


SURVEY PAPER

# Scalable algorithms for physics-informed neural and graph networks

Khemraj Shukla<sup>1</sup>, Mengjia Xu<sup>1,2</sup>, Nathaniel Trask<sup>3</sup> and George E. Karniadakis<sup>1,\*</sup> 

<sup>1</sup>Division of Applied Mathematics, Brown University, 182 George St, Providence, Rhode Island 02912, USA

<sup>2</sup>McGovern Institute for Brain Research, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, Massachusetts 02139, USA

<sup>3</sup>Center for Computing Research, Sandia National Laboratories, 1451 Innovation Pkwy SE #600, Albuquerque, New Mexico 87123, USA

\*Corresponding author. E-mail: [george\\_karniadakis@brown.edu](mailto:george_karniadakis@brown.edu)

**Received:** 28 March 2022; **Accepted:** 19 May 2022

**Keywords:** Graph calculus; graph neural networks; PINNs; scalability; scientific machine learning

## Abstract

Physics-informed machine learning (PIML) has emerged as a promising new approach for simulating complex physical and biological systems that are governed by complex multiscale processes for which some data are also available. In some instances, the objective is to discover part of the hidden physics from the available data, and PIML has been shown to be particularly effective for such problems for which conventional methods may fail. Unlike commercial machine learning where training of deep neural networks requires big data, in PIML big data are not available. Instead, we can train such networks from additional information obtained by employing the physical laws and evaluating them at random points in the space–time domain. Such PIML integrates multimodality and multi-fidelity data with mathematical models, and implements them using neural networks or graph networks. Here, we review some of the prevailing trends in embedding physics into machine learning, using physics-informed neural networks (PINNs) based primarily on feed-forward neural networks and automatic differentiation. For more complex systems or systems of systems and unstructured data, graph neural networks (GNNs) present some distinct advantages, and here we review how physics-informed learning can be accomplished with GNNs based on graph exterior calculus to construct differential operators; we refer to these architectures as physics-informed graph networks (PIGNs). We present representative examples for both forward and inverse problems and discuss what advances are needed to scale up PINNs, PIGNs and more broadly GNNs for large-scale engineering problems.

## Impact Statement


Many complex problems in computational engineering can be described by parametrized differential equations while boundary conditions or material properties may not be fully known, for example, in thermal-fluid or solid mechanics systems. These ill-defined problems cannot be solved with standard numerical methods but if some sparse data are available, progress can be made by recasting them as large-scale minimization problems. Here, we review physics-informed neural networks (PINNs) and physics-informed graph networks (PIGNs) that integrate seamlessly data and mathematical physics models, even in partially understood or uncertain contexts. Using automatic differentiation in PINNs and external graph calculus in PIGNs, the physical laws are enforced by penalizing the residuals on random points in the space–time domain for PINNs and on the nodes of a graph in PIGNs.

## 1. Introduction


Physics-informed machine learning (PIML) involves the use of neural networks, graph networks or Gaussian process regression to simulate physical and biomedical systems, using a combination of mathematical models and multimodality data (Raissi et al., 2018, 2019; Karniadakis et al., 2021; Gao et al., 2022).

Physics-informed neural networks (PINNs) (Raissi et al., 2019) can solve a partial differential equation (PDE) by directly incorporating the PDE into the loss function of the neural network (NN) and employing automatic differentiation to represent all the differential operators. Hence, PINNs do not require any mesh generation, thus avoiding a tremendous cost that has hindered progress in computational engineering, especially for moving and deformable domains. They can easily be applied to any known PDE, but also to different types of differential equations, including fractional PDEs (Pang et al., 2019), integro-differential equations (Lu et al., 2021a), or stochastic PDEs (Zhang et al., 2019, 2020a). PINNs are very simple to implement, and even the most elaborate codes contain less than 1,000 lines. Unlike traditional numerical approaches, the same exact code is able to treat both forward and inverse problems. Hence, different simulation scenarios, as shown in Figure 1, can be executed without any extra effort, other than preparing the input data. These advantages have been demonstrated across different fields, for example, in fluid mechanics (Mao et al., 2020; Raissi et al., 2020), optics (Chen et al., 2020; Lu et al., 2021b), systems biology (Yazdani et al., 2020), geophysics (Waheed et al., 2020), nondestructive evaluation of materials (Shukla et al., 2020, 2021a), and biomedicine (Sahli Costabal et al., 2020; Goswami et al., 2021; Yin et al., 2021).

Graph neural networks (GNNs) and specifically physics-informed graph networks (PIGNs) have emerged as an alternative method, as they can be particularly effective for composite physical systems where a state evolves as a function of its neighboring states, forming dynamic relational graphs instead of grids. In contrast to PINNs, these models apply message-passing between a small number of moving and interacting objects, which deviate from PDEs that are strictly differential functions. Another application of GNNs and PIGNs is the discovery of PDEs from sparse data, see for example, Iakovlev et al. (2020). The authors combined message passing neural networks with the method of lines and neural Ordinary Differential Equation (ODE) (Chen et al., 2018) and obtained good results for advection–diffusion, heat equation and Burgers equation in one dimension. Similar work was done in Kumar and Chakraborty (2021), where the authors designed a PIGN called GrADE for learning the system dynamics from data, for example, the one-dimensional and two-dimensional Burgers equations. GrADE combines a GNN that can handle unstructured data with neural ODE to represent the temporal domain. In addition, the authors used the graph attention (GAT) for embedding of nodes during aggregation. A somewhat different approach was proposed in Gao et al. (2022), where the authors introduced the graph convolution operation into physics-informed learning and combined it with finite elements in order to handle unstructured meshes for fluid mechanics and solid mechanics applications. GNNs have also been proposed for solving ordinary differential equations (GNODE) (Poli et al., 2019) in modeling continuous-time signals on graphs. This



**Figure 1.** *Physics versus Data. Simulating and forecasting the response of real-world problems requires both data and physical models. In many applications throughout physics, engineering and biomedicine we have some data and we can describe some but not all physical process. Physics-informed machine learning enables seamless integration of data and models. On the left, we have the classical paradigm of well posed problems. On the right, we depict black-box type system identification. Most real-life applications fall in the middle of this diagram.*



**Figure 2.** Recurrent architecture to incorporate physics in graph networks. The blue blocks contain learnable parameters while the orange blocks are objective functions. The middle core block can be repeated as many times as the required time steps ( $T$ ). Schematic adapted from Seo and Liu (2019).

work was extended in Seo and Liu (2019) for PDEs and for applications to climate, where a recurrent architecture was proposed to incorporate physics on graph networks, see Figure 2. We note that in this architecture the physics is not directly applied to the input observations but rather to the latent representations.

One of the major limitations of the original PINN algorithm is the large computational cost associated with the training of the neural networks, especially for forward multiscale problems. To reduce the computational cost, Jagtap et al. (2020) introduced a domain decomposition-based PINN for conservation laws, namely conservative PINN (cPINN), where continuity of the states as well as their fluxes across subdomain interfaces is enforced to connect the subdomains together. In subsequent work, Jagtap and Karniadakis (2020) applied domain decomposition to general PDEs using the so-called extended PINN (XPINN). Unlike cPINN, which offers space decomposition, XPINN offers both space–time domain decomposition for any irregular, nonconvex geometry thereby reducing the computational cost effectively. By exploiting the decomposition process of the cPINN/XPINN methods and its implicit mapping onto modern heterogeneous computing platforms, the training time of the network can be reduced to a great extent. Another domain decomposition method applied to the variational formulation of PINNs was proposed in Kharazmi et al. (2021). Referred to as the *hp*-VPINN method, it mimics the well-known spectral element method (Karniadakis et al., 2005), which exhibits dual *h*-*p* convergence. An effective parallel implementation of PINNs was reported in Shukla et al. (2021b). Similarly, drawing on existing parallel frameworks for standard GNNs, one can scale up PIGNs for simulating complex systems or systems.

This chapter is organized as follows. In the next section, we review the PINN formulation and the scalable extensions, and provide prototypical examples. In the subsequent section, we review popular variants of GNNs (spectral, spatial, and graph-attention) before discussing their relationship to traditional PDE discretizations and introducing PIGNs, which employ a graph exterior calculus; we also discuss scalability of GNNs.

