

# Evolutionary Optimization of Physics-Informed Neural Networks: Advancing Generalizability by the Baldwin Effect

Jian Cheng Wong, Chin Chun Ooi, Abhishek Gupta, *Senior Member, IEEE*, Pao-Hsiung Chiu, Joshua Shao Zheng Low, My Ha Dao, and Yew-Soon Ong, *Fellow, IEEE*

**Abstract**—Physics-informed neural networks (PINNs) are at the forefront of scientific machine learning, making possible the creation of machine intelligence that is cognizant of physical laws and able to accurately simulate them. However, today’s PINNs are often trained for a single physics task and require computationally expensive re-training for each new task, even for tasks from similar physics domains. To address this limitation, this paper proposes a pioneering approach to advance the generalizability of PINNs through the framework of Baldwinian evolution. Drawing inspiration from the neurodevelopment of precocial species that have evolved to learn, predict and react quickly to their environment, we envision PINNs that are pre-wired with connection strengths inducing strong biases towards efficient learning of physics. A novel two-stage stochastic programming formulation coupling evolutionary selection pressure (based on proficiency over a distribution of physics tasks) with lifetime learning (to specialize on a sampled subset of those tasks) is proposed to instantiate the Baldwin effect. The evolved Baldwinian-PINNs demonstrate fast and physics-compliant prediction capabilities across a range of empirically challenging problem instances with more than an order of magnitude improvement in prediction accuracy at a fraction of the computation cost compared to state-of-the-art gradient-based meta-learning methods. For example, when solving the diffusion-reaction equation, a 70x improvement in accuracy was obtained while taking 700x less computational time. This paper thus marks a leap forward in the evolutionary meta-learning of PINNs as generalizable physics solvers. Sample codes are available at <https://github.com/chiu/h/Baldwinian-PINN>.

**Index Terms**—Baldwin effect, evolutionary optimization, neuroevolution, meta-learning, physics-informed neural networks

## I. INTRODUCTION

**T**HE emerging field of scientific machine learning seeks to create more accurate, data-efficient, and explainable machine intelligence for science and engineering. The direct incorporation of mathematically expressible laws of nature into

learned models to ensure physically consistent predictions is an appealing proposition, as evidenced by the proliferation of physics-informed neural networks (PINNs) across multiple scientific domains since seminal work by *Raissi et al.* [1]. The key concept is to utilize physics-based mathematical relations or constraints as a regularization loss (*aka* physics-informed loss). This physics-informed loss is amenable to various forms of scientific knowledge and theories, including fundamental ordinary or partial differential equations (ODEs or PDEs). It flexibly incorporates scientific discoveries accumulated across centuries into the machine intelligence models of today across diverse scientific and engineering disciplines [2–6].

However, PINNs remain limited in their ability to generalize across physics scenarios. Contrary to its promise, a PINN does not guarantee compliance with physics when used for new scenarios unseen during training, e.g., variations in PDE parameters, initial conditions (ICs) or boundary conditions (BCs) that lie outside the confines of their training. Instead, these predictions remain physics-agnostic and may experience similar negative implications for reliability as typical data-driven models.

In principle, physics-compliant predictions for any new scenario can be achieved by performing physics-informed retraining—an attractive feature of PINNs—even without labelled data. However, the additional training can be cost-prohibitive as physics-informed learning is much more difficult than data-driven learning due to the ruggedness of the loss landscapes [7–10]. This has motivated the exploration of transfer learning techniques where connection strengths from similar (source) physics scenarios are used to facilitate accurate learning of solutions for new, harder problems [11–13]. The related notion of meta-learning seeks to discover an optimized initialization of a model to enable rapid adaptation to a new test task with minimal training [11]. Nonetheless, most transfer- and meta-learned PINNs proposed to date still require the aforementioned physics-informed retraining (with a substantial number of optimization iterations) to achieve more accurate solutions, and are therefore not ideal for applications that call for repeated, fast evaluations. A method to arrive at a generalizable PINN, one that can provide fast and accurate physics prediction/simulation across a varied set of unseen scenarios, remains elusive.

In search of machine intelligence that generalizes for physically-consistent simulations of varied processes in the natural world, this paper studies the meta-learning of PINNs

Jian Cheng Wong, Chin Chun Ooi, and Pao-Hsiung Chiu are with the Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore (e-mail: [wongj@a-star.edu.sg](mailto:wongj@a-star.edu.sg); [ooicc@a-star.edu.sg](mailto:ooicc@a-star.edu.sg); [chiuph@a-star.edu.sg](mailto:chiuph@a-star.edu.sg)).

Abhishek Gupta is with the School of Mechanical Sciences, Indian Institute of Technology Goa, India (e-mail: [abhishekgupta@iitgoa.ac.in](mailto:abhishekgupta@iitgoa.ac.in)).

Joshua Shao Zheng Low is with the College of Computing and Data Science, Nanyang Technological University, Singapore (e-mail: [joshualow188@gmail.com](mailto:joshualow188@gmail.com)).

My Ha Dao is with the Technology Centre for Offshore and Marine, Singapore (e-mail: [Dao\\_My\\_Ha@tcoms.sg](mailto:Dao_My_Ha@tcoms.sg)).

Yew-Soon Ong is with the Agency for Science, Technology and Research, Singapore, and is also with the College of Computing and Data Science, Nanyang Technological University, Singapore (e-mail: [Ong\\_Yew\\_Soon@a-star.edu.sg](mailto:Ong_Yew_Soon@a-star.edu.sg)).

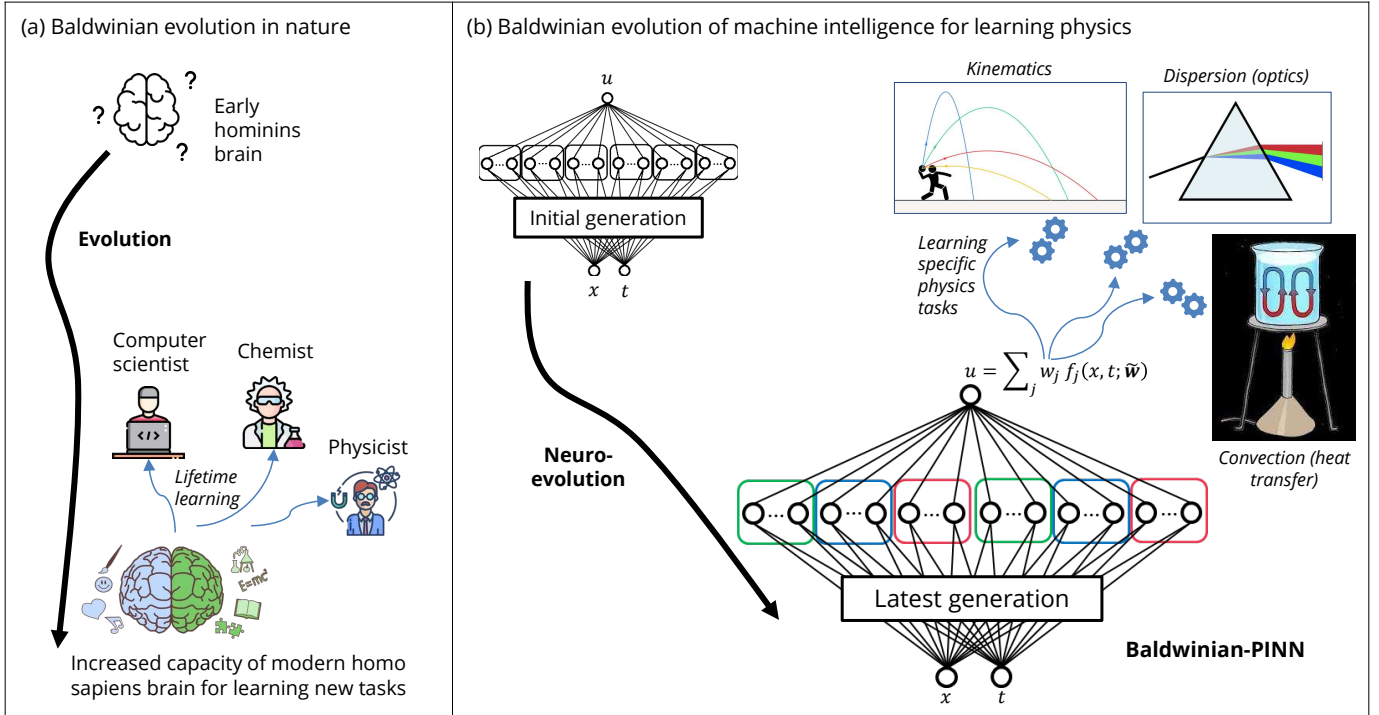


Figure 1. Schematic diagram of (a) Baldwinian evolution in nature and (b) evolving machine intelligence for learning physics with Baldwinian-PINNs. In nature, the Baldwin effect describes how learned traits are eventually reinforced in the genetic makeup of a population of organisms through natural selection. Equivalently, a population of Baldwinian-PINNs evolves over generations by being exposed to a broad distribution of physics tasks, gradually reinforcing traits promoting accurate physics learning into their genetic makeup. The evolved Baldwinian-PINNs are inherently equipped with strong learning biases to accurately solve any physics tasks over a broad task distribution.

through the lens of neural Baldwinism—an expression of the *Baldwin effect* in the evolution of brains and intelligence [14]. Inspiration is drawn from the transmission of knowledge and predispositions across generations in precocial species, whereby their young are “born ready” with strong learning biases to perform a wide range of tasks. In order to pre-wire such learning ability into the initial connection strengths of a neural network, we examine an algorithmic realization of neural Baldwinism in the context of PINNs. The essence of Baldwinism lies in the phenomenon that characters learned by individual organisms of a group during their lifetime may eventually, under selection pressure, get reinforced by associated hereditary characters [15]. Analogously, Baldwinian neuroevolution of PINNs consists of an outer evolutionary optimization loop in which populations of PINN models are collectively exposed to a wide range of physics tasks sampled from a probability distribution over tasks of interest. Models with higher propensity to perform a random subset of those tasks well in an inner lifetime learning loop are evaluated as being fitter for survival, thus inducing a selection pressure towards connection strengths that encode stronger learning biases. The parallels between neural Baldwinism in nature and that of PINNs is depicted in Figure 1. The mathematical formulation of the problem resembles a *two-stage stochastic program* [16, 17], where the solution to lifetime learning enables the PINN to rapidly specialize to a specific physics task at test time.

Harnessing evolutionary procedures to optimize neural networks lends much greater versatility relative to other meta-

learning approaches in terms of jointly crafting network architectures, initial network connection strengths, as well as learning algorithms, all through the use of potentially non-differentiable fitness functions [18]. Such biological “neuroevolution” [19, 20] precludes the need for explicit parameterization of tasks, facilitating generalization over task distributions comprising any broad mix of ODE/PDEs, ICs, and BCs. The physics-informed lifetime learning of the neural network can be accelerated by reduction to a least-squares learning problem in its final layer, making Baldwinian neuroevolution computationally feasible. Such a least-squares formulation yields a closed-form result by means of the Tikhonov regularization [21] and guarantees zero physics-informed loss given a sufficiently overparameterized network. The closed-form expression vastly reduces (or even eliminates) the need for iterative parameter updates to enable extremely fast lifetime learning of desired physics. Critically, the evolutionary search procedure is inherently highly parallelizable, thereby allowing for efficient meta-learning at scale by capitalizing on state-of-the-art advances in multi-CPU/GPU hardware infrastructure.

Neural Baldwinism thus makes it possible to achieve generalizable PINNs that are “genetically equipped” to perform well over a wide range of physics tasks. The evolved models (referred to as Baldwinian-PINNs) are demonstrated in this study to be broadly applicable to the simulation of families of linear and nonlinear ODE/PDEs encompassing diverse physical phenomena such as particle kinematics, heat and mass transfer, and reaction-diffusion. These Baldwinian-PINNs are capable of fast and accurate physics-aware predictions on

previously unseen tasks, demonstrating up to several orders of magnitude computation speedup along with an order of magnitude improvement in prediction accuracy relative to recent meta-learned PINNs [22, 23].

The key contributions of this paper are summarized below.

- This is the first study to unveil Baldwinian evolution as a compelling route towards discovering neural nets with high capacity to learn diverse physics tasks, thereby advancing generalizability in PINNs.
- An instantiation of the Baldwinian evolution framework is proposed through a novel two-stage stochastic programming formulation, wherein the first stage evolves the initial layers of a generalizable PINN model and the second stage trains its final layers to specialize to any new physics task (analogous to lifetime learning).
- Comprehensive numerical experimentation and analysis shows that this methodology produces PINNs that effectively learn and predict across multiple PDE problems spanning different physics domains, demonstrating significant advancements in terms of speed and accuracy over models meta-learned by gradient descent.

The remainder of the paper is organized as follows. The basic problem setup for a PINN, the notion of generalization over physics tasks, and related work are presented in Section II. The proposed methodology for evolving Baldwinian-PINNs is detailed in Section III. Extensive numerical assessment of the method is carried out in Section IV over a range of linear and nonlinear ODE/PDEs. The paper is concluded in Section V with a discussion of future research directions.

## II. PRELIMINARIES

### A. Problem setup

1) *Single PINN problem*: For simplicity of exposition, let us consider a problem with single spatial dimension  $x$ , time dimension  $t$ , and a single variable of interest  $u$ . In general, PINNs can learn a mapping between the input variables  $(x, t)$  and the output variable  $u$  while satisfying specified governing equations representing the physical phenomenon or dynamical process of interest:

$$\text{PDE: } \mathcal{N}_\vartheta[u(x, t)] = h(x, t), \quad x \in \Omega, t \in (0, T] \quad (1a)$$

$$\text{IC: } u(x, t = 0) = u_0(x), \quad x \in \Omega \quad (1b)$$

$$\text{BC: } \mathcal{B}[u(x, t)] = g(x, t), \quad x \in \partial\Omega, t \in (0, T] \quad (1c)$$

where the general differential operator  $\mathcal{N}_\vartheta[u(x, t)]$  can include linear and/or nonlinear combinations of temporal and spatial derivatives and PDE parameters  $\vartheta$ , and  $h(x, t)$  is an arbitrary source term in the domain  $x \in \Omega, t \in (0, T]$ . The IC (Eq. 1b) specifies the initial state,  $u_0(x)$ , at time  $t = 0$ , and the BC (Eq. 1c) specifies that  $\mathcal{B}[u(x, t)]$  equates to  $g(x, t)$  at the domain boundary  $\partial\Omega$ .

Crucially, individual PINN models arrive at an accurate and physics-compliant prediction  $u(x, t)$  for a single target scenario by minimizing the discrepancy between Eq. 1 and the model's prediction during training.

2) *Generalizable neural physics solver*: While most PINN models are trained to solve a specific physics task, there is increasing interest in generalizable neural physics solvers, i.e. models that can be flexibly applied to multiple physics problems once trained. In this context, we can consider a physical phenomena of interest that is represented by a set of training tasks belonging to some underlying task-distribution  $p(\mathcal{T})$ , e.g., a family of PDEs spanning different PDE parameters  $\vartheta$ , different ICs  $u_0(x)$  and/or different BCs  $g(x, t)$ .

Hence, the goal is to discover generalizable PINN models capable of fast, accurate, and physics-aware predictions on unseen scenarios, i.e., any new task from the distribution,  $\mathcal{T}_i \sim p(\mathcal{T})$ , by learning the underlying governing physics. In the context of meta-learning, the learning objective is to use training tasks from  $p(\mathcal{T})$  to find network initializations that are most amenable to a quick and accurate solution for the PINN loss, thereby accelerating the solution of multiple related physics problems at test time and providing a potential route to a generalizable neural physics solver.

### B. Related Work

PINN models are usually trained to solve a single, specific physics task. However, recent studies on meta-learning of PINNs to solve different physics scenarios as separate tasks have emerged, although no work has been reported from the neuroevolution perspective to our knowledge.

Most of the meta-learning PINN approaches reported in literature use weight interpolation as the basic framework. The simplest instantiation of this approach is to first train independent PINN models for each task and then interpolate across model weights for the new task [22, 24]. Several interpolation methods such as the Gaussian Process (GP) and Radial Basis Function (RBF) have been studied and shown to improve physics-informed learning on new tasks [22]. In other studies, interpolation is learned through the use of hypernetworks, whereby the hypernetworks and PINN are trained simultaneously from all the tasks [25, 26]. Other variants include encoding tasks as latent variables and passing them into the input layer of the PINN model [27].

However, there are two major drawbacks to such weight interpolation frameworks. Firstly, these methods rely heavily on task parameterization to perform interpolation and operate under the assumption of smoothness. This means that all the train tasks (and the new task) must adhere to this parameterization requirement. In addition, these methods do not capitalize on the principal characteristic of PINN, which is the potential for physics-informed learning (retraining or fine-tuning) for a new task, in their meta-learning formulation. The data-free nature of physics-informed fine-tuning is not exploited during the meta-learning phase, given that there is no guiding principle on how the fine-tuning towards any new task should be performed given the interpolated weights.

The model-agnostic meta-learning (MAML) [28] framework can theoretically overcome the limitations of weight interpolation methods. MAML and its variant, Reptile [29], aim to learn an optimal weight initialization that can be quickly fine-tuned on new tasks without the need for interpolation [23]. An illustration of MAML for PINN is shown in

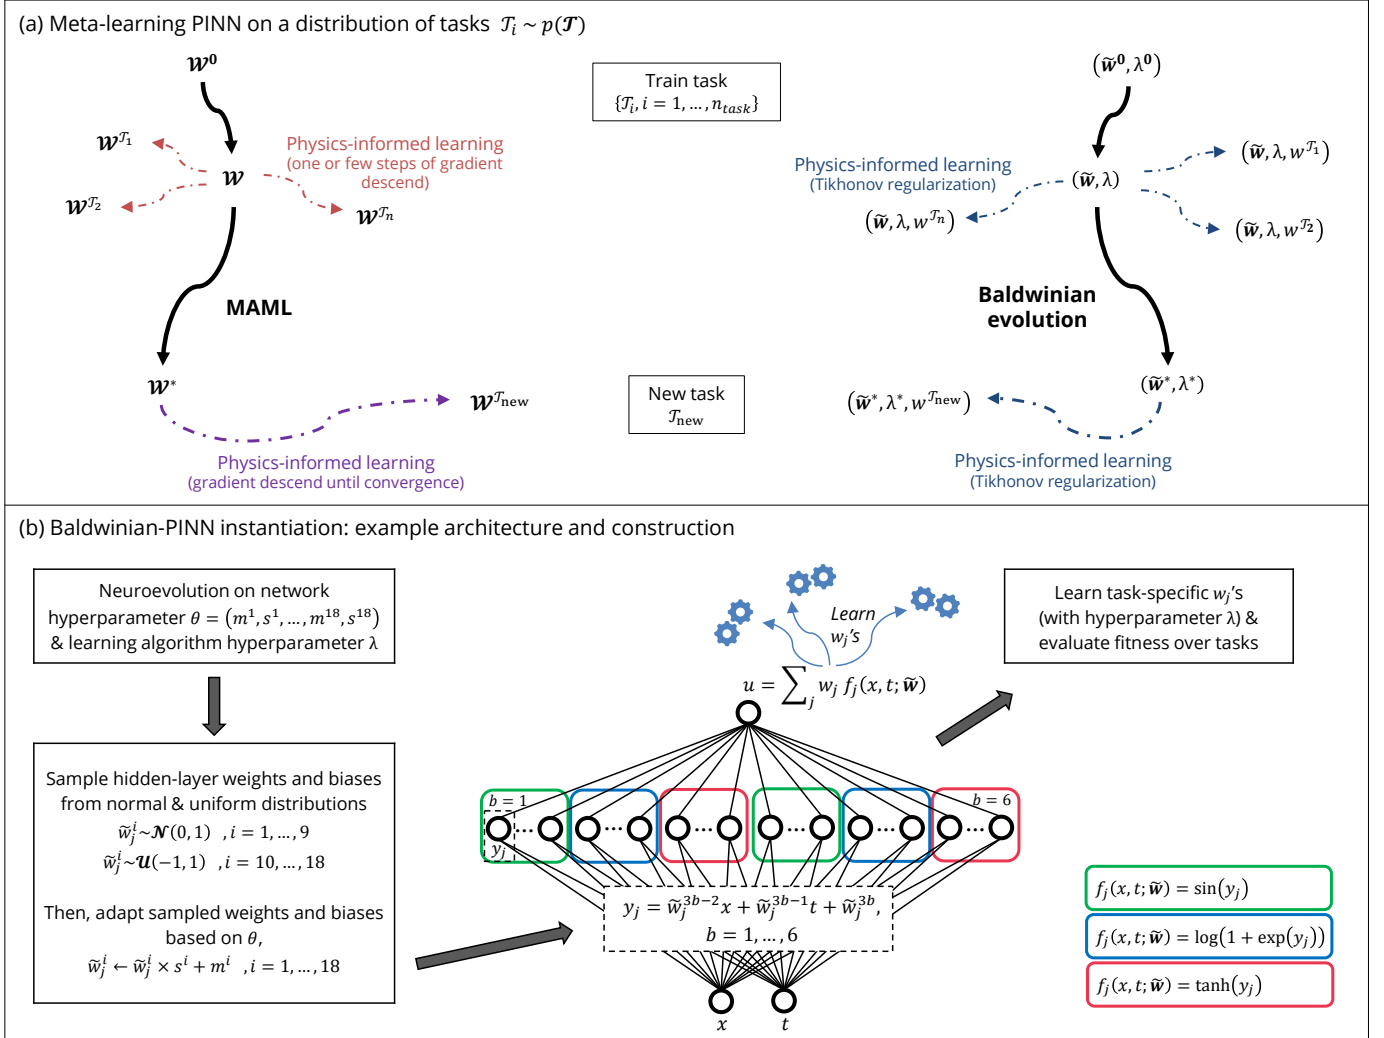


Figure 2. (a) Meta-learning PINN with Baldwinian neuroevolution (right) versus MAML (left). In MAML, the initial weights  $\mathbf{W}$  are learned using gradient-based method, such that they can be quickly fine-tuned (physics-informed learning) on new tasks. Although the task-specific fine-tuning is limited to one or a few gradient descent updates during training, such amount of fine-tuning is usually insufficient for a PINN at test time. In Baldwinian neuroevolution, the weight distribution in the pre-final nonlinear hidden layers  $\tilde{\mathbf{w}}$  and learning hyperparameters  $\lambda$  are jointly evolved. The task-specific physics-informed learning is performed on the final layer weights  $w$  (segregated from  $\tilde{\mathbf{w}}$ ) with a 1-step Tikhonov regularization operation at both training and test time (for linear ODE/PDEs). (b) Schematic of Baldwinian-PINNs architecture used in present study and procedure to obtain nonlinear hidden layers' weights  $\tilde{\mathbf{w}}$ 's from the evolved network hyperparameter  $\theta$  for learning task-specific outputs.

