape. The dashed line denotes the trajectory obtained by solving the lower-level problem to a small tolerance before updating the PDE parameters, corresponding to a reference gradient descent path. The black lines show BiLO trajectories from simultaneous gradient descent, which follow the reference path.

simultaneous gradient descent. Because the lower-level problem is not solved to convergence at each iteration, the resulting descent directions at each step contain small fluctuations; nevertheless, the BiLO trajectories closely follow the reference path.

### 3.1.2. Estimation under noisy data

We solve the inverse problem using different methods with different noise levels. Due to the presence of noise, the minimizer of the PDECO is no longer the ground truth parameters that generate the data. We evaluate three approaches: (1) BiLO, (2) PINNs with different  $w_{\text{data}}$ , and (3) Neural Operators with varying pretraining datasets. We show the mean and standard deviation of various metrics: the relative error of the inferred parameters  $D$  and  $\rho$  with respect to the GT, the relative L2 error of  $u_{\text{NN}}$  compared to  $u_{\text{FDM}}$ .

*Comparison with PINN.* Fig. 4 summarizes the performance of BiLO and PINN across varying noise levels and data loss weights. As expected, increasing the noise level generally leads to higher errors in the inferred parameters for all methods. The results for PINN are highly sensitive to the choice of the data loss weight  $w_{\text{data}}$  and depend non-monotonically on  $w_{\text{data}}$  when the noise level is small. When the noise level is high, the accuracy deteriorates significantly and smaller  $w_{\text{data}}$  yield better, but still limited accuracy. The weight  $w_{\text{data}} = 10^3$  and noise  $\sigma = 0.1$  result in unphysical solutions, such as negative values for  $D$ , and are therefore omitted from the plot. Across all noise levels, BiLO consistently outperforms PINN in terms of parameter accuracy, with especially pronounced improvements for low or zero noise-achieving up to an order of magnitude lower error (note the logarithmic scale on the y-axis). In contrast, BiLO demonstrates robust performance in both parameter inference and solution accuracy, maintaining low error even as the noise level increases. For PINN, the accuracy of the solution decreases as noise level increases. The deterioration becomes more pronounced as  $w_{\text{data}}$  increases. For the accuracy of the solution, BiLO is robust against noise levels.

*Comparison with neural operators.* For the NO, we use the DeepONet architecture [35] as an example, which is shown to have comparable performance with FNO [32,81]. In this experiment, we first train the NO using numerical PDE solutions corresponding to different values of  $D$  and  $\rho$ , and then we use the NO as a surrogate and use gradient-based optimization to infer the parameters of the PDE. We show that the quality of the inferred parameters depends on the training data used to train the NO.

We use the notation  $a : h : b$  to denote an array from  $a$  to  $b$  with step  $h$ . We consider the following 3 datasets for pretraining, where the PDE parameters are sampled with different ranges and different resolutions:

- Coarse:  $D = 0.8 : 01.05 : 3, \rho = 0.8 : 0.05 : 3.$
- Dense:  $D = 0.8 : 0.02 : 3, \rho = 0.8 : 0.02 : 3.$
- Out-of-distribution (OOD):  $D = 0.8 : 0.02 : 3, \rho = 0.8 : 0.02 : 1.8.$

In the ‘‘Coarse’’ dataset, the parameters are sampled with a coarse grid; In the ‘‘Dense’’ dataset, the parameters are sampled with a fine grid. In the ‘‘OOD’’ dataset, the parameters are sampled with a fine grid, but the ground truth  $\rho$  is outside the range of the training data.

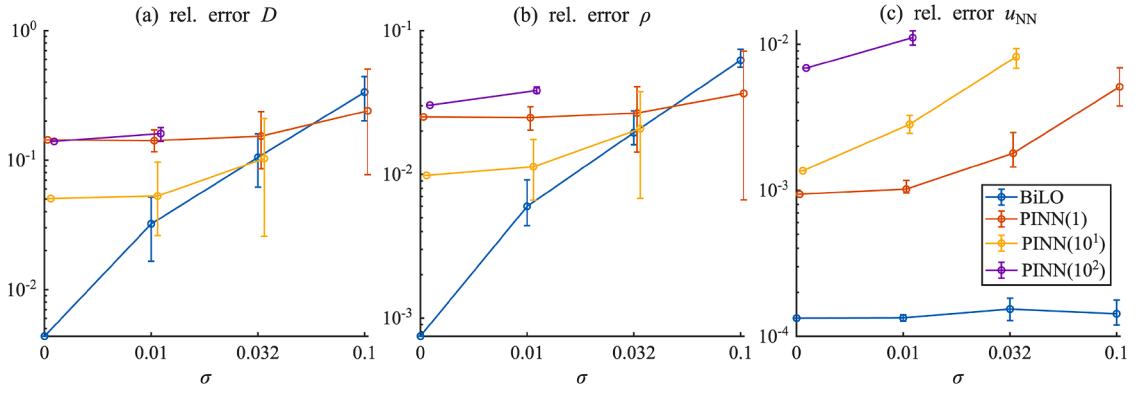


Fig. 4. Comparison of performance (a,b) relative error of inferred parameters  $D$  and  $\rho$ , (c) relative L2 error of  $u_{NN}$  compared to  $u_{FDM}$  for  $\sigma^2 = 0, 10^{-4}, 10^{-3}, 10^{-2}$  across different methods: BiLO and PINN (with  $w_{data} = 1, 10, 100$ ).

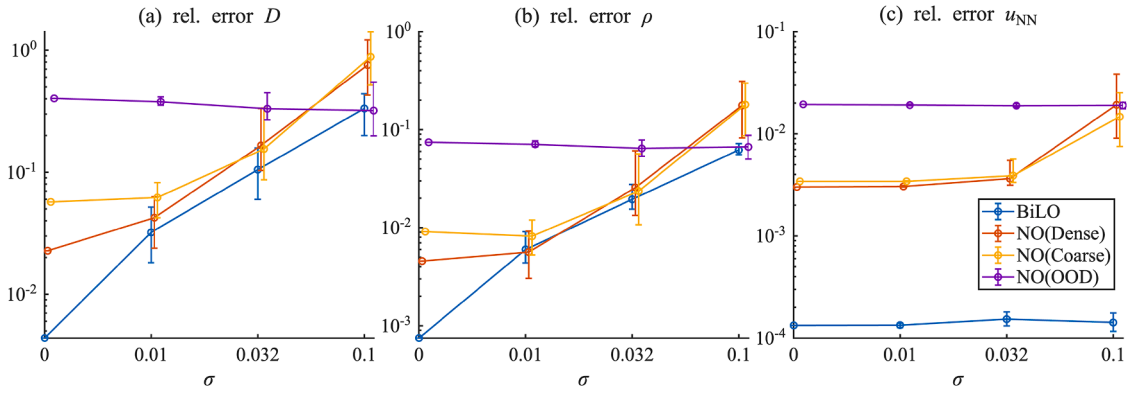


Fig. 5. Comparison of performance (a,b) relative error of inferred parameters  $D$  and  $\rho$ , (c) relative L2 error of  $u_{NN}$  compared to  $u_{FDM}$  for  $\sigma^2 = 0, 10^{-4}, 10^{-3}, 10^{-2}$  across different methods: BiLO and NO (with different pretrain datasets).

Fig. 5 illustrates that overall, BiLO achieves more accurate parameter estimation, better solution accuracy compared to NO-based methods. The performance of NO is dependent on the choice of pretraining dataset. In particular, the NO trained on out-of-distribution data exhibits degraded performance, as the inferred parameters fall outside the support of the training distribution, resulting in relatively large errors in both the estimated parameters and the reconstructed solution. The accuracy of NO shows less sensitivity to noise levels than PINN, consistent with its role as a surrogate solver.

### 3.2. Elliptic equation with singular forcing


We consider the following elliptic equation with singular forcing, which models the steady-state spatial distribution of mRNA molecules resulting from gene expression in a cell [82]:

$$\begin{cases} \Delta u + \lambda \delta(x - z) - \mu u = 0 & \text{in } \Omega \\ u = 0 & \text{on } \partial\Omega \end{cases} \quad (23)$$

Here,  $u(x)$  is the intensity measure of a spatial Poisson point process describing the location of mRNA particles in a simplified 1D domain, and  $\delta(x - z)$  denotes the Dirac delta function representing a point source at the gene site  $z$ . The parameter  $\lambda$  is the dimensionless transcription (birth) rate of mRNA,  $\mu$  is the degradation rate, and the boundary conditions correspond to nuclear export. This formulation is motivated by inferring the dynamics of gene expression from static images of single cells.

Given  $M$  snapshots of particle locations  $q_i^j$  for  $i = 1, \dots, N_j$  and  $j = 1, \dots, M$ , we aim to infer the parameters  $\lambda$  and  $\mu$ , by minimizing the negative log-likelihood function, which we refer to as the “data loss”:

$$\min_{\lambda, \mu} M \int_{\Omega} u(x) dx - \sum_{j=1}^M \sum_{i=1}^{N_j} \log u(q_i^j) \quad (24)$$



**Fig. 6.** Visualization of BiLO after pretraining at  $\lambda_0 = 250$ ,  $\mu_0 = 5$  and  $z = 0.5$ , with perturbations  $\delta\lambda = 100$  and  $\delta\mu = 1$ . Evaluating the pretrained operator  $u_{\text{NN}}$  at nearby parameter values yields approximations to the corresponding FDM solutions  $u_{\text{FDM}}$ .

To solve this example in 1D with  $\Omega = [0, 1]$ , the Eq. (23) can be written as a elliptic interface problem:

$$\begin{cases} \Delta u - \mu u = 0 \\ u(0) = u(1) = 0 \\ u^+(z) = u^-(z) \\ u_x^+(z) - u_x^-(z) = -\lambda \end{cases} \quad (25)$$

where the superscript + and - denote the limiting values from the right and left side of the interface at  $z$ , respectively. The solution is continuous, but its derivative is discontinuous at  $z$ .

We handle the singular forcing using the cusp-capturing PINN [83], which has been proposed to learn functions of the form  $\tilde{u}(x, \phi)$  such that  $u(x) = \tilde{u}(x, |x - z|)$ . The continuity condition is automatically satisfied, and the jump condition translates into an additional constraint in  $\mathcal{F}$ :

$$\partial_\phi \tilde{u}(z, 0) = -\lambda.$$

The cusp-capturing PINN is parameterized by  $\tilde{u}(x, \phi; W)$ , and to enforce the jump condition, we need the ‘‘jump loss’’:

$$\mathcal{L}_{\text{jump}}^{\text{PINN}}(W) = (\partial_\phi \tilde{u}(z, 0; W) + \lambda)^2 \quad (26)$$

The total loss for the PINN will include the residual loss, the jump loss, and the data loss.

In the BiLO framework, the local operator is parameterized as  $\tilde{u}(x, \phi, \theta; W)$ , where  $\theta = (\lambda, \mu)$ . The ‘‘jump loss’’ is defined as

$$\mathcal{L}_{\text{jump}}(\theta, W) = \partial_\phi \tilde{u}(z, 0; \theta; W) + \lambda)^2 \quad (27)$$

and we also need the ‘‘jump gradient loss’’:

$$\mathcal{L}_{\text{jgrad}}(\theta, W) = (\nabla_\theta \partial_\phi \tilde{u}(z, 0; \theta; W))^2 \quad (28)$$

The integral in the upper-level loss is approximated using the Simpson’s rule with 100 uniform grid points. The lower-level loss is given by  $\mathcal{L}_{\text{res}} + \mathcal{L}_{\text{jump}} + w_{\text{rgrad}}(\mathcal{L}_{\text{rgrad}} + \mathcal{L}_{\text{jgrad}})$ . The upper-level loss is given by the data loss (24) We visualize the local operator  $u(x, \lambda, \mu; W)$  in Fig. 6.

The objective function is more challenging than the MSE, as the values of  $u$  are not available at the particle locations –  $u$  is only proportional to the histogram of the particle locations. In this problem, the decay rate  $\mu$  is typically on the order of 10, while the birth rate  $\lambda$  is on the order of several hundreds, consistent with biologically plausible dynamics. To improve numerical conditioning during training, we reparameterize  $\lambda = 100\bar{\lambda}$  and learn the rescaled parameter  $\bar{\lambda}$ . In addition, we apply a final transformation  $\tau(m, x) = m^2 x(1 - x)$ , where  $m$  is the raw output of the neural network. This transformation enforces the boundary conditions and ensures non-negativity of the solution, which is necessary for evaluating  $\log u$  in the likelihood. The initial guesses are  $\lambda_0 = 250$  and  $\mu_0 = 5$ , with a relative error of 50% compared to the ground truth. The particle positions are sampled using ground truth parameters  $\lambda_{\text{GT}} = 500$  and  $\mu_{\text{GT}} = 2.5$  with  $M = 100$ .

For experiments using neural operator, we consider two pre-training datasets:

- In-distribution (ID):  $\lambda = 100 : 20 : 800$ ,  $\mu = 2 : 1 : 20$ .




Fig. 7. Comparison of performance metrics for BiLO and PINN ( $w_{\text{data}} = 1$  and 10) and FNO (with different pretraining datasets) for the elliptic equation with singular forcing.

- Out-of-distribution (OOD):  $\lambda = 100 : 20 : 800$ ,  $\mu = 3.5 : 0.5 : 10$ .

The  $\mu_{\text{GT}}$  are outside the range of the training data in the OOD dataset.

The results are shown in Fig. 7, and BiLO achieves significantly better performance than PINN ( $w_{\text{data}} = 1$  and 10) and FNO (with different pretraining datasets). For PINN, the results depend on the choice of  $w_{\text{data}}$ .  $w_{\text{data}} = 10$  leads to larger error in  $\lambda$  and overfitting, as indicated by the elevated relative error of  $u_{NN}$  compared to  $u_{\text{FDM}}$ . However,  $w_{\text{data}} = 1$  leads to larger relative error in  $\mu$ . For FNO, when pretrained on the ID dataset, although the relative errors in parameter estimation are larger than that of BiLO, the PDE solution accuracy remains comparable, with no signs of overfitting or underfitting. However, the OOD dataset results in larger parameter errors and a less accurate solution.

### 3.3. Poisson equation with variable diffusion

We consider the following 1D Poisson equation

$$\begin{cases} -(D(x)u'(x))' = f(x) & \text{in } [0, 1] \\ u(0) = u(1) = 0 \end{cases} \quad (29)$$

where  $f(x) = \pi^2 \sin(\pi x)$ . We aim to infer the variable diffusion coefficient  $D(x)$  such that  $D(0) = D(1) = 1$ . The ground truth  $D(x)$  is a ‘‘hat’’ function:

$$D(x) = \begin{cases} 1 + 0.5x, & \text{if } x \in [0, 0.5] \\ 1.5 - 0.5x, & \text{if } x \in [0.5, 1] \end{cases} \quad (30)$$

The initial guess is  $D(x) = 1$ . The relative error of the initial guess in  $L^\infty$  norm is around 25%. Moreover, the initial guess is overly smooth and misses the kink at  $x = 0.5$ , offering little information about the true profile.

**Visualizing BiLO.** In Fig. 8, we visualize the local operator  $u(x, z; W)$  after pre-training with  $D_0(x) = 1$ . We consider the variation  $\delta D_1(x) = -0.1$ , and  $\delta D_2(x) = 0.1x$  and evaluate the neural network at  $u(x, D_0(x) + \delta D_i(x); W)$  for  $i = 1, 2$ . The FDM solutions of the PDE corresponding to  $D_0(x) + \delta D_i(x)$  are also plotted. We can see that the neural network provides a good approximation to the solution corresponding to  $D_0(x) + \delta D_i(x)$ .