## 2. Physics-Informed Neural Networks

We consider the following PDE for the solution  $u(\mathbf{x}, t)$  parametrized by the parameters  $\lambda$  defined on a domain  $\Omega$ :

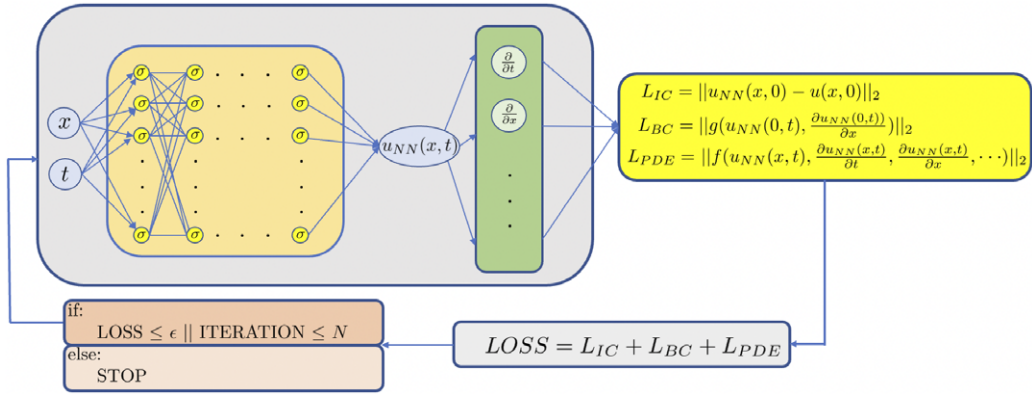
$$f\left(\mathbf{x}; \frac{\partial u}{\partial x_1}, \dots, \frac{\partial u}{\partial x_d}; \frac{\partial^2 u}{\partial x_1 \partial x_1}, \dots, \frac{\partial^2 u}{\partial x_1 \partial x_d}; \dots; \lambda\right) = 0, \quad \mathbf{x} = (x_1, \dots, x_d) \in \Omega, \quad (1)$$

with the boundary conditions

$$\mathcal{B}(u, \mathbf{x}) = 0 \quad \text{on} \quad \partial\Omega.$$

In PINNs, the initial conditions are treated similarly as the boundary conditions.

A schematic of a PINN for Equation (1) is shown in Figure 3. Specifically, in order to solve the parametrized PDE via PINNs, we construct a deep NN  $\hat{u}(\mathbf{x}; \theta)$  with the trainable parameters  $\theta$ . We incorporate the constraints implied by the PDE as well as the boundary and initial conditions in the loss



**Figure 3.** Schematic of PINN. The left part represents the data NN whereas the right part represents the physics-informed NN. All differential operators are obtained via automatic differentiation, hence no mesh generation is required to solve the PDE.

function. To evaluate the residual of the PDE and the boundary/initial conditions, we introduce a set of random points inside the domain ( $\mathcal{T}_f$ ) and another random set of points on the boundary ( $\mathcal{T}_b$ ). The loss function is then defined as Raissi et al. (2019) and Lu et al. (2021a)

$$\mathcal{L}(\theta; \mathcal{T}) = w_f \mathcal{L}_f(\theta; \mathcal{T}_f) + w_b \mathcal{L}_b(\theta; \mathcal{T}_b),$$

where

$$\mathcal{L}_f(\theta; \mathcal{T}_f) = \frac{1}{|\mathcal{T}_f|} \sum_{\mathbf{x} \in \mathcal{T}_f} \left| f\left(\mathbf{x}; \frac{\partial \hat{u}}{\partial x_1}, \dots, \frac{\partial \hat{u}}{\partial x_d}; \frac{\partial^2 \hat{u}}{\partial x_1 \partial x_1}, \dots, \frac{\partial^2 \hat{u}}{\partial x_1 \partial x_d}; \dots; \lambda \right) \right|^2, \tag{2}$$

$$\mathcal{L}_b(\theta; \mathcal{T}_b) = \frac{1}{|\mathcal{T}_b|} \sum_{\mathbf{x} \in \mathcal{T}_b} |\mathcal{B}(\hat{u}, \mathbf{x})|^2. \tag{3}$$

Here,  $w_f$  and  $w_b$  are the weights, and their choice is very important, especially for multiscale problems; below we show an effective algorithm on how to learn these weights, which can be a function of space-time, directly from the data.

One main advantage of PINNs is that the same formulation can be used not only for forward problems but also for inverse PDE-based problems. If the parameter  $\lambda$  in Equation (1) is unknown, and instead we have some extra measurements of  $u$  on the set of points  $\mathcal{T}_i$ , then we add an additional data loss (Raissi et al., 2019; Lu et al., 2021a) as


$$\mathcal{L}_i(\theta, \lambda; \mathcal{T}_i) = \frac{1}{|\mathcal{T}_i|} \sum_{\mathbf{x} \in \mathcal{T}_i} |\hat{u}(\mathbf{x}) - u(\mathbf{x})|^2$$

to learn the unknown parameters simultaneously with the solution  $u$ . We can recast the loss function as

$$\mathcal{L}(\theta, \lambda; \mathcal{T}) = w_f \mathcal{L}_f(\theta, \lambda; \mathcal{T}_f) + w_b \mathcal{L}_b(\theta, \lambda; \mathcal{T}_b) + w_i \mathcal{L}_i(\theta, \lambda; \mathcal{T}_i).$$

### 2.1. Self-adaptive loss weights

In some cases, it is possible to enforce the boundary conditions automatically by modifying the network architecture (Lagaris et al., 1998; Pang et al., 2019; Lagari et al., 2020; Lu et al., 2021b). A more general approach was proposed by McClenny and Braga-Neto (2020), which we explain next and give a simple



**Figure 4.** Self-adaptive weights. Solution of (4) using vanilla PINN (left) and self-adaptive PINNs(right). The self-adaptive PINN can capture the boundary layers whereas the vanilla PINN fails.

example of a boundary layer problem. Let us consider a boundary layer problem defined by an ODE as follows:

$$v \frac{d^2 u}{dx^2} - u = e^x, \tag{4}$$

where  $x \in [-1, 1]$ ,  $u[-1] = 1$ ,  $u[1] = 0$ , and viscosity  $v = 10^{-3}$ .

To compute the solution of (4), the loss function for self-adaptive PINN is defined as

$$\mathcal{L}(w, \lambda_r, \lambda_b, \lambda_0) = \mathcal{L}_r(w, \lambda_r) + \mathcal{L}_b(w, \lambda_b) + \mathcal{L}_0(w, \lambda_0), \tag{5}$$

where  $\lambda_r = (\lambda_r^1, \dots, \lambda_r^{N_r})$ ,  $\lambda_b = (\lambda_b^1, \dots, \lambda_b^{N_b})$ , and  $\lambda_n = (\lambda_0^1, \dots, \lambda_0^{N_0})$  are trainable self-adaptation weights for the initial, boundary, and collocation points, respectively, and

$$\begin{aligned} \mathcal{L}_r(w, \lambda_r) &= \frac{1}{N_r} \sum_{i=1}^{N_r} [\lambda_r^i r(x_r^i, t_r^1; w)]^2 \\ \mathcal{L}_b(w, \lambda_b) &= \frac{1}{N_b} \sum_{i=1}^{N_b} [\lambda_b^i (u(x_b^i, t_b^i; w) - g_b^i)]^2 \\ \mathcal{L}_0(w, \lambda_0) &= \frac{1}{N_0} \sum_{i=1}^{N_0} [\lambda_0^i (u(x_0^i, 0; w) - h_0^i)]^2. \end{aligned} \tag{6}$$

The key feature of self-adaptive PINNs is that the loss  $\mathcal{L}(w, \lambda_r, \lambda_b, \lambda_0)$  is minimized with respect to the network weights  $w$ , as usual, but is maximized with respect to the self-adaptation weights  $\lambda_r, \lambda_b, \lambda_0$ ; in other words, in training we seek a saddle point

$$\min_w \max_{r, \lambda_b, \lambda_0} \mathcal{L}(w, \lambda_r, \lambda_b, \lambda_0). \tag{7}$$