Figure 2(a). To reduce the computation cost of meta-learning, task-specific physics-informed learning is limited to one or a few gradient descent updates from the initialization. This essentially assumes that the model performance after one step (or a few steps of) gradient descent update from the initialization is already indicative of the learning performance. While this assumption may be appropriate in data-driven models, PINN models are much more challenging to optimize. Hence, they require more thorough training (e.g., a large number of updates, even with the use of 2nd-order gradient descent methods) to achieve a good convergence. Moreover, gradient-based MAML methods are prone to getting stuck in a local minimum and struggle to find a good initialization for a diverse set of tasks, as the gradients can be noisy or even deceptive, i.e., requiring going against it to reach the optimum. Hence, recent studies have shown poor MAML results, potentially due to non-convergence during meta-learning or insufficient

physics-informed fine-tuning performed at test time [22, 25].

### III. PROPOSED METHODOLOGY

We address the above limitations with a meta-learning PINN framework based on Baldwinian neuroevolution. A novel evolutionary algorithm is crafted to jointly optimize the weight distribution in the network's hidden layers alongside other hyperparameters that are essential for achieving optimal performance on downstream physics-informed learning tasks. We limit our task-specific learning to the final layer, allowing us to solve linear problems with a 1-step Tikhonov regularization operation (or a few steps for nonlinear problems) during both the meta-learning process and test time, in a consistent manner. Importantly, once given the nonlinear hidden layers' weights and learning hyperparameters, the Tikhonov regularization is exact and fast (does not require further gradient descent

updates). An illustration of the Baldwinian evolution of PINNs is shown in Figure 2(a).

#### A. The Baldwin effect as a two-stage stochastic program

Without loss of architectural generality, Baldwinian-PINNs can be assigned the form of a multilayer perceptron (MLP) with proven representation capacity [30, 31], thereby ensuring the ability to learn a diversity of tasks. Its output  $u(x, t)$  can be written as:

$$u(x, t) = \sum_j w_j f_j(x, t; \tilde{\mathbf{w}}) \quad (2)$$

where  $\mathbf{w} = [\dots w_j \dots]^T$  are the final layer weights, and  $f_j(x, t; \tilde{\mathbf{w}})$ 's represent nonlinear projections of the input variables with the hidden layers' weights  $\tilde{\mathbf{w}}$ . The connection strengths  $\tilde{\mathbf{w}}$  are typically learned during PINN training. However, in the proposed implementation of Baldwinian-PINNs as detailed below, they are drawn from a probability distribution defined by a dimensionally reduced set of network hyperparameters  $\theta$ —see example in Section III-D1 and Figure. 2(b)—that are assigned *at birth* as per Baldwinian evolution.

The model's reaction to a given environment—i.e., finding the best set of  $w_j$ 's such that the model's output satisfies Eq. 1 for a specific physics task  $\mathcal{T}_i$ —is the focus of a typical PINN described in Section II-A. This is reduced to a physics-informed least-squares problem in this work:

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} (\mathbf{A}_{\tilde{\mathbf{w}}}^{\mathcal{T}_i} \mathbf{w} - \mathbf{b}^{\mathcal{T}_i})^T (\mathbf{A}_{\tilde{\mathbf{w}}}^{\mathcal{T}_i} \mathbf{w} - \mathbf{b}^{\mathcal{T}_i}) + \lambda \mathbf{w}^T \mathbf{w} \quad (3)$$

where  $\mathbf{A}_{\tilde{\mathbf{w}}}^{\mathcal{T}_i} \mathbf{w}$  is obtained by substituting the model's output into the left hand side of Eq. 1 for a given set of collocation points, and  $\mathbf{b}^{\mathcal{T}_i}$  represents the corresponding right hand side of Eq. 1. Eq. 3 yields a closed-form solution  $\mathbf{w}^*$  when using the Moore-Penrose generalized inverse, permitting extremely fast learning (order of milli-seconds in our experiments) for physics-compliant prediction. Similar least-squares formulations studied in the literature have been shown to be competitive with, or faster than, widely-used numerical solvers such as the finite element method [32, 33]. The learning hyperparameter  $\lambda \geq 0$  reduces the  $L^2$ -norm of the least-squares solution, thereby improving the solution numerically. Detailed derivations, including information on the construction of matrix  $\mathbf{A}_{\tilde{\mathbf{w}}}^{\mathcal{T}_i}$  and vector  $\mathbf{b}^{\mathcal{T}_i}$ , are provided in Section III-C.

Going beyond specialization to a single physics task by means of Eq. 3, we formulate the search for Baldwinian-PINN models that can generalize to a whole family of PDEs as the following two-stage stochastic optimization problem [17]:

$$\begin{aligned} \min_{\theta, \lambda} \mathbb{E}_{\mathcal{T}_i \sim p(\mathcal{T})} \mathbb{E}_{\tilde{\mathbf{w}} \sim p_{\theta}(\tilde{\mathbf{w}})} [\kappa_{LSE} l_{LSE}(\mathbf{w}^*) + \kappa_{MSE} l_{MSE}(\mathbf{w}^*)] \\ \text{subject to } \kappa_{LSE} \geq 0, \kappa_{MSE} \geq 0 \end{aligned} \quad (4)$$

where  $\mathbf{w}^*$  is the solution of the second stage problem defined earlier in Eq. 3, which provides optimal final layer weights that allow the model to specialize to any realization of task  $\mathcal{T}_i \sim p(\mathcal{T})$  for the given network's  $\tilde{\mathbf{w}} \sim p_{\theta}(\tilde{\mathbf{w}})$ .

The first stage optimization objective is defined by the weighted sum of the physics learning proficiency, i.e., sum of squared residuals or least-squares error (LSE),

$$l_{LSE}(\mathbf{w}^*) = (\mathbf{A}_{\tilde{\mathbf{w}}}^{\mathcal{T}_i} \mathbf{w}^* - \mathbf{b}^{\mathcal{T}_i})^T (\mathbf{A}_{\tilde{\mathbf{w}}}^{\mathcal{T}_i} \mathbf{w}^* - \mathbf{b}^{\mathcal{T}_i}) \quad (5)$$

and the actual predictive performance, i.e., mean squared error (MSE),

$$l_{MSE}(\mathbf{w}^*) = \frac{1}{n} \sum_{s=1}^n (u_s^{label} - \sum_j w_j^* f_j(x_s, t_s; \tilde{\mathbf{w}}))^2 \quad (6)$$

given labelled data  $u_s^{label}, s = 1, \dots, n$ , over the task distribution  $p(\mathcal{T})$  and the network connections' distribution  $p_{\theta}(\tilde{\mathbf{w}})$ .

*In what follows, we present a novel procedure for solving the two-stage stochastic program via an evolutionary algorithm. The hyperparameters  $\theta, \lambda$  are evolved from one generation to the next, with their fitness evaluations for the first stage problem conditioned on lifetime learning to obtain the optimal  $\mathbf{w}^*$  from the second stage problem. Note that while the outcome of lifetime learning influences the selection pressure acting on the evolving hyperparameters, it does not directly alter the genetic makeup of the hyperparameters within a generation. As such, our overall method exhibits a clear connection with the evolutionary principles of Baldwinism.*

#### B. Baldwinian neuroevolution

The Baldwinian neuroevolution procedure to solve the two-stage stochastic programming problem defined in Section III-A is described in Algorithm 1 and Algorithm 2. Recall from Figure. 2(b) that a Baldwinian-PINN is represented by  $(\theta, \lambda)$ , i.e., distribution of network weights and biases and lifetime learning hyperparameters, during neuroevolution.

---

#### Algorithm 1 Baldwinian neuroevolution of PINNs

---

**INPUT:** training tasks distribution,  $p(\mathcal{T})$

**OUTPUT:** best solution found (i.e., Baldwinian-PINN model's weights and biases in nonlinear hidden-layers and lifetime learning hyperparameters  $(\theta, \lambda)$ )

**Require:**  $\mathcal{F}$ : procedure to return fitness based on lifetime learning performance of individual given a batch of tasks (described in detail in Algorithm 2)

- 1: Initialize population  $\mathcal{P}$
  - 2: **while** not done **do**
  - 3:   Sample a new batch of offspring  $(\theta^g, \lambda^g), g = 1, \dots, n_{pop}$  from the population  $\mathcal{P}$
  - 4:   Sample batch of tasks  $\mathcal{T}_i \sim p(\mathcal{T}), i = 1, \dots, n_{task}$
  - 5:   **for** all  $(\theta^g, \lambda^g)$  **do**
  - 6:      $f^g = \mathcal{F}(\theta^g, \lambda^g, \{\mathcal{T}_i, i = 1, \dots, n_{task}\})$
  - 7:   **end for**
  - 8:   Update population  $\mathcal{P}$  based on fitness of individuals:  $\{(\theta^g, \lambda^g, f^g), g = 1, \dots, n_{pop}\}$
  - 9: **end while**
  - 10: Return best individual  $(\theta, \lambda)$  found in population  $\mathcal{P}$
- 

The Baldwinian neuroevolution procedure described in Algorithm 1 is generic for evolutionary optimization methods. The algorithm initializes a population  $\mathcal{P}$  of Baldwinian-PINN

---

**Algorithm 2** Baldwinian-PINN lifetime learning and fitness calculation  $\mathcal{F}$ 


---

**INPUT:** network and lifetime learning hyperparameters

 $(\theta, \lambda)$ , batch of tasks  $\mathcal{T}_i, i = 1, \dots, n_{task}$ 
**OUTPUT:** fitness  $f$ 


---

**Require:**  $\mathcal{G}$ : procedure to sample weights and biases in nonlinear hidden-layers (exemplified in Section III-D1 and Figure. 2(b))

**Require:**  $\mathcal{C}$ : procedure to construct least squares problem on a set of collocation points based on underlying physics (PDEs, BCs, ICs) of the task (described in detail in Section III-C)

- 1: Sample hidden layers' weights and biases  $\tilde{\mathbf{w}} = \mathcal{G}(\theta)$
  - 2: **for** all  $\mathcal{T}_i$  **do**
  - 3:   Construct least squares matrix and vector:  
 $(\mathbf{A}, \mathbf{b}) = \mathcal{C}(\mathcal{T}_i, \tilde{\mathbf{w}})$
  - 4:   Compute least squares solution as per Eq. 8:  
 $\mathbf{w}^* = (\lambda \mathbf{I} + \mathbf{A}^T \mathbf{A})^{-1} \mathbf{A} \mathbf{b}$
  - 5:   Compute least squares error (LSE) as per Eq. 5:  
 $l_{LSE}^{\mathcal{T}_i} = (\mathbf{A} \mathbf{w}^* - \mathbf{b})^T (\mathbf{A} \mathbf{w}^* - \mathbf{b})$
  - 6:   Compute mean squared error (MSE) as per Eq. 6:  
 $l_{MSE}^{\mathcal{T}_i} = \frac{1}{n} \sum_s (u_s^{label} - \sum_j w_j^* f_j(x_s, t_s; \tilde{\mathbf{w}}_j))^2$
  - 7: **end for**
  - 8: Compute overall fitness:  $f = -(\kappa_{LSE} \sum_{\mathcal{T}_i} l_{LSE}^{\mathcal{T}_i} + \kappa_{MSE} \sum_{\mathcal{T}_i} l_{MSE}^{\mathcal{T}_i})$
- 

models given by different weights and biases in the nonlinear hidden-layers and lifetime learning hyperparameters  $(\theta, \lambda)$ . For a probabilistic model-based evolutionary algorithm,  $\mathcal{P}$  is commonly represented by a *search distribution*. In each generation,  $n_{pop}$  new offspring individuals are sampled from the search distribution (probabilistic model-based EAs) or through crossover/mutation (traditional EAs), and their fitness are evaluated for a batch of tasks randomly sampled from the training task distribution  $p(\mathcal{T})$ . The fitness evaluation procedure  $\mathcal{F}$  gives the lifetime learning outcome  $f^g$  of these offspring individuals  $(\theta^g, \lambda^g), g = 1, \dots, n_{pop}$  for the given tasks. In line with the essence of Baldwinism, the lifetime learning procedure (described in detail in Algorithm 2) does not alter the genetic makeup  $(\theta, \lambda)$  of the individuals. The outcome of the lifetime learning procedure specifies the fitness which creates the selection pressure influencing the evolution of the population, but is not directly inherited by the next population (unlike Lamarckian evolution). The Baldwinian neuroevolution algorithm iteratively adapts  $\mathcal{P}$  towards offspring with better fitness until the convergence criteria, e.g., a pre-determined number of generations or fitness value (trade-off between computation resource and convergence), is met.

Given an individual  $(\theta, \lambda)$  sampled from the search distribution and a batch of tasks sampled from the training distribution,  $\mathcal{T}_i \sim p(\mathcal{T}), i = 1, \dots, n_{task}$ , the procedure  $\mathcal{F}$  to return fitness is detailed in Algorithm 2. It starts with a procedure  $\mathcal{G}$  to populate nonlinear hidden-layers' weights and biases  $\tilde{\mathbf{w}}$  of a Baldwinian-PINN from the sampled individual  $\theta$ . In the present study, Baldwinian-PINNs are designed to have  $\tilde{\mathbf{w}}$  fixed *at birth* to random values drawn from a probability

distribution defined by a dimensionally reduced set of network hyperparameters  $\theta$ , akin to a randomized neural networks setup [17, 34]. This procedure is exemplified in Section III-D1. Then, lifetime learning of the Baldwinian-PINN is performed to obtain the optimal network weights  $\mathbf{w}^*$  in the final layer, for each of the sampled tasks  $\mathcal{T}_i$ . It involves procedure  $\mathcal{C}$  to construct least squares problem, i.e., matrix and vector  $(\mathbf{A}, \mathbf{b})$ , on a fixed set of collocation points based on underlying physics (PDEs, BCs, ICs) of the task. This procedure is detailed in Section III-C. We choose the collocation points to coincide with the location of labelled data for MSE computation, although this is not a prerequisite. The overall fitness  $f$  is computed by aggregating the LSE and MSE based on lifetime learning outcomes for all the sampled tasks.

In the spirit of the Baldwin effect, the task-specific final layer  $\mathbf{w}^*$  is the outcome of lifetime learning and not directly inherited by the next generation of offspring [35]; only the hyperparameters  $(\theta, \lambda)$  are subjected to evolutionary variation and inheritance.

### C. Baldwinian-PINN lifetime learning procedure

The lifetime learning procedure of the Baldwinian-PINNs occurs only in the linear final layer of the network, i.e., finding the best set of  $w_j$ 's such that the output  $u(x, t) = \sum_j w_j f_j(x, t; \tilde{\mathbf{w}})$  satisfies the governing equations in Eq. 1 for a specific task. Given a set of collocation points  $(x_i^{pde}, t_i^{pde}), i = 1, \dots, n_{pde}, (x_i^{ic}, 0), i = 1, \dots, n_{ic}, (x_i^{bc}, t_i^{bc}), i = 1, \dots, n_{bc}$  sampled from the respective domain, the following system of equations can be formed:

$$\begin{bmatrix} \dots & \mathcal{N}_\theta[f_j(x_1^{pde}, t_1^{pde}; \tilde{\mathbf{w}})] & \dots \\ & \vdots & \\ \dots & \mathcal{N}_\theta[f_j(x_{n_{pde}}^{pde}, t_{n_{pde}}^{pde}; \tilde{\mathbf{w}})] & \dots \\ & \vdots & \\ \dots & f_j(x_1^{ic}, 0; \tilde{\mathbf{w}}) & \dots \\ & \vdots & \\ \dots & f_j(x_{n_{ic}}^{ic}, 0; \tilde{\mathbf{w}}) & \dots \\ & \vdots & \\ \dots & \mathcal{B}[f_j(x_1^{bc}, t_1^{bc}; \tilde{\mathbf{w}})] & \dots \\ & \vdots & \\ \dots & \mathcal{B}[f_j(x_{n_{bc}}^{bc}, t_{n_{bc}}^{bc}; \tilde{\mathbf{w}})] & \dots \end{bmatrix} \begin{bmatrix} \vdots \\ w_j \\ \vdots \end{bmatrix} = \begin{bmatrix} h(x_1^{pde}, t_1^{pde}) \\ \vdots \\ h(x_{n_{pde}}^{pde}, t_{n_{pde}}^{pde}) \\ \vdots \\ u_0(x_1^{ic}) \\ \vdots \\ u_0(x_{n_{ic}}^{ic}) \\ \vdots \\ g(x_1^{bc}, t_1^{bc}) \\ \vdots \\ g(x_{n_{bc}}^{bc}, t_{n_{bc}}^{bc}) \end{bmatrix} \quad \mathbf{A} \mathbf{w} = \mathbf{b} \quad (7)$$

Note that the derivatives required to construct  $\mathbf{A}$  can be easily computed by automatic differentiation [36]. The best-fit solution to the above system of linear equations with unknown  $\mathbf{w} = [\dots w_j \dots]^T$  can be obtained by means of the Tikhonov regularization:

$$\mathbf{w}^* = (\lambda \mathbf{I} + \mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b} \quad (8)$$

where the learning hyperparameter  $\lambda \geq 0$  reduces the  $L^2$ -norm of the least-squares solution, thereby improving the solution numerically. Such a least-squares formulation yields a closed-form result in a single computation step for any linear PDE (i.e., the governing equations in Eq. 1 are all linear with respect to  $u$ ) which encapsulates a wide class of physics phenomena in the natural world.

For nonlinear PDEs, iterative methods can be used for arriving at optimized  $w_j$ 's in Baldwinian-PINNs (detailed in Suppl. S.II). Hence, the Baldwinian-PINNs' lifetime learning is broadly applicable to **both** linear and nonlinear PDEs, with the psuedoinverse formulation permitting extremely fast computation for physics-compliant prediction. It is worth emphasizing that each Baldwinian-PINN undergoes lifetime learning in order to produce accurate, physics-compliant predictions on a single, new physics task, hence, the fast nature of this procedure permits flexible and rapid prediction for each new task on-demand. In addition, there is no requirement on prior labelled data for any new task as the learning can be entirely physics-informed (as per conventional PINNs).

#### D. Implementation details

1) *Baldwinian-PINN architecture*: In the present study, Baldwinian-PINNs are designed to have weights and biases in the nonlinear hidden-layers  $\tilde{w}$  fixed to random values *at birth*. They are akin to a randomized neural networks setup [17, 34] and drawn from a probability distribution defined by a dimensionally reduced set of network hyperparameters  $\theta$ . Both normal and uniform distributions for the weights and biases are possible for the randomized PINNs [32, 33]. Similarly, smooth nonlinear activations such as *sin*, *softplus*, and *tanh* are common in PINN literature and have their own merits. For greater flexibility, we apply both distributions for setting weights and biases and all three activations to different hidden layer blocks in the Baldwinian-PINN models.

Our base Baldwinian-PINN model is depicted in Figure 2(b). It's hidden layer architecture is segmented into  $3 \times 2 = 6$  unique blocks, with the weights and biases of the first 3 blocks sampled from normal distributions and the weights and biases of the other 3 blocks sampled from uniform distributions. Each block has a fixed number of neurons ( $n_{neuron} = 150$  or  $200$ ). Assuming 2 input variables ( $x, t$ ), we can write the output  $f_j$  for all the neurons  $j = 1, \dots, n_{neuron}$  in a nonlinear hidden-layer block as:

$$y_j = \tilde{w}_j^{3b-2}x + \tilde{w}_j^{3b-1}t + \tilde{w}_j^{3b} \quad (9a)$$

$$f_j(x, t; \tilde{w}_j) = \varphi^b(y_j) \quad (9b)$$

for each of the blocks  $b = 1, \dots, 6$ , where the activation  $\varphi^b$  can be *sin* ( $b = 1, 4$ ), *softplus* ( $b = 2, 5$ ), or *tanh* ( $b = 3, 6$ ). The  $\tilde{w}_j^i$ 's are weights and biases with their own distributional mean  $m^i$  and spread  $s^i$ ,  $i = 1, \dots, 18$ . Their values can be obtained by the following sampling procedure:

$$\tilde{w}_j^i \sim \mathcal{N}(0, 1) \quad \text{or} \quad \tilde{w}_j^i \sim \mathcal{U}(-1, 1) \quad (10a)$$

$$\tilde{w}_j^i \leftarrow \tilde{w}_j^i \times s^i + m^i \quad (10b)$$

Given the Baldwinian-PINNs' configuration, a set of network hyperparameters  $\theta = (m^1, s^1, \dots, m^{18}, s^{18})$  control the distributional mean and spread of the weights and biases in different blocks. Since the neuroevolution only searches the distribution parameters for groups of weights instead of evolving each individual weight in the nonlinear hidden layer, the reduced dimensionality can be even more effectively searched by today's evolutionary algorithms.

While the above description is for a single hidden layer, we note that the Baldwinian-PINN framework is not restricted to a single hidden layer. We include additional examples in Suppl. S.IV.B to show that the methodology also works for deeper neural architectures.

2) *Evolutionary algorithm*: In the majority of this study, we employ the covariance matrix adaptation evolution strategy (CMA-ES) [37] for evolving  $(\theta, \lambda)$ , although results in Suppl. S.IV.B show the extensibility of this framework to other neuroevolution algorithms. As an instantiation of information-geometric optimization algorithms [38], CMA-ES represents the population of  $(\theta, \lambda)$  in  $\mathcal{P}$  using a multivariate normal search distribution, initialized with zero mean and standard deviation (std.) as tuning hyperparameter. It iteratively adapts the search distribution based on the rank-based fitness landscape until the convergence criteria is met. Our experiments show that the performance of Baldwinian neuroevolution is robust across a range of CMA-ES hyperparameters such as population size and initial standard deviation (std.) of search distribution, and the number of tasks sampled for fitness evaluation per iteration. Hence, a robust setting that shows good convergence in fitness across different types of problems is chosen based on initial experiments.

The network weights and biases can take any value from  $(-\infty, \infty)$ , hence there is no restriction to the continuous search space of  $\theta$  representing their distributional means and spreads. The learning hyperparameter  $\lambda \geq 0$  can be evolved in continuous search space and its absolute value is then used for computing the least-squares solution. In our implementation with CMA-ES,  $(\theta, \lambda)$  share the same initial standard deviation, and we scale the learning hyperparameter by a factor, i.e.,  $\lambda \leftarrow 1e-4 \times \text{abs}(\lambda)$ , to improve the performance of the Tikhonov regularization.

In our preliminary experiments, we found it helpful to have the  $l_{MSE}$  component in the optimization objective even though the learning of future test tasks remains solely physics-informed. We set  $\kappa_{LSE} = \kappa_{MSE} = 1$  as default for the computation of overall fitness unless there is a huge difference in magnitude between  $l_{LSE}$  and  $l_{MSE}$  during Baldwinian neuroevolution. In addition, we can multiply the BC/IC rows in both **A** and **b** to re-balance the importance between PDE and BC/IC errors in the least-squares solution.