**Results.** We next estimate  $D(x)$  in the presence of noise at levels  $\sigma = 0, 0.01, 0.03$  and use  $w_{\text{reg}} = 10^{-3}$ . We use BiLO, PINNs with  $w_{\text{data}} = 1, 10, 100$ , and the adjoint method. Fig. 9 shows that BiLO consistently produces more accurate estimates of the diffusion coefficient  $D(x)$  across all noise levels. In contrast, the performance of PINN is highly sensitive to the choice of  $w_{\text{data}}$ : larger values such as  $w_{\text{data}} = 10^3$  work well under noise-free conditions but lead to poor performance when the noise level increases; smaller values such as  $w_{\text{data}} = 10$  are more robust to noise but underfit when no noise is present. The reconstruction accuracy of the solution  $u_{NN}$  is comparable between BiLO and PINN. The adjoint method, which solves the PDE with high accuracy and is considered ‘‘exact’’ (hence no corresponding bar in (b)), reconstructs the diffusion coefficient less accurately than BiLO.

We also infer the variable diffusion coefficient  $D(x)$  in the Poisson equation using a DeepONet. The pretraining dataset is generated by solving the Poisson equation with 1000 samples of variable diffusion coefficients  $D(x)$ .  $D(x)$  is sampled from a Gaussian Random field on  $[0, 1]$ , conditioned on  $D(0) = D(1) = 1$ . The covariance function is the Gaussian kernel, with variance 0.05 and different length scales  $l = 0.2, 0.3, 0.4$ . As  $l$  increases, the samples of  $D(x)$  become smoother.




Fig. 8. Visualizing the operator  $u(x, D_0(x) + \delta D_i(x); W)$ , where  $\delta D_1(x) = -0.1$  and  $\delta D_2(x) = 0.1x$ , after pre-training with  $D_0(x) = 1$ . The FDM solutions corresponding to  $D_0(x) + \delta D_i(x)$  are also shown.

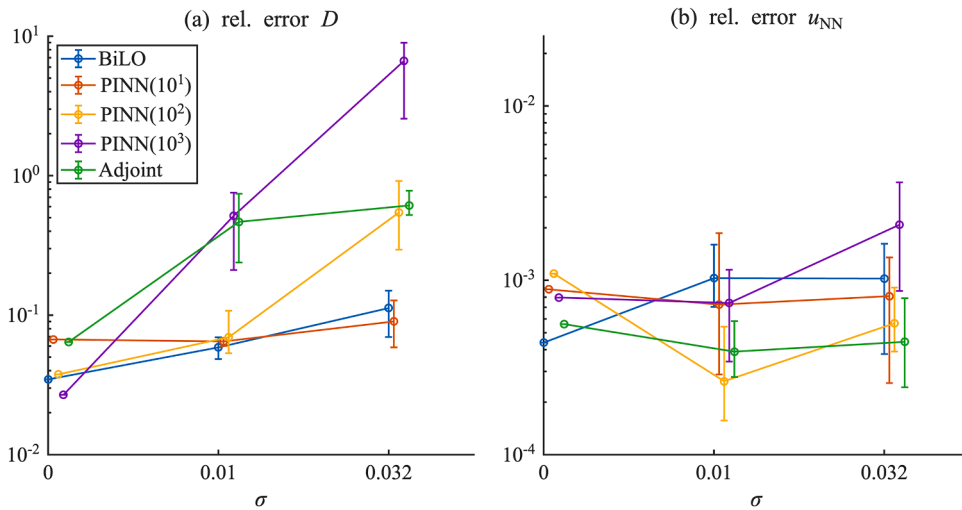


Fig. 9. Comparison of performance metrics – (a) relative error in  $D$  and (b) relative error in  $u$  – for BiLO, PINN (with various  $w_{\text{data}}$ ), and the adjoint method.

Fig. 10 shows that the performance of NO depends on the choice of pretraining dataset. Because of the ill-posedness, the inference of  $D(x)$  is similar across the different methods, although BiLO is more accurate at smaller noise levels and at least as accurate as NO at larger noise levels. While the accuracy of NO in approximating the solution remains relatively stable across different noise levels, it is less accurate than both PINN and BiLO. Among all methods, BiLO consistently achieves comparable or better inferences of the diffusion coefficient  $D(x)$ .

Fig. 11 illustrates qualitative differences in the inferred diffusion coefficient  $D(x)$  across the methods. BiLO best captures the kink in the ground truth, leading to a more accurate reconstruction. In contrast, PINN with low  $w_{\text{data}}$  produces overly smooth estimates, while large  $w_{\text{data}}$  leads to oscillatory artifacts due to overfitting. The NO reconstructions, regardless of the length scale used in pretraining, tend to be smooth and lack sharp features. The adjoint method yields slightly more oscillatory reconstructions compared to BiLO.

### 3.4. Inferring the initial condition of a heat equation

In this example, we aim to infer the initial condition of a 1D heat equation from the final state. Consider the heat equation

$$\begin{cases} u_t(x, t) = Du_{xx}(x, t) \\ u(x, 0) = f(x) \\ u(0, t) = u(1, t) = 0 \end{cases} \tag{31}$$

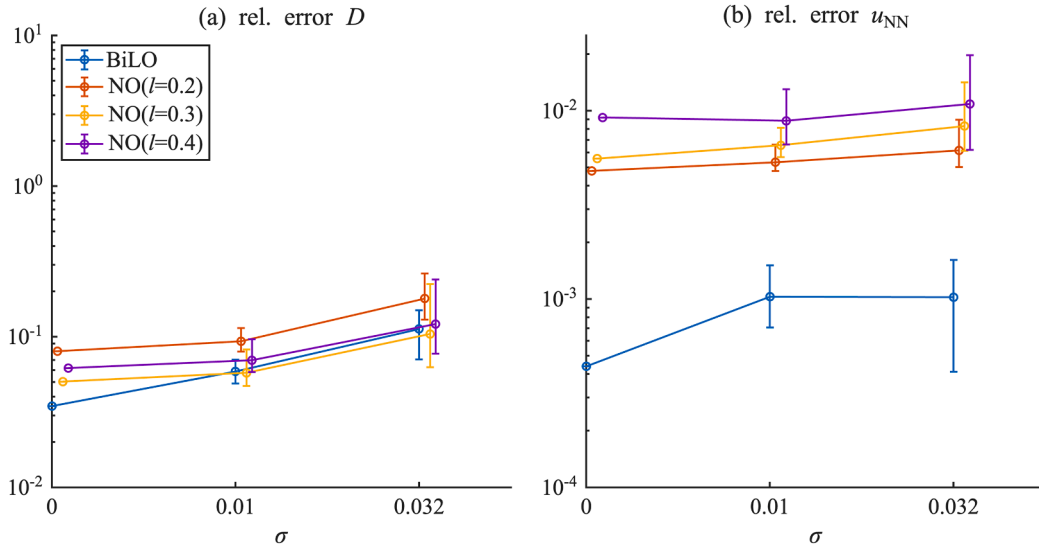


Fig. 10. Comparison of performance metrics for different methods: BiLO and NO (with different pretrain datasets).

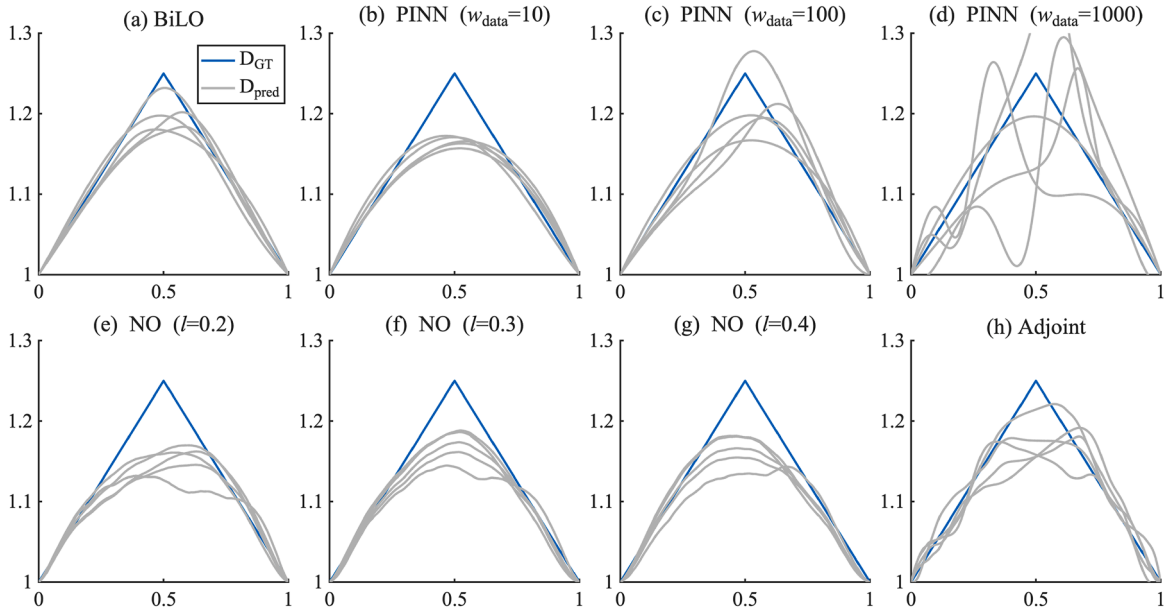


Fig. 11. Examples of the inferred diffusion coefficient  $D(x)$  from various methods with  $\sigma = 0.01$ .

on  $x \in [0, 1]$  and  $t \in [0, 1]$ , with fixed diffusion coefficient  $D = 0.01$ , and unknown initial condition  $f(x)$ , where  $f(0) = f(1) = 0$ . Our goal is to infer the initial condition  $f(x)$  from observation of the final state  $u(x, 1)$ . We set the ground truth initial condition  $f_{GT}$  to be the hat function

$$f_{GT}(x) = \begin{cases} 2x, & \text{if } x \in [0, 0.5] \\ 2 - 2x, & \text{if } x \in [0.5, 1] \end{cases} \quad (32)$$

We set the initial guess  $f_0(x) = \sin(\pi x)$ . The initial guess has a relative  $L^\infty$  error of about 21%. In addition, its smooth sinusoidal shape does not capture the nondifferentiability at  $x = 0.5$ . To evaluate the performance of the estimated initial condition  $f$ , we use the  $L^2$  norm of the estimated initial condition and the ground truth initial condition, which are evaluated at 1001 evenly spaced points in the spatial domain. We consider the case with noise  $\epsilon \sim N(0, 0.001)$ . Due to the ill-posedness of the inverse problem, we need to regularize the problem by penalizing the 2-norm of the derivative of the unknown function with  $w_{reg} = 10^{-2}$ . The results, shown in Fig. 12, show that BiLO outperforms PINNs with various values of  $w_{data}$ .

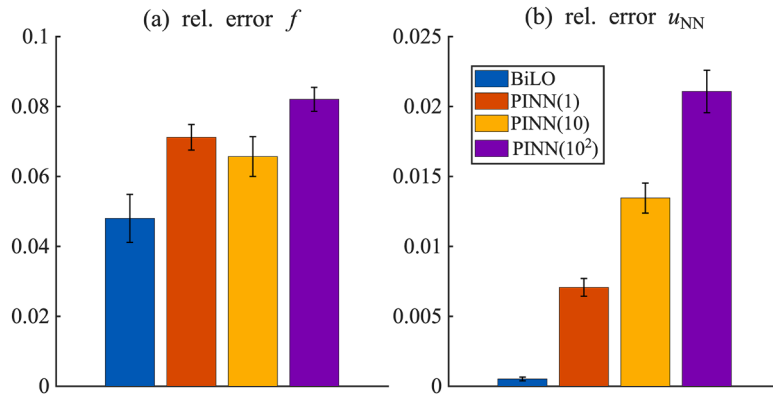


Fig. 12. Comparison of the BiLO and PINN (with various  $w_{\text{data}}$ ) for inference of the initial condition of the heat equation.

### 3.5. Inviscid burgers' equation

We consider an inverse problem governed by an inviscid Burgers' equation on the domain  $x \in [0, 1]$  and  $t \in [0, 1]$ .

$$\begin{cases} u_t + auu_x = 0, \\ u(x, 0) = f(x), \\ u(0, t) = u(1, t) = 0 \end{cases} \tag{33}$$

where  $a = 0.2$ . We infer the initial condition  $f$  from the observational data at  $t = 1$ . The numerical solutions are computed by using the Godunov scheme. The inviscid Burgers' equation is a hyperbolic PDE, and the solution can develop shocks and rarefaction waves.

In Fig. 13, we show the initial guess in the first column, the ground truth in the second column, and the inference results by BiLO ( $w_{\text{reg}} = 10^{-3}$ ) in the third column. The first row shows the initial condition  $f(x)$ , the second row shows the solution  $u(x, t)$  on the domain  $x \in [0, 1]$  and  $t \in [0, 1]$ , and the third row shows the solution  $u(x, 1)$ . For inference, only the solution at  $t = 1$  of the ground truth is provided. The solution of the initial guess remains smooth over the entire time interval  $t \in [0, 1]$ , with no discontinuity forming. In contrast, the ground-truth initial condition shown in the second column evolves into a solution with a discontinuity over time. The inferred initial condition and solution shown in the third column recover this non-smooth behavior, demonstrating that BiLO can capture discontinuity despite starting from a smooth initial guess. Additional examples are shown in Appendix C.5.