This can be accomplished by a gradient descent–ascent procedure, with updates given by

$$\begin{aligned} w^{k+1} &= w^k - \eta_k \nabla_w \mathcal{L}(w^k, \lambda_r^k, \lambda_b^k, \lambda_0^k), \\ \lambda_r^{k+1} &= \lambda_r^k + \eta_k \nabla_{\lambda_r} \mathcal{L}(w^k, \lambda_r^k, \lambda_b^k, \lambda_0^k), \\ \lambda_b^{k+1} &= \lambda_b^k + \eta_k \nabla_{\lambda_b} \mathcal{L}(w^k, \lambda_r^k, \lambda_b^k, \lambda_0^k), \\ \lambda_0^{k+1} &= \lambda_0^k + \eta_k \nabla_{\lambda_0} \mathcal{L}(w^k, \lambda_r^k, \lambda_b^k, \lambda_0^k). \end{aligned} \tag{8}$$

In the case of vanilla PINN, the  $\lambda$ 's are positive constant scalars, which weight the solution equally and do

not adapt to the regularity of the solution. However, in self-adaptive PINN, initial, boundary or collocation points in stiff regions of the solution automatically emphasize more these terms in the loss function, hence forcing the approximation to improve on those regions as shown in Figure 4.

2.2. Gradient-enhanced training of PINNs

In the standard PINN algorithm we minimize the residual  $f$  of the PDE in the  $L_2$ -norm but it may be beneficial to minimize it in the Sobolev norm, and this makes sense since the derivatives of  $f$  are also zero. This was proposed in Yu et al. (2021) in the gradient-enhanced PINNs (gPINNs), that is,

$$\nabla f(\mathbf{x}) = \left( \frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \dots, \frac{\partial f}{\partial x_d} \right) = \mathbf{0}, \quad \mathbf{x} \in \Omega.$$

Hence, the loss function of gPINNs is:

$$\mathcal{L} = w_f \mathcal{L}_f + w_b \mathcal{L}_b + w_i \mathcal{L}_i + \sum_{i=1}^d w_{g_i} \mathcal{L}_{g_i}(\boldsymbol{\theta}; \mathcal{T}_{g_i}),$$

where

$$\mathcal{L}_{g_i}(\boldsymbol{\theta}; \mathcal{T}_{g_i}) = \frac{1}{|\mathcal{T}_{g_i}|} \sum_{\mathbf{x} \in \mathcal{T}_{g_i}} \left| \frac{\partial f}{\partial x_i} \right|^2. \tag{9}$$

Here,  $\mathcal{T}_{g_i}$  is the set of residual points for the derivative  $\frac{\partial f}{\partial x_i}$ , and in general,  $\mathcal{T}_f$  and  $\mathcal{T}_{g_i}$  ( $i = 1, \dots, d$ ) can be different.

For example, for the Poisson’s equation  $\Delta u = f$  in 1D, the additional loss term is

$$\mathcal{L}_g = w_g \frac{1}{|\mathcal{T}_g|} \sum_{x \in \mathcal{T}_g} \left| \frac{d^3 u}{dx^3} - \frac{df}{dx} \right|^2.$$

Next, we present a comparison between PINNs and gPINNs based on the work of Yu et al. (2021) for the following diffusion equation

$$\frac{\partial u}{\partial t} = D \frac{\partial^2 u}{\partial x^2} + R(x, t), \quad x \in [-\pi, \pi], t \in [0, 1],$$

where  $D = 1$ . Here,  $R$  is given by:

$$R(x, t) = e^{-t} \left[ \frac{3}{2} \sin(2x) + \frac{8}{3} \sin(3x) + \frac{15}{4} \sin(4x) + \frac{63}{8} \sin(8x) \right].$$

The initial and boundary conditions are:

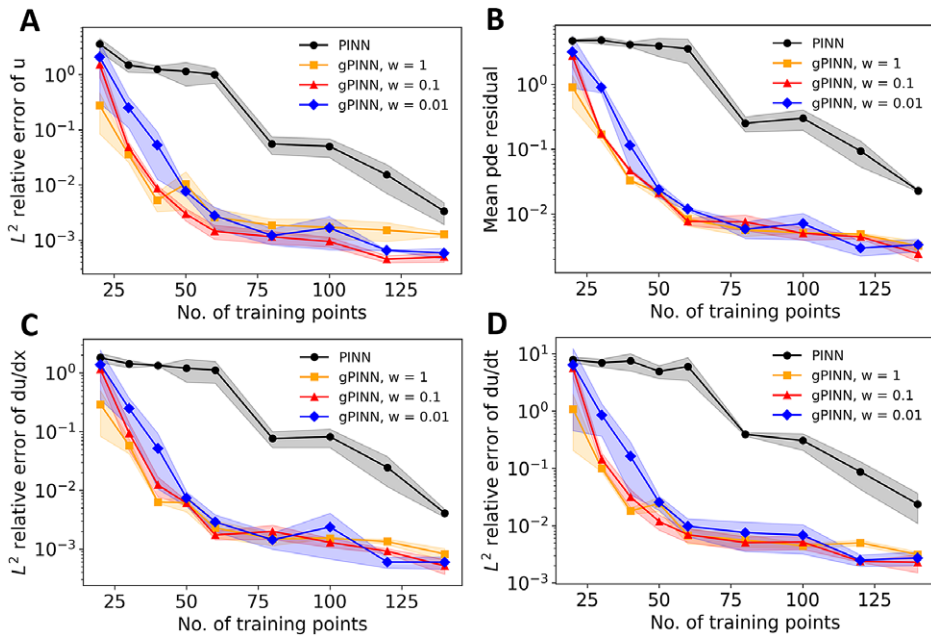
$$u(x, 0) = \sum_{i=1}^4 \frac{\sin(ix)}{i} + \frac{\sin(8x)}{8},$$

$$u(-\pi, t) = u(\pi, t) = 0,$$

corresponding to the analytical solution for  $u$ :

$$u(x, t) = e^{-t} \left[ \sum_{i=1}^4 \frac{\sin(ix)}{i} + \frac{\sin(8x)}{8} \right]. \tag{10}$$

We have two loss terms of the gradient, and the total loss function is



**Figure 5.** Comparison between PINN and gPINN. (a)  $L^2$  relative error of  $u$  for PINN and gPINN with  $w = 1, 0.1,$  and  $0.01$ . (b) Mean absolute value of the PDE residual. (c)  $L^2$  relative error of  $\frac{du}{dx}$  (d)  $L^2$  relative error of  $\frac{du}{dt}$  (figure is from Yu et al., 2021).

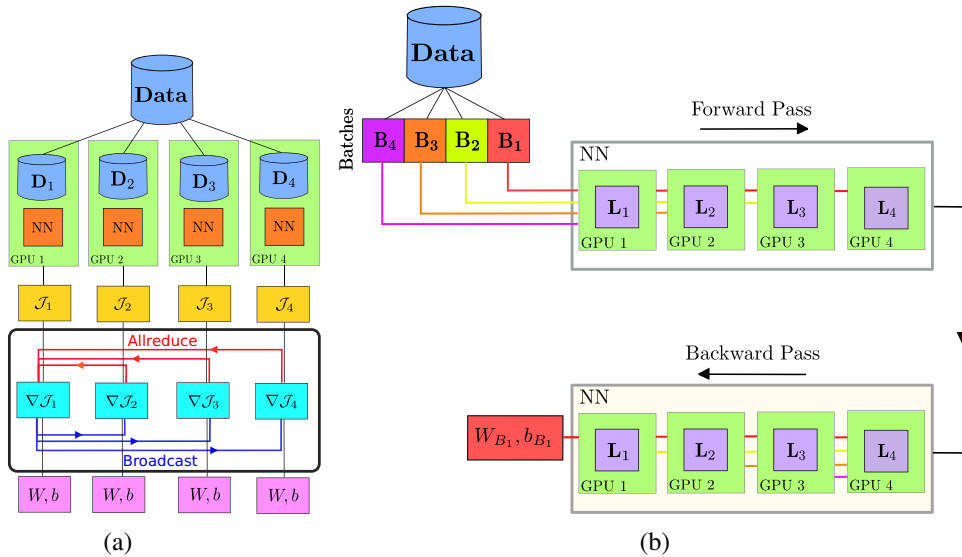
$$\mathcal{L} = \mathcal{L}_f + w\mathcal{L}_{g_x} + w\mathcal{L}_{g_t}.$$

We see in Figure 5 that gPINN results with the values of  $w = 0.01, 0.1,$  and  $1$  all outperform PINN by up to two orders of magnitude. However, we note that while gPINN reaches 1%  $L^2$  relative error of  $u$  by using only 40 training points, and PINN requires more than 100 points to reach the same accuracy, gPINN is almost twice as expensive as PINN due to extra expensive automatic differentiation.