As this study focuses on the paradigm of Baldwinian neuroevolution as a pathway towards generalizable neural physics solvers, other combinations of evolutionary optimization or lifetime learning algorithms may be used. One key advantage of evolutionary optimization is that the fitness evaluations (population size  $n_{pop} \times$  number of random tasks  $n_{task}$ ) required each iteration can be easily parallelized across multiple GPUs to fully harvest any hardware advantage. In particular, we utilized the JAX framework to harness previously reported performance improvements for automatic differentiation and linear algebra operations [39, 40]. The experimental study is performed on a workstation with an Intel Xeon W-2275 Processor and 2 NVIDIA GeForce RTX 3090 GPUs.

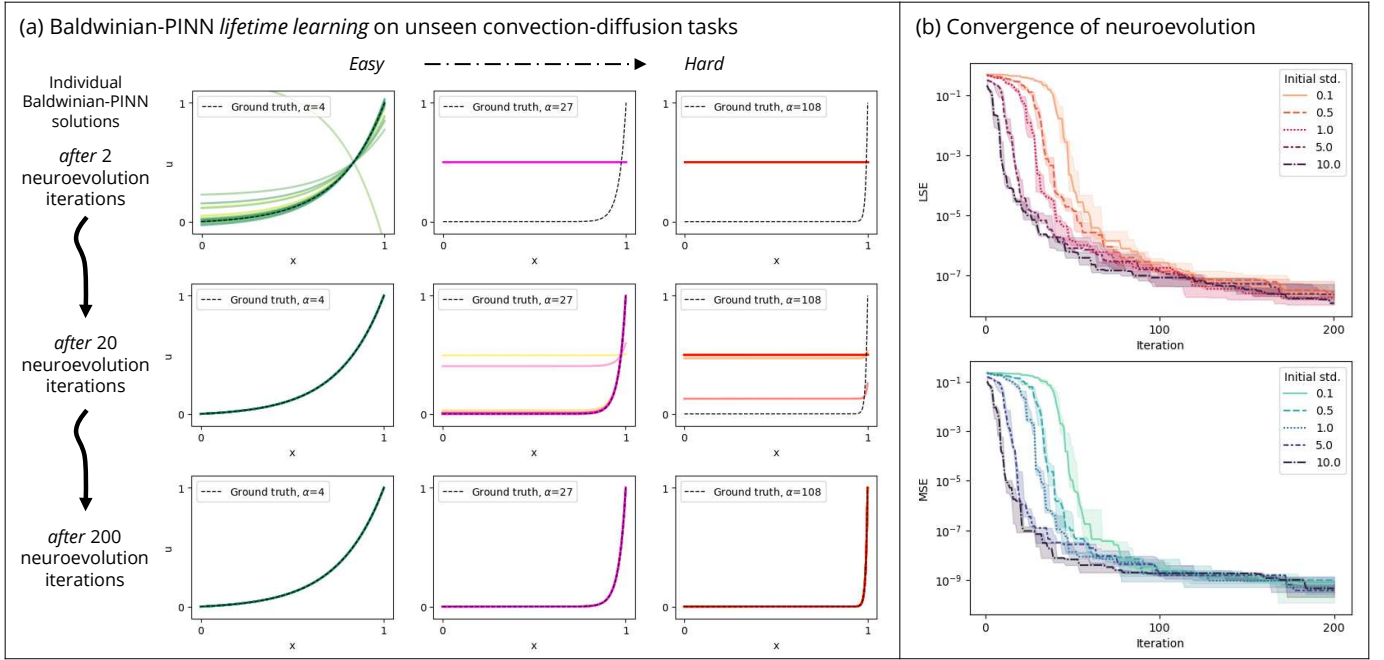


Figure 3. (a) Solution of 20 individual Baldwinian-PINN models sampled from the CMA-ES search distribution (initial std. = 1), for unseen convection-diffusion tasks  $\alpha = \{4, 27, 108\}$ . The task is more challenging with increasing  $\alpha$ . Baldwinian neuroevolution is effective for evolving good Baldwinian-PINN models which can generalize across different difficulties. (b) Baldwinian neuroevolution demonstrates effective LSE and MSE convergences on convection-diffusion problem, for different CMA-ES initial std. values (best std. = 1). The bold lines indicate their median convergence path from 5 individual runs, and the shaded areas indicate their interquartile ranges.

#### IV. EXPERIMENTAL STUDIES

We examine the efficacy of Baldwinian neuroevolution for learning physics (Algorithm 1 and Algorithm 2 in Section III-B) as formalized in a two-stage stochastic programming problem in Section III-A. Several ODE/PDE problems which are representative of real-world phenomena are used to demonstrate Baldwinian neuroevolution for physics in the following sections. Table I summarizes neuroevolution configurations and Baldwinian-PINN lifetime learning performance on their respective test tasks, including computational time for both neuroevolution and lifetime learning phases.

##### A. Learning to solve and generalize linear ODE/PDEs

1) *Convection-diffusion*: The steady-state convection-diffusion equation is a ubiquitous physics model that describes the final distribution of a scalar quantity (e.g. mass, energy, or temperature) in the presence of convective transport and diffusion [41]. Solutions to this physics are key to characterization and design of many systems, including microfluidic chip cooling in electronics [42]. The 1D equation is defined as:

$$\text{(Problem 1)} \quad \alpha \frac{du}{dx} - \frac{d^2u}{dx^2} = 0, \quad x \in [0, 1] \quad (11)$$

subject to BCs  $u(x=0) = 0$ ;  $u(x=1) = 1$ . These real-world problems have characteristic physics that vary with non-dimensional constants such as the Peclet number  $Pe$  (ratio of convection to diffusion) [43]. Hence, it is helpful to learn a PINN model that can return  $u(x)$  for a diverse range of  $Pe$ -related problems (determined by  $\alpha$  here). The training tasks consist of  $\alpha = \{5, 10, \dots, 100\}$ , encompassing both smoother

output patterns at lower  $\alpha$  and very high gradient patterns at higher  $\alpha$ , with the latter being challenging for PINNs to learn by both stochastic gradient descent (SGD) [8, 12] and classical numerical methods [44].

The predictive performance of the learned model is evaluated for an unseen range of test tasks, i.e.,  $\alpha = \{1, 2, \dots, 110\}$ . The efficacy of Baldwinian neuroevolution is demonstrated in Figure 3, with the successful evolution of Baldwinian-PINNs which can learn an extremely accurate solution on new test tasks in milli-seconds. The learned solutions can achieve an average MSE of  $5.8e-9 \pm 9.9e-9$  ( $n = 110$  tasks  $\times$  5 individual runs) after 200 neuroevolution iterations.

2) *Family of linear PDEs (convection, diffusion, and dispersion)*: Next, we extend Baldwinian-PINNs to a family of linear PDEs. This PDE family is a further generalization of the convection-diffusion equation:

$$\text{(Problem 2)} \quad \frac{du}{dt} + \alpha \frac{du}{dx} - \gamma \frac{d^2u}{dx^2} + \delta \frac{d^3u}{dx^3} = q(x, t), \\ x \in [0, 1], \quad t \in (0, 2] \quad (12a)$$

$$q(x, t) = \sum_{j=1}^J A_j \sin \left( \omega_j t + \frac{2\pi l_j x}{L} + \varphi_j \right) \quad (12b)$$

with IC  $u(x, 0) = q(x, 0)$  and periodic BC  $u(0, t) = u(1, t)$ . Eq. 12 models the time evolution of a scalar quantity in the presence of physics phenomena such as convection ( $\frac{du}{dx}$  component), diffusion ( $\frac{d^2u}{dx^2}$  component), and dispersion ( $\frac{d^3u}{dx^3}$  component), and a rich diversity of dynamical processes can be generated from different PDE and IC combinations [45, 46].

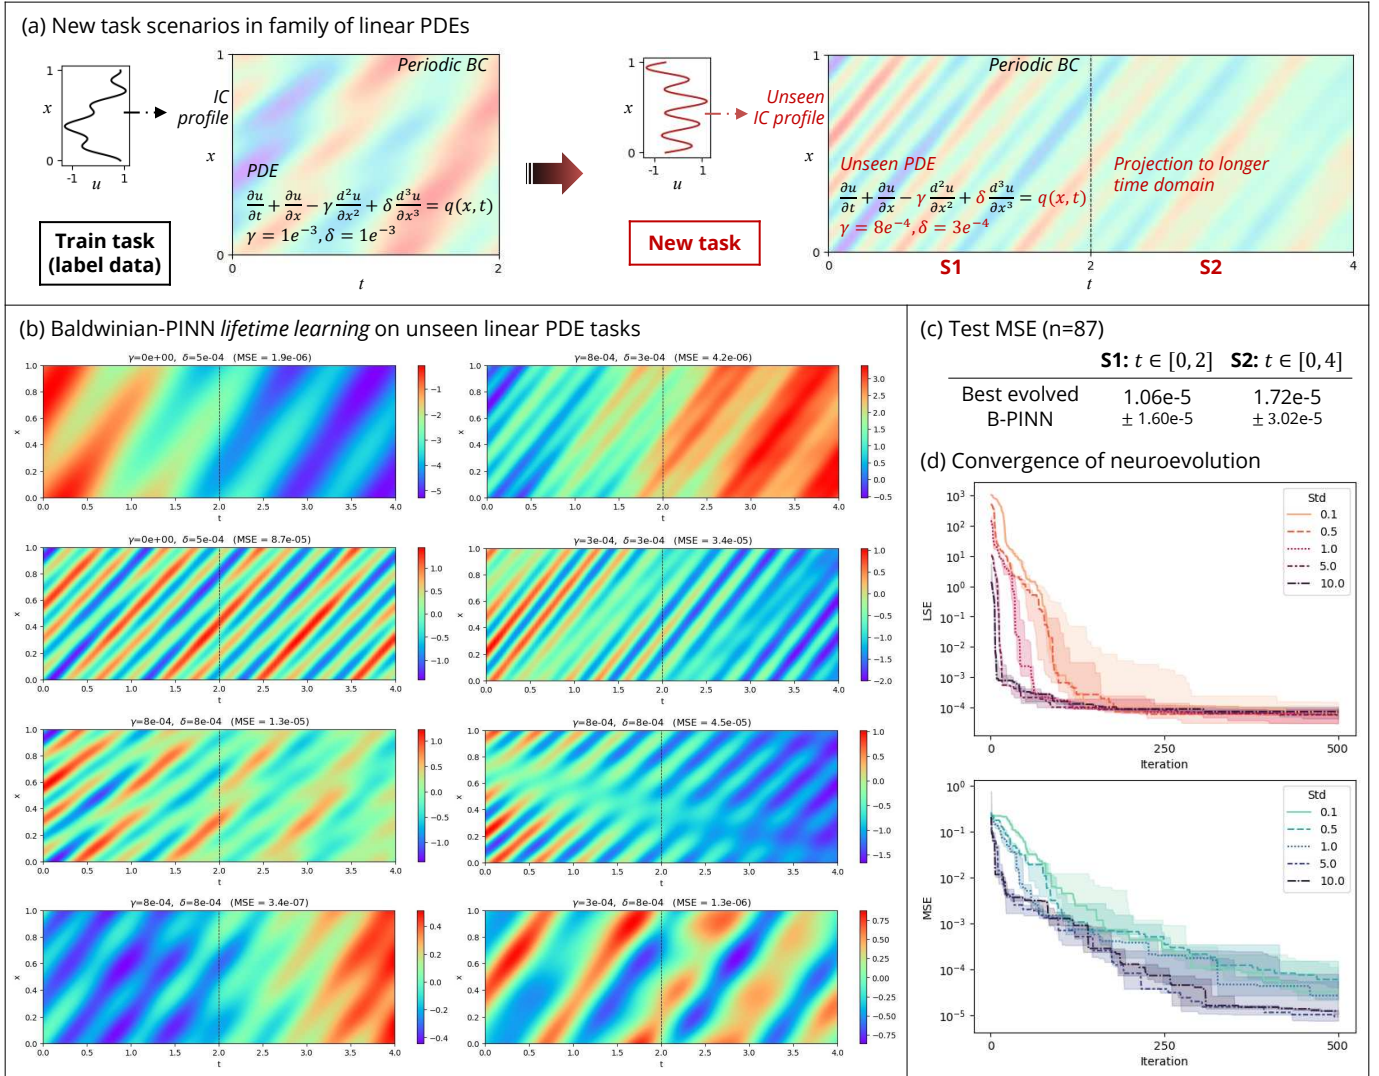


Figure 4. (a) Schematic to illustrate new tasks arising from family of linear PDEs problem: **S1** change to new PDE and IC profile for  $t \in [0, 2]$  (same time domain as train tasks), and **S2** projection to longer time domain  $t \in [0, 4]$ . (b) Solution for unseen linear PDE tasks obtained by best evolved Baldwinian-PINN sampled from the center of CMA-ES search distribution after 500 iterations with initial std. = 5 (they are visually indistinguishable from the ground truth). (c) The mean MSE over  $n = 87$  test tasks for 2 test scenarios described in (a) are below  $5e-5$ . (d) Baldwinian neuroevolution demonstrates effective LSE and MSE convergences for different CMA-ES initial std. values, with a superior performance given by std. = 5 and 10. The bold lines indicate their median convergence path from 5 individual runs, and the shaded areas indicate their interquartile ranges.

$q(x, t)$  is a source term, with different  $q(x, t = 0)$  being the corresponding IC profiles. We consider the following PDE scenarios:  $\alpha = 1$ ,  $\gamma = \{0, 5e-4, 1e-3\}$ ,  $\delta = \{0, 5e-4, 1e-3\}$ . The ratio between  $\alpha$ ,  $\gamma$ , and  $\delta$  determine non-dimensional constants (e.g.  $Pe$ ), and consequently, the systems' characteristic physics.  $q(x, t)$  comprises scenarios with  $J = 5$ ,  $L = 6$  and coefficients sampled uniformly from  $A_j \in [-0.8, 0.8]$ ,  $\omega_j \in [-2, 2]$ ,  $l_j \in [0, 1, 2, 3, 4]$ ,  $\varphi_j \in [-\pi, \pi]$ . The training set comprises 108 tasks with different PDE and IC combinations.

The effective generalization of a Baldwinian-PINN to an entire PDE family with diverse output patterns is demonstrated on 2 task scenarios: **S1** shows successful learning of  $u(x, t)$  for  $t \in [0, 2]$  on unseen set of PDEs, which include changes to PDE parameters ( $\gamma$  and  $\delta$ ) and the source term / IC  $q$ ; while **S2** shows effective extrapolation of solution to a longer time domain, i.e.,  $t \in [0, 4]$ . For both scenarios, the Baldwinian-PINNs learn the solution accurately in milli-seconds. The

average MSE given by a successfully evolved Baldwinian-PINN over all test tasks ( $n = 87$ ) for **S1** and **S2** are  $1.06e-5 \pm 1.60e-5$  and  $1.72e-5 \pm 3.02e-5$ , respectively. Illustrative results are in Figure 4.

Interestingly, the Baldwinian-PINNs maintain good accuracy on tasks from **S2**, whereby the evolved Baldwinian-PINN model first predicts  $u(x, t)$  for  $t \in [0, 2]$ , before using the solution  $u(x, t = 2)$  as new IC for  $t \in [2, 4]$ . This is achievable because Baldwinian neuroevolution does not require parameterization for the tasks, and Baldwinian-PINNs can generalize to new ICs, BCs, and PDE source terms. This is tricky for existing meta-PINN methods as they require interpolation across a potentially infinitely large distribution of tasks (e.g. possible BCs or ICs), in contrast to the lifetime learning encapsulated in the Baldwinian paradigm. Additional results illustrating the ability of Baldwinian-PINNs to perform well despite variations in the lifetime learning task objectives

(e.g., solving for different time durations) are in Suppl. S.III.A, further emphasizing the merits of Baldwinian neuroevolution for physics.

3) *Additional linear ODE/PDE problems:* The Baldwinian neuroevolution of physics is further demonstrated on 3 linear problems (Problems 3-5), namely 1D Poisson’s equation, 2D Poisson’s equation, and Helmholtz equation (see Suppl. S.IV). Suppl. Table S1 enumerates the accuracy advantages of Baldwinian-PINN relative to results reported by recent meta-learning PINN works [23, 25]. Suppl. S.IV also presents additional visualization results from Baldwinian-PINN, other instantiations of the proposed Baldwinian neuroevolution framework (such as natural evolution strategies (NES) variant [47]), and empirical study of Baldwinian-PINN performance when training-task labels are noisy or scarce, demonstrating the Baldwinian-PINNs’ versatility and generalizability.

### B. Learning to solve and generalize nonlinear ODE/PDEs

1) *Kinematics:* Extending beyond linear ODE/PDEs, Baldwinian-PINNs are applied to nonlinear kinematics equations. Assuming a ball is thrown at specific launch angle  $a_0$  and initial velocity  $vel_0$ , the following 2D kinematics equations describe the ball’s motion under the influence of gravity  $g$  and air resistance  $R$ :

$$\text{(Problem 6)} \quad \frac{d^2x}{dt^2} + R\frac{dx}{dt} = 0, \quad t \in (0, T] \quad (13a)$$

$$\frac{d^2y}{dt^2} + R\frac{dy}{dt} = -g, \quad t \in (0, T] \quad (13b)$$

subject to ICs  $x(t = 0) = 0$ ,  $\frac{dx}{dt}(t = 0) = vel_0 \times \cos(\frac{a_0\pi}{180})$ ;  $y(t = 0) = 0$ ,  $\frac{dy}{dt}(t = 0) = vel_0 \times \sin(\frac{a_0\pi}{180})$ . The air resistance  $R = \frac{1}{2} \frac{\rho C_d A}{m} V$  is related to air density  $\rho$ , object properties (drag coefficient  $C_d$ , cross-sectional area  $A$ , and mass  $m$ ), and object velocity  $V = \sqrt{(dx/dt)^2 + (dy/dt)^2}$ , hence the equations are nonlinear with respect to  $x$  and  $y$ . The 150 training tasks comprise different launch angles  $a_0 \in [15, 85]$ , initial velocity  $vel_0 \in [10, 110]$ , and object properties  $C_d \in [0.2, 0.7]$ ,  $A \in [0.00145, 0.045]$ ,  $m \in [0.046, 0.6]$  as may be representative of different projectiles (e.g., baseball, basketball).  $g$  and  $\rho$  are assumed to be 9.8 and 1.3 respectively.

The Baldwinian-PINN learns the horizontal and vertical position  $x(t)$  and  $y(t)$  of the object via iterative least-squares computation (see Section III-C) with a fixed number of nonlinear iterations  $N = 15$ . 100 test tasks encompassing different  $a_0 \in [5, 90]$ ,  $vel_0 \in [8, 98]$ , and  $C_d$  are constructed to assess generalizability. Results presented in Figure 5a-b show that Baldwinian-PINNs learn very accurate solutions (MSE =  $2.2e-8 \pm_{1.3e-7}$ ) in milli-seconds.

2) *Set of nonlinear PDEs:* The Baldwinian neuroevolution of physics is further demonstrated on 5 nonlinear PDE problems (Problems 7-11) such as the Burger’s, nonlinear Allen-Cahn and nonlinear reaction-diffusion equations, as per recent meta-learning PINN study [22] and described in Suppl. S.V. That study compares several meta-learning PINNs based on weight interpolation methods and MAML.

A key difference between this study and meta-learning of PINNs in [22] is the availability of data for training tasks. On the test tasks, the evolved Baldwinian-PINNs can achieve  $\sim 1$  order of magnitude lower relative error compared to [22]. Crucially, the computation time for Baldwinian-PINN for new predictions is at most two seconds whereas existing meta-learning approaches may take more than 500 seconds (2 orders of magnitude acceleration). For example, when solving the diffusion-reaction equation, a 70x improvement in accuracy was achieved relative to other state-of-the-art approaches, while taking 700x less computational time. Complete quantitative results are summarized in Suppl. Table S2.

The versatility of Baldwinian neuroevolution is also key here as we can accelerate the Baldwinian-PINN’s lifetime learning by using a much coarser discretization and smaller number of nonlinear iterations during training while still learning a good solution on test tasks with finer (e.g.,  $16\times$ ) discretization and more (e.g.,  $2\times$ ) nonlinear iterations.

3) *Navier-Stokes (N-S) equations for fluid flow:* We explore a more challenging physics scenario in computational fluid dynamics, the lid-driven cavity flow problem. The fluid flow inside a 2D unit square,  $x \in [0, 1]$ ,  $y \in [0, 1]$ , is driven by the top lid velocity ( $u_{lid} = 1$ ), and governed by the steady-state incompressible N-S equations for dependent variables velocity  $u, v$  and pressure  $p$ :

$$\nabla \cdot \vec{u} = 0 \quad (14a)$$

$$(\vec{u} \cdot \nabla)u = \frac{1}{Re} \nabla^2 \vec{u} - \nabla p \quad (14b)$$

Complex flow structures emerge at moderate to high Reynolds numbers, making this a challenging benchmark for PINNs, which may require hours of training with additional techniques such as transfer or curriculum learning to achieve convergence.

The Baldwinian neuroevolution of physics is performed on 7 training tasks  $Re \in [1, 800]$ . After 400 neuroevolution iterations, the evolved Baldwinian-PINNs can generalize to new solutions  $Re \in [5, 1000]$  with an average MSE of  $2.1e-4 \pm_{2.4e-4}$  ( $n = 18$  tasks  $\times$  5 individual runs). Results presented in Figure 5c-d and Suppl. S.VI Figure S9 demonstrate excellent agreement between Baldwinian-PINN predictions and the numerical benchmark from Ghia *et al.* [48]. Baldwinian-PINNs learn accurate solutions to the N-S equations in under one second, highlighting the potential to address difficult physics problems characterized by strong nonlinearity and strong coupling among multiple variables.

### C. Analysis on effectiveness of Baldwinian neuroevolution

We further investigate the advantages of Baldwinian neuroevolution via an ablation study using the *convection-diffusion* and *kinematics* examples.