### 3.6. Darcy flow in 2D

The setup of this experiment is similar to the steady state Darcy flow inverse problem in [33]. We consider the following 2D Poisson equation with variable diffusion coefficient  $A(\mathbf{x})$  in the unit square domain  $\Omega = [0, 1] \times [0, 1]$  with Dirichlet boundary condition:

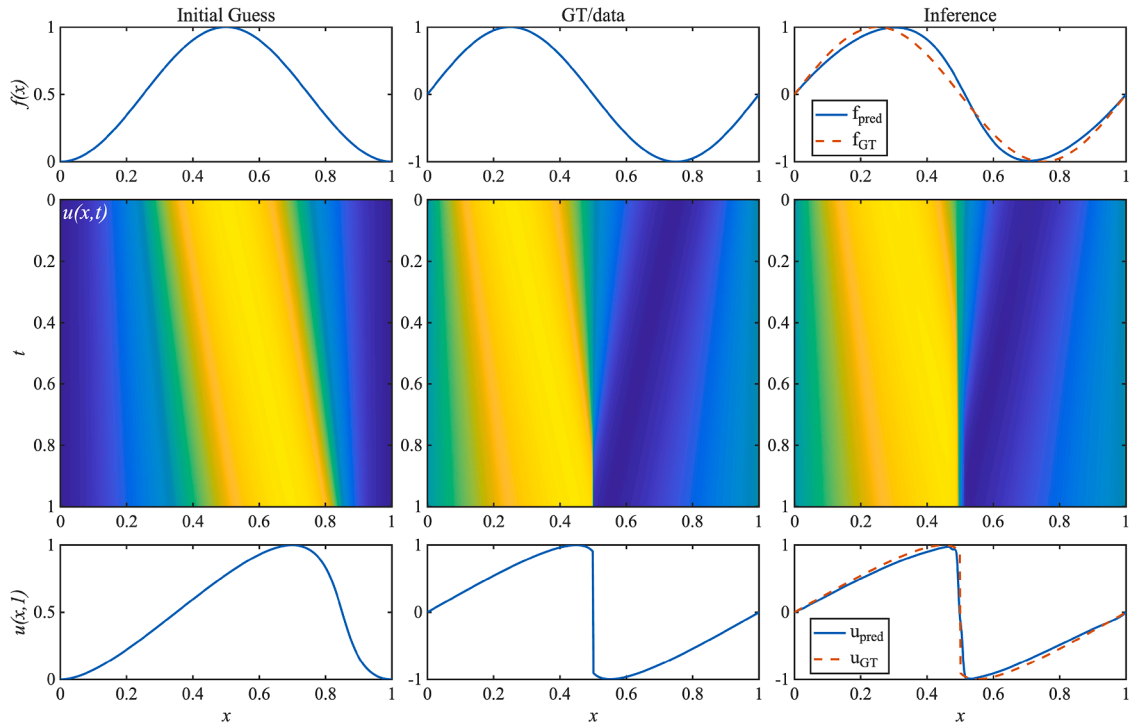
$$\begin{cases} -\nabla \cdot (A(\mathbf{x})\nabla u(\mathbf{x})) = f(\mathbf{x}) & \text{in } \Omega \\ u(\mathbf{x}) = 0, & \text{on } \partial\Omega \end{cases} \tag{34}$$

Our goal is to infer the variable diffusion coefficient  $A(\mathbf{x})$  from the solution  $u(\mathbf{x})$ . Let  $\phi(\mathbf{x})$  be samples of a Gaussian random field (GRF) with mean 0 and squared exponential (Gaussian) covariance structure

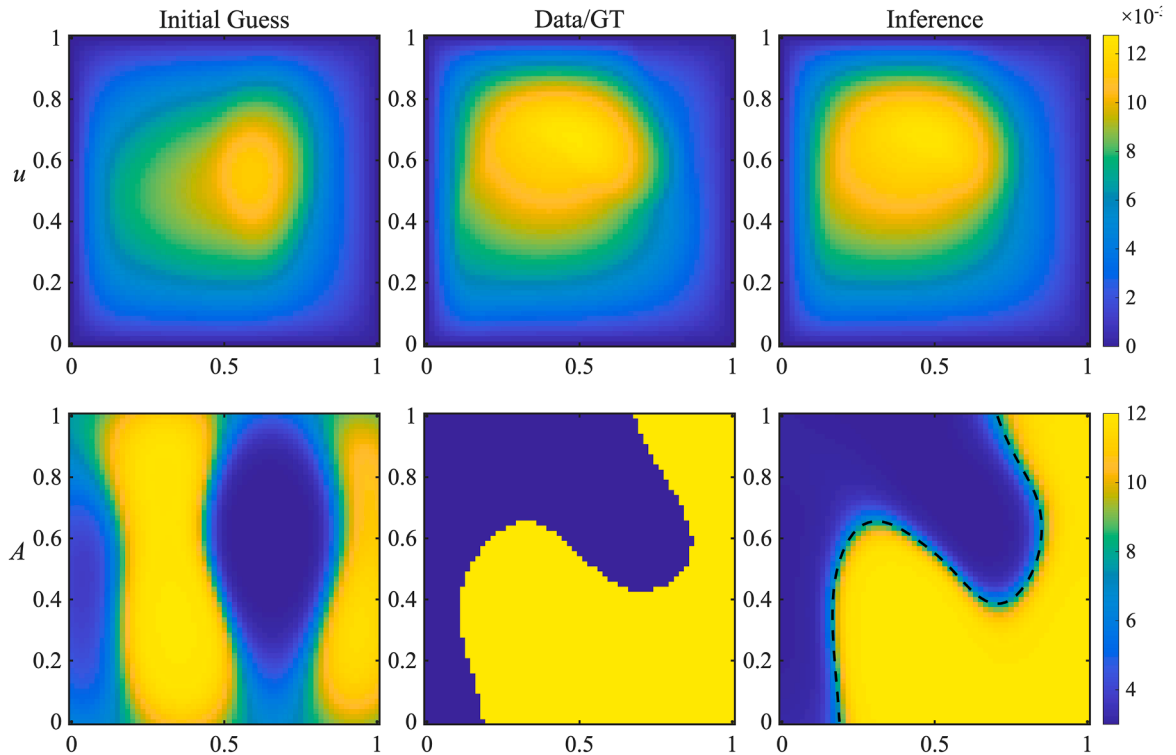
$$C(\mathbf{x}, \mathbf{y}) = \sigma \exp(-\|\mathbf{x} - \mathbf{y}\|^2 / \lambda^2),$$

where the marginal standard deviation  $\sigma = \sqrt{10}$  and the correlation length  $\lambda = 0.01$  [84]. We generate the initial guess  $A_0(\mathbf{x}) = s(\phi_0(\mathbf{x})) \times 9 + 3$ , where  $\phi_0(\mathbf{x})$  is a sample of the GRF, and  $s(x) = 1/(1 + e^{-x})$  is the logistic function. We consider the ground truth diffusion coefficient to be a piecewise constant function with high and low diffusivity:  $A_{\text{GT}}(\mathbf{x}) = 12$  if  $\phi_{\text{GT}}(\mathbf{x}) > 0$  and  $A_{\text{GT}}(\mathbf{x}) = 3$  otherwise, where  $\phi_{\text{GT}}$  is another sample of the GRF. The neural network representation of  $A(\mathbf{x})$  is smooth by construction; hence, we approximate the discontinuous ground truth diffusion coefficient by thresholding the inferred field at 7.5 [33]. The corresponding solution of  $A_0$  and  $A_{\text{GT}}$  are denoted as  $u_0$  and  $u_{\text{GT}}$ . Following [33], we use the total variation regularization  $|\nabla A|$  with weight  $w_{\text{reg}} = 10^{-9}$ .

Fig. 14 shows an example of the results of the inverse problem. The first and the second row show  $u$  and  $A$ , respectively. The first column shows the initial guess, the second column shows the ground truth, and the third column shows the inference results by BiLO. The dashed line marks the 7.5 level used to classify regions of high and low diffusivity, yielding a classification accuracy of 96.1%, which is comparable to the results (97.1% classification accuracy) from the Physics-informed Neural Operator (PINO) in [33], which requires pretraining a FNO with a synthetic dataset, and instance-wise fine-tuning with physics-informed loss. In our method, we only need to pretrain the BiLO with a single initial guess, which is visually different from the ground truth. Additional examples are shown in Appendix C.6.



**Fig. 13.** Example of inferring the initial condition of the inviscid Burgers' equation using data at  $t = 1$ . The initial guess is used to pre-train the network. The solution at  $t = 1$  of the GT is the data for inference. First column: initial guess, second column: ground truth, third column: inferred initial condition. First row: initial condition, second row: solution  $u(x, t)$ , third row: solution  $u(x, 1)$ .



**Fig. 14.** Example of inferring the variable diffusion coefficient in the 2D Darcy flow problem. The network prediction achieves a relative MSE of 1.7% for  $u$ . Thresholding the inferred  $A(x)$  at 7.5 (dashed line) recovers the discontinuous ground-truth structure with a classification accuracy of 96.1%.



**Table 1**

Results of the GBM inverse problem: Average DICE scores for the WT and TC regions based on predicted segmentations from the neural network solution  $u_{\text{NN}}$  and the numerical solution  $u_{\text{FDM}}$ , along with the relative mean squared error (MSE) of  $u_{\text{NN}}$  at  $t = 1$ .

Methods	DICE <sub>NN</sub>	DICE <sub>FDM</sub>	rel.MSE(%)
BiLO	0.84	0.84	0.04
PINN( $10^{-3}$ )	0.98	0.75	28
PINN( $10^{-5}$ )	0.87	0.83	1
PINN( $10^{-7}$ )	0.83	0.83	0.03

### 3.7. Glioblastoma (GBM) inverse problem

In this section, we consider a real-world application of BiLO for a patient-specific parameter estimation of GBM growth models using patient MRI data in 2D. The challenge lies in the high noise levels and the potential model mis-specification, as the Fisher-KPP PDE likely does not fully capture the complexities present in the tumor MRI data. The setup of the problem follows [7,24,85,86].

*Tumor growth and imaging model.* Let  $\Omega$  be the brain region in 2D based on MRI images. The normalized tumor cell density is  $u(\mathbf{x}, t)$ .

$$\begin{cases} \frac{\partial u}{\partial t} = D\bar{D}\nabla \cdot (P(\mathbf{x})\nabla u) + \rho\bar{\rho}u(1-u) & \text{in } \Omega \\ \nabla u \cdot \mathbf{n} = 0 & \text{on } \partial\Omega \end{cases} \quad (35)$$

where  $P$  depends on the tissue distribution (e.g., white and grey matter) and is obtained from the MRI data,  $\bar{D}$  and  $\bar{\rho}$  are known patient-specific characteristic parameters based on the data, and  $D$  and  $\rho$  are the unknown nondimensionalized parameters that we aim to infer from the data.

We consider two regions of interest in the tumor, the whole tumor (WT) region and the tumor core (TC) region. Let  $\hat{y}^s$ ,  $s \in \{\text{WT}, \text{TC}\}$  be indicator function of the WT and TC regions, which can be obtained from the MRI data and serves as the observational data in the inverse problem. We assume that the segmentations are the regions in which the tumor cell density  $u$  at nondimensional  $t = 1$  lies above a certain threshold  $u_c^s$ :

$$\mathbf{y}^s(\mathbf{x}) = s(20(u(\mathbf{x}, 1) - u_c^s)),$$

where  $s$  is the logistic function. The predicted segmentation depends on the solution of the PDE, and thus on the parameters  $D$ ,  $\rho$  and  $u_c^s$ . In the end, we aim to minimize the relative error between the predicted and the observed segmentations:

$$\min_{D, \rho, u_c^{\text{WT}}, u_c^{\text{TC}}} \sum_{s \in \{\text{WT}, \text{TC}\}} \|\mathbf{y}^s - \hat{\mathbf{y}}^s\|_2^2 / \|\hat{\mathbf{y}}^s\|_2^2 \quad (36)$$


*Results.* In this inference problem, the ground truth values for the parameters  $D$ ,  $\rho$ ,  $u_c^{\text{WT}}$ , and  $u_c^{\text{TC}}$  are not available. Therefore, we evaluate the quality of the inferred parameters by comparing the predicted segmentations  $\hat{\mathbf{y}}$  with the observed segmentation  $\mathbf{y}$  data using the DICE score, which is defined as  $2\langle \mathbf{y}, \hat{\mathbf{y}} \rangle / (\|\mathbf{y}\|_1 + \|\hat{\mathbf{y}}\|_1)$  [87]. DICE is a standard metric in medical image segmentation that quantifies the overlap between two binary masks [6,7].

The predicted segmentations are obtained by thresholding the tumor cell density at the inferred thresholds  $u_c^{\text{WT}}$  and  $u_c^{\text{TC}}$ . These densities can be computed either from the numerical PDE solution  $u_{\text{FDM}}$  or the neural network surrogate  $u_{\text{NN}}$ , resulting in predicted masks denoted by  $\mathbf{y}_{\text{FDM}}^s$  and  $\mathbf{y}_{\text{NN}}^s$ , respectively. We define DICE <sub>$m$</sub> ,  $m \in \{\text{NN}, \text{FDM}\}$ , as the average DICE score across the WT and TC regions using the corresponding predicted segmentation.

Table 1 reports the relative errors of  $u_{\text{NN}}$  and  $u_{\text{FDM}}$  at  $t = 1$ , as well as the DICE scores DICE <sub>$m$</sub> . Fig. 15 visualizes the predicted segmentations using BiLO and PINN for different values of  $w_{\text{data}}$ . For the PINNs, we observe that the DICE score based on  $u_{\text{NN}}$  is generally higher than that based on  $u_{\text{FDM}}$ , indicating a tendency to overfit the data while compromising the fidelity of the PDE solution. This behavior is reflected in the larger relative errors of  $u_{\text{NN}}$  and is visually apparent in Fig. 15, where the contours from  $u_{\text{NN}}$  track the noisy segmentation data more closely than those from  $u_{\text{FDM}}$ .

Reducing  $w_{\text{data}}$  helps mitigate this overfitting by regularizing the data fitting. In contrast, BiLO achieves both accurate PDE solutions and well-performing parameters without the need to tune  $w_{\text{data}}$ , leading to better segmentations in this case. Interestingly, even when  $u_{\text{NN}}$  is not accurate in the PINN setting, the inferred parameters can still yield reasonable segmentation when evaluated using  $u_{\text{FDM}}$ , as evidenced by the corresponding DICE scores.

*Summary of additional results.* For each numerical experiments, the training details and some additional results are provided in Appendix C. In Appendix D, we demonstrate the robustness of BiLO with respect to the learning rate and the residual-gradient weight  $w_{\text{grad}}$ . These findings highlight the practical reliability of BiLO without requiring extensive hyperparameter tuning. The computational costs, including seconds-per-step and peak memory usage, are provided in Appendix E. While BiLO is more costly per step than PINN, overall it requires fewer steps to converge, leading to a lower total computational cost.



**Fig. 15.** Predicted segmentation using PINN with  $w_{\text{data}} = 10^{-3}, 10^{-5}, 10^{-7}$  and BiLO. The solid and dashed contours are the predicted segmentation based on  $u_{\text{NN}}$  and  $u_{\text{FDM}}$ . BiLO gives almost overlapping contours, suggesting high accuracy of  $u_{\text{NN}}$ .

#### 4. Conclusion

In this work, we presented a Bi-level Local Operator (BiLO) learning framework for solving PDE-constrained optimization problems. We minimize the data loss with respect to the PDE parameters at the upper level, and learn the local solution operator of the PDE at the lower level. The bi-level optimization problem is solved using simultaneous gradient descent, leading to an efficient algorithm. Empirical results demonstrate more accurate parameter recovery and stronger fidelity to the underlying PDEs under sparse and noisy data, compared with the soft PDE-constraint formulation, which faces the delicate trade-off between adhering to the PDE constraints and accurately fitting the data. Future work includes a deeper theoretical analysis of the simultaneous gradient descent dynamics, reducing computational cost via Gaussian Process representations of the PDE solution and Gauss–Newton updates at the lower level [55,88], and extending BiLO to more complex settings such as full 3D tumor inverse problems [7,24].

#### Code availability

The code for the numerical experiments is available at <https://github.com/Rayzhangzirui/BiLO>.

#### CRedit authorship contribution statement

**Ray Zirui Zhang:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Christopher E. Miles:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Xiaohui Xie:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Conceptualization; **John S. Lowengrub:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization.

#### Data availability

The code is available at <https://github.com/Rayzhangzirui/BiLO>.

## Declaration of competing interest

John S. Lowengrub reports financial support was provided by National Science Foundation. Christopher E. Miles reports financial support was provided by National Science Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

R.Z.Z and J.S.L thank Babak Shahbaba for GPU resources. J.S.L acknowledges partial support from the [National Science Foundation](#) through grants [DMS-2309800](#) and [DMS-1763272](#) and the [Simons Foundation \(594598QN\)](#) for an NSF-Simons Center for Multiscale Cell Fate Research. C.E.M was partially supported by a NSF CAREER grant DMS-2339241.

## Appendix A. Theoretical analysis

In the main text, we describe a simultaneous gradient descent algorithm for solving the bi-level optimization problem. In this section, we provide a theoretical justification for the algorithm under certain simplifying assumptions.