### 2.3. Scalable PINNs

As discussed above, PINNs can efficiently tackle both forward problems, where the solution of governing physical law is inferred, as well as ill-posed inverse problems, where the unknown physics and/or free parameters in the governing equations are to be inferred from the available multi-modal measurements. However, one of the major concerns with PINNs is the large computational cost associated with the training of the neural networks, especially for forward multiscale problems. To improve the scalability and reduce the computational cost, Jagtap et al. (2020) introduced a domain decomposition-based (in-space) PINN for conservation laws, namely conservative PINN (cPINN), where continuity of the state variables as well as their fluxes across subdomain interfaces is used to obtain the global solution from the local solutions in the subdomains. In subsequent work, Jagtap and Karniadakis (2020) applied domain decomposition to general PDEs using the so-called extended PINN (XPINN). Unlike cPINN, which offers space decomposition, XPINN offers both space–time domain decomposition for any irregular, nonconvex geometry thereby reducing the computational cost effectively. By exploiting the decomposition process of the cPINN/XPINN methods and its implicit mapping on the modern heterogeneous computing platforms, the training time of the network can be reduced to a great extent.

There are currently two existing approaches for distributed training of neural networks, namely, the data-parallel approach (Sergeev and Del Balso, 2018) and the model parallel approach, which are agnostic to physics-based priors. The data-parallel approach is based on the single instruction and multiple data



**Figure 6.** Schematic of the implementation of data and model parallel algorithms. (a) Method for the data-parallel approach, where the same neural network model, represented by NN, is loaded by each processor but works on a different chunk of input data. Synchronization of training (gradient of loss) is performed after the computation of loss on each processor via “allreduce” and “broadcast” operations represented by horizontal red and blue lines. (b) Represents the model parallel approach, where each layer of the model (represented by  $L_1 \dots L_4$ ) is loaded on a processor and each processor works on a batch of data ( $B_1 \dots B_4$ ). Forward and backward passes are performed by using a pipeline approach. Adapted from Shukla et al. (2021b).

(SIMD) parallel programming model, which results in a simple performance model driven by weak scaling. A block diagram showing the basic building blocks of the data-parallel approach is shown in Figure 6a for a four processors or coprocessors system. The programming model used for the data-parallel approach falls into the regime of message passing interface (MPI) +  $X$  system, where  $X$  is chosen as CPU(s) or GPU(s), depending on the size of the input data. In the data-parallel approach, the data is uniformly split into a number of chunks ( $D_1, \dots, D_4$  in Figure 6a), equal to the number of processors. The neural network (NN) model is initialized with the same parameters on all the processes as shown in Figure 6a. These neural networks are working on different chunks of the data, and therefore, work on different loss functions as shown by  $\mathcal{J}_1, \dots, \mathcal{J}_4$  in Figure 6a.

NVIDIA also introduced the parallel code MODULUS (Hennigh et al., 2020), which implements the standard PINN-based multiphysics simulation framework. The underlying idea of MODULUS is to solve the differential equation by modeling the mass balance condition as a hard constraint as well as a global constraint. MODULUS provides the functionality for multi-GPU and multinode implementation based on the data-parallel approach (Figure 6a).

To ensure a consistent neural network model (defined with same weights and biases) across all the processes and during each epoch or iteration of training, a distributed optimizer is used, which averages out the gradient of loss values ( $\nabla \mathcal{J}_1, \dots, \nabla \mathcal{J}_4$  in Figure 6a) on the root processor through an “allreduce” operation. Subsequently, the averaged gradient of the loss function is sent to all the processors in the communicator world through a broadcast operation (collective communication). Then, the parameters of the model are updated through an optimizer. An additional component, which arises in the data-parallel approach, is the increase in global batch size with the number of processes. Goyal et al. (2017) addressed this issue by multiplying the learning rate with the number of processes in the communicator. We note that in the data-parallel approach the model size (size of neural network parameters) remains uniform on each

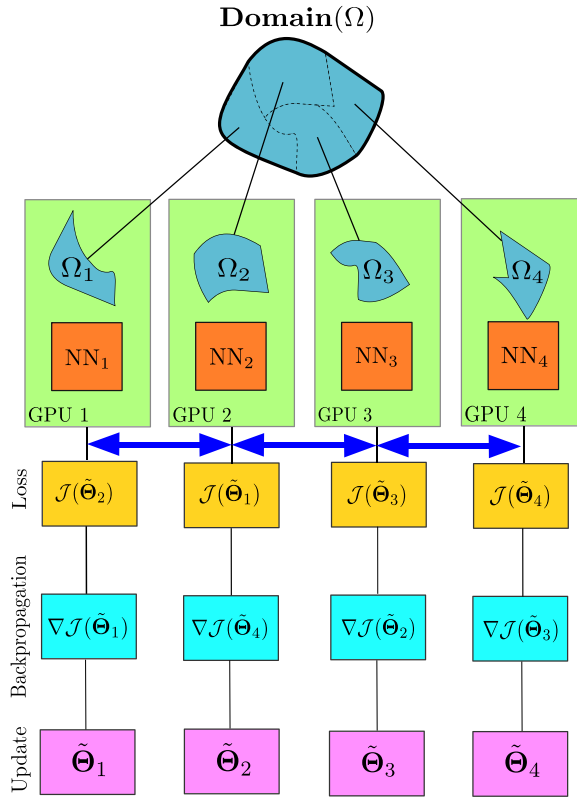


processor or GPU and that imposes a problem for large models to be trained on GPUs as they have fixed memory.


To circumvent this issue, another distributed algorithm approach namely, the model parallel approach is proposed. A block diagram of the algorithm in the model parallel approach is shown in Figure 6b, which can be interpreted as a classic example of pipeline algorithm (Rajbhandari et al., 2020). In the model parallel approach, data is first divided into small batches ( $\mathbf{B}_1, \dots, \mathbf{B}_4$  in Figure 6b) and each hidden layer ( $\mathbf{L}_1, \dots, \mathbf{L}_4$  in Figure 6b) is distributed to a processor (or GPU). During the forward pass, once  $\mathbf{B}_1$  is processed by  $\mathbf{L}_1$ , the output of  $\mathbf{L}_1$  is passed as the input to  $\mathbf{L}_2$  and  $\mathbf{L}_1$  will start working on  $\mathbf{B}_2$  and so on. Once  $\mathbf{L}_4$  finishes working,  $\mathbf{B}_1$  the backward pass (optimization process) will kickoff, thus, completing epochs of training in a pipeline mode. We note that the implementation of both algorithms is problem agnostic and does not incorporate any prior information on solutions to be approximated, which makes the performance of these algorithms to be dependent on the data size and model parameters. In the literature, the implementation of data and model parallel approaches is primarily carried out for problems pertaining to the classification and natural language processing (Goyal et al., 2017; Rasley et al., 2020), which are based on large amounts of training data. Therefore, the efficiency of data and the model parallel approach for scientific machine learning is not explored, which is primarily dominated by the high-dimensional and sparse data set. Apart from these two classical approaches, recently Xu et al. (2020a) deployed the topological concurrency on data structures of neural network parameters. In brief, this implementation could be comprehended as task-based parallelism and it is rooted in the idea of model-based parallelism. Additionally, it also provides interactivity with other discretization-based solvers such as Fenics (Alnæs et al., 2015).