Briefly, we explore deep and shallow MLP architectures for baseline SGD-trained DNN and PINN models for comparison with the Baldwinian-PINNs: **1.** the deep architecture consists of similar total number of network weights as the corresponding Baldwinian-PINN models, but distributed across multiple nonlinear hidden layers with smaller number of nodes; **2.** the shallow architecture has single nonlinear hidden layer and

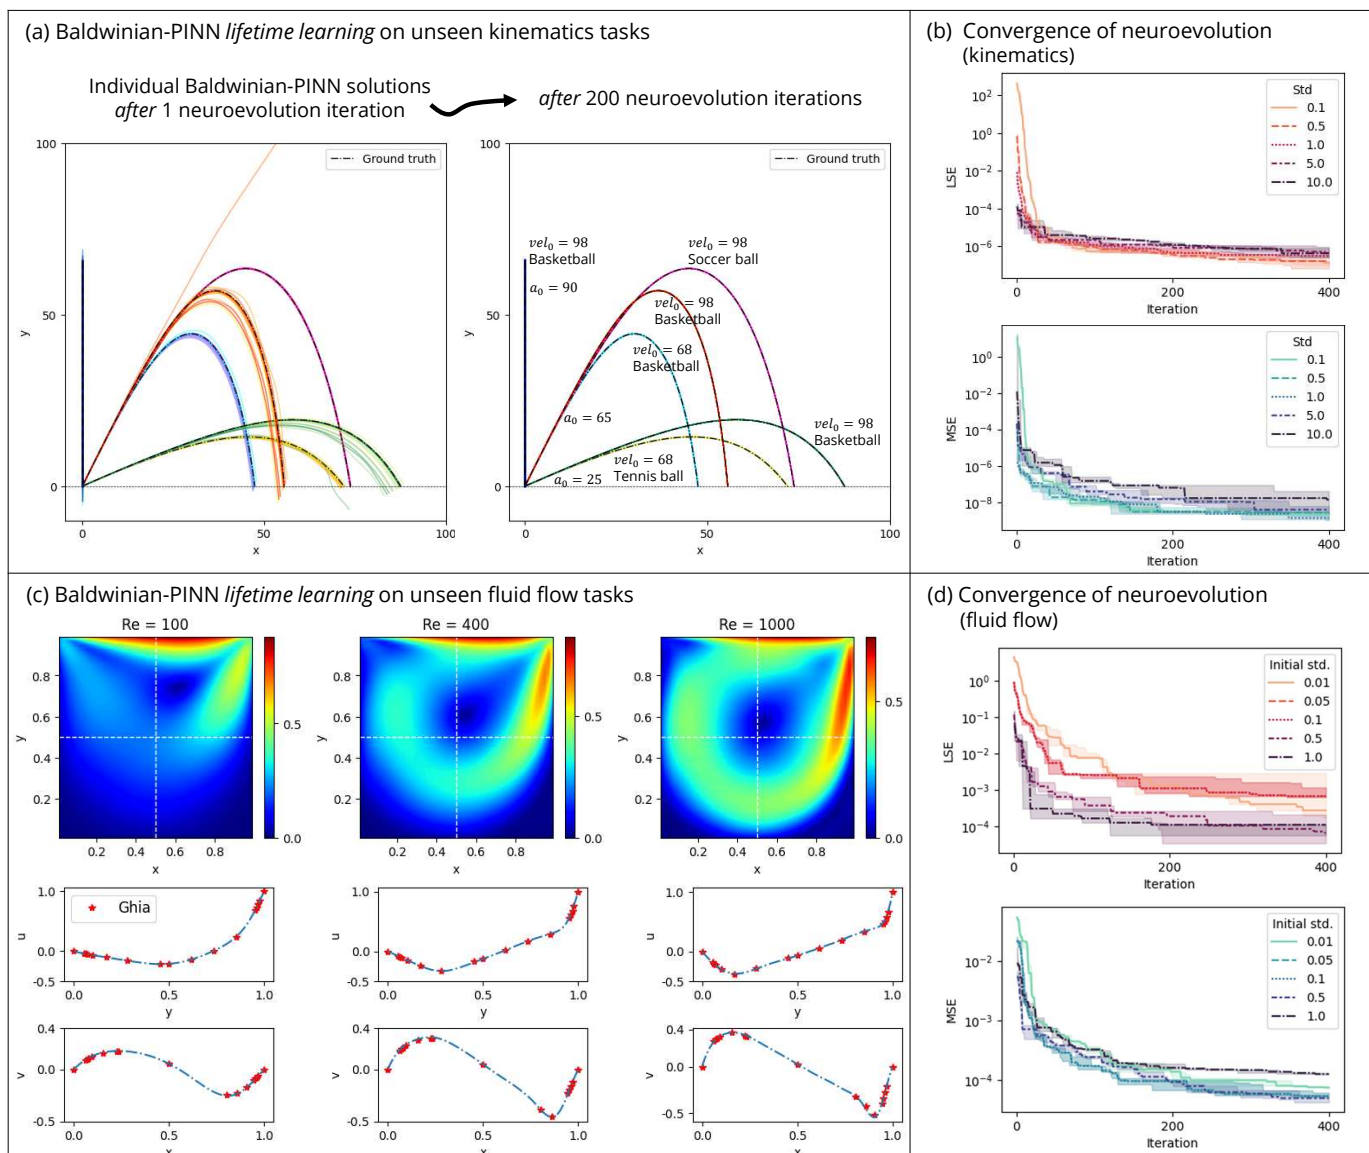


Figure 5. (a) Solution of 20 individual Baldwinian-PINN models sampled from the CMA-ES search distribution (initial std. = 0.5), for unseen kinematics tasks. (c) Solution of Baldwinian-PINN model sampled from the center of NES search distribution after 400 iterations (initial std. = 0.5), for unseen fluid flow (N-S equations) tasks. The velocity magnitude predictions are visually indistinguishable from the ground truth, even as Reynolds number (Re) increases. Their mid-section  $u$  and  $v$ -velocity profiles (dashed lines) show excellent agreement with the classic benchmark results from Ghia [48]. Baldwinian neuroevolution demonstrates effective LSE and MSE convergence on both (b) kinematics and (d) fluid flow problem, for different CMA-ES initial std. values. The bold lines indicate their median convergence path from 5 individual runs, and the shaded areas indicate their interquartile ranges.

same number of nodes as the corresponding Baldwinian-PINN model but this necessitates more total network weights because of the additional input variables.

Each architecture also consists of 2 variants: **a.**  $\tanh$  activation for the nonlinear hidden layers whereby the network weights are initialized by *Xavier* method; **b.**  $\sin$  activation for the nonlinear hidden layers whereby the network weights are initialized by *He* method.

This notation is maintained when referencing the respective model performance in Figure 6. For example, the models labelled as DNN-1a and PINN-2b refer to the corresponding baseline DNN model with a deep architecture and  $\tanh$  activation layer and baseline PINN model with a shallow architecture (single hidden layer) and  $\sin$  activation layer respectively. Additional model descriptions are in Suppl. S.VII.

1) *Direct prediction across tasks with parametric DNNs / PINNs:* As a baseline, DNN and PINN models are trained with SGD (ADAM) based on data-driven loss and PINN (data and physics) loss respectively, and applied to predictions for unseen tasks based on interpolation. We explore different MLP configurations as per Suppl. Table S3 in order to more fairly compare the best performance across different models.

Critically, when making a direct prediction for a new task, the respective baseline DNN and PINN models do not explicitly incorporate (and guarantee compliance with) the known physics prior. Although the parametric PINN can learn a more physically-consistent prediction from the physics loss during training, compliance is not guaranteed for a new task.

In addition, the new tasks must follow an *a priori* determined input parameterization for interpolation. For example,

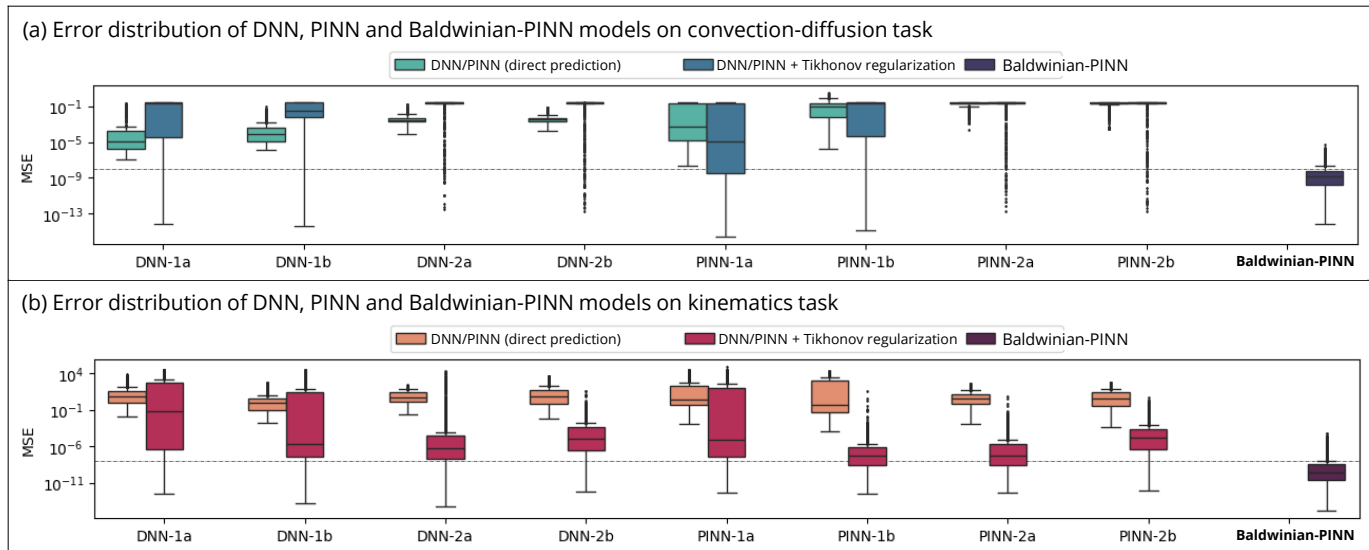


Figure 6. The generalization performance of the DNN, PINN, and Baldwinian-PINN models are compared, across (a) 110 convection-diffusion and (b) 100 kinematics test tasks. Model-1a/2a: deep/shallow architecture with  $\tanh$  activations (network weights initialized by *Xavier* method); Model-1b/2b: deep/shallow architecture with  $\sin$  activations (weights initialized by *He* method). The MSE results from each of the DNN and PINN models are pooled from 3 initial learning rate configurations  $\times$  5 individual runs. The MSE results from Baldwinian-PINN model are pooled from 5 initial std. values  $\times$  5 individual runs.

the DNN and PINN models for convection-diffusion need to be *a priori* set-up with  $(x, \alpha)$  as inputs to enable predictions across different  $\alpha$ 's. In contrast, the Baldwinian-PINN is parameterization-agnostic, and does not require any input parameterization for predictions of new tasks.

The *direct prediction* performance of DNNs and PINNs with different model configurations are summarized in Figure 6 and Suppl. Table S4. These baseline DNN and PINN predictions have high generalization errors (several orders of magnitude higher than Baldwinian-PINN). As explained above, PINNs do not necessarily outperform DNNs for direct prediction on test tasks (e.g., MSE for convection-diffusion for different model configurations range from  $1.7e-4$  to  $1.2e-2$  for DNN, and from  $1.9e-3$  to  $3.4e-1$  for PINN), suggesting that PINN training by itself does not always guarantee generalization in the absence of additional physics-informed fine-tuning.

2) *Prediction across tasks with parametric DNNs / PINNs and physics-informed fine-tuning*: In order to assess the utility of a physics-informed fine-tuning step, we optimize the final layer weights of the SGD-trained baseline DNNs and PINNs with the Tikhonov regularization at test time. This fine-tuning procedure is almost identical to the Baldwinian-PINN's lifetime learning.

The key difference is that the pre-final nonlinear hidden layers weight distribution and learning hyperparameter  $\lambda$  in Baldwinian-PINNs are jointly learned by evolution to share among tasks. In contrast, nonlinear hidden layers' weights for the DNNs and PINNs are SGD-learned from train tasks. As the Tikhonov regularization solution is highly dependent on  $\lambda$ , we need to perform a grid search across  $\lambda = \{1e-2, 1e-4, 1e-6, 1e-8, 1e-10, 1e-12, 0\}$  to determine the best  $\lambda$  during test time and present the solution associated with the lowest LSE. The results *after Tikhonov regularization* are compared in Figure 6 and Suppl. Table S4.

Results show that physics-informed fine-tuning via

Tikhonov regularization can improve the accuracy of both data-trained DNN and physics-trained PINN models at test time. However, their results are not as generalizable (i.e., the improvements are isolated to certain model and task instances) as the Baldwinian-PINN (which has been optimized for generating accurate solution over task distribution).

For example, we observe a significant improvement for the nonlinear kinematics problem (i.e., the best MSE for DNN and PINN improves from  $3.0$  to  $4.4e-5$  and  $1.3$  to  $7.9e-6$  respectively after Tikhonov regularization), potentially because the training distribution is sparser relative to the larger variation in output patterns (see Suppl. S.VIII).

In contrast, fine-tuning via Tikhonov regularization doesn't produce better results than direct prediction for the convection-diffusion problem. The MSE with and without Tikhonov regularization are  $2.3e-4$  and  $2.8e-2$ , and  $1.9e-3$  and  $7.2e-3$  for the best-configured DNN and PINN respectively. Tikhonov regularization cannot jointly minimize both PDE and IC/BC errors to sufficiently small levels, leading to inferior performance than direct prediction for some tasks.

While the best DNN model (after Tikhonov regularization) can be more accurate than the corresponding Baldwinian-PINN on some convection-diffusion tasks at lower  $\alpha$ , solution quality deteriorates quickly at larger  $\alpha$ . Interestingly, PINNs with the same configuration (after Tikhonov regularization) show the opposite trend, highlighting the challenge of obtaining a generalizable model for entire task distributions through SGD, in contrast to Baldwinian neuroevolution.

Overall, these results suggest that an additional step of physics-informed learning (i.e. Tikhonov regularization) on any new task can be beneficial. However, SGD-trained PINN and DNN models learn pre-final nonlinear hidden layers that are highly variable in their suitability for the physics-informed Tikhonov regularization-based fine-tuning across new tasks, in contrast to the proposed Baldwinian-PINNs. Hence, these

results highlight the advantages of Baldwinian-PINN framework as a complete solution to the challenges encountered in generalizing with physics-informed learning.

## V. CONCLUSION

In this paper, we study the Baldwin effect as a novel means to advancing the generalizability of PINNs over a family of governing differential equations. The Baldwin effect is instantiated through a two-stage stochastic programming formulation, wherein the first stage evolves the initial layers of a generalizable PINN model, and the second stage trains its final layer to specialize to any new physics task (analogous to lifetime learning). The Baldwinian neuroevolution biases pre-final hidden layers' weight distribution so that final layer specialization is both fast and accurate, preserving (near) constant-time solves while allowing capacity to scale. Although the pre-final hidden layers are not fine-tuned per task, the Baldwinian-PINN's capacity can be effectively scaled—by increasing width and/or depth—to handle more complex physics (e.g., fluid flow) while preserving fast test-time learning.

Our method is demonstrated to be broadly applicable to the learning of different linear and nonlinear ODE/PDEs encompassing diverse physical phenomena such as convection-diffusion, particle kinematics, heat and mass transfer, and fluid flow. Relative to recent meta-learned PINNs, Baldwinian-PINNs can accelerate the physics-aware predictions by several orders of magnitude, while improving the prediction accuracy by up to one order. A Baldwinian-PINN is thus in the image of a precocious species with accelerated learning ability at birth.

The lifetime learning encapsulated in the Baldwinian paradigm does not require *a priori* parameterization for the task scenarios, permitting both flexible generalization to new ICs, BCs, and PDE source terms and variations in the lifetime learning task objectives (e.g., different domain and/or sample size). This allows Baldwinian-PINNs to be useful in applications when large number of evaluations with *a priori* unknown input conditions are sought, e.g. generative design or what-if analysis. In addition, results in Section IV-A2 suggest that Baldwinian-PINNs could be suited to the continual modelling of dynamical systems through their versatility in handling ICs and ability to rapidly model and stitch time windows together with minimal error [49].

In the context of recent interest in *foundation models for scientific machine learning* [50], specifically through the use of neural operators or neural PDE solvers, our experiments showing accurate and fast generalization across families of linear and nonlinear ODE/PDEs suggest that Baldwinian learning can be an alternate route to such flexible and generalizable machine intelligence models. It will be interesting to test the limits to which Baldwinian-PINNs can learn across broad classes of physics phenomena and/or differential operators in future work. Hidden-layer fine-tuning adds flexibility but forfeits the closed-form update and re-introduces slow, nonconvex optimization. For large task shifts (out-of-distribution), a light hidden-layer adaptation (e.g., a few gradient-descent steps on a low-rank adapter) are a potential practical extension that trades some speed for additional flexibility.

Lastly, while the experiments in this work focus on optimizing the center and spread of the probability distributions that sample the initial weights in the neural network layers, this can be easily extended to incorporate other state-of-the-art neural architecture search approaches, which directly optimize the graph structure of the node connections. Our experiments indicate that significant improvement relative to other recent meta-learning PINN works can already be observed even when we only optimize the center and spread of the weight sampling distributions under the Baldwinian-PINN framework. The proposed Baldwinian-PINN framework can be seamlessly extended in future work, e.g. via integration with other state-of-the-art neural architecture search algorithms.

## ACKNOWLEDGMENT

This research was in part supported by the National Research Foundation, Singapore through the AI Singapore Programme, under the project “AI-based urban cooling technology development” (Award No. AISG3-TC-2024-014-SGKR). Abhishek Gupta was supported by the Ramanujan Fellowship from the Anusandhan National Research Foundation, Government of India (Grant No. RJF/2022/000115).

## REFERENCES

- [1] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational physics*, vol. 378, pp. 686–707, 2019.
- [2] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, “Scientific machine learning through physics-informed neural networks: Where we are and what’s next,” *Journal of Scientific Computing*, vol. 92, no. 3, p. 88, 2022.
- [3] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, “Physics-informed machine learning,” *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.
- [4] M. Raissi, A. Yazdani, and G. E. Karniadakis, “Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations,” *Science*, vol. 367, no. 6481, pp. 1026–1030, 2020.
- [5] B. Huang and J. Wang, “Applications of physics-informed neural networks in power systems—a review,” *IEEE Transactions on Power Systems*, 2022.
- [6] T. de Wolff, H. Carrillo, L. Martí, and N. Sanchez-Pi, “Assessing physics informed neural networks in ocean modelling and climate change applications,” in *AI: Modeling Oceans and Climate Change Workshop at ICLR 2021*, 2021.
- [7] S. Wang, S. Sankaran, H. Wang, and P. Perdikaris, “An expert’s guide to training physics-informed neural networks,” *arXiv preprint arXiv:2308.08468*, 2023.
- [8] J. C. Wong, C. Ooi, A. Gupta, and Y.-S. Ong, “Learning in sinusoidal spaces with physics-informed neural networks,” *IEEE Transactions on Artificial Intelligence*, 2022.
- [9] P.-H. Chiu, J. C. Wong, C. Ooi, M. H. Dao, and Y.-S. Ong, “Can-pinn: A fast physics-informed neural network based on coupled-automatic-numerical differentiation method,” *Computer Methods in Applied Mechanics and Engineering*, vol. 395, p. 114909, 2022.
- [10] J. C. Wong, P.-H. Chiu, C. Ooi, M. H. Dao, and Y.-S. Ong, “Lsa-pinn: Linear boundary connectivity loss for solving pdes on complex geometry,” in *2023 International Joint Conference on Neural Networks (IJCNN)*, 2023, pp. 1–10.
- [11] J. C. Wong, A. Gupta, C. C. Ooi, P.-H. Chiu, J. Liu, and Y.-S. Ong, “Evolutionary optimization of physics-informed

- neural networks: Evo-pinn frontiers and opportunities,” 2025. [Online]. Available: <https://arxiv.org/abs/2501.06572>
- [12] J. C. Wong, A. Gupta, and Y.-S. Ong, “Can transfer neuroevolution tractably solve your differential equations?” *IEEE Computational Intelligence Magazine*, vol. 16, no. 2, pp. 14–30, 2021.
  - [13] A. Krishnapriyan, A. Gholami, S. Zhe, R. Kirby, and M. W. Mahoney, “Characterizing possible failure modes in physics-informed neural networks,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 26 548–26 560, 2021.
  - [14] K. L. Downing, “Heterochronous neural baldwinism,” in *Artificial Life Conference Proceedings*. Citeseer, 2012, pp. 37–44.
  - [15] G. G. Simpson, “The baldwin effect,” *Evolution*, vol. 7, no. 2, pp. 110–117, 1953.
  - [16] W. B. Powell, “Clearing the jungle of stochastic optimization,” in *Bridging data and decisions*. Catonsville, Maryland, USA: Informs, 2014, pp. 109–137.
  - [17] H. Bakker, F. Dunke, and S. Nickel, “A structuring review on multi-stage optimization under uncertainty: Aligning concepts from theory and practice,” *Omega*, vol. 96, p. 102080, 2020.
  - [18] C. Fernando, J. Sygnowski, S. Osindero, J. Wang, T. Schaul, D. Teplyashin, P. Sprechmann, A. Pritzel, and A. Rusu, “Meta-learning by the baldwin effect,” in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, ser. GECCO ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1313–1320. [Online]. Available: <https://doi.org/10.1145/3205651.3208249>
  - [19] K. O. Stanley, J. Clune, J. Lehman, and R. Miikkulainen, “Designing neural networks through neuroevolution,” *Nature Machine Intelligence*, vol. 1, no. 1, pp. 24–35, 2019.
  - [20] R. Miikkulainen and S. Forrest, “A biological perspective on evolutionary computation,” *Nature Machine Intelligence*, vol. 3, no. 1, pp. 9–15, 2021.
  - [21] G. H. Golub, P. C. Hansen, and D. P. O’Leary, “Tikhonov regularization and total least squares,” *SIAM journal on matrix analysis and applications*, vol. 21, no. 1, pp. 185–194, 1999.
  - [22] M. Penwarden, S. Zhe, A. Narayan, and R. M. Kirby, “A metalearning approach for physics-informed neural networks (pinns): Application to parameterized pdes,” *Journal of Computational Physics*, vol. 477, p. 111912, 2023.
  - [23] X. Liu, X. Zhang, W. Peng, W. Zhou, and W. Yao, “A novel meta-learning initialization method for physics-informed neural networks,” *Neural Computing and Applications*, vol. 34, no. 17, pp. 14 511–14 534, 2022.
  - [24] Y. Chen and S. Koohy, “Gpt-pinn: Generative pre-trained physics-informed neural networks toward non-intrusive meta-learning of parametric pdes,” *Finite Elements in Analysis and Design*, vol. 228, p. 104047, 2024.
  - [25] W. Cho, K. Lee, D. Rim, and N. Park, “Hypernetwork-based meta-learning for low-rank physics-informed neural networks,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
  - [26] M. Toloubidokhti, Y. Ye, R. Missel, X. Jiang, N. Kumar, R. Shrestha, and L. Wang, “Dats: Difficulty-aware task sampler for meta-learning physics-informed neural networks,” in *The Twelfth International Conference on Learning Representations*, 2023.
  - [27] X. Huang, Z. Ye, H. Liu, S. Ji, Z. Wang, K. Yang, Y. Li, M. Wang, H. Chu, F. Yu *et al.*, “Meta-auto-decoder for solving parametric partial differential equations,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 23 426–23 438, 2022.
  - [28] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *International conference on machine learning*. PMLR, 2017, pp. 1126–1135.
  - [29] A. Nichol and J. Schulman, “Reptile: a scalable metalearning algorithm,” *arXiv preprint arXiv:1803.02999*, vol. 2, no. 3, p. 4, 2018.
  - [30] K. Hornik, M. Stinchcombe, and H. White, “Multilayer feedforward networks are universal approximators,” *Neural networks*, vol. 2, no. 5, pp. 359–366, 1989.
  - [31] R. Zhang, Y. Lan, G.-b. Huang, and Z.-B. Xu, “Universal approximation of extreme learning machine with adaptive growth of hidden nodes,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 2, pp. 365–371, 2012.
  - [32] S. Dong and Z. Li, “Local extreme learning machines and domain decomposition for solving linear and nonlinear partial differential equations,” *Computer Methods in Applied Mechanics and Engineering*, vol. 387, p. 114129, 2021.
  - [33] S. Dong and J. Yang, “On computing the hyperparameter of extreme learning machines: Algorithm and application to computational pdes, and comparison with classical and high-order finite elements,” *Journal of Computational Physics*, vol. 463, p. 111290, 2022.
  - [34] P. N. Suganthan and R. Katuwal, “On the origins of randomization-based feedforward neural networks,” *Applied Soft Computing*, vol. 105, p. 107239, 2021.
  - [35] A. Gupta and Y.-S. Ong, *Memetic computation: the main-spring of knowledge transfer in a data-driven optimization era*. Switzerland: Springer, 2018, vol. 21.
  - [36] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, “Automatic differentiation in machine learning: a survey,” *Journal of Machine Learning Research*, vol. 18, pp. 1–43, 2018.
  - [37] N. Hansen, “The cma evolution strategy: A tutorial,” *arXiv preprint arXiv:1604.00772*, 2016.
  - [38] Y. Ollivier, L. Arnold, A. Auger, and N. Hansen, “Information-geometric optimization algorithms: A unifying picture via invariance principles,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 564–628, 2017.
  - [39] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne *et al.*, “Jax: composable transformations of python+ numpy programs,” 2018.
  - [40] Y. Tang, Y. Tian, and D. Ha, “Evojax: Hardware-accelerated neuroevolution,” *arXiv preprint arXiv:2202.05008*, 2022.
  - [41] M. Stynes, “Steady-state convection-diffusion problems,” *Acta Numerica*, vol. 14, pp. 445–508, 2005.
  - [42] R. Van Erp, R. Soleimanzadeh, L. Nela, G. Kampitsis, and E. Matioli, “Co-designing electronics with microfluidics for more sustainable cooling,” *Nature*, vol. 585, no. 7824, pp. 211–216, 2020.
  - [43] S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*. New York, NY, USA: Hemisphere Publishing Corporation, 1980.
  - [44] A. Gupta, “Numerical modelling and optimization of non-isothermal, rigid tool liquid composite moulding processes,” Ph.D. dissertation, ResearchSpace@ Auckland, 2013.
  - [45] Y. Bar-Sinai, S. Hoyer, J. Hickey, and M. P. Brenner, “Learning data-driven discretizations for partial differential equations,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 31, pp. 15 344–15 349, 2019.
  - [46] J. Brandstetter, D. Worrall, and M. Welling, “Message passing neural pde solvers,” *arXiv preprint arXiv:2202.03376*, 2022.
  - [47] M. Nomura and I. Ono, “Fast moving natural evolution strategy for high-dimensional problems,” in *2022 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2022, pp. 1–8.
  - [48] U. Ghia, K. N. Ghia, and C. Shin, “High-re solutions for incompressible flow using the navier-stokes equations and a multigrid method,” *Journal of computational physics*, vol. 48, no. 3, pp. 387–411, 1982.
  - [49] X. Meng, Z. Li, D. Zhang, and G. E. Karniadakis, “Ppinn: Parareal physics-informed neural network for time-dependent pdes,” *Computer Methods in Applied Mechanics and Engineering*, vol. 370, p. 113250, 2020.
  - [50] S. Subramanian, P. Harrington, K. Keutzer, W. Bhimji, D. Morozov, M. Mahoney, and A. Gholami, “Towards foundation models for scientific machine learning: Characterizing scaling and transfer behavior,” *arXiv preprint arXiv:2306.00258*, 2023.