This appendix provides a rigorous bound on the error of the approximate hypergradient used in the BiLO framework, with results that apply to both PDE-constrained optimization setting in Part I and Bayesian inference in Part II [13]. While the proofs rely on highly idealized conditions, the resulting theory offers valuable insights into the theoretical behavior of BiLO. Numerical experiments in the main text further illustrate that BiLO remains effective even when some restrictive assumptions are relaxed.

We present our theoretical analysis through two theorems organized into distinct subsections to enhance clarity and accessibility. The first theorem ([Theorem 1](#) in [Appendix A.1](#)) focuses on the consistency of the approximate gradient, showing that it is exact when the lower-level problem is solved to optimality. It illustrates the key calculations in a simplified setting, specifically under the idealized condition of exact minimization at the lower level, and assuming a linear PDE operator with established stability properties. This special case offers a transparent view into why BiLO's approximate gradient can be exact under ideal circumstances. The second theorem ([Theorem 2](#), in [Appendix A.2](#)) extends this analysis into a more general and abstract setting, focusing on the approximation error introduced by the inexact minimization of the lower-level problem. Although [Theorem 1](#) can technically be viewed as a special case or a corollary of [Theorem 2](#), explicitly stating and proving it separately aids in highlighting the fundamental mechanism underlying BiLO and facilitates understanding of the broader and more general error analysis that follows.

### A.1. Consistency

*Setup.* We consider the minimization problem

$$\min_{\theta} \|u - \hat{u}\|$$

where  $\hat{u}$  is the observation, and  $u$  is the solution of the PDE

$$\begin{cases} \mathbf{L}u = f & \text{in } \Omega \\ u = g & \text{on } \partial\Omega, \end{cases} \quad (\text{A.1})$$

where  $\Omega$  is a connected, open and bounded subset of  $\mathbb{R}^d$ .  $\mathbf{L}$  denotes a second-order partial differential operator:

$$\mathbf{L}u = \sum_{i,j=1}^d a_{ij} \partial_{ij} u + \sum_{i=1}^d b_i \partial_i u + cu \quad (\text{A.2})$$

where the coefficients  $a_{ij}$ ,  $b_i$ ,  $c$  are collectively denoted as  $\theta$ . We denote  $\mathbf{L}_\theta$  as the Fréchet derivative of  $\mathbf{L}$  with respect to  $\theta$ , which is also a differential operator.

We define the **ideal optimal weight** as the exact minimizer of the following problem:

$$W^*(\theta) = \arg \min \int_{\Omega} (\mathbf{L}u(\theta, W) - f)^2 dx$$

That is, for all  $\theta$

$$\mathbf{L}u(\theta, W^*(\theta)) = f. \quad (\text{A.3})$$

Denote  $u^* = u(\theta, W^*)$  and define

$$v := d_\theta u(\theta, W^*(\theta)) = \nabla_\theta u^* + \nabla_W u^* \nabla_\theta W^* \quad (\text{A.4})$$

The exact hypergradient is given by

$$g_{\text{true}}(\theta) = \int_{\Omega} (u^* - \hat{u}) v dx \quad (\text{A.5})$$

In BiLO, the gradient of the upper level objective with respect to  $\theta$  is given by

$$g_{\text{approx}}(\bar{W}, \theta) = \int_{\Omega} \left( u(\bar{W}, \theta) - \hat{u} \right) \nabla_{\theta} u(\bar{W}, \theta) d\mathbf{x} \quad (\text{A.6})$$

where  $\bar{W}$  is the minimizer of the lower level problem

$$\bar{W} = \arg \min \int_{\Omega} (\mathbf{L}u - f)^2 + w_{\text{reg}} (\mathbf{L}_{\theta}u + \mathbf{L}\nabla_{\theta}u)^2 d\mathbf{x}$$

**Theorem 1** (Consistency of the approximate gradient). *Assuming (i) The maximum principle holds for the PDE operator; (ii) The parametrized local operator  $u(\theta, W)$  satisfies the boundary condition for all  $\theta$  and  $W$ ; (iii) The lower-level optimization is solved exactly. then the approximate gradient of the upper level objective is exact.*

**Proof.**

Taking the total derivative of (A.3) with respect to  $\theta$ , we have

$$2\mathbf{L}_{\theta}u^* + \mathbf{L}[\nabla_{\theta}u^* + \nabla_W u^* \nabla_{\theta}W^*] = \mathbf{L}_{\theta}u^* + \mathbf{L}v = 0 \quad (\text{A.7})$$

By assumption (iii), the residual loss and the residual-gradient loss vanish exactly at the solution  $\bar{u}$ :

$$\mathbf{L}\bar{u} = f \quad (\text{A.8})$$

$$\mathbf{L}_{\theta}\bar{u} + \mathbf{L}\nabla_{\theta}\bar{u} = 0 \quad (\text{A.9})$$

By the uniqueness of the solution to the PDE, we have  $u^* = \bar{u}$ . Subtract (A.7) from (A.9), we obtain

$$\mathbf{L}[\nabla_{\theta}\bar{u} - v] = 0 \quad (\text{A.10})$$

Since  $u(\theta, W) = g$  on  $\partial\Omega$  for all  $W$  and  $\theta$ , we have  $\nabla_{\theta}\bar{u} - v = 0$  on  $\partial\Omega$ . If the maximum principle holds for the operator  $\mathbf{L}$  (for example, when  $\mathbf{L}$  uniformly elliptic and  $c \geq 0$  [89]), then  $v - \nabla_{\theta}\bar{u} = 0$ .

The difference between the exact gradient and the approximate gradient, which we denote as  $\Delta g$ , is given by

$$\begin{aligned} \Delta g &= g_{\text{approx}}(\bar{W}, \theta) - g_{\text{true}}(\theta) \\ &= \int_{\Omega} (u^* - \hat{u})(\nabla_{\theta}\bar{u} - v) d\mathbf{x} \end{aligned} \quad (\text{A.11})$$

By the Cauchy-Schwarz inequality, we have

$$|\Delta g| \leq \|u^* - \hat{u}\| \|\nabla_{\theta}\bar{u} - v\| = 0 \quad (\text{A.12})$$

That is, the approximate gradient at  $\bar{W}$  is exact.  $\square$

These assumptions are more restrictive than the numerical experiments. For example, in the Fisher-KPP example, the PDE operator is nonlinear. A more general analysis is shown in the next section, where we bound the error of the approximate gradient by the lower-level optimization error.

## A.2. Approximation error

*Setup.* We consider the following PDE-constrained optimization problem:

$$\begin{aligned} \min_{\theta} \quad & \mathcal{J}[u] \\ \text{s.t.} \quad & \mathcal{F}[u, \theta] = \mathbf{0} \end{aligned} \quad (\text{A.13})$$

on an open and bounded domain  $\Omega \subset \mathbb{R}^d$ . The PDE parameter  $\theta \in \mathbb{R}^m$ . The PDE residual operator  $\mathcal{F} : H_0^1(\Omega) \times \mathbb{R}^m \rightarrow H^{-1}(\Omega)$ . The objective function  $\mathcal{J} : H_0^1(\Omega) \rightarrow [0, \infty)$ . Denote the functional derivative of  $\mathcal{J}$  by  $\mathcal{J}'[\cdot]$ . We consider the PDE solution map, parameterized by weights  $W \in \mathbb{R}^n$ ,  $u : \mathbb{R}^m \times \mathbb{R}^n \rightarrow H_0^1(\Omega)$ . The residual as a function of  $\theta$  and  $W$  is defined as

$$r(\theta, W) := \mathcal{F}[u(\theta, W), \theta] \quad (\text{A.14})$$

Denote the (partial) Fréchet derivative of  $\mathcal{F}$  by  $\mathcal{F}_u$  and  $\mathcal{F}_{\theta}$ . The residual-gradient is defined as

$$\nabla_{\theta} r(\theta, W) := \mathcal{F}_u(u, \theta)[\nabla_{\theta}u(\theta, W)] + \mathcal{F}_{\theta}(u(\theta, W), \theta) \quad (\text{A.15})$$

Denote the linearized PDE operator at  $u$  as

$$\mathbf{L}_u[\cdot] := \mathcal{F}_u(u, \theta)[\cdot]$$

The **ideal optimal weights**  $W^*(\theta)$  satisfies the PDE for all  $\theta$ :

$$\mathcal{F}[u(\theta, W^*(\theta)), \theta] = \mathbf{0} \quad (\text{A.16})$$

This assumes that both the neural network's intrinsic approximation error and the error due to using a finite set of collocation points are negligible.

Its practical **approximation** is denoted by  $\overline{W}$ , which is obtained by terminating the optimization once the lower-level loss is within a specified tolerance,  $\mathcal{L}_{\text{LO}}(\theta, \overline{W}) \leq \epsilon$ . We also denote  $u^* = u(\theta, W^*)$  the solution at the ideal optimal weights  $W^*$ .  $\hat{u} = u(\theta, \overline{W})$  is the solution at the approximate weights  $\overline{W}$ .

For BiLO, since the residual-gradient weight  $w_{\text{rgrad}} < 1$  is fixed, we may assume without loss of generality that both the residual loss and the residual-gradient loss are within a specified tolerance  $\epsilon$ :

$$\|\mathcal{F}(\hat{u}, \theta)\| \leq \epsilon \quad (\text{A.17})$$

$$\|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] + F_{\theta}(\hat{u}, \theta)\| \leq \epsilon \quad (\text{A.18})$$

Denote:

$$v := \nabla_{\theta} u(W^*, \theta) + \nabla_W u(W^*, \theta) \nabla_{\theta} W^*(\theta) \quad (\text{A.19})$$

which is the sensitivity of the ideal PDE operator  $u^*$  with respect to  $\theta$

The true hypergradient of the upper level objective  $\mathcal{J}$  is given by

$$g_{\text{true}}(\theta) := \nabla_{\theta} \mathcal{J}[u(\theta, W^*(\theta))] = \mathcal{J}'[u^*]v \quad (\text{A.20})$$

The approximate gradient in BiLO is given by

$$g_a(\theta) := \nabla_{\theta} \mathcal{J}[u(\theta, \overline{W})] = \mathcal{J}'[\hat{u}] \nabla_{\theta} \hat{u} \quad (\text{A.21})$$

We use  $C$  to denote a generic constant that includes the stability constant, lipschitz constant, etc., which may vary from line to line.

**Theorem 2** (Approximate Gradient Error Bound). *Assuming (1) The PDE operator  $\mathcal{F}$  is Lipschitz continuous with respect to  $u$  and  $\theta$ , (2) The PDE operator  $\mathcal{F}(\cdot, \theta)$  is stable; that is, if  $\mathcal{F}(u, \theta) = 0$  and  $\|\mathcal{F}(v, \theta)\| \leq \epsilon$ , then  $\|v - u\| \leq C\epsilon$  for some constant  $C$ . (3) The linearized PDE operator  $\mathbf{L}_u$  is Lipschitz continuous with respect to  $u$  and is stable, (4) The objective functional  $\mathcal{J}$  is Lipschitz continuous and has bounded a derivative. Then*

$$\|g_a - g_{\text{true}}\| = O(\epsilon)$$

**Proof.**

Taking the total derivative of (A.16) with respect to  $\theta$ , we have

$$\mathbf{L}_{u^*}[v] + F_{\theta}(u^*, \theta) = 0 \quad (\text{A.22})$$

Because  $\|\mathcal{F}(\hat{u}, \theta)\| \leq \epsilon$  and  $\mathcal{F}(u^*, \theta) = 0$ , by the stability assumption on  $\mathcal{F}$ , we have  $\|\hat{u} - u^*\| = O(\epsilon)$ . With Lipschitz continuity of  $F_{\theta}$ , we have  $\|\mathcal{F}_{\theta}(\hat{u}, \theta) - F_{\theta}(u^*, \theta)\| = O(\epsilon)$ . From (A.18) and the definition of  $v$ , we have:

$$\|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[v] + \mathcal{F}_{\theta}(\hat{u}, \theta) - \mathcal{F}_{\theta}(u^*, \theta)\| \leq \epsilon \quad (\text{A.23})$$

Since  $\|\mathcal{F}_{\theta}(\hat{u}, \theta) - \mathcal{F}_{\theta}(u^*, \theta)\|$  is  $O(\epsilon)$ , by the triangle inequality, this implies:

$$\|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[v]\| \leq \epsilon + \|\mathcal{F}_{\theta}(\hat{u}, \theta) - \mathcal{F}_{\theta}(u^*, \theta)\| = O(\epsilon) \quad (\text{A.24})$$

By the Lipschitz continuity of  $\mathbf{L}_u$  with respect to  $u$ , we have

$$\|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[\nabla_{\theta} \hat{u}]\| \leq \|\mathbf{L}_{\hat{u}} - \mathbf{L}_{u^*}\| \|\nabla_{\theta} \hat{u}\| \leq C \|\hat{u} - u^*\| = O(\epsilon) \quad (\text{A.25})$$

By (A.22), we have

$$\|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[v]\| = \|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[\nabla_{\theta} \hat{u}] + \mathbf{L}_{u^*}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[v]\| \quad (\text{A.26})$$

By the triangle inequality, we have

$$\|\mathbf{L}_{u^*}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[v]\| \leq \|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[\nabla_{\theta} \hat{u}]\| + \|\mathbf{L}_{\hat{u}}[\nabla_{\theta} \hat{u}] - \mathbf{L}_{u^*}[v]\| = O(\epsilon) \quad (\text{A.27})$$

By the stability of the operator  $\mathbf{L}_{u^*}$ ,

$$\|\nabla_{\theta} \hat{u} - v\| = O(\epsilon) \quad (\text{A.28})$$

By the Lipschitz continuity of  $J'$ , we have

$$\|J'[\bar{u}] - J'[u^*]\| \leq C\|\bar{u} - u^*\| = O(\epsilon) \tag{A.29}$$

By the boundedness of  $J'$  and  $\|\nabla_{\theta}\bar{u}\|$ , we have the difference

$$\begin{aligned} & \|g_a - g_{true}\| \\ &= \|J'[\bar{u}]\nabla_{\theta}\bar{u} - J'[u^*]v\| \\ &\leq \|J'[\bar{u}] - J'[u^*]\| \|\nabla_{\theta}\bar{u}\| + \|J'[u^*]\| \|\nabla_{\theta}\bar{u} - v\| \\ &= O(\epsilon) \end{aligned} \tag{A.30}$$

In summary, the approximation error of the hypergradient is bounded by a constant multiple of  $\epsilon$ , where this constant depends on problem-specific parameters such as the smoothness and stability of the PDE operator  $\mathcal{F}$ , the functional  $J$ , and the parameterization of the solution  $u$ .  $\square$

### Appendix B. Comparison with BPN

BPN [75] is motivated by the same concern about the trade-off between the data loss and the PDE loss in a penalty-like formulation in PINN. BPN follows the PINN framework: the solution of the PDE is represented by a neural network  $u(x; W)$  ( $\theta$  is not input to the network). The definition of data loss and the residual loss is the same as in PINN. However, in BPN, the residual loss is separate from the data loss, leading to the bilevel optimization problem

$$\begin{aligned} & \min_{\theta} \mathcal{L}_{data}^{PINN}(W^*(\theta)) \\ \text{s.t. } & W^*(\theta) = \arg \min_W \mathcal{L}_{res}^{PINN}(W, \theta). \end{aligned}$$

The gradient of the data loss with respect to the PDE parameters is given by the chain rule

$$\frac{d\mathcal{L}_{data}}{d\theta} = \frac{d\mathcal{L}_{data}(W^*(\theta))}{dW} \frac{dW^*(\theta)}{d\theta},$$

where the hypergradient is given by

$$\frac{dW^*(\theta)}{d\theta} = - \left[ \frac{\partial^2 \mathcal{L}_{res}}{\partial W \partial W^T} \right]^{-1} \cdot \frac{\partial^2 \mathcal{L}_{res}}{\partial W \partial \theta^T}.$$

Broyden’s method [90] is used to compute the hyper-gradient, which is based on the low-rank approximation of the inverse Hessian. In BPN, the bilevel optimization problem is solved iteratively. At each step, gradient descent is performed at the lower level for a fixed number of iterations,  $N_f$ . Following this, the hypergradient is computed using Broyden’s method, which requires  $r$  iterations to approximate the inverse vector-Hessian product. This hypergradient is then used to perform a single step of gradient descent at the upper level.