The main advantage of cPINN and XPINN over data and model parallel approaches (including the work of Xu et al., 2020a) is their adaptivity for all hyperparameters of the neural network in each subdomain. In the vanilla data-parallel approach, the training across processors is synchronized by averaging the gradient of loss function across and subsequently broadcasting, enforcing the same network architectures on each processor. In PIML, the solution of the underlying physical laws is inferred based on a small amount of data, where the chosen neural network architecture depends on the complexity and regularity of the solutions to be recovered. To address this issue, the synchronization of the training process across all the processors has to be achieved by the physics of the problem, and therefore cPINNs and XPINNs come to the rescue. The convergence of the solution based on a single domain is constrained by its global approximation, which is relatively slow. On the other hand, computation of the solution based on local approximation by using the local neural networks can result in fast convergence. With regards to computational aspects, the domain decomposition-based approach requires a point-to-point communication protocol with a very small size of the buffer. However, for the data-parallel approach, it requires one *allreduce* operation and one *broadcast* operation with each having the buffer size equal to the number of parameters of the neural network, which makes communication across the processor slower. Moreover, for problems involving heterogeneous physics the data/model parallel approaches are not adequate, whereas cPINNs/XPINNs can easily handle such multiphysics problems. Therefore, domain decomposition approaches paired with physics-based synchronization help more than the vanilla data-parallel approaches.

The implementation of cPINN and XPINN on distributed and heterogeneous computing architectures are discussed in detail by Shukla et al. (2021b). The core idea of implementation is based on the domain decomposition approach, where in each subdomain we compute the corresponding solution, and the residuals and solutions at shared edges or planes are passed using MPI as shown in Figure 7 to be averaged. Figure 8 shows the weak scaling of cPINN and XPINN for up to 24 GPUs measured on the supercomputer Summit of Oak Ridge National Laboratory. The weak scaling of cPINN and XPINN shows a very good agreement between observed and theoretical performance. The weak scaling also shows that XPINN has less throughput than the cPINN, which could be easily justified as computation of flux typically requires one order less auto-differentiation computation than computation of residuals for



**Figure 7.** Parallel PINNs. Building blocks of distributed cPINN and XPINN methodologies deployed on a heterogeneous computing platform. The domain  $\Omega$  is decomposed into several subdomains  $\Omega_1, \dots, \Omega_4$  equal to the number of processors, and individual neural networks ( $NN_1, \dots, NN_4$ ) are employed in each subdomain, which gives separate loss functions  $\mathcal{J}(\Theta), q = 1, \dots, 4$  coupled through the interface conditions (shown by the double-headed blue arrow). The trainable parameters are updated by finding the gradient of loss function individually for each network and penalizing the continuity of interface conditions.



**Figure 8.** Two-dimensional incompressible Navier–Stokes equations: Weak GPU scaling for the distributed cPINN (left) and XPINN (right) algorithms.

XPINN. Shukla et al. (2021b) carried out a detailed strong and weak scaling along with communication and computation times required for both methods.

#### 2.4. Inverse problem with parallel PINNs

To validate the scalability of cPINN and XPINN for a very large inverse problem, we solve the steady-state heat conduction problem over a domain represented by the USA map. The governing equation for the steady-state heat conduction problem is written as

$$\partial_x(K(x,y)T_x) + \partial_y(K(x,y)T_y) = f(x,y) \quad (11)$$

with Dirichlet boundary conditions for temperature  $T$  and thermal conductivity  $K$ , obtained from the exact solution. The forcing term  $f(x,y)$  is obtained from the exact solution for variable thermal conductivity and temperature, which is assumed to be of the form

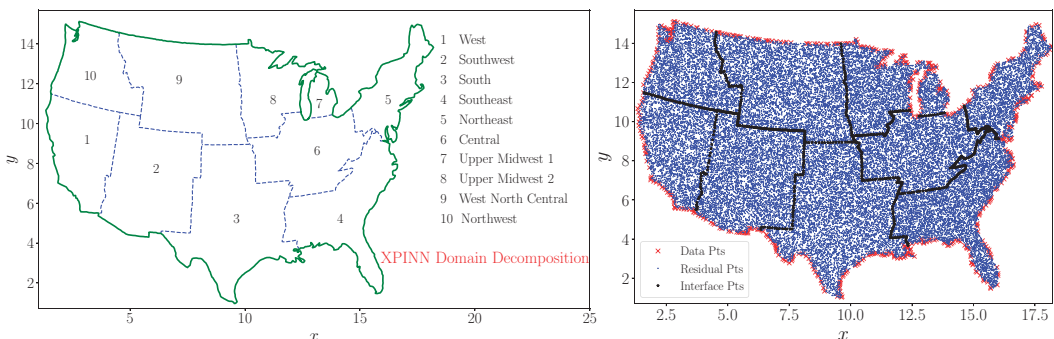
$$\begin{aligned} T(x,y) &= 20 \exp(-0.1y) \\ K(x,y) &= 20 + \exp(0.1y) \sin(0.5x). \end{aligned}$$

Materials whose thermal conductivity is a function of space are the so-called *functionally graded materials*, which are a particular type of composite materials. The material domain is chosen to be a map of the United States divided into 10 subdomains as shown in Figure 9. The interfaces between these subdomains are shown by the dashed blue line whereas the domain boundary is shown by the solid green line. The boundary and the training data is available for every subdomain. Unlike cPINN, XPINN can easily handle such complex subdomains with arbitrarily shaped interfaces, thus, providing greater flexibility in terms of domain subdivision. Figure 9(right) shows the residual, training data, and interface points in blue dot, red cross, and black asterisk, respectively, over the entire domain.

In this inverse problem, the temperature is assumed to be known in the domain but the variable thermal conductivity is unknown. The aim is to infer the unknown thermal conductivity of the material from a few data points for  $K$  available along the boundary line and temperature values available inside the domain. We employed a single PINN in each subdomain, and the details about the network's architecture are given in Table 1. In each subdomain, we used 3 hidden-layers with 80 neurons in each layer, and the learning rate is  $6e-3$ , which is fixed for all networks.

Figure 10 (top row) shows the predicted values of temperature and thermal conductivity. The absolute point-wise errors are given in the bottom row. The XPINN method accurately inferred the thermal conductivity, which shows that XPINN can easily handle highly irregular and nonconvex interfaces.

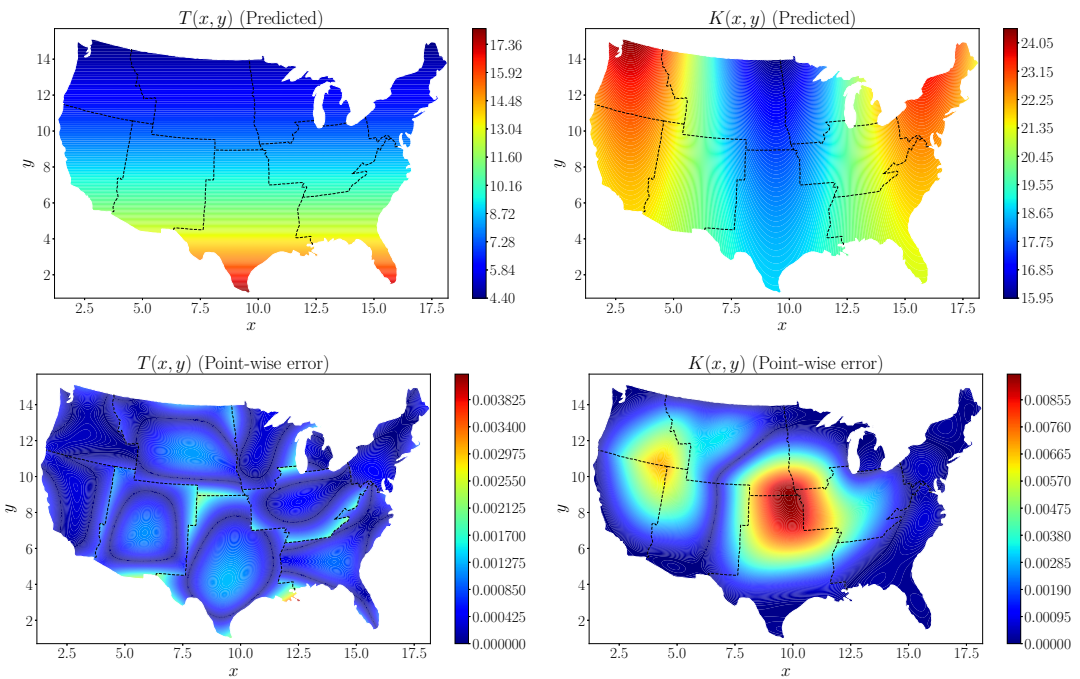
Figure 11 represents the performance of the parallel XPINN algorithm deployed for solving the inverse heat conduction problem. The performance is measured for the data points given in Table 1. The left



**Figure 9.** Steady-state heat conduction with variable conductivity: Domain decomposition of the US map into 10 regions (left) and the corresponding data, residual, and interface points in these regions (right).