Table I

SUMMARY OF BALDWINIAN NEUROEVOLUTION AND BALDWINIAN-PINN’S LIFETIME LEARNING CONFIGURATIONS, AND GENERALIZATION PERFORMANCE ON PINN PROBLEMS IN EXPERIMENTAL STUDY.

Task distribution		Baldwinian neuroevolution w/ CMA-ES or NES						Baldwinian-PINN lifetime learning			Predictive performance			
Problem	No. train tasks	Batch size for task, $n_{task}$	Population size, $n_{pop}$	Max. iteration	Initial std.	Search dimension for $(\theta, \lambda)$	Time <sup>1</sup> (s)	dim( $w$ )	Sample size $(n_{pde}, n_{ic}, n_{bc})$	nonlinear iterations, $N$	No. test tasks	Time <sup>2</sup> (s)	MSE <sup>3</sup>	
1	Convection-diffusion $x \rightarrow u$	20	10	20	200	1	25	250	900	1001, -, 2	-	110	0.005	5.8e-9 $\pm$ 9.9e-9
2	Family of linear PDEs $(x, t) \rightarrow u$	108	15	20	500	5	37	10830	1200	5151, 101, 50 <small>(10201, 101, 100)<sup>4</sup></small>	-	87	0.122	7.1e-4 $\pm$ 3.1e-3
3	1D Poisson’s equation $x \rightarrow u$	60	10	20	100	1	25	120	900	1001, -, 2	-	60	0.019	1.0e-9 $\pm$ 5.8e-9
4	2D Poisson’s equation $(x, y) \rightarrow u$	100	20	20	100	1	37	270	900	1089, -, 128 <small>(16641, -, 512)<sup>5</sup></small>	-	100	0.132	8.9e-12 $\pm$ 1.0e-11
5	Helmholtz equation $(x, y) \rightarrow u$	20	10	20	400	5	38	520	900	1024, -, 124 <small>(16384, -, 508)<sup>5</sup></small>	-	60	0.131	1.6e-5 $\pm$ 5.1e-5
6	Kinematics $t \rightarrow (x, y)$	150	30	20	400	0.5	25	1600	1800	101, 1, -	15	100	0.039	2.2e-8 $\pm$ 1.3e-7
7	Burgers’ equation $(x, t) \rightarrow u$	16	5	20	200	1	37	5240	900	13107, 257, 100 <small>(25957, 257, 200)<sup>5</sup></small>	5 <sub>(10)<sup>6</sup></sub>	32	1.96	2.3e-7 $\pm$ 9.6e-7
8	Nonlinear heat $(x, t) \rightarrow u$	13	5	20	100	5	37	440	900	1600, 64, 50 <small>(25600, 256, 200)<sup>5</sup></small>	5	64	0.87	1.3e-7 $\pm$ 4.1e-7
9	Allen-Cahn equation $(x, y) \rightarrow u$	16	8	20	100	0.5	37	510	900	1024, -, 128 <small>(16384, -, 512)<sup>5</sup></small>	5 <sub>(10)<sup>6</sup></sub>	32	1.19	2.0e-7 $\pm$ 6.1e-7
10	Diffusion-reaction $(x, y) \rightarrow u$	22	10	20	100	1	37	640	900	1024, -, 128 <small>(16384, -, 512)<sup>5</sup></small>	5	64	0.58	1.2e-8 $\pm$ 8.1e-8
11	6D diffusion-reaction $(x, y) \rightarrow u$	17	8	20	100	1	37	510	900	1024, -, 128 <small>(16384, -, 512)<sup>5</sup></small>	5	100	0.60	1.6e-8 $\pm$ 2.5e-8
12	Fluid flow $(x, y) \rightarrow (u, v, p)$	7	4	20	400	0.1	66	15868	1152	2401, -, 196	5	18	0.903	2.1e-4 $\pm$ 2.4e-4

<sup>1</sup>Baldwinian neuroevolution time cost, on 2 GPUs (NVIDIA GeForce RTX 3090).<sup>2</sup>Computation time per task, on single GPU (NVIDIA GeForce RTX 3090).<sup>3</sup>MSE results are aggregated from 5 individual runs.<sup>4</sup>Prediction of test task includes projection to a longer (2 $\times$ ) time domain.<sup>5</sup>Prediction of test task on denser sample points.<sup>6</sup>Prediction of test task with more nonlinear iterations.

# Supplementary Material for Evolutionary Optimization of Physics-Informed Neural Networks: Advancing Generalizability by the Baldwin Effect

Jian Cheng Wong, Chin Chun Ooi, Abhishek Gupta, *Senior Member, IEEE*, Pao-Hsiung Chiu,  
Joshua Shao Zheng Low, My Ha Dao, and Yew-Soon Ong, *Fellow, IEEE*

## S.I. DATA GENERATION

In this study, the majority of the problems have corresponding analytical solutions as described in the problem setups. In addition, the ground truth for the PDE family (Problem 2) and Burgers' equation (Problem 7) is obtained by a high-resolution finite volume scheme. To alleviate convection instability, the dispersion-relation-preserving (DRP) finite volume scheme with a universal limiter has been utilized [S1, S2]. Other spatial derivative terms are discretized by central difference. For the temporal term, the second-order TVD Runge Kutta scheme [S3] is employed. For the incompressible Navier-Stokes (N-S) equations (Problem 12), the dispersion-relation-preserving (DRP) finite volume scheme, together with improved divergence-free-condition compensated framework [S1], is employed to produce high-fidelity solutions as well as ensure the velocity-pressure coupling.

For the PDE family equation, the spatial resolution,  $\Delta x$ , is  $1/400$ , while temporal resolution,  $\Delta t$ , is  $5e-5\Delta x$ ; For the Burgers' equation, the spatial resolution,  $\Delta x$ , is  $1/512$ , while temporal resolution  $\Delta t$ , is  $1e-2\Delta x$ ; for the incompressible N-S equations, the spatial resolution,  $\Delta x$  and  $\Delta y$ , are  $1/200$ .

## S.II. LAGGING OF COEFFICIENT METHOD FOR NONLINEAR PDES

For nonlinear PDEs, iterative methods can be used for arriving at optimized Baldwinian-PINNs. In this work, we use a *lagging of coefficient approach* which is common in numerical methods [S4]. Briefly, we approximately linearize the nonlinear term(s) in Eq. 1 by substituting the output  $u(x, t) = \sum_j w_j f_j(x, t; \tilde{\mathbf{w}})$  obtained from previous step, and

Jian Cheng Wong, Chin Chun Ooi, and Pao-Hsiung Chiu are with the Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore (e-mail: wongj@a-star.edu.sg; oiccc@a-star.edu.sg; chiuph@a-star.edu.sg).

Abhishek Gupta is with the School of Mechanical Sciences, Indian Institute of Technology Goa, India (e-mail: abhishekgupta@iitgoa.ac.in).

Joshua Shao Zheng Low is with the College of Computing and Data Science, Nanyang Technological University, Singapore (e-mail: joshualow188@gmail.com).

My Ha Dao is with the Technology Centre for Offshore and Marine, Singapore (e-mail: Dao\_My\_Ha@tcoms.sg).

Yew-Soon Ong is with the Agency for Science, Technology and Research, Singapore, and is also with the College of Computing and Data Science, Nanyang Technological University, Singapore (e-mail: Ong\_Yew\_Soon@a-star.edu.sg).

iteratively solve Eq. 8 to update  $w_j$ 's for a fixed number of steps,  $N$ , or until a convergence criterion is reached.

The nonlinear equation  $(\frac{d^2u}{dx^2} + \frac{d^2u}{dy^2}) + u(1 - u^2) = f$  is used to demonstrate the *lagging of coefficient* method for computing the  $w_j$ 's in Baldwinian-PINNs. We approximately linearize the nonlinear term  $u(1 - u^2)$  as  $u(1 - \tilde{u}^2)$  with  $\tilde{u} = 0$  being the initial guess solution at first iteration. The ( $i$ -th PDE sample,  $j$ -th neuron) entry of the least squares matrix  $\mathbf{A}$  in Eq. 7 (Section III-C) now becomes:

$$\begin{aligned} \mathcal{N}_\theta[f_j(x_i^{pde}, y_i^{pde}; \tilde{\mathbf{w}}_j)] &= \frac{d^2 f_j(x_i^{pde}, y_i^{pde}; \tilde{\mathbf{w}}_j)}{dx^2} \\ &+ \frac{d^2 f_j(x_i^{pde}, y_i^{pde}; \tilde{\mathbf{w}}_j)}{dy^2} \\ &+ f_j(x_i^{pde}, y_i^{pde}; \tilde{\mathbf{w}}_j)(1 - \tilde{u}^2) \end{aligned} \quad (S1)$$

Hence, the Tikhonov regularization solution  $\mathbf{w}$  can be obtained from the linearized version of the equation. The solution obtained from past iteration is used to compute  $\tilde{u}(x, y) = \sum_j w_j f_j(x, y; \tilde{\mathbf{w}}_j)$  for the next iteration, until the solution  $\mathbf{w}$  reaches a specified convergence criterion or reaches a pre-determined number of iterations,  $N$ .

## S.III. ADDITIONAL RESULTS FOR FAMILY OF LINEAR PDE PROBLEM

### A. Prediction on new time interval

The experimental results in Section IV-A2 show that Baldwinian-PINNs trained on a set of linear PDE tasks for  $t \in [0, 2]$ , are capable of learning time-dependent solution  $u(x, t)$  on a set of test tasks for unseen PDEs and ICs, as well as for a longer ( $2 \times$ ) time domain by performing the learning twice, i.e., for  $t \in [0, 2]$  and then using the learned solution  $u(x, t = 2)$  as new IC for  $t \in [2, 4]$ . Recall that the average MSE given by the best evolved Baldwinian-PINN model (sampled from the center of CMA-ES search distribution from the best run) over all test tasks for  $t \in [0, 2]$  and  $t \in [0, 4]$  are  $1.06e-5 \pm 1.60e-5$  and  $1.72e-5 \pm 3.02e-5$ , respectively.

We further demonstrate the versatility of Baldwinian-PINNs by changing the time domain of interest in the test tasks to  $t \in [0, 3]$ . From the best evolved Baldwinian-PINN model, we can obtain the prediction on new tasks in the following ways: **P0** the solution for the original time domain  $t \in [0, 2]$

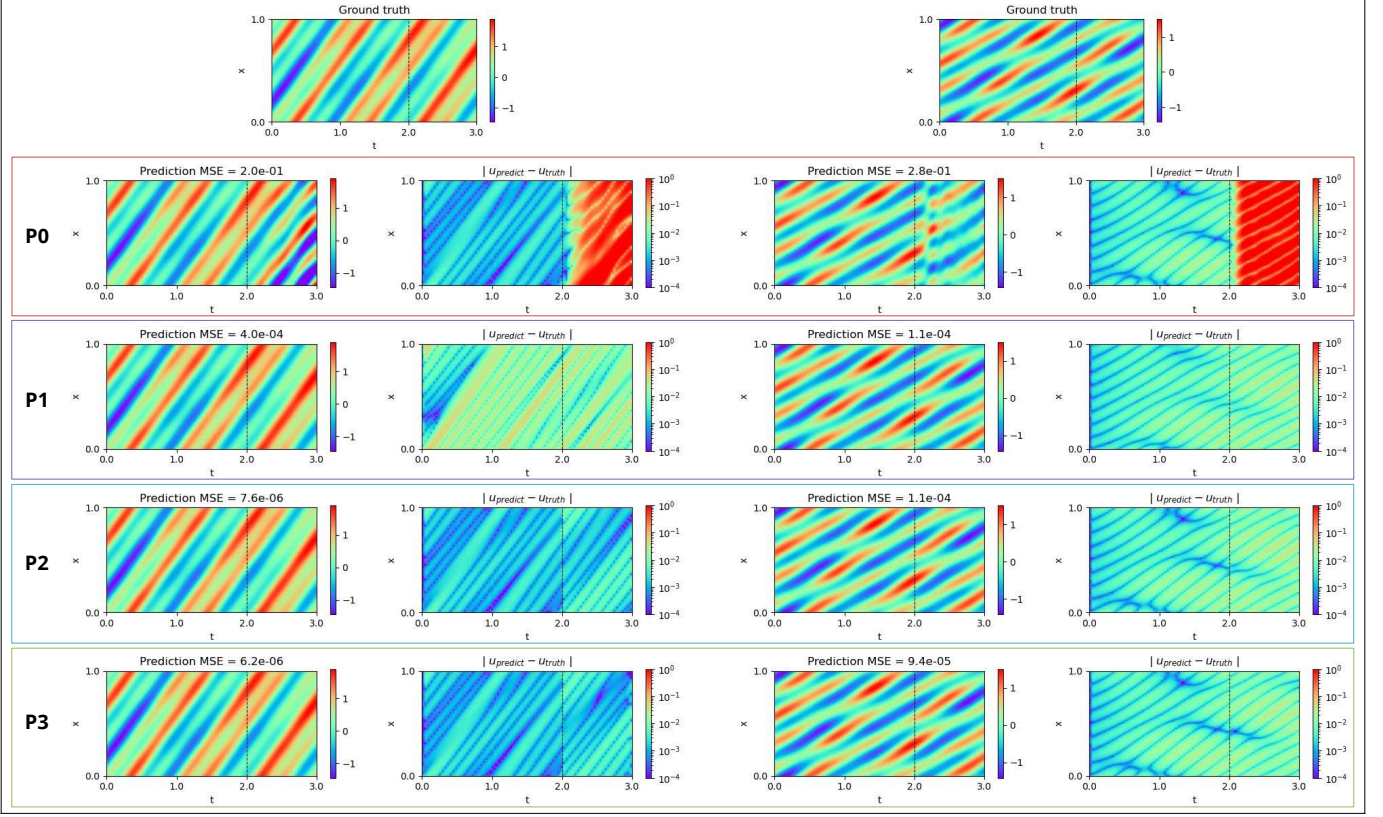


Figure S1. Solution for unseen family of linear PDE tasks on  $t \in [0, 3]$  obtained by best evolved Baldwinian-PINN (sampled from the center of CMA-ES search distribution after 500 iterations with initial std. = 5) using different ways **P0-P3**.

via physics-based lifetime learning, and the extended time  $t \in [2, 3]$  based on neural network interpolation; **P1** the solution for original and extended time domain  $t \in [0, 3]$  altogether via physics-based lifetime learning; **P2** physics-based learning of the solution for the first 2s time  $t \in [0, 2]$ , before using the learned solution  $u(x, t = 2)$  as new IC for next 2s time  $t \in [2, 4]$ ; **P3** physics-based learning of the solution for the first 2s time  $t \in [0, 2]$ , before using the learned solution  $u(x, t = 2)$  as new IC for next time domain  $t \in [2, 3]$ .

Figure S1 shows the Baldwinian-PINN’s solution obtained using different ways **P0-P3** on selected test tasks. Their MSE results over all test tasks for  $t \in [0, 3]$  are  $5.35e-2 \pm 6.08e-2$  (**P0**),  $4.24e-5 \pm 1.04e-5$  (**P1**),  $1.47e-5 \pm 2.32e-5$  (**P2**), and  $1.43e-5 \pm 2.22e-5$  (**P3**), respectively. The performance of **P0** is significantly worse than **P1-P4**, as expected, because of the physics-agnostic extrapolation. The Baldwinian-PINNs can flexibly learn the solutions on extended time domain in a single Tikhonov regularization solve (**P1**), although the most accurate solutions are given by **P2** and **P3**.

### B. Performance variation across runs

The Baldwinian neuroevolution outcome has the most variation in accuracy across individual runs, on the family of linear PDEs problem relative to the other problems in the experimental studies. The Baldwinian-PINN MSE results over all test tasks for  $t \in [0, 4]$  obtained by 5 different runs are  $5.3e-5 \pm 1.8e-4$ ,  $1.7e-5 \pm 3.0e-5$ ,  $2.0e-3 \pm 5.6e-3$ ,  $1.5e-3$

$\pm 3.5e-3$ , and  $2.8e-5 \pm 9.1e-5$ , respectively. Figures S2 and S3 compare the solutions of 2 Baldwinian-PINNs obtained from separate Baldwinian neuroevolution runs, on the same 5 test tasks selected at different levels of accuracy along the MSE spectrum (pooled from  $n = 87$  test tasks  $\times$  5 individual runs with initial std. = 5). In all our other experiments, the overall variation across runs remains fairly small.

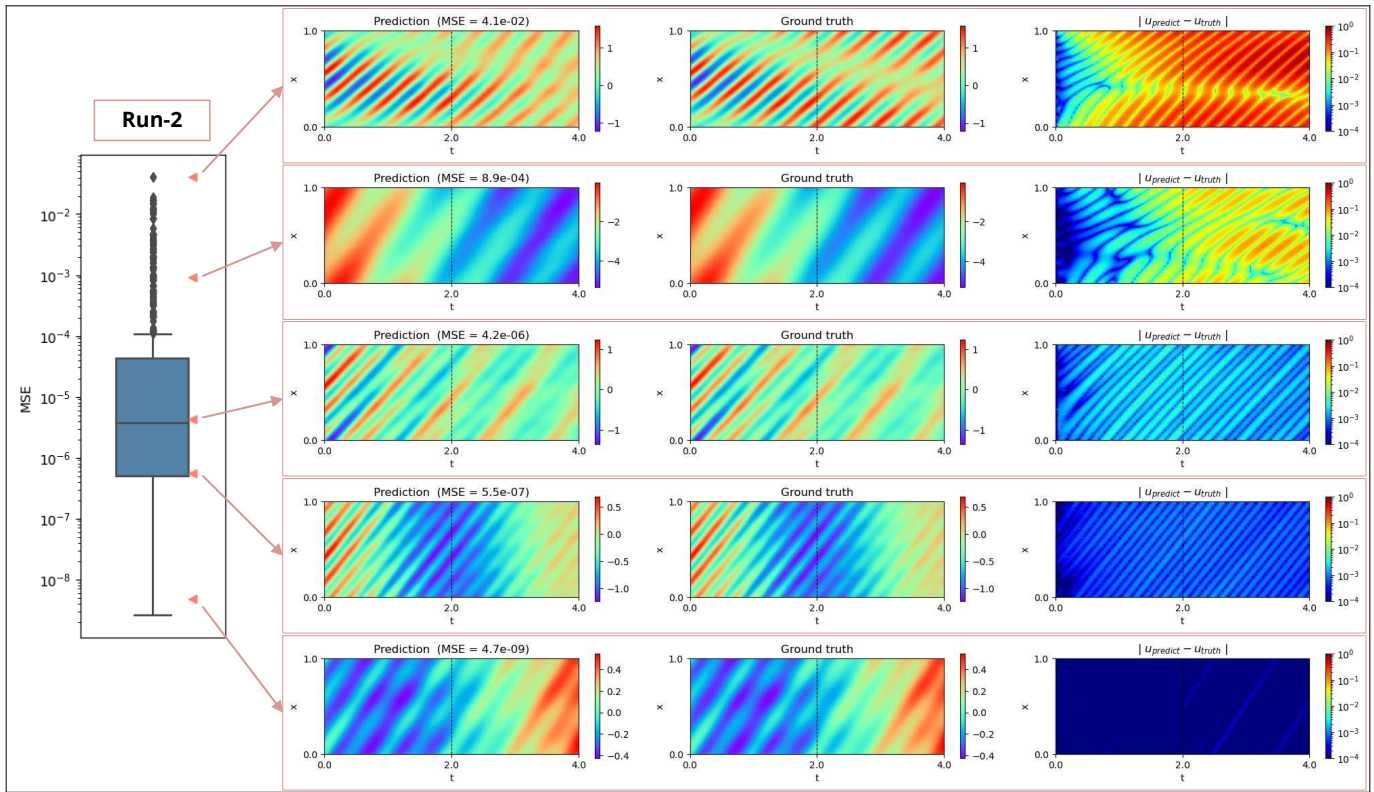


Figure S2. Baldwinian-PINN's solution (run no. 2) vs. ground truth on 5 selected family of linear PDE tasks, and the position of their accuracy along the MSE spectrum (pooled from  $n = 87$  test tasks  $\times$  5 individual runs with initial std. = 5).