Fig. B.16. Comparison of BPN and BiLO methods. x-axis is the number of lower level optimization steps. Top: Parameter error  $\|\theta - \theta_{GT}\|^2$  versus iterations. Middle: PDE loss  $\log_{10}(\mathcal{L}_{res})$ . Bottom: Data loss  $\mathcal{L}_{data}$ .

The BiLO approach differs significantly. Instead of representing the PDE solution, BiLO represents the local PDE operator, leading to a different lower level problem that includes the residual-gradient loss. This enables direct computation of gradients for  $Ldata$  with respect to  $\theta$ , eliminating the need for specialized algorithms to approximate the hypergradient. This formulation also allows us to perform simultaneous gradient descent at the upper and lower levels, which is more efficient than the iterative approach in BPN. Our method is specialized for PDE-constrained optimization, leveraging the structure of the PDE constraint for efficiency (see the theorem in Appendix A). In contrast, BPN adopts a more general bilevel optimization framework, which, while broadly applicable, does not fully exploit the unique characteristics of PDE problems.

To compare BiLO with BPN, we adopted the problem (B.1) and the setup from [75], using the same residual points (64), neural network architecture (4 hidden layers with 50 units), upper-level optimizer (Adam with learning rate 0.05), lower-level optimizer (Adam with learning rate 0.001), and initial guess ( $\theta_0 = 0, \theta_1 = 1$ ). Both methods included 1000 pretraining steps to approximate the PDE solution at the initial parameters. In BPN, 64 lower iterations are performed for each upper iteration, with 32 Broyden iterations to compute the hypergradient. By contrast, BiLO performs simultaneous gradient descent at the upper and lower levels, where each iteration updates both levels concurrently.

$$\begin{aligned} \min_{\theta_0, \theta_1} J &= \int_0^1 (y - x^2)^2 dx \\ \text{s.t.} \quad \frac{d^2 y}{dx^2} &= 2, \quad y(0) = \theta_0, \quad y(1) = \theta_1 \end{aligned} \tag{B.1}$$

Fig. B.16 presents the loss and the error of the PDE parameters for both methods versus the number of lower-level iterations. BiLO achieves a parameter error below 0.01 in fewer than 80 iterations and just 6.4 seconds, while BPN requires 27 upper iterations (1728 lower iterations) and 231 seconds to reach the same accuracy. While this highlights BiLO’s efficiency, we note that both methods may benefit from further hyperparameter tuning, and the comparison is made under the settings reported in [75].

### C. Training details and additional results

In all the numerical experiments, we use the tanh activation function and 2 hidden layers, each with 128 neurons, for both PINN and BiLO. The collocation points are evenly spaced as a grid in the domain. For all the optimization problems, we use the Adam optimizer with learning rate 0.001 and run a fixed number of steps.

#### C.1. Fisher-KPP equation

Our local operator takes the form

$$u(x, t, D, \rho; W) = u(x, 0) + \mathcal{N}(x, t, D, \rho; W)x(1 - x)t$$

so that the initial condition and the boundary condition are satisfied. Let  $X_r, X_d$  be the spatial coordinates evenly spaced in  $[0, 1]$ , and  $T_r$  be temporal coordinates evenly spaced in  $[0, 1]$ . We set  $\mathcal{T}_{data} = \mathcal{T}_{res} = X_r \times T_r$  and  $|X_r| = |T_r| = 51$ . Both BiLO and PINN are pretrained with the initial guess for 10,000 steps, and fine-tuned for 50,000 steps. In Fig. C.17, we show the training history of the inferred parameters and the inferred parameters during the fine-tuning stage for solving inverse problems using BiLO for the Fisher-KPP equation (Section 3.1.1, with no noise and  $w_{rgrad} = 0.1$ ).

##### C.1.1. Comparison with neural operators

We use the following architecture for the DeepONet [35] in Section C.1.1

$$G_W(D, \rho, \mathbf{x}) = \sum_{i=1}^k b_k(D, \rho)t_k(\mathbf{x}) \tag{C.1}$$

where  $b_k(D, \rho)$  is the  $k$ -th output of the “branch net”, and  $t_k(\mathbf{x})$  is the  $k$ -th output of the “trunk net”. Both the branch net and the trunk net are parameterized by fully neural networks with 2 hidden layers, each with 128 neurons, so that the total number of parameters (46179) are comparable to the network used by BiLO (42051). The weights of the DeepONet are denoted as  $W$ . A final transformation on the output  $G_W$  is used to enforce the boundary condition. We pre-train multiple DeepONets with 10,000 steps using each dataset.

Given a pretrain dataset with collections of  $\{D^j, \rho^j\}$  and their corresponding solutions  $u^j$  for  $j = 1, \dots, m$ , we first train the DeepONet with the following operator data loss:

$$\min_W \sum_{j=1}^m \sum_{\mathbf{x} \in \mathcal{T}_{data}} |G_W(D^j, \rho^j, \mathbf{x}) - u^j(\mathbf{x})|^2 \tag{C.2}$$

where  $\mathcal{T}_{data}$  is the same as those used in the BiLO and PINN. For the inverse problem, we fix the weights  $W$  and treat the  $D$  and  $\rho$  as unknown variables. We minimize the data loss:

$$\min_{D, \rho} \frac{1}{|\mathcal{T}_{data}|} \sum_{\mathbf{x} \in \mathcal{T}_{data}} |G_W(D, \rho, \mathbf{x}) - \hat{u}(\mathbf{x})|^2 \tag{C.3}$$

where  $\hat{u}$  is the noisy data.

## C.2. Variable-diffusion coefficient poisson equation

The local operator takes the form of  $u(x, z; W) = \mathcal{N}_1(x, z; W)x(1-x)$  to enforce the boundary condition, where the fully connected neural network  $\mathcal{N}_1$  has 2 hidden layers, each with 128 neurons. The unknown function is parameterized by  $D(x; V) = \mathcal{N}_2(x, V)x(1-x) + 1$ , where  $\mathcal{N}_2$  has 2 hidden layers, each with 64 neurons. For pre-training, we set  $|\mathcal{T}_{\text{res}}| = |\mathcal{T}_{\text{reg}}| = |\mathcal{T}_{\text{data}}| = 101$ , and train 10,000 steps. For fine-tuning, we set  $|\mathcal{T}_{\text{res}}| = |\mathcal{T}_{\text{reg}}| = 101$  and  $|\mathcal{T}_{\text{data}}| = 51$ , and train 10,000 steps.

### C.2.1. Implementation of the adjoint methods

For the numerical example on learning the variable diffusion coefficient of the Poisson Equation, we implement the adjoint method following [91]. The domain is discretized with uniformly spaced grid points:  $x_i = hi$  for  $i = 0, \dots, n, n+1$ , where  $h$  is the spacing of the grid points and  $n$  is the number of intervals. We use the finite element discretization with linear basis functions  $\phi_i$ . Let  $\mathbf{u}$  be the nodal value of the solution  $u$  at  $x_i$  for  $i = 1, \dots, n$  and similar for  $\mathbf{D}$ . We have  $\mathbf{u}_0 = \mathbf{u}_{n+1} = 0$  and  $\mathbf{D}_0 = \mathbf{D}_{n+1} = 1$ . The stiffness matrix  $A(\mathbf{D})$  is given by

$$A(\mathbf{D})_{ij} = \frac{1}{2} \begin{cases} \mathbf{D}_{i-1} + 2\mathbf{D}_i + \mathbf{D}_{i+1} & \text{if } i = j \\ -(\mathbf{D}_i + \mathbf{D}_j) & \text{if } |i - j| = 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{C.4})$$

The load vector  $\mathbf{f}$  is given by  $\mathbf{f}_i = f(x_i)$ . Suppose the observed data is located at some subset of the grid points of size  $m$ . Then  $\hat{\mathbf{v}}\mathbf{u} = C\mathbf{u} + \eta$ , where  $\eta$  is the noise, and  $C \in \mathbb{R}^{n \times m}$  is the observation operator. After discretization, the minimization problem is

$$\min_{\mathbf{D}} \|\mathbf{C}\mathbf{u} - \hat{\mathbf{u}}\|_2^2 + \frac{w_{\text{reg}}}{2} \sum_{i=1}^N (\mathbf{D}_{i+1} - \mathbf{D}_i)^2$$

$$\text{s.t } A(\mathbf{D})\mathbf{u} = \mathbf{f}$$

The gradient of the loss function with respect to the diffusion coefficient is given by

$$\mathbf{g}_i = \left\langle \frac{\partial A}{\partial \mathbf{D}_i} \mathbf{u}, \mathbf{z} \right\rangle + w_{\text{reg}} (\mathbf{D}_{i+1} - 2\mathbf{D}_i + \mathbf{D}_{i-1})$$

where  $\mathbf{z}$  is the solution of the adjoint equation  $A^T \mathbf{z} = C^T (\mathbf{C}\mathbf{u} - \hat{\mathbf{u}})$ . Gradient descent with step size 0.1 is used to update  $D$ , and is stopped when the norm of the gradient is less than  $10^{-6}$ .

### C.2.2. Comparison with DeepONet

Fig. C.18 shows the samples of  $D(x)$  with various length scale  $l$  and their corresponding solutions  $u$ .

The DeepONet has the following architecture:

$$G_W(\mathbf{D}, \mathbf{x}) = \sum_{i=1}^k b_k(\mathbf{D}) t_k(\mathbf{x}) \quad (\text{C.5})$$

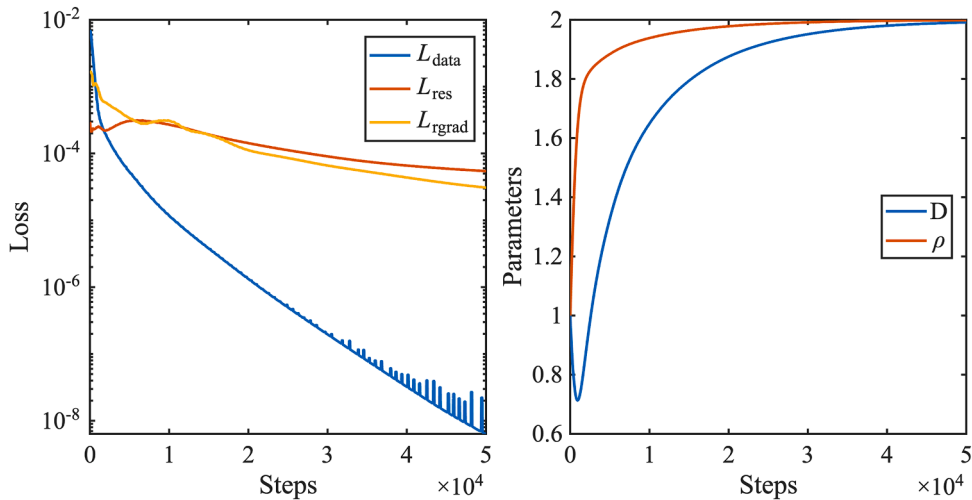


Fig. C.17. Training history of the unweighted losses ( $\mathcal{L}_{\text{res}}$ ,  $\mathcal{L}_{\text{data}}$ , and  $\mathcal{L}_{\text{rgrad}}$ ) and the PDE parameters ( $D$  and  $\rho$ ) during the fine-tuning stage for solving inverse problems using BiLO for the Fisher-KPP equation.



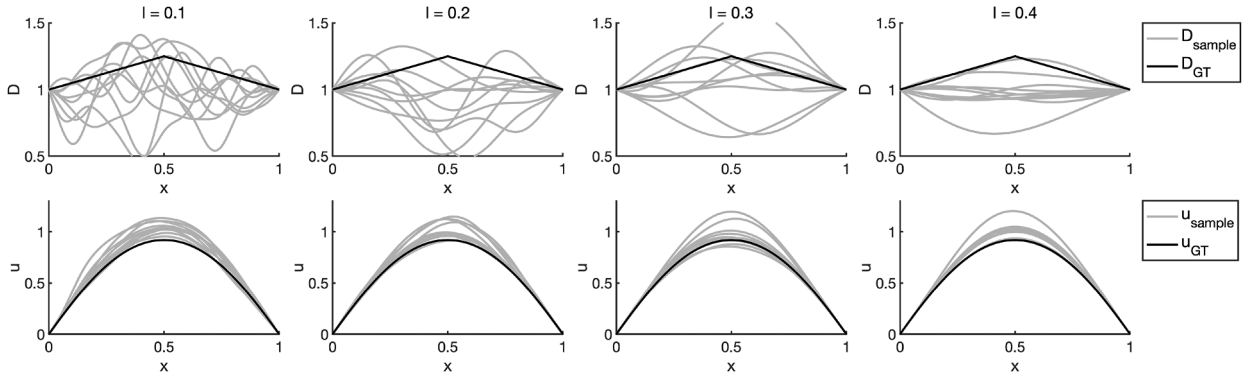


Fig. C.18. Samples (gray lines) of  $D(x)$  with various length scale  $l$  and their corresponding solutions. Black line is the ground truth  $D$  and  $u$ .

where the vector  $\mathbf{D}$  represent the values of  $D(x)$  at the collocation points. A final transformation on the output  $G_W$  is used to enforce the boundary condition. In this experiment, both  $D$  and  $u$  are evaluated at 101 points in  $[0, 1]$ . Let  $x_i$  be the collocation points in  $[0, 1]$  for  $i = 1, \dots, N$ . Let  $\{D^j(x_i), u^j(x_i)\}$  be the samples of  $D$  and the corresponding solutions  $u$  at  $x_i$  for  $j = 1, \dots, m$ . We denote  $\mathbf{D}^j$  as the vector of  $D^j(x_i)$  for  $i = 1, \dots, N$ . In the pre-training step, we solve the following minimization problem

$$\min_W \sum_{j=1}^m \sum_{i=1}^N |G_W(\mathbf{D}^j, x_i) - u^j(x_i)|^2 \tag{C.6}$$

For the inverse problem, we fix the weights  $W$  and treat the  $\mathbf{D}$  as an unknown variable. We minimize the data loss and a finite difference discretization of the regularization term  $|D(x)|^2$ :

$$\min_{\mathbf{D}} \frac{1}{N} \sum_{i=1}^N |G_W(\mathbf{D}, x_i) - \hat{u}(x_i)|^2 + w_{\text{reg}} \sum_{i=0}^N |(\mathbf{D}_{i+1} - \mathbf{D}_i)/h|^2 \tag{C.7}$$

where  $h$  is the spacing of the collocation points,  $\mathbf{D}_0 = \mathbf{D}_N = 1$ . Here we work with the vector  $\mathbf{D}$  for simplicity. Alternatively, we can represent  $D(x)$  as a neural network as in PINN and BiLO experiments.