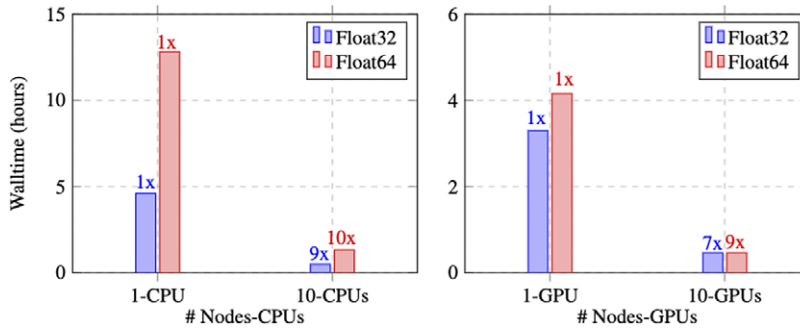
**Table 1.** Steady-state heat conduction with variable conductivity: Neural network architecture in each subdomain.

Subdomain number	1	2	3	4	5
# Residual points	3,000	4,000	5,000	4,000	3,000
Adaptive activation function	tanh	sin	cos	tanh	sin
Subdomain number	6	7	8	9	10
# Residual points	4,000	800	3,000	5,000	4,000
Adaptive activation function	cos	tanh	sin	cos	tanh



**Figure 10.** Steady-state heat conduction with variable conductivity: The first row shows the contour plots for the predicted temperature  $T(x,y)$  and thermal conductivity  $K(x,y)$  while the second row shows the corresponding absolute point-wise errors. Adapted from Shukla et al. (2021b).

subfigure in Figure 11 represents the wall time and scaling of the parallel XPINN algorithm on CPUs by computing the solution in each domain using one CPUs. The continuity of solution on the shared boundaries of the subdomain is imposed by passing the solution on each domain boundary to the neighboring subdomains through MPI protocols. Here, each CPU corresponds to a rank mapped by node, and Intel’s Xeon(R) Gold 6126 CPUs are used for measurement. First, we measured the wall time of the algorithm on one CPU and considered it as baseline data (1X) for the scaling. Thereafter, we computed the wall time for 10 CPUs, leading to a (9X) and (10X) for single (32-bit float) and double-precision (64-bit float) computation, respectively. The scaling for double-precision is relatively better as the communication process is shadowed by the computation time due to more time being spent on double-precision



**Figure 11.** Steady-state heat conduction with variable conductivity: Walltime and speedup of parallel XPINN algorithm on CPUs and GPUs implemented for the inverse heat conduction problem in (11); (a) speedup and wall time for the parallel XPINN code on Intel's Xeon(R) Gold 6126 CPU. The speed and wall time is measured for computations performed with single (Float32) and double-precision numbers (Float64); (b) speedup and wall time, measured for single- and double-precision operations, on Nvidia's V100 GPUs. Adapted from Shukla et al. (2021b)).

arithmetic. However, double-precision-based computation increases wall time by a factor of 2.5 for single and multiple CPU-based implementations.

Next, we report the walltime and GPU-based implementation. The algorithm is implemented on Nvidia V100 GPUs with 16 GB memories. The hardware architecture is similar to that reported in the example of incompressible Navier–Stokes equations. The right subfigure of Figure 11 represents the walltime with 1 and 10 GPUs for single- and double-precision arithmetic. On a single GPU, considered as (1X) model the wall time for double-precision arithmetic 30% more than that received for single-precision. In the multi-GPU implementation, one GPU is used for each subdomain, and here one rank mapped by the node corresponds to the combination of 1CPU + 1GPU. The walltime for 10-GPUs (10 Nodes or 10 ranks) yields a scaling (7X) and (9X) for single- and double-precision arithmetic. We note that the residual points in each subdomain are not enough to saturate the GPUs and therefore more time is spent on fetching the data to memory and inter-GPU communication. To provide more intuition in the purview of the current implementation, for a typical V100 (16 GB) memory, the single-precision performance is 14 TFLOPs with a memory bandwidth of 900 GB/s and therefore for each byte of transfer (memory to GPU core) 15.6 instruction (FLOPs/Bandwidth) needs to be issued to occupy the GPUs completely. In the context of the current problem, the partition and therefore the load on GPUs (or CPUs) are static and does not change throughout the computation; thus, subdomain 7, endowed with only 800 residual points, has to wait until all the GPUs (or CPUs) complete their work for the respective subdomain. This also results in slight sublinear performance. However, the performance presented in Figure 11 shows a very good linear scaling for CPUs on a heterogeneous (CPU + GPU) architecture. Additionally, the scaling of the presented algorithm on CPUs only architecture concurs with the idea of Daghighi et al. (2021), where the authors try to revisit the algorithms used in machine learning to make them faster on CPUs rather than getting fixated with one-dimensional development of specific hardware to run the matrix multiplication faster. In this test case, the partition of the domain is performed by manually choosing the interface points. This was done to show the efficacy of the XPINN algorithm for nonconvex or irregular subdomains. In such complex subdomains, the distribution system often faces the problem of load imbalance, which can seriously degrade the performance of the system. A more efficient approach could be utilized to decompose the domain such that the subdomains are packed optimally for internode or interprocess communication. A suitable point cloud (Rusu and Cousins, 2011) or K-way partition (Karypis and Kumar, 2009) based approach will result in further optimizing the communication.

### 3. Graph Neural Networks

In contrast to more traditional machine learning (ML) data handled in image recognition or natural language processing, scientific data is fundamentally unstructured, requiring processing of polygonal finite element meshes, LIDAR point clouds, Lagrangian drifters, and other data formats involving scattered collections of differential forms. This poses both challenges and opportunities. Convolutional networks serve as a foundation for image processing, and are able to employ weight-sharing by exploiting the Cartesian structure of pixels, but cannot be applied in unstructured settings where stencils vary in size and shape from node to node. Unstructured data do, however, encode nontrivial topological information regarding connectivity, allowing application of ideas from topological data analysis and combinatorial Hodge theory (Carlsson, 2009; Jiang et al., 2011; Bubenik, 2015; Wasserman, 2018). In contrast, GNNs form a class of deep neural network methods designed to handle unstructured data. Graphs serve as a flexible topological structure, which may be used for learning useful latent “graph-level,” “node-level,” or “edge-level” representations of data for diverse graph prediction tasks and analytics (Monti et al., 2017).

This unstructured nature enables GNNs to naturally handle a wide range of graph analytics problems (i.e., node classification, link prediction, data visualization, graph clustering community detection, anomaly detection) and have been applied effectively across a diverse range of domains, for example, protein structure prediction (Jumper et al., 2021), untangling the mathematics of knots (Davies et al., 2021), brain networks (Rosenthal et al., 2018; Xu et al., 2020b, 2021a) in brain imaging, molecular networks (Liu et al., 2019) in drug discovery, protein–protein interaction networks (Kashyap et al., 2018) in genetics, social networks (Wang et al., 2019b) in social media, bank-asset networks (Zhou and Li, 2019) in finance, and publication networks (West et al., 2016) in scientific collaborations.