Figure S3. Baldwinian-PINN's solution (run no. 4) vs. ground truth on 5 selected family of linear PDE tasks, and the position of their accuracy along the MSE spectrum (pooled from  $n = 87$  test tasks  $\times$  5 individual runs with initial std. = 5).

Table S1  
COMPARISON WITH BASELINE META-LEARNING PINN MODELS ON SINGLE TEST TASK FOR LINEAR ODE/PDE PROBLEMS.

Problem	Method / Model						Remarks
	NRPINN		Hyper-LR-PINN		Baldwinian-PINN		
	MAE	MSE	MAE	MSE	MAE	MSE	
3 1D Poisson's equation	5.1e-4 <sup>1</sup>	-	-	-	2.5e-7	8.4e-14	<p>Test task: <math>\alpha_1 = 1, \alpha_2 = 1, \alpha_3 = 0.1, \alpha_4 = 0, \omega_1 = 0.7, \omega_2 = 1.5, x \in [-10, 10]</math></p> <p>NRPINN: 60 train tasks from <math>\alpha_1 = 1, \alpha_2 = 1, \alpha_3 = 0.1, \alpha_4 = 0, \omega_1 \in [0, 1], \omega_2 \in [0.2]</math> as per [S5]</p> <p>Baldwinian-PINN: 60 train tasks from <math>\alpha</math>'s <math>\in [0, 4], \omega</math>'s <math>\in [0, 4]</math></p>
4 2D Poisson's equation	6.4e-7 <sup>2</sup>	-	-	-	3.0e-7	2.0e-13	<p>Test task: <math>a_1 = 0.15, a_2 = 0.18, a_3 = 0.20, a_4 = 0.31, a_5 = 0.43, a_6 = 0.56, a_7 = 0.70, a_8 = 0.80, b_1 = 0.34, b_2 = 0.31, b_3 = 0.65, b_4 = 0.86, b_5 = 0.65, b_6 = 0.38, b_7 = 0.64, b_8 = 0.12, c_1 = 0.84, c_2 = 1.07, c_3 = 1.12, c_4 = 0.83, c_5 = 1.12, c_6 = 1.11, c_7 = 0.99, c_8 = 0.91, d_j = 0.01, x \in [0, 1], y \in [0, 1]</math></p> <p>NRPINN: 100 train tasks from <math>J \in \{1, 5, 10\}, a_j \in [0.1, 0.9], b_j \in [0.1, 0.9], c_j \in [0.8, 1.2], d_j = 0.01</math> as per [S5]</p> <p>Baldwinian-PINN: 100 train tasks from <math>J \in [1, 10], a_j \in [-0.6, 0.6], b_j \in [-0.6, 0.6], c_j \in [0.5, 2], d_j \in [0.005, 0.02]</math></p>
5 Helmholtz equation	-	-	2.8e-2 <sup>3</sup>	-	1.5e-4	4.7e-8	<p>Test task: <math>\alpha_1 = 2.5, \alpha_2 = 2.5, x \in [-1, 1], y \in [-1, 1]</math></p> <p>Hyper-LR-PINN: train tasks from <math>\alpha \in [2, 3]</math> with interval 0.1 (<math>\alpha = \alpha_1 = \alpha_2</math>) as per [S6]</p> <p>Baldwinian-PINN: 20 train tasks from <math>\alpha_1 \in [0.1, 6], \alpha_2 \in [0.1, 6]</math></p>

<sup>1</sup>MAE obtained after fine-tuning with 900 training iterations; result extracted from [S5].

<sup>2</sup>MAE obtained after fine-tuning with 4000 training iterations; result extracted from [S5].

<sup>3</sup>MAE obtained after fine-tuning with 10 training epochs; result extracted from [S6].

#### S.IV. STUDIES ON SET OF LINEAR ODE/PDE PROBLEMS

The additional linear ODE/PDE problems (Problems 3-5) are described below.

1) *1D Poisson's equation*: The 1D Poisson's equation consists of 60 randomly sampled train tasks  $\alpha$ 's  $\in [0, 4]$ ,  $\omega$ 's  $\in [0, 4]$  and 60 randomly sampled test tasks  $\alpha$ 's  $\in [-5, 5]$ ,  $\omega$ 's  $\in [-5, 5]$  for the PDE/BC parameters  $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \omega_1, \omega_2)$ :

$$\text{(Problem 4)} \quad \frac{d^2u}{dx^2} = q, \quad x \in [-10, 10] \quad (\text{S2})$$

where the exact solution  $u(x; \alpha_1, \alpha_2, \alpha_3, \alpha_4, \omega_1, \omega_2) = \alpha_1 \sin(\omega_1 x) + \alpha_2 \sin(\omega_2 x) - \alpha_3 x + \alpha_4$  is used to derive the corresponding BCs and source term  $q$ .

We further test the evolved Baldwinian-PINN on a single test task ( $\alpha_1 = 1, \alpha_2 = 1, \alpha_3 = 0.1, \alpha_4 = 0, \omega_1 = 0.7, \omega_2 = 1.5$ ) as described in [S5].

2) *2D Poisson's equation*: The 2D Poisson's equation is represented by:

$$\text{(Problem 5)} \quad - \left( \frac{d^2u}{dx^2} + \frac{d^2u}{dy^2} \right) = q(x, y), \\ x \in [-1, 1], \quad y \in [-1, 1] \quad (\text{S3a})$$

$$q(x, y) = \sum_{j=1}^J c_j \exp \left( \frac{(x - a_j)^2 + (y - b_j)^2}{d_j} \right) \quad (\text{S3b})$$

subject to BCs  $u(x = -1) = 0, u(x = 1) = 0, u(y = -1) = 0, u(y = 1) = 0$ . The heat source  $q(x, y)$  is generated by the following scenarios: the number of heat source  $J \in [1, 10]$ , and coefficients sampled uniformly from  $a_j \in [-0.6, 0.6]$ ,  $b_j \in [-0.6, 0.6]$ ,  $c_j \in [0.5, 2]$ , and  $d_j \in [0.005, 0.02]$ . The training and test sets both comprise 100 tasks with different source  $q(x, y)$  scenarios.

We further test the evolved Baldwinian-PINN on a single test task as described in [S5], which is subject to a different domain  $x \in [0, 1], y \in [0, 1]$  and BCs  $u(x = 0) = 0, u(x = 1) = 0, u(y = 0) = 0, u(y = 1) = 0$ , with  $J = 8$  and  $a_j \in [0.1, 0.9]$ ,  $b_j \in [0.1, 0.9]$ ,  $c_j \in [0.8, 1.2]$ , and  $d_j = 0.01$ .

3) *Helmholtz equation*: The Helmholtz equation consists of 20 randomly sampled train tasks and 20 randomly sampled test tasks from the PDE/BC parameters  $\alpha_1 \in (0, 6], \alpha_2 \in (0, 6]$ :

$$\text{(Problem 6)} \quad \left( \frac{d^2u}{dx^2} + \frac{d^2u}{dy^2} \right) + u^2 = q, \\ x \in [-1, 1], \quad y \in [-1, 1] \quad (\text{S4})$$

where the exact solution  $u(x, y; \alpha_1, \alpha_2) = \sin(\alpha_1 \pi x) \sin(\alpha_2 \pi y)$  is used to derive the corresponding BCs and source term  $q$  [S6, S7].

We further test the evolved Baldwinian-PINN on a single test task ( $\alpha_1 = 2.5, \alpha_2 = 2.5$ ) as described in [S6].

#### A. Results for set of linear ODE/PDE problems

Table S1 compares the generalization performance between Baldwinian-PINNs and baseline meta-learning PINN models (NRPINN [S5] and Hyper-LR-PINN [S6]) on linear benchmark problems.

Figure S4 and Figure S5 provide visualization of the Baldwinian-PINN solutions and errors for the *2D Poisson's* and *Helmholtz* tasks, respectively. The visualization results show good performance on the diverse set of PDE tasks, as exhibited by source patterns and frequencies. Note that the train and test task distributions as used in this work have much greater diversity than those studied in prior meta-learning PINN works such as [S5–S7]. Our PDE parameters' range for the Helmholtz problem is 6 times larger than [S6, S7], thereby encapsulating a much broader frequency spectrum, including tasks in the higher frequency range which are also more challenging to learn.

To demonstrate the diversity of the tasks, we compare the performance of different Baldwinian-PINN models, with the first one learning from  $n = 10$  lower frequency train tasks  $\alpha_1, \alpha_2 \in (0, 1]$  for 200 neuroevolution iterations (train MSE  $< 1e-11$ ); while the second one is learning from  $n = 10$  higher frequency train tasks  $\alpha_1, \alpha_2 \in [5, 6]$  for 200 neuroevolution iterations (train MSE  $< 1e-6$ ). Both low- and high-frequency-learned models are applied to the same set of  $n = 60$  test tasks drawing from the full frequency range  $\alpha_1, \alpha_2 \in (0, 6]$ , and their test MSE results are  $4.8e-2 \pm 1.5e-1$  and  $3.8e-2 \pm 8.7e-2$ , respectively.

Their test MSE results are 2-3 orders of magnitudes higher than the  $1.6e-5 \pm 5.1e-5$  achieved by the original Baldwinian-PINN. Note that the original model is learned from  $n = 20$  train tasks  $\alpha_1, \alpha_2 \in (0, 6]$  for 400 neuroevolution iterations (train MSE  $< 1e-5$ ).

Figure S6 provides the comparison of test MSE distributions and visualization of selected (worst to best along the MSE spectrum) Baldwinian-PINN solutions and errors for the low- and high-frequency-learned Baldwinian-PINN models. From the results, we can observe that the low-frequency-learned model tends to give inaccurate physics-informed predictions for the test tasks from the other side of the frequency spectrum (i.e., high and mixed frequencies), and vice versa.

These results highlight the effectiveness of the Baldwinian neuroevolution for generating Baldwinian-PINN models that are "genetically equipped" to perform well over diverse task distribution pertaining to the training environment, e.g., low frequency, high frequency, or full frequency spectrum, as represented in the population.

#### B. Baldwinian-PINN with different evolutionary algorithms and neural architectures

The CMA-ES algorithm used in our experimental study is merely an instantiation of the proposed Baldwinian neuroevolution framework for meta-learning PINN. The outer-loop evolution procedure is agnostic to, and can be seamlessly switched to other state-of-the-art evolutionary search algorithms.

In this section, we present the results when Baldwinian neuroevolution is carried out by 3 alternate algorithms: simple

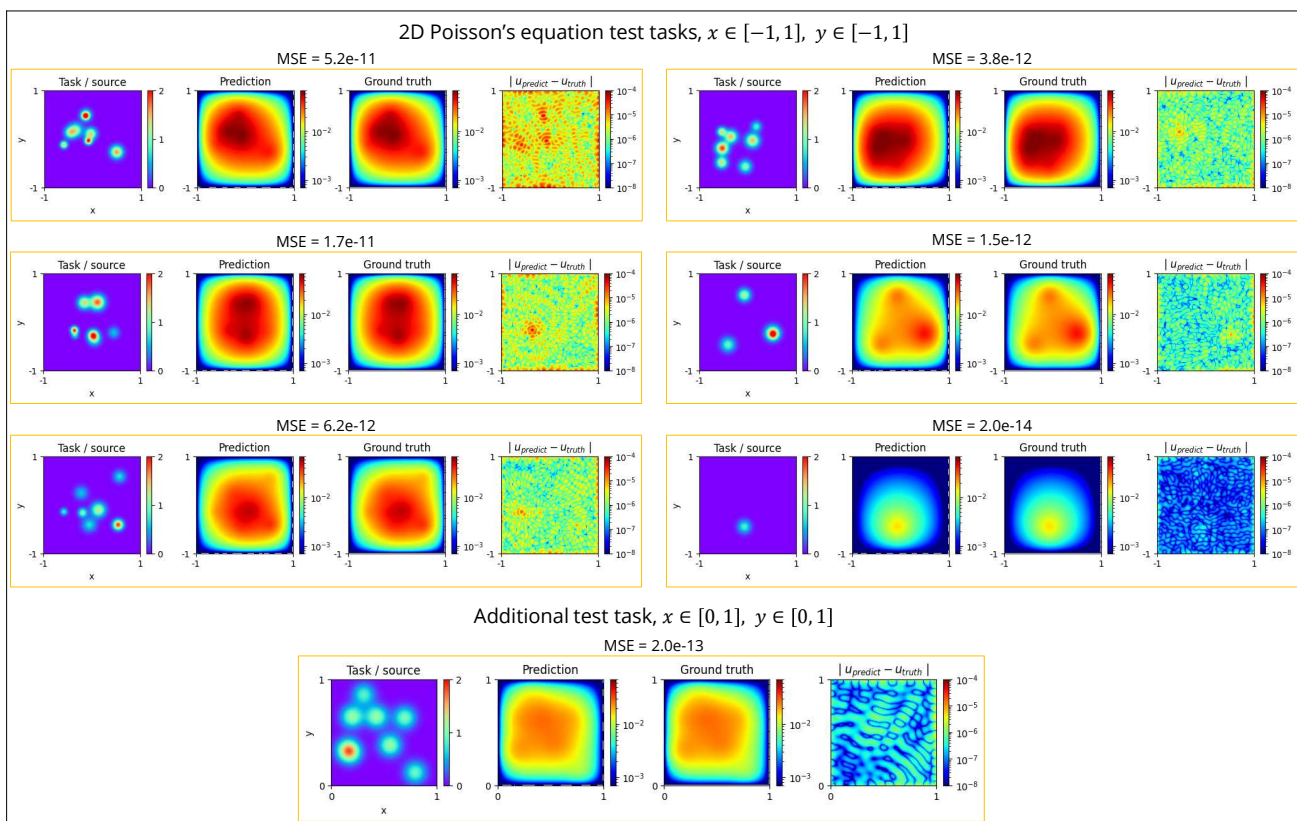


Figure S4. Baldwinian-PINN solutions selected (worst to best along the MSE spectrum) from the  $n = 60$  2D Poisson's test tasks, and an additional test task (subject to a different domain  $x \in [0, 1]$ ,  $y \in [0, 1]$ ) described in [S5], showing good performance on the diverse set of tasks exhibited by the source patterns.

genetic algorithm (GA) [S8], policy gradients with parameter-based exploration (PGPE) [S9], and natural evolution strategies (NES) variant [S10], using the *Helmholtz* example. Note that the same optimization settings (e.g., initial standard deviation, population size, neuroevolution iterations) as the original (CAM-ES-evolved) Baldwinian-PINN are used (i.e., we do not fine-tune the settings for each algorithm). The search dimension consists of 36 network hyperparameters controlling the distributional mean and spread of the weights and biases, the learning hyperparameter  $\lambda$ , and the loss importance hyperparameter  $\lambda_{pde}$  for re-balancing the relative importance between PDE and BC/IC errors in the least-squares solution.

Their test MSE values are  $1.1 \pm 8.1e-1$ ,  $3.0e-5 \pm 1.5e-4$ , and  $2.9e-8 \pm 9.8e-8$ , respectively. The simple GA performs much worse than the original Baldwinian-PINN. On the other hand, the NES-evolved Baldwinian-PINN significantly outperforms the CMA-ES-evolved Baldwinian-PINN, with more than 2 orders of magnitude MSE improvement on test tasks. Their MSE distributions on  $n = 60$  test tasks are compared in Figure S6, together with visualization of the selected (worst to best along the MSE spectrum) solutions and errors obtained by the NES-evolved Baldwinian-PINN.

Similarly, the Baldwinian-PINN neural architecture described in Section III-D1 is merely an instantiation of the proposed Baldwinian neuroevolution framework for meta-learning PINN. The extension from a simple, single hidden layer neural architecture to a deeper, more complex neural architecture is straightforward.

To demonstrate this, we construct a MLP with multiple hidden layers with a mix of *sin* and *softplus* activation functions, and additional skip connections from the early hidden layers to the final layer. The total number of weight parameters in this deeper Baldwinian-PINN is 68,480 (including 1280 weights in the final layer). Note that the learned final layer is the outcome of lifetime learning through the Tikhonov regularization computation, whereas these 67,200 weight parameters before the final layer are sampled from 32 normal distributions. There are 64 distributional hyperparameters, 1 learning hyperparameter  $\lambda$ , and 1 loss importance hyperparameter  $\lambda_{pde}$  evolved by the NES algorithm, using the same optimization settings (e.g., initial standard deviation, population size, neuroevolution iterations) as the original Baldwinian-PINN. For this study, the original Baldwinian-PINN architecture consists of 3600 weight parameters (including 900 weights in the final layer).

Despite having a much larger number of weight parameters in the deeper Baldwinian-PINN model, Baldwinian neuroevolution managed to arrive at a good set of distributional hyperparameters for the model to make a fast and accurate physics-based prediction on the test tasks (test MSE= $2.6e-7 \pm 6.7e-7$ ), in the *Helmholtz* problem.

The results described in this section underscore the promise of the Baldwinian neuroevolution framework for meta-learning PINN. Further improvements to the search algorithm and neural architecture design may boost the Baldwinian-PINN's learning speed and accuracy on other complex physics problems.

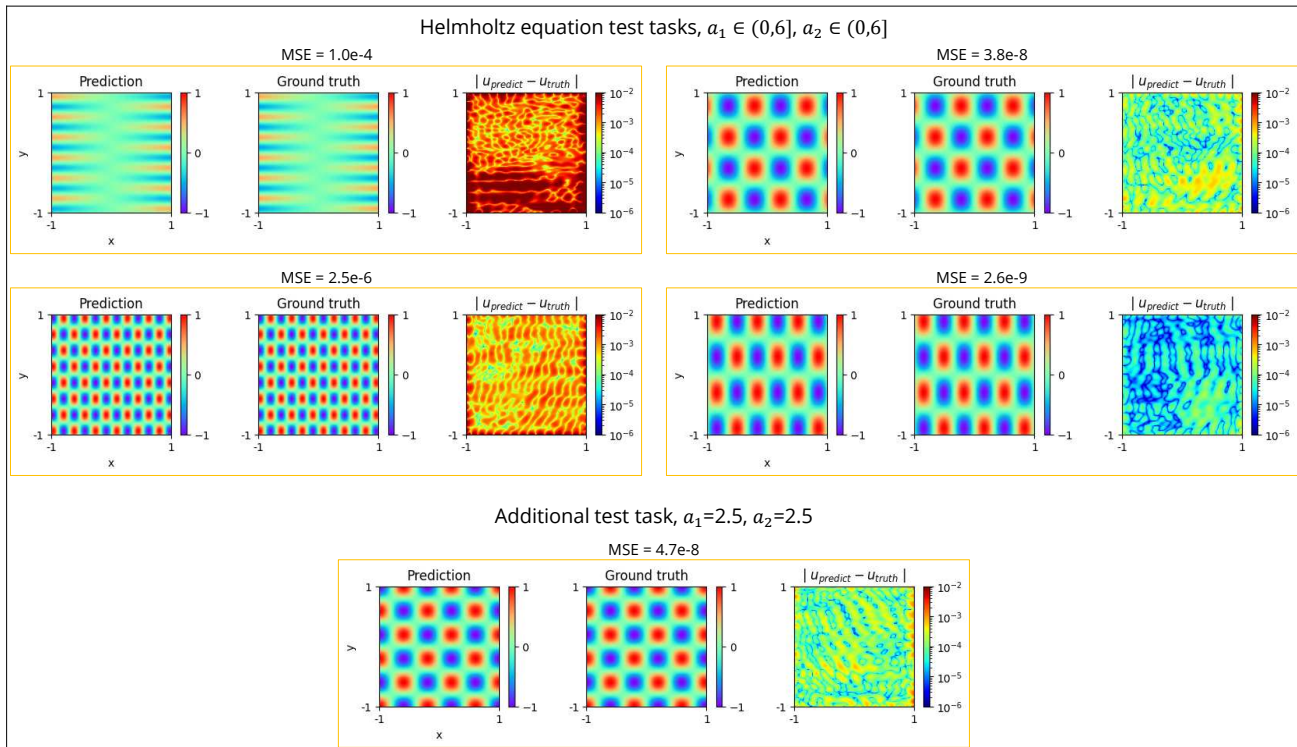


Figure S5. Baldwinian-PINN solutions selected (worst to best along the MSE spectrum) from the  $n = 60$  Helmholtz test tasks, and an additional test task ( $\alpha_1 = 2.5, \alpha_2 = 2.5$ ) described in [S6]. Baldwinian-PINN shows good performance on the diverse set of PDE tasks exhibited by the frequencies.

### C. Noisy or scarce training-task labels

Our lifetime-learning step is entirely physics-based, requires no labels at test time, and is solved in closed form with Tikhonov regularization; this inherently resists noisy or scarce data. When labels in the training tasks are noisy or sparse, we can rebalance the physics-data trade-off in the loss function by tuning the loss weights  $\kappa_{LSE}$  and  $\kappa_{MSE}$  to further strengthen the physics prior.

Empirically, Figure S7 evaluates Baldwinian-PINNs' lifetime-learning performance on a set of 1D Poisson's equation tasks ( $n = 60$ ) when the training-task labels are noisy or scarce. As observation noise increases from 0 to 100% under both zero-mean Gaussian noise and positively biased noise scenarios, the test error rises smoothly rather than catastrophically. Meanwhile, test accuracy remains low even when only 1–10% of points from the training tasks are observed, because the physics residual and BC/IC rows constrain the solution and meta-learning can be effectively guided by a few quality measurements, demonstrating that Baldwinian-PINNs tolerate non-ideal data and can operate under severe data scarcity.