Table C.2 presents the relative errors in the inferred diffusion coefficient  $D$  and the solution  $u_{\text{NN}}$  across different methods and hyperparameter settings, under a fixed noise level  $\sigma = 0.01$ . For Neural Operator (NO), we vary the length scale  $l = 0.2, 0.3, 0.4$  of the GRF in the pretrain dataset and regularization weight  $w_{\text{reg}} = 10^{-5}, 10^{-4}, 10^{-3}$ . The best performance is achieved with  $l = 0.3$  and  $w_{\text{reg}} = 10^{-4}$ . For PINN, BiLO, and the Adjoint method, we consider  $w_{\text{data}} = 10, 100, 1000$ , and evaluate each under  $w_{\text{reg}} = 10^{-4}, 10^{-3}, 10^{-2}$ . For PINN, we observe that both smaller  $w_{\text{data}}$  and larger  $w_{\text{reg}}$  tend to promote smoother solutions, necessitating tuning of both hyperparameters. The combination  $w_{\text{data}} = 100$  and  $w_{\text{reg}} = 10^{-3}$  gives the most accurate reconstruction among PINN variants. The BiLO method achieves the best overall performance with  $w_{\text{reg}} = 10^{-3}$ , outperforming both NO and PINN. For reference, the adjoint method solves the inverse problem numerically on a fine grid and is treated as the ground truth; hence, no relative error in  $u_{\text{NN}}$  is reported.

### C.3. Infer the initial condition of a heat equation

We can represent the unknown function  $f(x; V) = s(\mathcal{N}(x; V))x(1-x)$ , where  $\mathcal{N}_f$  is a fully connected neural network with 2 hidden layers and width 64, and  $s$  is the softplus activation function (i.e.,  $s(x) = \log(1 + \exp(x))$ ). The transformation ensures that the initial condition satisfies the boundary condition and is non-negative. For BiLO, the neural network is represented as  $u(x, t, z) = N_u(x, t, z; W)x(1-x)t + z$ , where  $N_u$  is a fully connected neural network with 2 hidden layers and width 128. For the PINN, we have  $u(x, t; W, V) = N_u(x, t; W)x(1-x)t + f(x; V)$ . These transformations ensure that the networks satisfy the boundary and initial conditions.

Let  $X_r, X_d$  be spatial coordinates evenly spaced in  $[0, 1]$  and  $T_r$  be temporal coordinates evenly spaced in  $[0, 1]$  (both including the boundary). We set  $\mathcal{T}_{\text{res}} = X_r \times T_r$  and  $|X_r| = |T_r| = 51$ . That is, the residual collocation points is a uniform grid in space and time. We set  $\mathcal{T}_{\text{data}} = X_d \times \{1\}$  and  $|X_d| = 11$ . That is, the data collocation points is a uniform grid in space at the final time  $t = 1$ . We set the collocation point for the regularization loss of the unknown function  $\mathcal{T}_{\text{reg}}$  to be 101 evenly spaced points in the spatial domain. To

**Table C.2**

Relative errors in the inferred diffusion coefficient  $D$  and neural network solution  $u_{\text{NN}}$  for different methods under noise level  $\sigma = 0.01$ . We vary the regularization strength  $w_{\text{reg}}$  across all methods. For NO, we additionally vary the GRF length scale  $l$  used in the pretraining data. For PINN, we vary the number of training data  $w_{\text{data}} = 10, 100, 1000$ . The adjoint method serves as the reference solution and does not report error in  $u_{\text{NN}}$ .

Method	$w_{\text{reg}}$	Rel Err. $D$	Rel Err. $u_{\text{NN}}$
BiLO	$10^{-4}$	$2.70 \times 10^{-2} \pm 6.74 \times 10^{-3}$	$7.33 \times 10^{-4} \pm 3.00 \times 10^{-4}$
	$10^{-3}$	$1.87 \times 10^{-2} \pm 8.25 \times 10^{-3}$	$1.62 \times 10^{-3} \pm 9.87 \times 10^{-4}$
	$10^{-2}$	$4.81 \times 10^{-2} \pm 2.13 \times 10^{-3}$	$7.32 \times 10^{-4} \pm 7.29 \times 10^{-4}$
PINN(10)	$10^{-4}$	$2.94 \times 10^{-2} \pm 5.51 \times 10^{-3}$	$4.76 \times 10^{-4} \pm 1.18 \times 10^{-4}$
	$10^{-3}$	$3.20 \times 10^{-2} \pm 3.08 \times 10^{-3}$	$7.25 \times 10^{-4} \pm 8.30 \times 10^{-4}$
	$10^{-2}$	$4.84 \times 10^{-2} \pm 3.38 \times 10^{-3}$	$5.20 \times 10^{-4} \pm 3.67 \times 10^{-4}$
PINN(100)	$10^{-4}$	$3.33 \times 10^{-2} \pm 1.32 \times 10^{-2}$	$2.48 \times 10^{-4} \pm 1.89 \times 10^{-4}$
	$10^{-3}$	$2.88 \times 10^{-2} \pm 8.10 \times 10^{-3}$	$2.64 \times 10^{-4} \pm 2.13 \times 10^{-4}$
	$10^{-2}$	$3.04 \times 10^{-2} \pm 7.48 \times 10^{-3}$	$7.05 \times 10^{-4} \pm 4.80 \times 10^{-4}$
PINN(1000)	$10^{-4}$	$6.99 \times 10^{-2} \pm 3.66 \times 10^{-2}$	$1.03 \times 10^{-3} \pm 4.51 \times 10^{-4}$
	$10^{-3}$	$5.71 \times 10^{-2} \pm 2.57 \times 10^{-2}$	$7.41 \times 10^{-4} \pm 5.14 \times 10^{-4}$
	$10^{-2}$	$3.64 \times 10^{-2} \pm 1.32 \times 10^{-2}$	$8.73 \times 10^{-4} \pm 5.71 \times 10^{-4}$
Adjoint	$10^{-4}$	$5.23 \times 10^{-2} \pm 1.36 \times 10^{-2}$	-
	$10^{-3}$	$3.04 \times 10^{-2} \pm 9.77 \times 10^{-3}$	-
	$10^{-2}$	$4.27 \times 10^{-2} \pm 4.35 \times 10^{-3}$	-
NO (0.2)	$10^{-5}$	$4.36 \times 10^{-2} \pm 1.01 \times 10^{-2}$	$5.75 \times 10^{-3} \pm 1.05 \times 10^{-3}$
	$10^{-4}$	$4.12 \times 10^{-2} \pm 8.69 \times 10^{-3}$	$5.32 \times 10^{-3} \pm 9.58 \times 10^{-4}$
	$10^{-3}$	$5.63 \times 10^{-2} \pm 2.89 \times 10^{-3}$	$6.04 \times 10^{-3} \pm 4.12 \times 10^{-4}$
NO (0.3)	$10^{-5}$	$3.13 \times 10^{-2} \pm 1.05 \times 10^{-2}$	$1.69 \times 10^{-2} \pm 1.20 \times 10^{-3}$
	$10^{-4}$	$3.09 \times 10^{-2} \pm 9.68 \times 10^{-3}$	$6.56 \times 10^{-3} \pm 1.59 \times 10^{-3}$
	$10^{-3}$	$5.63 \times 10^{-2} \pm 3.27 \times 10^{-3}$	$8.73 \times 10^{-3} \pm 9.74 \times 10^{-4}$
NO (0.4)	$10^{-5}$	$4.61 \times 10^{-2} \pm 8.54 \times 10^{-3}$	$2.72 \times 10^{-2} \pm 4.34 \times 10^{-3}$
	$10^{-4}$	$3.39 \times 10^{-2} \pm 9.10 \times 10^{-3}$	$8.85 \times 10^{-3} \pm 3.69 \times 10^{-3}$
	$10^{-3}$	$5.96 \times 10^{-2} \pm 3.77 \times 10^{-3}$	$1.28 \times 10^{-2} \pm 1.73 \times 10^{-3}$

evaluate the performance of the inferred initial condition  $f$ , we use the relative  $L^2$  norm of inferred initial condition and the ground truth initial condition, which are computed using 1001 evenly spaced points in the spatial domain.

#### C.4. Elliptic equation with singular forcing

For BiLO, the neural network is represented as

$$u(x, \phi, \mu, \lambda; W) = N_u(x, \phi, \mu, \lambda; W)x(1-x),$$

where  $N_u$  is a fully connected neural network with 2 hidden layers and width 128. The FNO has 4 layer, 32 modes, and 32 channels. The FNO has 82,785 parameters in total, the BiLO have 82,691 parameters.

#### C.5. 1D burgers' equation

Fig. C.19 shows an additional example of the experiment in Section 3.5 in the main text.

#### C.6. 2D darcy flow

The setup of this experiment is similar to the steady state Darcy flow inverse problem in Li et al. [33]. We pretrain the BiLO with  $A_0(\mathbf{x})$  and its corresponding solution  $u_0(\mathbf{x})$  for 10,000 steps. And we fine-tune the BiLO for 5000 steps using  $u_{\text{GT}}(\mathbf{x})$  to infer  $A_{\text{GT}}$ .

The unknown function is represented as  $A(\mathbf{x}; V) = s(\mathcal{N}(\mathbf{x}; V)) \times 9 + 3$ , where  $N_f$  is a fully connected neural network with 2 hidden layers and width 64, and  $s$  is the logistic function. The transformation is a smoothed approximation of the piece-wise constant function. For BiLO, the neural network is represented as  $u(\mathbf{x}, z) = N_u(\mathbf{x}, z; W)\mathbf{x}_1(1-\mathbf{x}_1)\mathbf{x}_2(1-\mathbf{x}_2)$ , where  $N_u$  is a fully connected neural network with 2 hidden layers and width 128, and  $z$  is our auxiliary variable such that  $z = A(\mathbf{x}; V)$ . Fig. C.20 shows an additional example of the experiment.

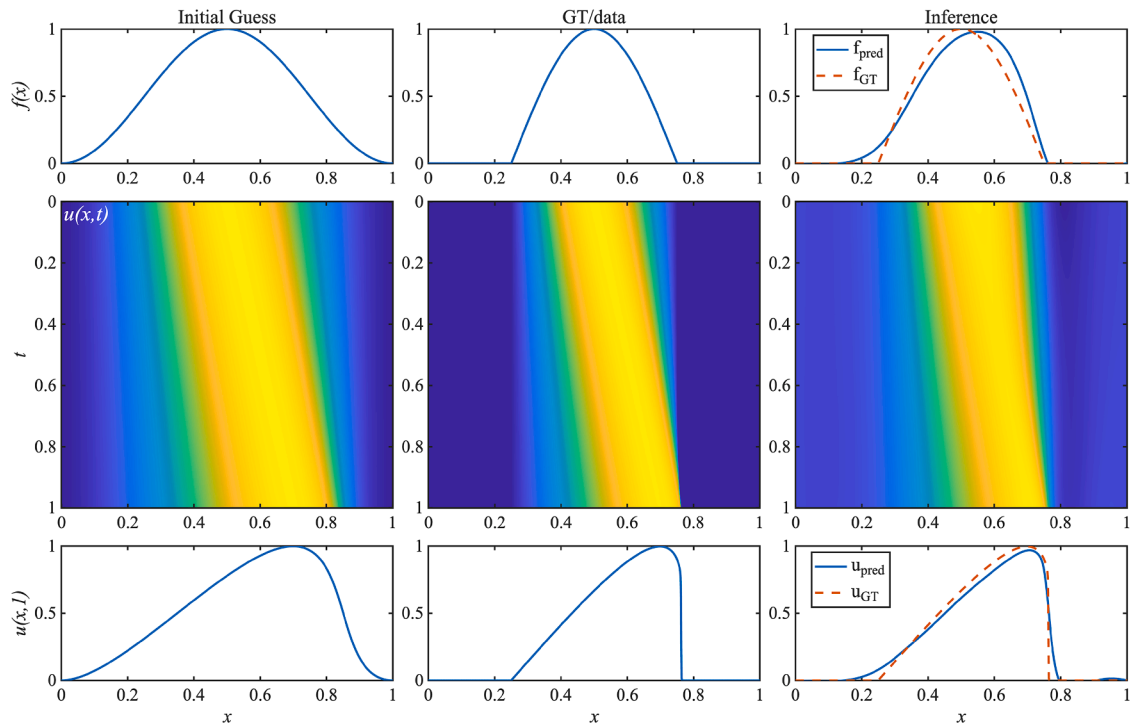


Fig. C.19. Example 2 of inferring the initial condition of the Burgers' equation. The initial guess is used to pre-train the network. The solution at  $t = 1$  of the GT is the data for inference. First column: initial guess, second column: ground truth, third column: inferred initial condition. First row: initial condition, second row: solution  $u(x, t)$ , third row: solution  $u(x, 1)$ .

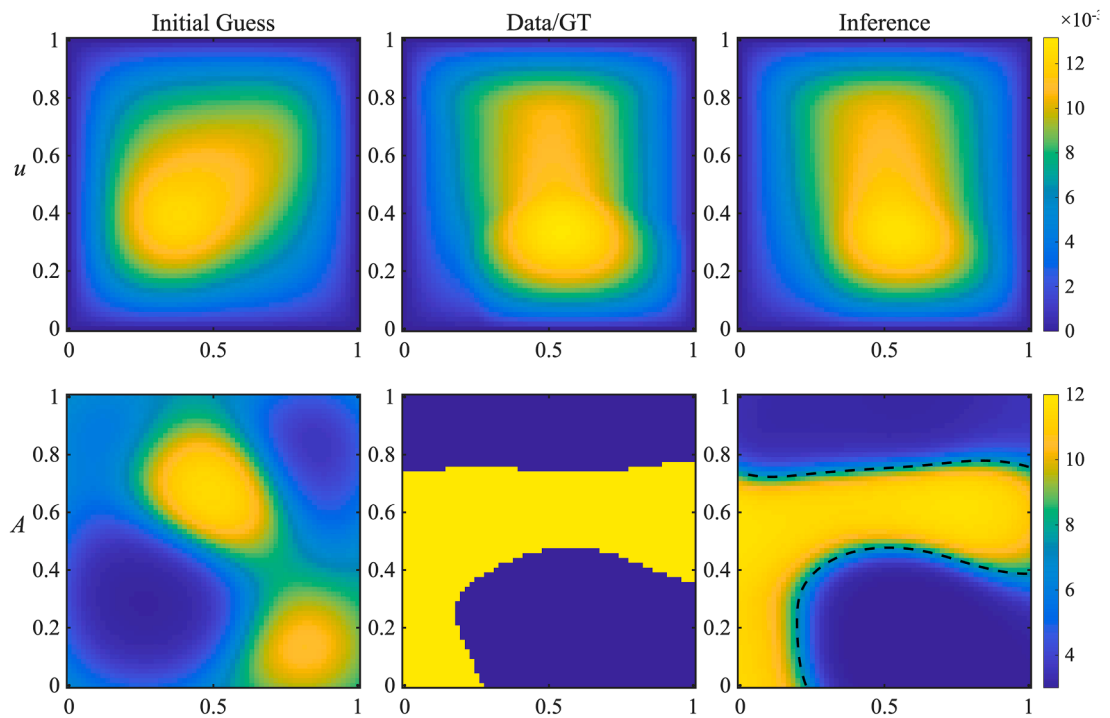


Fig. C.20. Example 2 of inferring the variable diffusion coefficient. The relative l2 error of  $u_{NN}$  against  $u_{GT}$  is 1.7%. The thresholded (at the dashed line) inferred diffusion coefficient has classification accuracy of 96.7% .