In this section, we consider the task of fitting dynamics to data defined on a set of nodes  $\mathcal{N}$ , often associated with an embedding as vertices  $\mathcal{V} \subset \mathbb{R}^d$ . We will consider the graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{E}$  represents the edge set denoted by  $\mathcal{E} \subseteq \binom{\mathcal{V}}{2}$ . We will also consider higher-order  $k$ -cliques consisting of ordered  $k$ -tuples of nodes, where for example, an edge  $e \in \mathcal{E}$  corresponds to a 2-clique. We associate scalar values  $\mathbf{x}$  with nodes, edges, or both, and aim to either learn dynamics ( $\dot{\mathbf{x}} = L(\mathbf{x}; \theta)$ ) or boundary value problems ( $L(\mathbf{x}; \theta) = f$ ) for learnable parameters  $\theta$ .

We present three approaches: (a) methods which aim to learn operators via finite difference-like stencils, (b) GNNs which represent operators by learning update and aggregation maps, and (c) graph calculus methods, which relate learnable operators to graph analogs of the familiar vector calculus div/grad/curl. First, we survey prevailing strategies and architectures for traditional GNNs.

#### 3.1. Basics of GNNs

Broadly, GNNs may be classified into three major categories: (a) Spectral GNNs, which locally aggregate connected node information by solving eigenproblems in the *spectral domain*. (b) Spatial GNNs, which perform graph convolution via aggregating node neighbor information from the first-hop neighbors in the *spatial graph domain*. (c) Graph attention networks (GATs), which leverage the *self-attention* mechanism for hidden feature aggregation. We recall first some mathematical definitions before presenting each in turn.

##### 3.1.1. Notation and mathematical preliminaries

Consider a node  $v_i \in \mathcal{V}$  and  $e_{ij} = \{v_i, v_j\} \in \mathcal{E}$  indicating either a directed or undirected edge between nodes  $v_i$  and  $v_j$ . By  $j \sim i$ , we denote neighbors  $j$  with corresponding edges  $(i, j)$  coincident upon node  $i$ .  $\mathbf{A} \in \mathbb{R}^{N \times N}$  corresponds to the adjacency matrix of  $\mathcal{G}$ . When  $\mathcal{G}$  is binary (or “unweighted”),  $A_{ij} = 1$  if there is an edge between nodes  $v_i$  and  $v_j$ , else 0. If  $\mathcal{G}$  is a weighted graph,  $A_{ij} = w_{ij}$  when there is an edge between nodes  $v_i$  and  $v_j$ . A graph is *homogeneous* if all nodes and edges are of a single type, and vice versa, the graph is *heterogeneous* if it consists of multiple types of nodes and edges. Additionally, each node may be associated with features (or attributes)  $\mathbf{x} \in \mathbb{R}^{N \times F}$  referring to the node attribute matrix of  $\mathcal{G}$ . The entry  $x_i \in \mathbb{R}$

corresponds to the node feature vector of node  $v_i$ . The graph Laplacian matrix is an  $N \times N$  matrix defined as  $\mathbf{L} = \mathbf{D} - \mathbf{A}$ , where  $\mathbf{D}$  is a  $N \times N$  diagonal node degree matrix whose  $i$ th element is the degree of the node  $i$ , that is,  $\mathbf{D}(i, i) = \sum_{j=1}^N A_{ij}$ .

### 3.1.2. Spectral methods

The main idea of spectral GNN approaches is to use Eigen-decompositions of the graph Laplacian matrix  $\mathbf{L}$  to generalize spatial convolutions to graph structured data. This allows simultaneous access to information over short and long time scales (Bruna et al., 2013). Spectral graph convolutions are defined in the spectral domain based on the graph Fourier transform.

Specifically, the normalized self-adjoint positive-semidefinite graph Laplacian matrix  $\mathbf{L}$  describes diffusion over a graph (Chung and Graham, 1997), and can be factorized by  $\mathbf{L} = U\Lambda U^T$ .  $U$  is  $N \times N$  matrix whose  $l$ th column is the eigenvector ( $u_l$ ) of  $\mathbf{L}$ ;  $\Lambda$  is a diagonal matrix whose diagonal elements  $\lambda_l = \Lambda(l, l)$  ( $l \in \mathcal{R}^N$ ) are the corresponding eigenvalues. We can apply the operation  $U^T \phi$  to project data  $\phi$  from graph nodes into the spectral domain, where features can be decomposed into a sum of orthogonal graphlets or motifs (i.e., eigenvectors) that make up  $\mathcal{G}$ . Typically, eigenvectors with the smallest eigenvalues may be used to distinguish vertices that will be slowest to share information under diffusion or message passing. Projecting vertex features onto these eigenvectors and processing the projected graph signal is equivalent to augmenting message passing with a nonlinear low-pass filter.

Recently, many spectral GNN variants have been proposed, for example, Bruna et al. (2013) and Kipf and Welling (2016). The main workflow of spectral GNNs follows four main steps: (a) transform the graph to the spectral domain using the graph Laplacian eigenfunctions (see Equation (12)); (b) perform the same transformation on the graph convolutional filters; (c) apply convolutions in the spectral domain; (d) transform from the spectral domain back to the original graph domain (see Equation (13)):

$$\hat{\mathbf{x}}(\lambda_l) = \langle \mathbf{x}, \mathbf{u}_l \rangle = \sum_{i=1}^N \mathbf{x}(i) \mathbf{u}_l^T(i), \quad (12)$$

$$\mathbf{x}(i) = \sum_{l=1}^N \hat{\mathbf{x}}(\lambda_l) \mathbf{u}_l(i). \quad (13)$$

The detailed spectral graph convolution operation is described in Equation (14), where  $U$  is the eigenvector matrix,  $X^{(k)}$  represents the node feature map at the  $k$ th layer, and  $W^{(k)}$  denotes the learnable spectral graph convolutional filters at the  $k$ th layer.

$$X^{(k+1)} = U \left( U^T X^{(k)} \odot U^T W^{(k)} \right). \quad (14)$$

However, the aforementioned spectral GNN methods share a few significant weaknesses: the Laplacian eigenfunctions are inconsistent across different domains, and therefore the method generalizes poorly across different geometries. Moreover, the spectral convolution filter is applied for the whole graph without considering the local graph structure property. In addition, the graph Fourier transform has high cost in computation. To solve the locality and efficiency problem, ChebNets (Defferrard et al., 2016) employs the Chebyshev polynomial basis to represent spectral convolution filter instead of graph Laplacian eigenvectors.

### 3.1.3. Spatial methods

Spatial methods involve a form of learnable message-passing that propagates information over the graph via a local aggregation (or diffusion) process. Graph convolutions are defined directly on the graph topology in spatial methods (Atwood and Towsley, 2016; Niepert et al., 2016; Gilmer et al., 2017). Spatial Graph Convolutional Networks (GCNs) directly ignore the “edge attributes” as well as the crucial “messages” sent by nodes along edges in the graph, however, the Message Passing Neural Networks (MPNNs) can

effectively make use of this information to update node states. MPNN consists of three main phases: *message passing, readout, and classification*. Message passing consists of: (a) A node feature transformation based on some sort of projection. (b) Node feature aggregation with a permutation-invariant function, for example, averaging, sum, or concatenation. (c) A node feature update via the current states and representations aggregated from each node’s neighborhood. Specifically, the message passing runs for  $T$  time steps, and includes two key components: message functions  $M_t$  and node update functions  $U_t$ . The hidden states  $h_i^t$  at each node in the graph are updated based on Equation (15) with the current state ( $h_i^t$ ) and the aggregated messages  $m_i^{t+1}$ , which are computed from neighboring nodes’ feature ( $h_j^t$ ) and edge features ( $e_{ij}$ ) according to Equation (16).

$$h_i^{t+1} = U_t(h_i^t, m_i^{t+1}), \tag{15}$$

$$m_i^{t+1} = \sum_{j \in \mathcal{N}_i} M_t(h_i^t, h_j^t, e_{ij}), \tag{16}$$

where  $\mathcal{N}_i$  denotes the neighbors of node  $i$ . The readout phase computes a feature vector for the whole graph using some readout function  $R$  according to

$$\hat{y} = R(\{h_i^T \mid i \in \mathcal{G}\}).$$

The message functions  $M_t$ , node update functions  $U_t$  and readout function ( $R$ ) are all learned differentiable functions. Additionally,  $R$  must be invariant to node permutation (i.e., graph isomorphism) to guarantee identical output for equivalent graphs.