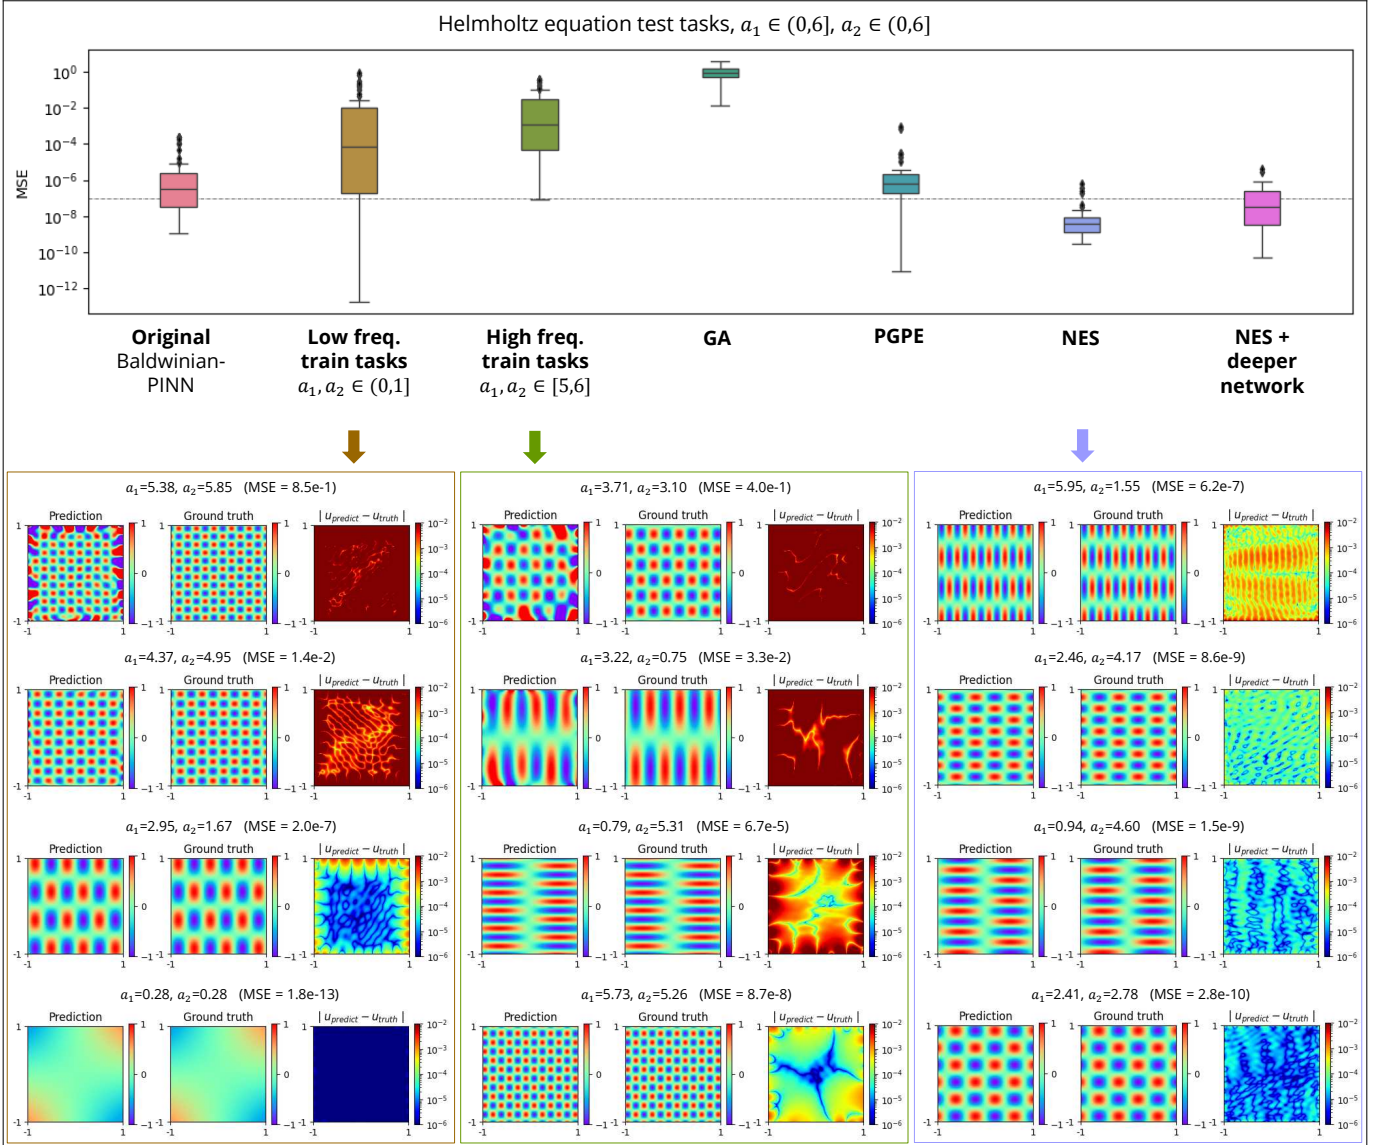


Figure S6. The MSE distributions on  $n = 60$  Helmholtz test tasks for the original Baldwinian-PINN (evolved by NES to meta-learn from  $n = 20$  train tasks  $\alpha_1 \in (0, 6]$ ,  $\alpha_2 \in (0, 6]$ ) and its variants, i.e., using a different set of train tasks, or a different evolutionary algorithm, or both different evolutionary algorithm and neural architecture. The solutions and errors for selected (worst to best along the MSE spectrum) Helmholtz test tasks are visualized for both low- and high-frequency-learned Baldwinian-PINN models, and also the NES-evolved Baldwinian-PINN.

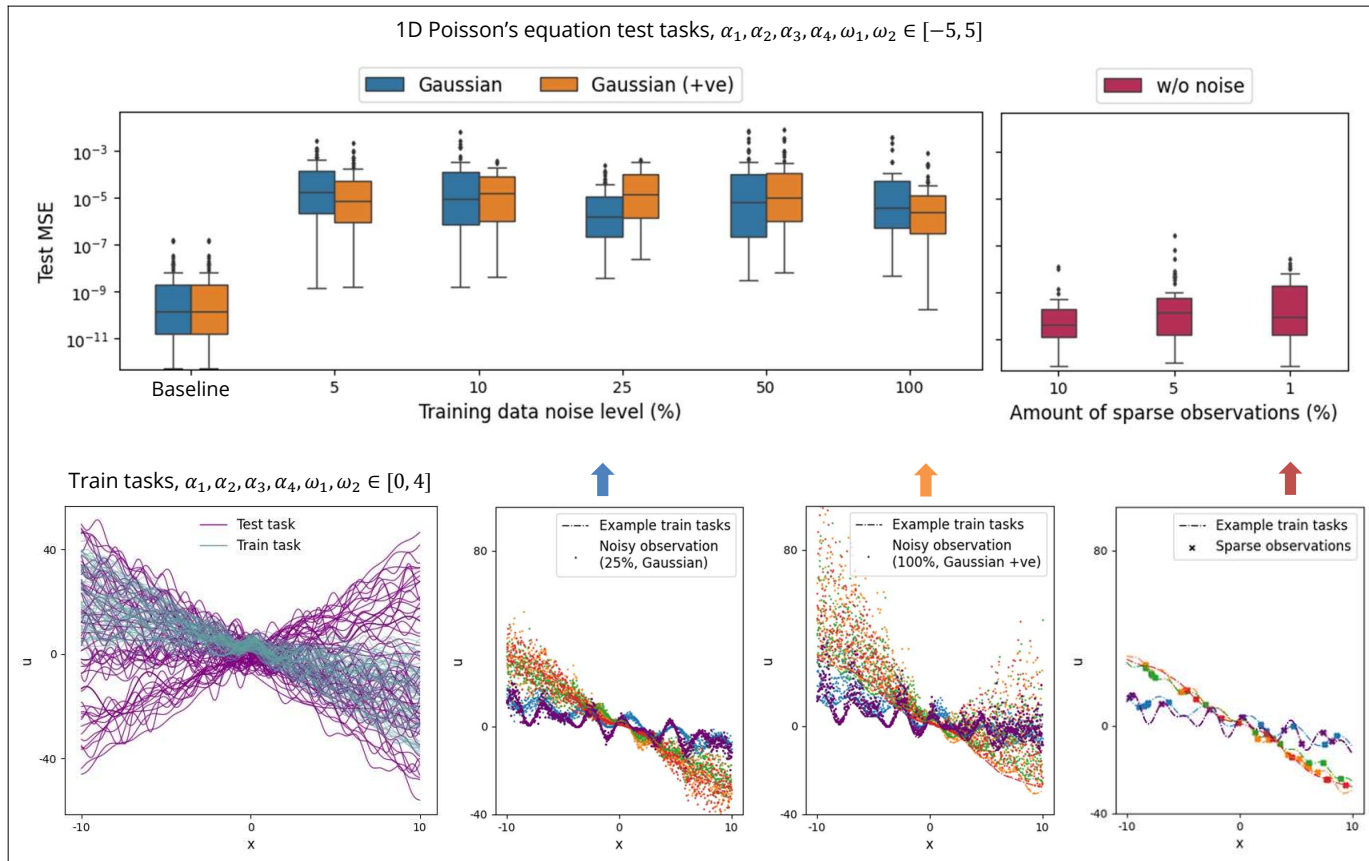


Figure S7. Robustness to noisy and sparse observations on 1D Poisson's problem. Top: distribution of test MSE obtained by Baldwinian-PINNs when noisy or sparse training-task labels are used. The training data noise level ranges from 0% (baseline) to 100% for both zero-mean “Gaussian” noise and positively biased “Gaussian (+ve)” noise scenarios, whereas the number of sparse observations ranges from 100% (baseline) to 1%. Bottom: examples of training tasks with corresponding noisy (25% Gaussian; 100% Gaussian +ve) and sparse (1%) observation scenarios. Baldwinian-PINNs degrade gracefully with noise and exhibit low error with very limited observations.

Table S2  
COMPARISON WITH BASELINE META-LEARNING PINN MODELS ON NONLINEAR BENCHMARK PROBLEMS.

Problem (no. test task)	Method / Model															
	Random		MAML		Center		Multitask		LMC		RBF (multiquadric)		Polynomial		Baldwinian-PINN	
	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>1</sup>	Time <sup>2</sup>	Error <sup>3</sup>	Time <sup>4</sup>
7 Burgers' equation ( $n = 32$ )	1.2e-3 $\pm 1.8e-3$	105 $\pm 29$	1.7e-3 $\pm 3.2e-3$	125 $\pm 38$	9e-4 $\pm 5e-4$	56 $\pm 48$	7e-4 $\pm 2e-4$	24 $\pm 21$	8e-4 $\pm 4e-4$	40 $\pm 27$	8e-4 $\pm 4e-4$	7 $\pm 8$	7e-4 $\pm 3e-4$	7 $\pm 16$	<b>3.8e-4</b> $\pm 6.9e-4$	<b>1.96</b> $\pm 0.01$
8 nonlinear heat ( $n = 64$ )	5.2e-3 $\pm 3.9e-3$	156 $\pm 42$	4.5e-3 $\pm 3.1e-3$	188 $\pm 43$	4.5e-3 $\pm 2.4e-3$	96 $\pm 41$	4.8e-3 $\pm 3.6e-3$	49 $\pm 30$	4.9e-3 $\pm 3.3e-3$	59 $\pm 28$	4.4e-3 $\pm 2.2e-3$	35 $\pm 30$	4.6e-3 $\pm 3.2e-3$	38 $\pm 30$	<b>6.0e-4</b> $\pm 7.6e-4$	<b>0.87</b> $\pm 0.02$
9 Allen-Cahn equation ( $n = 32$ )	1.5e-2 $\pm 1.2e-2$	496 $\pm 161$	-	-	1.2e-2 $\pm 4.6e-3$	201 $\pm 83$	1.2e-2 $\pm 4.7e-3$	120 $\pm 44$	1.1e-2 $\pm 4.0e-3$	120 $\pm 21$	1.2e-2 $\pm 5.3e-3$	68 $\pm 46$	1.2e-2 $\pm 4.2e-3$	44 $\pm 19$	<b>9.5e-4</b> $\pm 1.5e-3$	<b>1.19</b> $\pm 0.01$
10 Diffusion-reaction ( $n = 64$ )	1.1e-2 $\pm 6.2e-3$	1073 $\pm 206$	-	-	9.5e-3 $\pm 5.9e-3$	426 $\pm 159$	9.0e-3 $\pm 5.4e-3$	243 $\pm 93$	9.1e-3 $\pm 5.8e-3$	302 $\pm 127$	8.1e-3 $\pm 4.6e-3$	280 $\pm 161$	8.5e-3 $\pm 5.1e-3$	249 $\pm 129$	<b>1.3e-4</b> $\pm 2.0e-4$	<b>0.58</b> $\pm 0.02$
11 6D diffusion-reaction ( $n = 100$ )	2.2e-3 $\pm 1.3e-3$	612 $\pm 255$	-	-	1.8e-3 $\pm 1.7e-3$	494 $\pm 222$	1.7e-3 $\pm 1.9e-3$	431 $\pm 200$	1.5e-3 $\pm 1.8e-3$	428 $\pm 201$	1.7e-3 $\pm 2.0e-3$	375 $\pm 168$	1.8e-3 $\pm 2.0e-3$	496 $\pm 213$	<b>1.9e-4</b> $\pm 1.5e-4$	<b>0.60</b> $\pm 0.03$

<sup>0</sup>Benchmark results from meta-learning PINN method / model (i.e., Random, MAML, Center, Multitask, LMC, RBF (multiquadric), Polynomial) are extracted from [S11].

<sup>1</sup>Relative L2 errors are obtained from a meta-learned PINN after fine-tuning with ADAM (500 iterations) + L-BFGS optimizations.

<sup>2</sup>Only for the L-BFGS optimization time (s) of the fine-tuning phase (computation time of meta-learning phase, and ADAM optimization prior to the L-BFGS during fine-tuning phase not reported in [S11])

<sup>3</sup>Relative L2 errors are aggregated from 5 individual runs.

<sup>4</sup>Computation time (s) per task during lifetime-learning (fine-tuning) phase, on single GPU (NVIDIA GeForce RTX 3090). For neuroevolution time please refer to Table I

## S.V. STUDIES ON SET OF NONLINEAR PDE PROBLEMS

The nonlinear PDE problems (Problems 7-11) used for demonstrating the Baldwinian neuroevolution of physics are described below.

1) *Burgers' equation*: The Burgers' equation [S12] consists of 16 randomly sampled train tasks and 32 uniformly sampled test tasks from the PDE parameter  $\gamma \in [5e-3, 5e-2]$ :

$$\text{(Problem 7)} \quad \frac{du}{dt} + u \frac{du}{dx} - \gamma \frac{d^2u}{dx^2} = 0, \\ x \in [-1, 1], \quad t \in (0, 1] \quad (\text{S5})$$

subject to IC  $u(x, t = 0) = -\sin(\pi x)$ .

2) *Nonlinear heat equation*: The nonlinear heat equation consists of 13 randomly sampled train tasks and 64 uniformly sampled test tasks from the PDE parameters  $\gamma \in [1, \pi], k \in [1, \pi]$ :

$$\text{(Problem 8)} \quad \frac{du}{dt} - \gamma \frac{d^2u}{dx^2} + k \tanh(u) = q, \\ x \in [-1, 1], \quad t \in (0, 1] \quad (\text{S6})$$

where the exact solution  $u(x, t; \gamma, k) = k \sin(\pi x) \exp(-\pi k x^2) \exp(-\pi t^2)$  is used to derive the corresponding IC, BCs and source term  $q$ .

3) *Nonlinear Allen-Cahn equation*: The nonlinear Allen-Cahn equation consists of 16 randomly sampled train tasks and 32 uniformly sampled test tasks from the PDE parameter  $\gamma \in (0, \pi]$ :

$$\text{(Problem 9)} \quad \gamma \left( \frac{d^2u}{dx^2} + \frac{d^2u}{dy^2} \right) + u(u^2 - 1) = q, \\ x \in [-1, 1], \quad y \in [-1, 1] \quad (\text{S7})$$

where the exact solution  $u(x, y; \gamma) = \exp(-\gamma(x + 0.7)) \sin(\pi x) \sin(\pi y)$  is used to derive the corresponding BCs and source term  $q$ .

4) *Nonlinear diffusion-reaction equation*: The nonlinear diffusion-reaction equation [S13] consists of 22 randomly sampled train tasks and 64 uniformly sampled test tasks from the PDE parameter  $\gamma \in [1, \pi], k \in [1, \pi]$ :

$$\text{(Problem 10)} \quad \gamma \left( \frac{d^2u}{dx^2} + \frac{d^2u}{dy^2} \right) + k u^2 = q, \\ x \in [-1, 1], \quad y \in [-1, 1] \quad (\text{S8})$$

where the exact solution  $u(x, y; \gamma, k) = k \sin(\pi x) \sin(\pi y) \exp(-\gamma \sqrt{k x^2 + y^2})$  is used to derive the corresponding BCs and source term  $q$ .

5) *6D parametric diffusion-reaction*: The 6D parametric diffusion-reaction problem consists of 17 randomly sampled train tasks and 100 randomly sampled test tasks from the

PDE/BC parameters  $(\alpha_1, \alpha_2, \omega_1, \omega_2, \omega_3, \omega_4)$ ,  $\alpha$ 's  $\in [0.1, 1]$ ,  $\omega$ 's  $\in [1, 5]$ :

$$\text{(Problem 11)} \quad \left( \frac{d^2u}{dx^2} + \frac{d^2u}{dy^2} \right) + u(1 - u^2) = q, \\ x \in [-1, 1], \quad y \in [-1, 1] \quad (\text{S9})$$

where the exact solution  $u(x, y; \alpha_1, \alpha_2, \omega_1, \omega_2, \omega_3, \omega_4) = \alpha_1 \tanh(\omega_1 x) \tanh(\omega_2 y) + \alpha_2 \sin(\omega_3 x) \sin(\omega_4 y)$  is used to derive the corresponding BCs and source term  $q$ .

### A. Results for set of nonlinear PDE problems

As per [S11], the spatio-temporal domain in Problems 7-8 is uniformly discretized into  $256 \times 100$ , and the 2D spatial domain in Problems 9-11 is uniformly discretized into  $128 \times 128$ , for the test tasks. The comparison of generalization performance and compute cost (i.e., lifetime learning for Baldwinian-PINNs or gradient fine-tuning for baselines) between Baldwinian-PINN and several meta-learned PINN models (based on results reported in [S11]) are summarized in Table S2.

Figure S8 provides additional visualization results for the 5 nonlinear PDE problems described above, showing the corresponding Baldwinian-PINN solutions with the worst and median accuracy along the MSE spectrum (pooled from all test tasks  $\times$  5 individual runs) for each problem.

## S.VI. ADDITIONAL RESULTS FOR NAVIER-STOKES (N-S) EQUATIONS FOR FLUID FLOW

Our Baldwinian-PINN model for the fluid flow problem consists of multiple hidden layers with a mix of *sin*, *tanh*, and *softplus* activation functions, with skip connections from the early hidden layers to the final layer. The total number of weight parameters is 133,504 (including 1152 weights in the final layer). Figure S9 provides additional visualization comparing Baldwinian-PINN predictions and ground truth for the lid-driven cavity flow across unseen Reynolds numbers  $Re = \{5, 25, 100, 400, 1000\}$ , showing  $u$ - and  $v$ -velocity fields alongside absolute error maps. Errors remain uniformly small across the domain and over two orders of magnitude in  $Re$ , indicating strong generalization of the Baldwinian-PINN model to markedly different flow regimes.

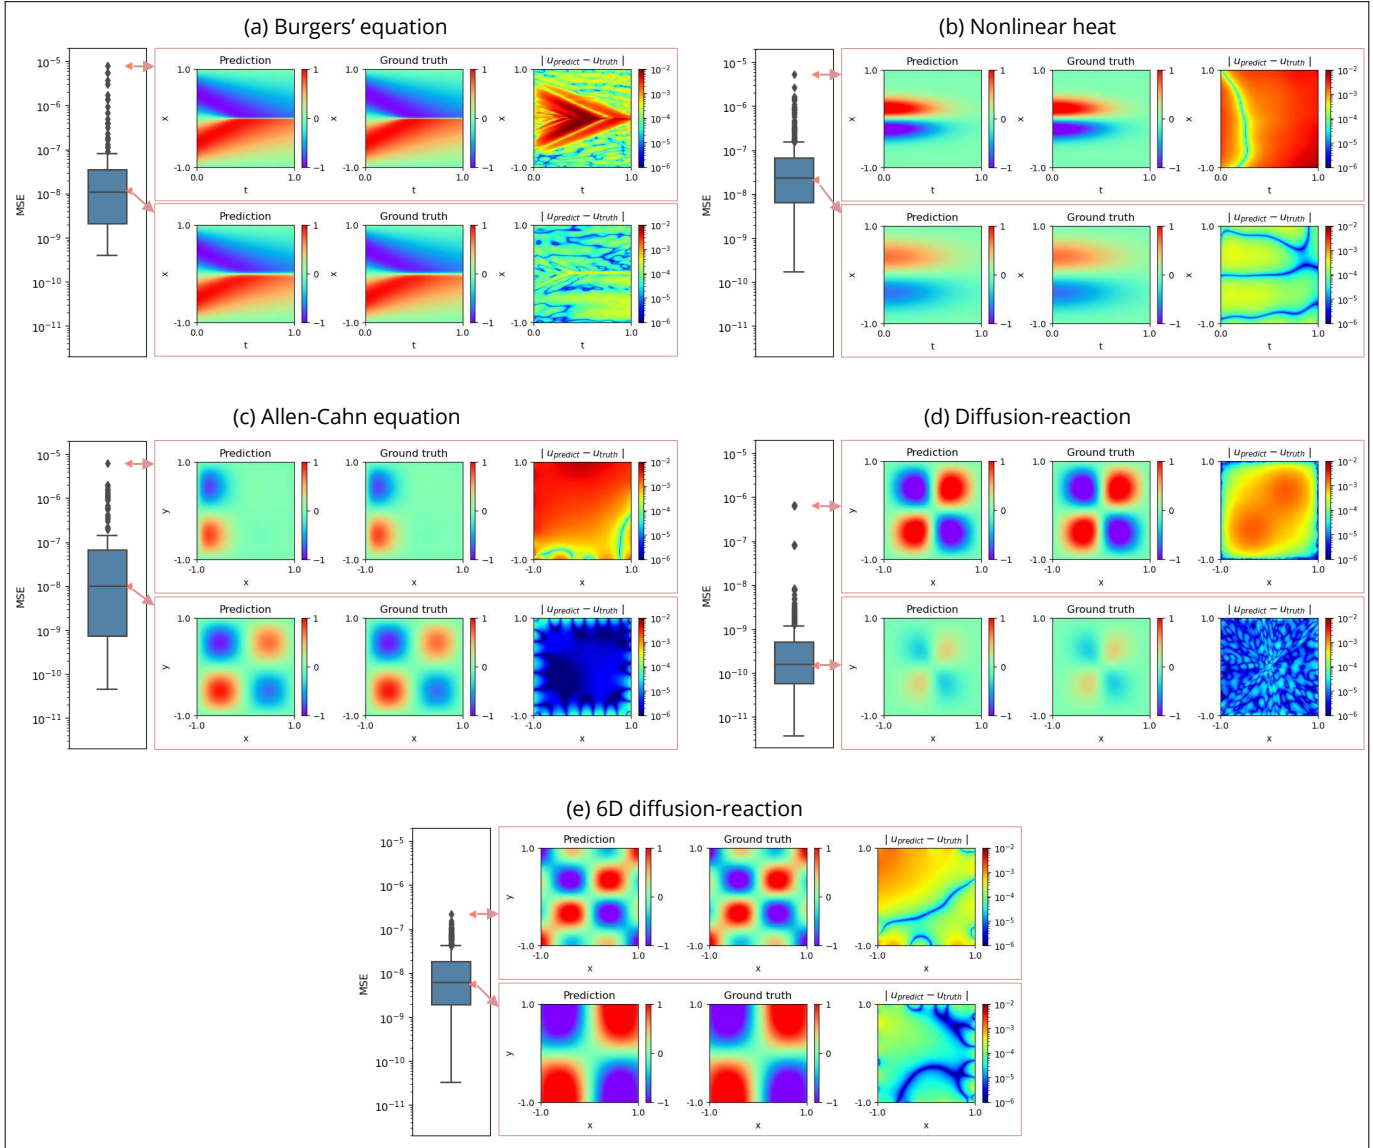


Figure S8. Baldwinian-PINN solutions with the worst and (approximately) average accuracy along the MSE spectrum (pooled from all test tasks  $\times$  5 individual runs), for 5 nonlinear PDE problems.