### D. Sensitivity to hyperparameters

We evaluate the sensitivity of BiLO to various hyperparameter choices in the noise-free FKPP inverse problem and compare it with PINN using the best-performing  $w_{\text{data}} = 10$ .

In both BiLO and PINN frameworks, it is possible to specify distinct learning rates for the PDE parameters  $\theta$  and the neural network weights  $W$ , denoted as  $\alpha_\theta$  and  $\alpha_W$ , respectively, as in Eq. (19) and (21). In our experiments, we fixed  $\alpha_W = 10^{-3}$ , which is a common choice for training neural networks, and varied the learning rate  $\alpha_\theta$  for both methods. As shown in Fig. D.21, BiLO consistently achieves more accurate parameter recovery than PINN for any fixed choice of  $\alpha_\theta$ . Moreover, BiLO remains robust across a wide range of  $\alpha_\theta$  values, including large values up to 0.1, and converges significantly faster without sacrificing accuracy. We attribute BiLO's superior performance to its more accurate descent direction for  $\theta$ . Although we use  $\alpha_\theta = 10^{-3}$  throughout the main text for consistent comparison, even better performance could likely be obtained through cross-validation over  $\alpha_\theta$ .




Fig. D.21. Trajectory of PDE parameters  $D$  and  $\rho$  with different  $\alpha_\theta$ . BiLO is robust to the choice of  $\alpha_\theta$  and consistently outperforms.

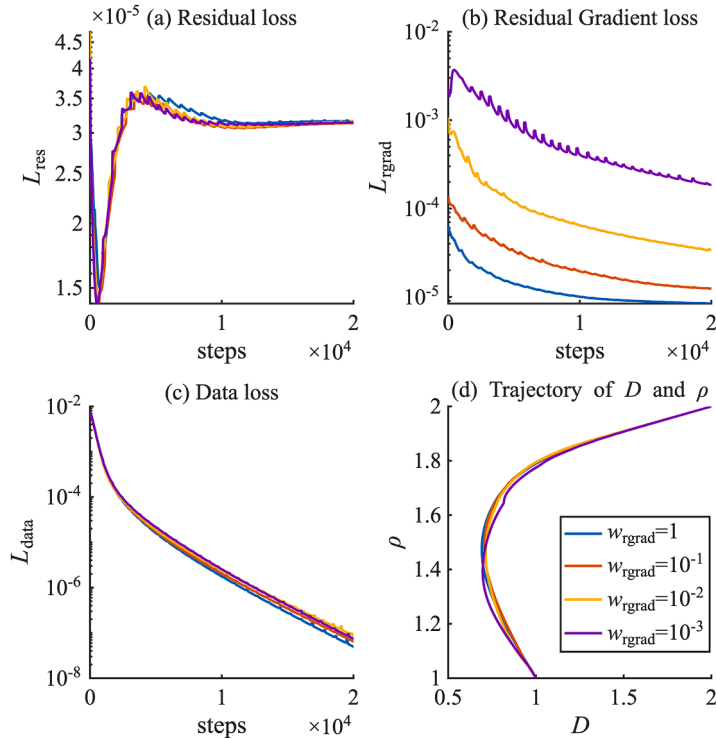


Fig. D.22. History of the losses-(a)  $\mathcal{L}_{\text{res}}$ , (b)  $\mathcal{L}_{\text{data}}$ , and (c)  $\mathcal{L}_{\text{rgrad}}$ -and (d) the trajectories of the parameters  $D$  and  $\rho$  for different values of the residual-gradient weight  $w_{\text{rgrad}}$ . The initial guess of  $(D, \rho)$  is  $(1, 1)$ , and the ground truth is  $(2, 2)$ .

Fig. D.22 illustrates that BiLO is robust with respect to the weight of the residual-gradient loss  $w_{\text{rgrad}}$ . Across different values of  $w_{\text{rgrad}}$ , the trajectories of the PDE parameters remain similar. As  $w_{\text{rgrad}}$  increases, the residual-gradient loss  $\mathcal{L}_{\text{rgrad}}$  decreases, but the residual loss  $\mathcal{L}_{\text{res}}$  remains small, confirming that the learned parameters still accurately solve the PDE. This suggests that  $\mathcal{L}_{\text{rgrad}}$  and  $\mathcal{L}_{\text{res}}$  exhibit a tradeoff distinct from that between  $\mathcal{L}_{\text{data}}$  and  $\mathcal{L}_{\text{res}}$ : while minimizing  $\mathcal{L}_{\text{data}}$  alone may lead to functions that do not satisfy the PDE, both  $\mathcal{L}_{\text{rgrad}}$  and  $\mathcal{L}_{\text{res}}$  can, in principle, be minimized simultaneously if the local operator can be sufficiently approximated.

**E. Computational cost**

Compared with PINN, BiLO involve computing a higher order derivative term in the residual-gradient loss. This increases the memory cost and computation time per step. In Table E.3, we show the seconds-per-step and the maximum memory allocation of 1 run of BiLO and PINN for the various problems. The seconds per step is computed by total training time divided by the number of steps. The maximum memory allocation is the peak memory usage during the training. For all the experiments, we use Quadro RTX 8000 GPU. We note that the measured seconds-per-step is not subject to rigorous control as the GPU is shared with other users and many runs are performed simultaneously.

While BiLO requires more computation per step, it achieves higher accuracy in less total time, as shown in Fig. E.23, where we plot the trajectory of the PDE parameters with respect to wall time in seconds. We attribute this to the more accurate descent direction of the PDE parameters, which leads to faster convergence.




Fig. E.23. Trajectory of PDE parameters  $D$  and  $\rho$  during the fine-tuning stage for solving inverse problems using BiLO. The left panel shows the trajectory of  $D$  and the right panel shows the trajectory of  $\rho$ . The x axis is the wall time in seconds. While BiLO takes more time per step than PINN, it converges faster to the optimal parameters.

**Table E.3**

Comparison of BiLO and PINN in terms of wall time per training step (in seconds) and maximum memory usage (in MB) across various PDE problems. Ratio is computed as BiLO / PINN. We note that the measured seconds-per-step is not subject to rigorous control as the GPU (Quadro RTX 8000) is shared.

Problem	Metric	BiLO	PINN	Ratio
Fisher-KPP	sec/step	0.11	0.06	1.69
	max-mem	200	65.2	3.07
Singular	sec/step	0.21	0.11	2.00
	max-mem	67.4	44.9	1.50
Poisson	sec/step	0.11	0.09	1.15
	max-mem	23.2	20.2	1.15
Burgers	sec/step	0.17	0.08	2.24
	max-mem	122	73.9	1.65
Heat	sec/step	0.11	0.07	1.57
	max-mem	211	109	1.93
Darcy	sec/step	0.13	0.09	1.55
	max-mem	418	152	2.75
GBM	sec/step	1.38	0.43	3.19
	max-mem	40,313	10,376	3.89

## References

- [1] C. Deng, S. Feng, H. Wang, X. Zhang, P. Jin, Y. Feng, Q. Zeng, Y. Chen, Y. Lin, et al., OpenFWI: Large-Scale Multi-Structural Benchmark Datasets for Seismic Full Waveform Inversion, 2023, 2111.02926 <https://doi.org/10.48550/arXiv.2111.02926>.
- [2] J. Martin, L.C. Wilcox, C. Burstedde, O. Ghattas, et al., A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion, *SIAM J. Scient. Comput.* 34 (3) (2012) A1460–A1487. <https://doi.org/10.1137/110845598>
- [3] Y. Yang, A.F. Gao, J.C. Castellanos, Z.E. Ross, K. Azizzadenesheli, R.W. Clayton, et al., Seismic wave propagation and inversion with neural operators, (*The Seismic Record* 1 (3) (2021) 126–134. <https://doi.org/10.1785/0320210026>
- [4] G. Uhlmann, Electrical impedance tomography and calderón's problem, *Inverse Probl.* 25 (12) (2009) 123011. <https://doi.org/10.1088/0266-5611/25/12/123011>
- [5] R. Molinaro, Y. Yang, B. Engquist, S. Mishra, et al., Neural Inverse Operators for Solving PDE Inverse Problems, 2023, 2301.11167 <https://doi.org/10.48550/arXiv.2301.11167>.
- [6] J. Lipkova, P. Angelikopoulos, S. Wu, E. Alberts, B. Wiestler, C. Diehl, C. Preibisch, T. Pyka, S.E. Combs, P. Hadjidoukas, K. Van Leemput, P. Koumoutsakos, J. Lowengrub, B. Menze, et al., Personalized radiotherapy design for glioblastoma: integrating mathematical tumor models, multimodal scans, and Bayesian inference, *IEEE Trans. Med. Imaging* 38 (8) (2019) 1875–1884. <https://doi.org/10.1109/TMI.2019.2902044>
- [7] R.Z. Zhang, I. Ezhov, M. Balcerak, A. Zhu, B. Wiestler, B. Menze, J.S. Lowengrub, Personalized predictions of glioblastoma infiltration: mathematical models, physics-informed neural networks and multimodal scans, *Med. Image Anal.* 101 (2025) 103423. <https://doi.org/10.1016/j.media.2024.103423>
- [8] A. Schäfer, M. Peirlinck, K. Linka, E. Kuhl, et al., Bayesian physics-based modeling of tau propagation in Alzheimer's disease, *Front. Physiol.* 12 (2021) 702975. <https://doi.org/10.3389/fphys.2021.702975>
- [9] S. Subramanian, A. Ghafouri, K.M. Scheufele, N. Himthani, C. Davatzikos, G. Biros, et al., Ensemble inversion for brain tumor growth models with mass effect, *IEEE Trans. Med. Imaging* 42 (4) (2023) 982–995. <https://doi.org/10.1109/TMI.2022.3221913>
- [10] K. Sen Mrinal, P.L. Stoffa, *Global Optimization Methods in Geophysical Inversion*, Cambridge University Press, Cambridge, Cambridge, 2013. <https://doi.org/10.1017/CBO9780511997570>
- [11] M. Hinze, R. Pinnau, M. Ulbrich, S. Ulbrich, *Optimization with PDE Constraints*, Springer Science & Business Media, 2008.
- [12] A.M. Stuart, Inverse problems: a Bayesian perspective, *Acta Numerica* 19 (2010) 451–559. <https://doi.org/10.1017/S0962492910000061>
- [13] R.Z. Zhang, C.E. Miles, X. Xie, J.S. Lowengrub, et al., BiLO: Bilevel Local Operator Learning for PDE Inverse Problems. Part II: Efficient Uncertainty Quantification with Low-Rank Adaptation, 2025, 2507.17019 <https://doi.org/10.48550/arXiv.2507.17019>.
- [14] R.E. Plessix, A review of the adjoint-state method for computing the gradient of a functional with geophysical applications, *Geophys. J. Int.* 167 (2) (2006) 495–503. <https://doi.org/10.1111/j.1365-246X.2006.02978.x>
- [15] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, et al., Physics-informed machine learning, *Nature Rev. Phys.* 3 (6) (2021) 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
- [16] M. Raissi, P. Perdikaris, G.E. Karniadakis, et al., Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* 378 (2019) 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [17] A.D. Jagtap, D. Mitsotakis, G.E. Karniadakis, et al., Deep learning of inverse water waves problems using multi-fidelity data: application to Serre–green–naghdi equations, *Ocean Eng.* 248 (2022) 110775. <https://doi.org/10.1016/j.oceaneng.2022.110775>
- [18] A.D. Jagtap, Z. Mao, N. Adams, G.E. Karniadakis, et al., Physics-informed neural networks for inverse problems in supersonic flows, *J. Comput. Phys.* 466 (2022) 111402. <https://doi.org/10.1016/j.jcp.2022.111402>
- [19] Y. Chen, L. Lu, G.E. Karniadakis, L.D. Negro, et al., Physics-informed neural networks for inverse problems in nano-optics and metamaterials, *Opt. Express* 28 (8) (2020) 11618–11633. <https://doi.org/10.1364/OE.384875>
- [20] L. Yang, X. Meng, G.E. Karniadakis, et al., B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data, *J. Comput. Phys.* 425 (2021) 109913. <https://doi.org/10.1016/j.jcp.2020.109913>
- [21] T. Kapoor, H. Wang, A. Nunez, R. Dollevoet, et al., Physics-informed neural networks for solving forward and inverse problems in complex beam systems, *IEEE Trans. Neural Netw. Learn. Syst.* (2024) 1–15. 2303.01055 <https://doi.org/10.1109/TNNLS.2023.3310585>
- [22] Y. Zong, Q. He, A.M. Tartakovsky, et al., Improved training of physics-informed neural networks for parabolic differential equations with sharply perturbed initial conditions, *Comput. Methods Appl. Mech. Eng.* 414 (2023) 116125. <https://doi.org/10.1016/j.cma.2023.116125>
- [23] P. Karnakov, S. Litvinov, P. Koumoutsakos, et al., Optimizing a Discrete Loss (ODL) to Solve Forward and Inverse Problems for Partial Differential Equations Using Machine Learning Tools, 2022, 2205.04611 <https://doi.org/10.48550/arXiv.2205.04611>.
- [24] M. Balcerak, J. Weidner, P. Karnakov, I. Ezhov, S. Litvinov, P. Koumoutsakos, R.Z. Zhang, J.S. Lowengrub, B. Wiestler, B. Menze, et al., Individualizing Glioma Radiotherapy Planning by Optimization of Data and Physics-Informed Discrete Loss, 2024, 2312.05063.
- [25] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, A. Anandkumar, et al., Neural Operator: Learning Maps Between Function Spaces, 2022, 2108.08481 <https://doi.org/10.48550/arXiv.2108.08481>.
- [26] O.D. Akyildiz, M. Girolami, A.M. Stuart, A. Vadeboncoeur, et al., Efficient prior calibration from indirect data, *SIAM J. Scient. Comput.* 47 (4) (2025) C932–C958. <https://doi.org/10.1137/24M166485X>
- [27] N.H. Nelsen, Y. Yang, Operator Learning Meets Inverse Problems: A Probabilistic Perspective, 2025, (n.d.). 2508.20207 <https://doi.org/10.48550/arXiv.2508.20207>.
- [28] T. Zhou, X. Wan, D.Z. Huang, Z. Li, Z. Peng, A. Anandkumar, J.F. Brady, P.W. Sternberg, C. Daraio, et al., AI-aided geometric design of anti-infection catheters, *Sci. Adv.* 10 (1) (2024) eadj1741. <https://doi.org/10.1126/sciadv.adj1741>
- [29] J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath, A. Anandkumar, et al., FourCastNet: A Global Data-driven High-resolution Weather Model Using Adaptive Fourier Neural Operators, 2022, 2202.11214 <https://doi.org/10.48550/arXiv.2202.11214>.
- [30] L. Lu, R. Pestourie, S.G. Johnson, G. Romano, et al., Multifidelity deep neural operators for efficient learning of partial differential equations with application to fast inverse design of nanoscale heat transport, *Phys. Rev. Res.* 4 (2) (2022) 023210. <https://doi.org/10.1103/PhysRevResearch.4.023210>
- [31] S. Mao, R. Dong, L. Lu, K.M. Yi, S. Wang, P. Perdikaris, et al., PPDONet: deep operator networks for fast prediction of steady-state solutions in disk–planet systems, *Astrophys. J. Lett.* 950 (2) (2023) L12. <https://doi.org/10.3847/2041-8213/acd77f>
- [32] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, A. Anandkumar, et al., Fourier Neural Operator for Parametric Partial Differential Equations, 2021, 2010.08895 <https://doi.org/10.48550/arXiv.2010.08895>.
- [33] Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Azizzadenesheli, A. Anandkumar, et al., Physics-informed neural operator for learning partial differential equations, *ACM / IMS J. Data Sci.* 1 (3) (2024) 9:1–9:27. <https://doi.org/10.1145/3648506>
- [34] C. White, J. Berner, J. Kossaifi, M. Elleithy, D. Pitt, D. Leibovici, Z. Li, K. Azizzadenesheli, A. Anandkumar, et al., Physics-informed neural operators with exact differentiation on arbitrary geometries, in: *The Symbiosis of Deep Learning and Differential Equations III*, 2023.
- [35] L. Lu, P. Jin, G. Pang, Z. Zhang, G.E. Karniadakis, et al., Learning nonlinear operators via deepONet based on the universal approximation theorem of operators, *Nature Mach. Intell.* 3 (3) (2021) 218–229. <https://doi.org/10.1038/s42256-021-00302-5>
- [36] S. Wang, H. Wang, P. Perdikaris, et al., Learning the solution operator of parametric partial differential equations with physics-informed DeepONets, *Sci. Adv.* 7 (40) (2021) eabi8605. <https://doi.org/10.1126/sciadv.abi8605>
- [37] L. Yang, S. Liu, T. Meng, S.J. Osher, et al., In-context operator learning with data prompts for differential equation problems, *Proc. Nat. Acad. Sci.* 120 (39) (2023) e2310142120. <https://doi.org/10.1073/pnas.2310142120>
- [38] T. O'Leary-Roseberry, P. Chen, U. Villa, O. Ghattas, et al., Derivative-informed neural operator: an efficient framework for high-dimensional parametric derivative learning, *J. Comput. Phys.* 496 (2024) 112555. <https://doi.org/10.1016/j.jcp.2023.112555>

- [39] D. Nganyu Tanyu, J. Ning, T. Freudenberg, N. Heilenkötter, A. Rademacher, U. Iben, P. Maass, et al., Deep learning methods for partial differential equations and related parameter identification problems, *Inverse Probl.* 39 (10) (2023) 103001. <https://doi.org/10.1088/1361-6420/ace9d4>
- [40] L. Herrmann, S. Kollmannsberger, Deep learning in computational mechanics: a review, *Comput. Mech.* (2024). <https://doi.org/10.1007/s00466-023-02434-4>
- [41] S.L. Brunton, J.N. Kutz, *Machine Learning for Partial Differential Equations*, 2023, 2303.17078 <https://doi.org/10.48550/arXiv.2303.17078>.
- [42] S. Dong, N. Ni, A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks, *J. Comput. Phys.* 435 (2021) 110242. <https://doi.org/10.1016/j.jcp.2021.110242>
- [43] L. Lu, R. Pestourie, W. Yao, Z. Wang, F. Verdugo, S.G. Johnson, et al., Physics-informed neural networks with hard constraints for inverse design, *SIAM J. Scient. Comput.* 43 (6) (2021) B1105–B1132. <https://doi.org/10.1137/21M1397908>
- [44] N. Sukumar, A. Srivastava, Exact imposition of boundary conditions with distance functions in physics-informed deep neural networks, *Comput. Methods Appl. Mech. Eng.* 389 (2022) 114333. <https://doi.org/10.1016/j.cma.2021.114333>
- [45] A.S. Krishnapriyan, A. Gholami, S. Zhe, R.M. Kirby, M.W. Mahoney, et al., Characterizing Possible Failure Modes in Physics-Informed Neural Networks, 2021, 2109.01050 <https://doi.org/10.48550/arXiv.2109.01050>.
- [46] K. He, X. Zhang, S. Ren, J. Sun, et al., Deep Residual Learning for Image Recognition, 2015, (n.d.). 1512.03385 <https://doi.org/10.48550/arXiv.1512.03385>.
- [47] S. Wang, H. Wang, P. Perdikaris, et al., On the eigenvector bias of Fourier feature networks: from regression to solving multi-Scale PDEs with physics-informed neural networks, *Comput. Methods Appl. Mech. Eng.* 384 (2021) 113938. <https://doi.org/10.1016/j.cma.2021.113938>
- [48] S. Wang, Y. Teng, P. Perdikaris, et al., Understanding and mitigating gradient flow pathologies in physics-informed neural networks, *SIAM J. Scient. Comput.* 43 (5) (2021) A3055–A3081. <https://doi.org/10.1137/20M1318043>
- [49] Y. Zhang, P. Khanduri, I. Tsaknakis, Y. Yao, M. Hong, S. Liu, et al., An Introduction to Bi-level Optimization: Foundations and Applications in Signal Processing and Machine Learning, 2023, 2308.00788 <https://doi.org/10.48550/arXiv.2308.00788>.
- [50] P. Khanduri, I. Tsaknakis, Y. Zhang, J. Liu, S. Liu, J. Zhang, M. Hong, et al., Linearly constrained bilevel optimization: a smoothed implicit gradient approach, in: *Proceedings of the 40th International Conference on Machine Learning*, PMLR, 2023, pp. 16291–16325.
- [51] M. Ye, B. Liu, S. Wright, P. Stone, Q. Liu, et al., BOME! Bilevel Optimization Made Easy: A Simple First-Order Approach, 2022, 2209.08709 <https://doi.org/10.48550/arXiv.2209.08709>.
- [52] H. Shen, Q. Xiao, T. Chen, et al., On Penalty-based Bilevel Gradient Descent Method, 2023, 2302.05185 <https://doi.org/10.48550/arXiv.2302.05185>.
- [53] A. Shaban, C.-A. Cheng, N. Hatch, B. Boots, et al., Truncated back-propagation for bilevel optimization, in: *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, PMLR, 2019, pp. 1723–1732.
- [54] M. Hong, H.-T. Wai, Z. Wang, Z. Yang, et al., A Two-Timescale Framework for Bilevel Optimization: Complexity Analysis and Application to Actor-Critic, 2022, 2007.05170 <https://doi.org/10.48550/arXiv.2007.05170>.
- [55] N.H. Nelsen, H. Owhadi, A.M. Stuart, X. Yang, Z. Zou, et al., Bilevel Optimization for Learning Hyperparameters: Application to Solving PDEs and Inverse Problems with Gaussian Processes, 2025, 2510.05568 <https://doi.org/10.48550/arXiv.2510.05568>.
- [56] S. Basir, I. Senocak, Critical Investigation of Failure Modes in Physics-Informed Neural Networks, 2022, 2206.09961 <https://doi.org/10.48550/arXiv.2206.09961>.
- [57] D.P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, 2017, 1412.6980 <https://doi.org/10.48550/arXiv.1412.6980>.
- [58] S. Kaltenbach, P. Perdikaris, P.-S. Koutsourelakis, et al., Semi-Supervised Invertible Neural Operators for Bayesian Inverse Problems, 2023, 2209.02772 <https://doi.org/10.48550/arXiv.2209.02772>.
- [59] Y. Yang, A.F. Gao, K. Azizzadenesheli, R.W. Clayton, Z.E. Ross, et al., Rapid seismic waveform modeling and inversion with neural operators, *IEEE Trans. Geosci. Remote Sens.* 61 (2023) 1–12. <https://doi.org/10.1109/TGRS.2023.3264210>
- [60] M.V. de Hoop, D.Z. Huang, E. Qian, A.M. Stuart, et al., The Cost-Accuracy Trade-Off In Operator Learning With Neural Networks, 2022, 2203.13181 <https://doi.org/10.48550/arXiv.2203.13181>.
- [61] L. Lu, X. Meng, Z. Mao, G.E. Karniadakis, et al., DeepXDE: a deep learning library for solving differential equations, *SIAM Rev.* 63 (1) (2021) 208–228. <https://doi.org/10.1137/19M1274067>
- [62] M.A. Nabian, R.J. Gladstone, H. Meidani, et al., Efficient training of physics-informed neural networks via importance sampling, *Comput.-Aided Civ. Infrastruct. Eng.* 36 (8) (2021) 962–977. <https://doi.org/10.1111/mice.12685>
- [63] C. Wu, M. Zhu, Q. Tan, Y. Kartha, L. Lu, et al., A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks, *Comput. Methods Appl. Mech. Eng.* 403 (2023) 115671. <https://doi.org/10.1016/j.cma.2022.115671>
- [64] S.J. Anagnostopoulos, J.D. Toscano, N. Stergiopoulos, G.E. Karniadakis, Residual-based attention in physics-informed neural networks, *Comput. Methods Appl. Mech. Eng.* 421 (2024) 116805. <https://doi.org/10.1016/j.cma.2024.116805>
- [65] A.D. Jagtap, G.E. Karniadakis, Adaptive activation functions accelerate convergence in deep and physics-informed neural networks, *J. Comput. Phys.* 404 (2020) 109136. 1906.01170 <https://doi.org/10.1016/j.jcp.2019.109136>
- [66] S. Wang, B. Li, Y. Chen, P. Perdikaris, et al., PirateNets: Physics-informed Deep Learning with Residual Adaptive Networks, 2024, 2402.00326 <https://doi.org/10.48550/arXiv.2402.00326>.
- [67] B. Moseley, A. Markham, T. Nissen-Meyer, et al., Finite basis physics-informed neural networks (FBPINNs): a scalable domain decomposition approach for solving differential equations, *Adv. Comput. Math.* 49 (4) (2023) 62. <https://doi.org/10.1007/s10444-023-10065-9>
- [68] S. Basir, I. Senocak, Physics and equality constrained artificial neural networks: application to forward and inverse problems with multi-fidelity data fusion, *J. Comput. Phys.* 463 (2022) 111301. <https://doi.org/10.1016/j.jcp.2022.111301>
- [69] C. Wang, S. Li, D. He, L. Wang, et al., Is  $\mathcal{L}^2$  physics informed loss always suitable for training physics informed neural network?, in: *Advances in Neural Information Processing Systems*, 2022.
- [70] J. Yu, L. Lu, X. Meng, G.E. Karniadakis, et al., Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems, *Comput. Methods Appl. Mech. Eng.* 393 (2022) 114823. <https://doi.org/10.1016/j.cma.2022.114823>
- [71] H. Son, J.W. Jang, W.J. Han, H.J. Hwang, et al., Sobolev Training for Physics Informed Neural Networks, 2021, 2101.08932 <https://doi.org/10.48550/arXiv.2101.08932>.
- [72] S. Maddu, D. Sturm, C.L. Müller, I.F. Sbalzarini, et al., Inverse dirichlet weighting enables reliable training of physics informed neural networks, *Mach. Learn. Sci. Technol.* 3 (1) (2022) 015026. <https://doi.org/10.1088/2632-2153/ac3712>
- [73] L. McClenny, U. Braga-Neto, Self-Adaptive Physics-Informed Neural Networks Using a Soft Attention Mechanism, 2022, 2009.04544 <https://doi.org/10.48550/arXiv.2009.04544>.
- [74] S. Wang, S. Sankaran, H. Wang, P. Perdikaris, et al., An Expert's Guide to Training Physics-informed Neural Networks, 2023, 2308.08468 <https://doi.org/10.48550/arXiv.2308.08468>.
- [75] Z. Hao, C. Ying, H. Su, J. Zhu, J. Song, Z. Cheng, et al., Bi-Level Physics-Informed Neural Networks for PDE Constrained Optimization Using Broyden's Hypergradients, 2023, 2209.07075 <https://doi.org/10.48550/arXiv.2209.07075>.
- [76] Z. Zou, X. Meng, G.E. Karniadakis, Correcting model misspecification in physics-informed neural networks (PINNs), *J. Comput. Phys.* (2024) 112918. <https://doi.org/10.1016/j.jcp.2024.112918>
- [77] K.R. Swanson, A.E.C. Jr., J.D. Murray, et al., A quantitative model for differential motility of gliomas in grey and white matter, *Cell Prolif.* 33 (5) (2000) 317–329. <https://doi.org/10.1046/j.1365-2184.2000.00177.x>
- [78] H.L.P. Harpold, E.C. Alvord, Jr., K.R. Swanson, et al., The evolution of mathematical modeling of glioma proliferation and invasion, *J. Neuropathol. Exp. Neurol.* 66 (1) (2007) 1–9. <https://doi.org/10.1097/nen.0b013e31802d9000>
- [79] A. Schäfer, E.C. Mormino, E. Kuhl, Network diffusion modeling explains longitudinal tau PET data, *Front. Neurosci.* 14 (2020).
- [80] Z. Zhang, Z. Zou, E. Kuhl, G.E. Karniadakis, et al., Discovering a reaction–diffusion model for Alzheimer's disease by combining PINNs with symbolic regression, *Comput. Methods Appl. Mech. Eng.* 419 (2024) 116647. <https://doi.org/10.1016/j.cma.2023.116647>
- [81] L. Lu, X. Meng, S. Cai, Z. Mao, S. Goswami, Z. Zhang, G.E. Karniadakis, et al., A comprehensive and fair comparison of two neural operators (with practical extensions) based on FAIR data, *Comput. Methods Appl. Mech. Eng.* 393 (2022) 114778. <https://doi.org/10.1016/j.cma.2022.114778>.

- [82] C.E. Miles, S.A. McKinley, F. Ding, R.B. Lehoucq, et al., Inferring stochastic rates from heterogeneous snapshots of particle positions, *Bull. Math. Biol.* 86 (6) (2024) 74. <https://doi.org/10.1007/s11538-024-01301-4>
- [83] Y.-H. Tseng, T.-S. Lin, W.-F. Hu, M.-C. Lai, et al., A cusp-capturing PINN for elliptic interface problems, *J. Comput. Phys.* 491 (2023) 112359. <https://doi.org/10.1016/j.jcp.2023.112359>
- [84] P. Constantine, *Random Field Simulation*, 2024, (MATLAB Central File Exchange).
- [85] I. Ezhov, K. Scibilia, K. Franitza, F. Steinbauer, S. Shit, L. Zimmer, J. Lipkova, F. Kofler, J.C. Paetzold, L. Canalini, D. Waldmannstetter, M.J. Menten, M. Metz, B. Wiestler, B. Menze, et al., Learn-morph-infer: a new way of solving the inverse problem for brain tumor modeling, *Med. Image Anal.* 83 (2023) 102672. <https://doi.org/10.1016/j.media.2022.102672>
- [86] K. Scheufele, S. Subramanian, G. Biros, et al., Fully automatic calibration of tumor-Growth models using a single mpMRI scan, *IEEE Trans. Med. Imaging* 40 (1) (2021) 193–204. <https://doi.org/10.1109/TMI.2020.3024264>
- [87] L.R. Dice, Measures of the amount of ecologic association between species, *Ecology* 26 (3) (1945) 297–302. 1932409 <https://doi.org/10.2307/1932409>
- [88] Y. Chen, B. Hosseini, H. Owhadi, A.M. Stuart, et al., Solving and learning nonlinear PDEs with Gaussian processes, *J. Comput. Phys.* 447 (2021) 110668. <https://doi.org/10.1016/j.jcp.2021.110668>
- [89] L. C. Evans, *Partial Differential Equations*, American Mathematical Soc (2010).
- [90] C. G Broyden, A class of methods for solving nonlinear simultaneous equations, *Math Comput.* 19 (92) (1965) 577–593. <https://doi.org/10.1090/S0025-5718-1965-0198670-6>.
- [91] C. R. Vogel, *Computational Methods for Inverse Problems*, Soc. Industr. Appl. Math. (2002). arXiv, <https://epubs.siam.org/doi/pdf/10.1137/1.9780898717570> <https://doi.org/10.1137/1.9780898717570>.