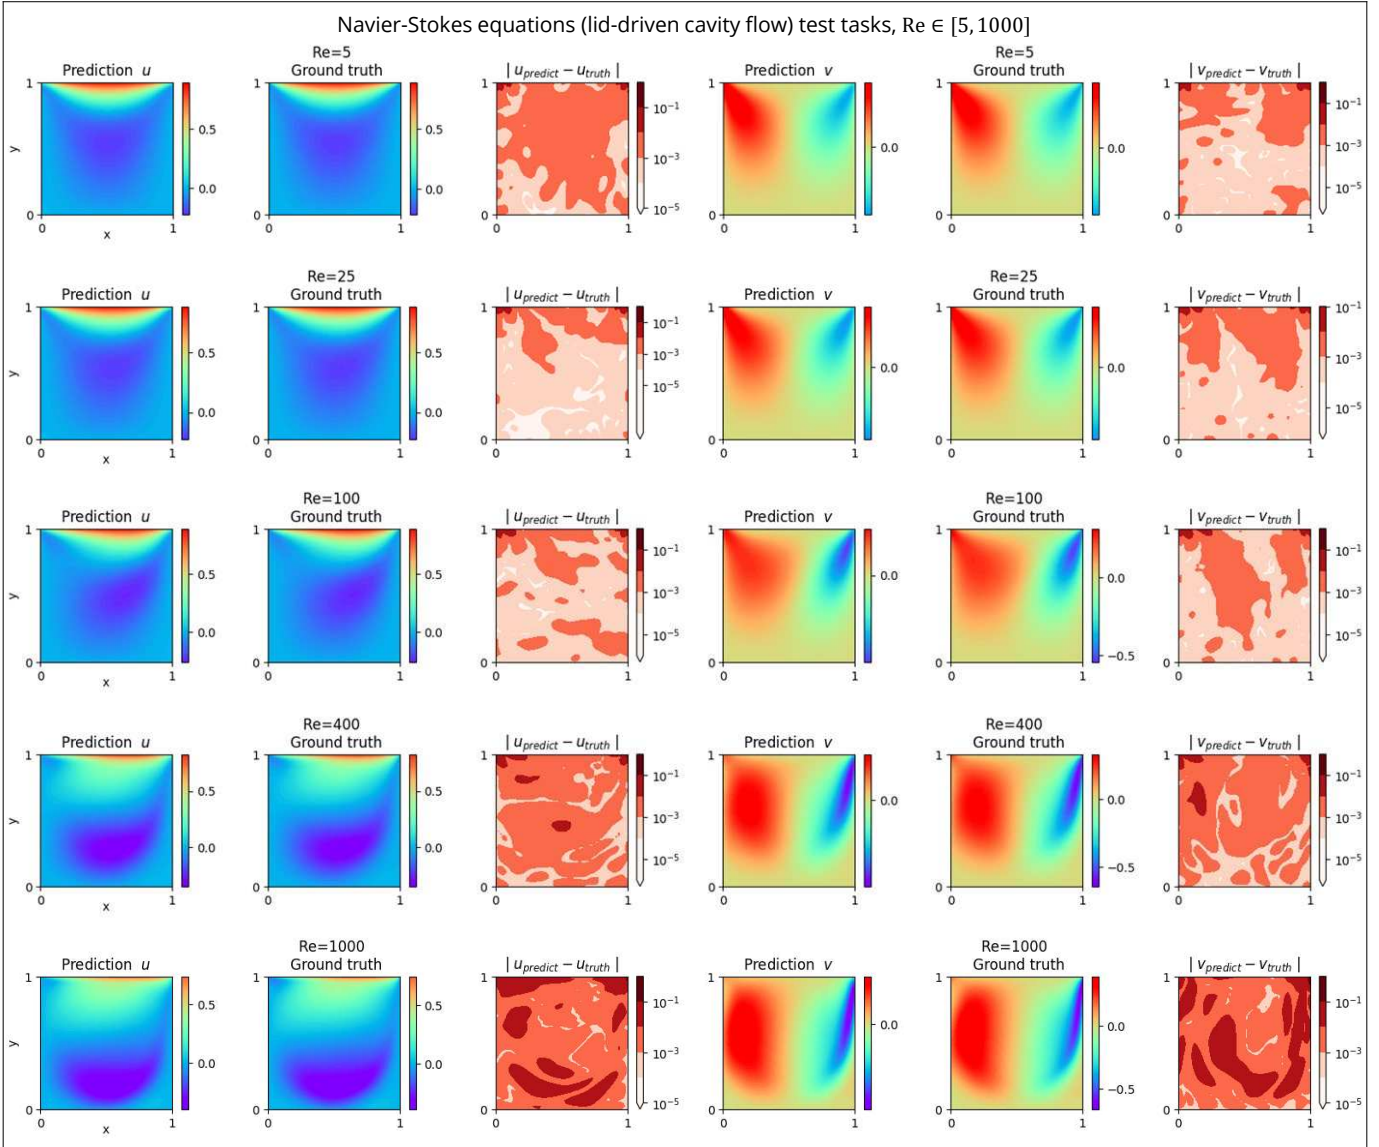


Figure S9. Baldwinian-PINN prediction vs. ground truth and their absolute error maps for unseen fluid flow (N-S equations) tasks on  $Re \in [5, 1000]$ . The MSE of the solutions obtained by the best evolved Baldwinian-PINN (sampled from the center of NES search distribution after 400 iterations with initial std. = 0.5) compared to numerical simulation ground truths are  $4.4e-5$  ( $Re=5$ ),  $3.4e-5$  ( $Re=25$ ),  $3.5e-5$  ( $Re=100$ ),  $5.1e-5$  ( $Re=400$ ), and  $1.8e-4$  ( $Re=1000$ ), respectively.

Table S3  
DNN / PINN MODEL AND TRAINING CONFIGURATIONS USED IN ABLATION STUDY.

	Problem	Model	Architecture	Activation	Initialization	No. weights (final layer)	Training / ADAM optimizer		
							Batch size for task	Max. iteration	Learning rate <sup>1</sup>
1	Convection-diffusion $(x, \alpha) \rightarrow u$	1a	Deep	tanh	Xavier	2750 (50)	10	50,000	5e-4 / 5e-3 / 5e-2
		1b	Deep	sin	He	2750 (50)	10	50,000	5e-4 / 5e-3 / 5e-2
		2a	Shallow	tanh	Xavier	3600 (900)	10	50,000	5e-4 / 5e-3 / 5e-2
		2b	Shallow	sin	He	3600 (900)	10	50,000	5e-4 / 5e-3 / 5e-2
6	Kinematics $(t, vel_0, a_0, C_d, A, m) \rightarrow (x, y)$	1a	Deep	tanh	Xavier	3930 (60)	30	150,000	5e-4 / 5e-3 / 5e-2
		1b	Deep	sin	He	3930 (60)	30	150,000	5e-4 / 5e-3 / 5e-2
		2a	Shallow	tanh	Xavier	6300 (1800)	30	150,000	5e-4 / 5e-3 / 5e-2
		2b	Shallow	sin	He	6300 (1800)	30	150,000	5e-4 / 5e-3 / 5e-2

<sup>1</sup>Flat learning rate for first 40% of iterations, then cosine decay towards 1e-6.

Table S4  
GENERALIZATION PERFORMANCE OF DNN / PINN MODELS IN ABLATION STUDY.

Problem (no. test task)	Model	Lr <sup>1</sup> Mode <sup>2</sup>	DNN						PINN					
			5e-4		5e-3		5e-2		5e-4		5e-3		5e-2	
			DNN	Tik-Reg	DNN	Tik-Reg	DNN	Tik-Reg	PINN	Tik-Reg	PINN	Tik-Reg	PINN	Tik-Reg
1 Convection-diffusion ( $x, \alpha$ ) $\rightarrow$ $u$ ( $n = 110$ )	1a		3.5e-4 $\pm 2.2e-3$	1.5e-1 $\pm 1.1e-1$	2.3e-4 $\pm 1.8e-3$	3.8e-2 $\pm 7.6e-2$	1.2e-2 $\pm 2.6e-2$	2.3e-1 $\pm 2.8e-2$	4.6e-3 $\pm 2.0e-2$	1.4e-2 $\pm 4.3e-2$	1.9e-3 $\pm 1.1e-2$	7.2e-3 $\pm 2.5e-2$	2.2e-1 $\pm 5.6e-2$	2.3e-1 $\pm 2.8e-2$
	1b		6.9e-4 $\pm 2.2e-3$	1.8e-1 $\pm 1.0e-1$	1.7e-4 $\pm 1.0e-3$	1.3e-1 $\pm 1.2e-1$	1.9e-3 $\pm 7.4e-3$	2.2e-2 $\pm 3.6e-2$	1.3e-1 $\pm 9.3e-2$	1.6e-1 $\pm 1.1e-1$	2.1e-2 $\pm 3.5e-2$	1.1e-1 $\pm 1.1e-1$	3.4e-1 $\pm 3.7e-1$	1.9e-1 $\pm 9.6e-2$
	2a		7.1e-3 $\pm 6.8e-3$	2.3e-1 $\pm 4.7e-2$	4.5e-3 $\pm 1.3e-2$	2.2e-1 $\pm 5.3e-2$	2.3e-3 $\pm 4.0e-3$	2.1e-1 $\pm 7.2e-2$	2.2e-1 $\pm 4.5e-2$	2.3e-1 $\pm 4.9e-2$	1.9e-1 $\pm 7.0e-2$	2.2e-1 $\pm 5.6e-2$	2.0e-1 $\pm 6.5e-2$	2.1e-1 $\pm 7.8e-2$
	2b		6.0e-3 $\pm 7.1e-3$	2.2e-1 $\pm 5.7e-2$	3.8e-3 $\pm 5.0e-3$	2.2e-1 $\pm 6.1e-2$	2.5e-3 $\pm 5.2e-3$	2.1e-1 $\pm 7.7e-2$	2.2e-1 $\pm 5.1e-2$	2.2e-1 $\pm 5.8e-2$	2.2e-1 $\pm 5.2e-2$	2.2e-1 $\pm 5.9e-2$	2.1e-1 $\pm 5.8e-2$	2.2e-1 $\pm 6.1e-2$
6 Kinematics ( $t, vel_0, a_0, C_d, A, m$ ) $\rightarrow$ ( $x, y$ ) ( $n = 100$ )	1a		8.5e+0 $\pm 2.8e+1$	2.9e-2 $\pm 4.2e-1$	3.7e+1 $\pm 8.8e+1$	9.4e+2 $\pm 3.5e+3$	2.0e+2 $\pm 6.2e+2$	2.6e+3 $\pm 4.6e+3$	1.1e+1 $\pm 3.7e+1$	2.5e+2 $\pm 3.3e+3$	1.9e+1 $\pm 6.0e+1$	4.8e+1 $\pm 6.3e+2$	2.5e+3 $\pm 4.5e+3$	2.6e+3 $\pm 4.6e+3$
	1b		3.5e+0 $\pm 1.4e+1$	3.0e-3 $\pm 4.5e-2$	1.9e+1 $\pm 6.7e+1$	3.2e-4 $\pm 2.9e-3$	5.2e+0 $\pm 9.1e+0$	2.6e+3 $\pm 4.6e+3$	1.3e+0 $\pm 5.7e+0$	7.9e-6 $\pm 1.1e-4$	3.5e+0 $\pm 1.4e+1$	1.6e-5 $\pm 1.1e-4$	3.0e+3 $\pm 3.5e+3$	8.1e-2 $\pm 1.6e+0$
	2a		3.0e+0 $\pm 4.3e+1$	4.4e-5 $\pm 2.5e-4$	7.9e+0 $\pm 1.6e+1$	2.4e+1 $\pm 3.8e+2$	1.8e+1 $\pm 2.7e+1$	2.3e+2 $\pm 1.6e+3$	2.8e+1 $\pm 5.5e+1$	9.4e-3 $\pm 1.7e-1$	3.3e+0 $\pm 6.5e+0$	3.1e-4 $\pm 2.8e-3$	1.5e+1 $\pm 2.9e+1$	1.2e-2 $\pm 2.7e-1$
	2b		1.5e+1 $\pm 2.7e+1$	5.2e-2 $\pm 5.4e-1$	2.1e+0 $\pm 3.5e+0$	2.0e-1 $\pm 2.1e+0$	5.0e+2 $\pm 7.7e+2$	1.4e-3 $\pm 7.9e-3$	2.2e+1 $\pm 4.4e+1$	4.7e-2 $\pm 3.1e-1$	4.0e+0 $\pm 8.7e+0$	5.2e-3 $\pm 3.5e-2$	7.8e+1 $\pm 1.1e+2$	5.9e-3 $\pm 3.9e-2$

<sup>0</sup>MSE results are aggregated from 5 individual runs, across 110 convection-diffusion and 100 kinematics test tasks. The DNN and PINN model configurations associated with a lower test MSE are highlighted.

<sup>1</sup>Initial learning rate.

<sup>2</sup>Prediction mode, i.e., direct DNN/PINN prediction vs. additional Tikhonov regularization (Tik-Reg) step.

### S.VII. BASELINE SGD-TRAINED DNN AND PINN MODELS USED IN ABLATION STUDY

Consider the general inputs  $(x, t, \vartheta)$  for the spatial-temporal domain and task parameter  $\vartheta$ . The data-driven loss function of a SGD-trained DNN model computes the MSE between the DNN output  $u_{\text{DNN}}(x_i, t_i, \vartheta_i)$  against the target  $u_i^{\text{label}}$  over  $i = 1, \dots, n$  labelled data pooled from a batch of training tasks:

$$l_{\text{DNN}} = l_{\text{data}} = \frac{1}{n} \sum_{i=1} (u_i^{\text{label}} - u_{\text{DNN}}(x_i, t_i, \vartheta_i))^2 \quad (\text{S10})$$

The loss function of a (baseline) SGD-trained PINN model is defined as:

$$l_{\text{PINN}} = \lambda_{\text{data}} l_{\text{data}} + \lambda_{\text{pde}} l_{\text{pde}} + \lambda_{\text{ic}} l_{\text{ic}} + \lambda_{\text{bc}} l_{\text{bc}} \quad (\text{S11a})$$

$$l_{\text{pde}} = \frac{1}{n_{\text{pde}}} \sum_{i=1} (\mathcal{N}_{\vartheta}[u_{\text{PINN}}(x_i^{\text{pde}}, t_i^{\text{pde}}, \vartheta_i)] - h(x_i^{\text{pde}}, t_i^{\text{pde}}))^2 \quad (\text{S11b})$$

$$l_{\text{ic}} = \frac{1}{n_{\text{ic}}} \sum_{i=1} (u_{\text{PINN}}(x_i^{\text{ic}}, 0, \theta_i) - u_0(x_i^{\text{ic}}))^2 \quad (\text{S11c})$$

$$l_{\text{bc}} = \frac{1}{n_{\text{bc}}} \sum_{i=1} (\mathcal{B}[u_{\text{PINN}}(x_i^{\text{bc}}, t_i^{\text{bc}}, \theta_i)] - g(x_i^{\text{bc}}, t_i^{\text{bc}}))^2 \quad (\text{S11d})$$

such that the PINN output  $u_{\text{PINN}}(x_i, t_i, \vartheta_i)$  satisfies PDE, IC, and BC for a set of training samples  $(x_i^{\text{pde}}, t_i^{\text{pde}}, \vartheta_i), i = 1, \dots, n_{\text{pde}}, (x_i^{\text{ic}}, 0, \vartheta_i), i = 1, \dots, n_{\text{ic}}, (x_i^{\text{bc}}, t_i^{\text{bc}}, \vartheta_i), i = 1, \dots, n_{\text{bc}}$  from the respective domain and task, in addition to minimizing the MSE from the labelled data. SGD-trained PINNs typically converge much slower than DNN because of the additional loss terms. We perform a coarse search for the loss balancing parameters  $(\lambda_{\text{data}}, \lambda_{\text{pde}}, \lambda_{\text{ic}}, \lambda_{\text{bc}})$  to improve the convergence of the PINN loss.

We tested a range of learning rate schedules, e.g., an initial learning rate =  $\{5e-4, 1e-3, 1e-2\}$  for the first 40% of training iterations followed by cosine decay towards  $1e-6$ . Table S3 and Table S4 give the DNN / PINN model and training configurations and their subsequent performance on test tasks for the ablation study.

### S.VIII. LINEAR AND NONLINEAR ODE: TASK DIVERSITY

Figure S10 provides additional visualization results for the convection-diffusion (Problem 1), 1D Poisson's (Problem 3), and nonlinear kinematics (Problem 6) tasks, showing the diversity of task distributions on which we have demonstrated the utility of the current study's Baldwinian-PINN.

#### REFERENCES

- [S1] P.-H. Chiu, "An improved divergence-free-condition compensated method for solving incompressible flows on collocated grids," *Computers & Fluids*, vol. 162, pp. 39–54, 2018.
- [S2] B. P. Leonard, "The ultimate conservative difference scheme applied to unsteady one-dimensional advection," *Computer Methods in Applied Mechanics and Engineering*, vol. 88, pp. 17–74, 1991.
- [S3] S. Gottlieb and C.-W. Shu, "Total variation diminishing runge-kutta schemes," *Mathematics of Computation*, vol. 67, pp. 73–85, 1998.
- [S4] D. Anderson, J. C. Tannehill, R. H. Pletcher, R. Munipalli, and V. Shankar, *Computational fluid mechanics and heat transfer*. Boca Raton, Florida, USA: CRC press, 2020.
- [S5] X. Liu, X. Zhang, W. Peng, W. Zhou, and W. Yao, "A novel meta-learning initialization method for physics-informed neural networks," *Neural Computing and Applications*, vol. 34, no. 17, pp. 14 511–14 534, 2022.
- [S6] W. Cho, K. Lee, D. Rim, and N. Park, "Hypernetwork-based meta-learning for low-rank physics-informed neural networks," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [S7] M. Toloubidokhti, Y. Ye, R. Missel, X. Jiang, N. Kumar, R. Shrestha, and L. Wang, "Dats: Difficulty-aware task sampler for meta-learning physics-informed neural networks," in *The Twelfth International Conference on Learning Representations*, 2023.
- [S8] Y. Tang, Y. Tian, and D. Ha, "Evojax: Hardware-accelerated neuroevolution," *arXiv preprint arXiv:2202.05008*, 2022.
- [S9] F. Sehnke, C. Osendorfer, T. Rückstieß, A. Graves, J. Peters, and J. Schmidhuber, "Parameter-exploring policy gradients," *Neural Networks*, vol. 23, no. 4, pp. 551–559, 2010.
- [S10] M. Nomura and I. Ono, "Fast moving natural evolution strategy for high-dimensional problems," in *2022 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2022, pp. 1–8.
- [S11] M. Penwarden, S. Zhe, A. Narayan, and R. M. Kirby, "A metalearning approach for physics-informed neural networks (pinns): Application to parameterized pdes," *Journal of Computational Physics*, vol. 477, p. 111912, 2023.
- [S12] J. Bec and K. Khanin, "Burgers turbulence," *Physics reports*, vol. 447, no. 1-2, pp. 1–66, 2007.
- [S13] C. Rao, P. Ren, Q. Wang, O. Buyukozturk, H. Sun, and Y. Liu, "Encoding physics to learn reaction-diffusion processes," *Nature Machine Intelligence*, vol. 5, no. 7, pp. 765–779, 2023.

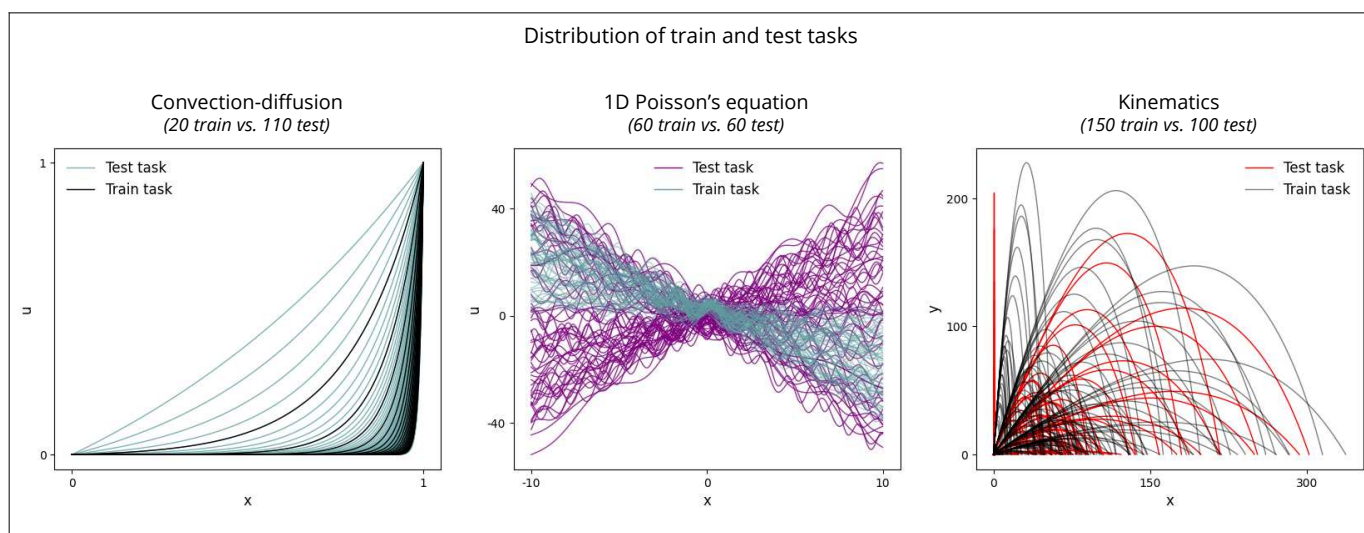


Figure S10. The distribution of convection-diffusion, 1D Poisson's, and nonlinear kinematics tasks for training and test.