### 3.1.4. Graph attention networks

In spectral and spatial GNN models, the message aggregation operations from nodes’ neighborhoods are mostly guided by the graph structure, which weights contributions from neighboring nodes equally. In contrast, GATs (Veličković et al., 2017) assign during aggregation a learnable propagation weight via a “self-attention” function ( $a$ ), defined in Equation (17). Specifically, given the node pair feature vectors ( $h_i$  and  $h_j$ ), the attention coefficients ( $e_{ij}$ ) between nodes  $i$  and  $j$  can be computed by using a self-attention mechanism  $a : \mathbb{R}^{F'} \times \mathbb{R}^{F'} \rightarrow \mathbb{R}$  ( $F'$  is the node feature size), which is usually implemented by a single-layer feed-forward neural network. To this end, the attention mechanism can help identify the importance of interactions between each node neighbor. Here,  $c_{ij}$  reflects the importance of node  $j$ ’s feature to node  $i$ . For better comparability across all nodes in the graph, the normalized attention coefficients ( $\alpha_{ij}$ ) between node pair ( $i, j$ ) can be defined as a softmax function (see Equation (17)), where  $\mathcal{N}_i$  denote the neighborhood of node  $i$ , that is, first-order neighbors.

$$\alpha_{ij} = \frac{\exp(c_{ij})}{\sum_{k \in \mathcal{N}_i} \exp(c_{ik})}, c_{ij} = a(h_i, h_j). \tag{17}$$

In practice, a multihead attention mechanism has been widely used in GATs which can stabilize the process by using  $K$  independent attention mechanisms to achieve the transformation. The corresponding learned node feature for node  $i$  with multihead attention mechanism (Vaswani et al., 2017) is shown in Equation (18).

$$h_i' = \parallel_{k=1}^K \sigma \left( \sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}^k h_j \right), \tag{18}$$

where  $\alpha_{ij}^k$  denotes the attention coefficients derived by the  $k$ th attention head;  $\mathbf{W}^k$  is the weight matrix specifying the linear transformation of the  $k$ th attention head;  $\sigma$  is the nonlinear activation function;  $\parallel$  represents ultimate concatenation, sum or average operation for the features learned by  $K$  independent attention heads.



In summary, GCNs are a special case of GATs which use the spectral normalized graph adjacency matrix to serve the role of the attention function. GATs are a special case of MPNNs with hidden feature aggregation by using self-attention mechanism as message passing rule. All methods share common inductive bias and node/edge permutation invariance properties. Lastly, there are also some works (DyRep [Trivedi et al., 2019], EvolveGCN [Pareja et al., 2020], TGN [Rossi et al., 2020], DynG2G [Xu et al., 2021b], etc.) for *dynamic* graph structure data inference (climate network, traffic network, etc.).

### 3.2. Learning stencil operators

In this section, we consider learning operators expressed via a discrete stencil

$$L(\mathbf{x}; \theta) = \sum_{i \in \mathcal{V}} \theta_i \mathbf{x}_i, \quad (19)$$

where the learnable parameters  $\theta_i$  correspond to stencil coefficients. An attractive feature of this simple setup is that differential operators may be encoded as a finite difference stencil. However, unlike traditional finite difference methods, often we consider a fixed size graph and do not consider a sequence of refined grids which converge to some continuum operator. Therefore, we are more interested in recovering an effective discrete operator in the traditional finite difference sense.

In Bar-Sinai et al. (2019), the authors repurpose a convolutional network to this end. We note that this corresponds to a graph in which  $\mathcal{V}$  consists of a Cartesian grid of nodes, while  $\mathcal{E} = \{(i, j) \mid |x_i - x_j|_1 < \varepsilon\}$  consists of neighbor nodes in an  $\varepsilon$  ball in the taxicab norm; see also Figure 12. While simple, this work is extremely interpretable, and the authors demonstrate an ability to learn stencils for Burgers, Korteweg–de Vries, and Kuramoto–Sivashinsky equations with improved error compared to a naive application of the finite difference method. We recall from introductory finite difference methods (Gustafsson et al., 1995) that schemes such as the Lax–Wendroff or Macormack schemes introduce higher-order dissipation or dispersion into a discrete operator to achieve improved accuracy or stability. A simple Taylor series analysis reveals such schemes incorporate higher-order differential correction terms which depend on the spatiotemporal resolution, hence learning finite difference stencils may also learn similar higher-order corrections from data.

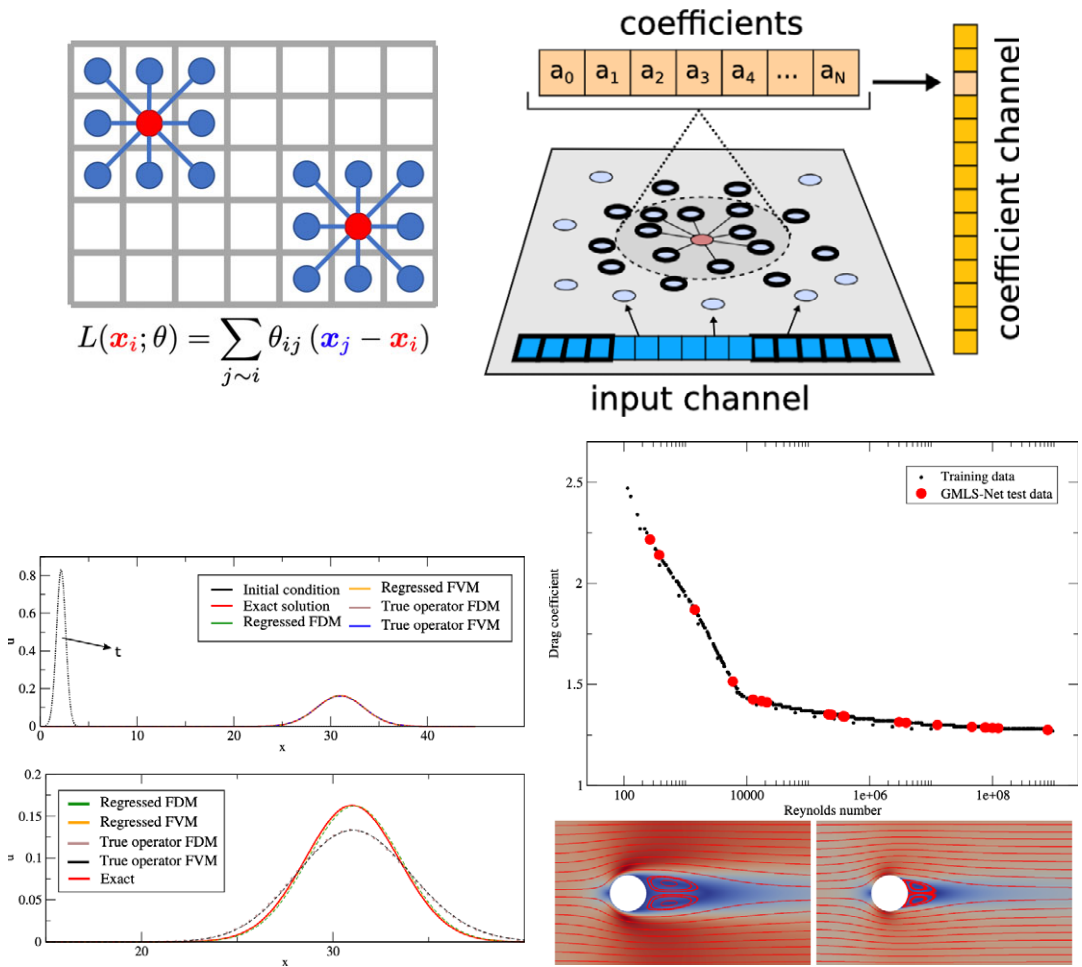
The requirement of a Cartesian node set is, however, restrictive and precludes application to complex geometries. In Trask et al. (2019) unstructured data is handled by finding a polynomial fit to data at each node of an unstructure  $\varepsilon$ -ball graph and then providing the resulting coefficient weights to a multilayer perceptron.

$$\min_{\mathbf{c}_i} \sum_{j \in \mathcal{V}} (u_j - \mathbf{c}_i^T P(x_j))^2 W(|x_i - x_j|) \quad L(\mathbf{x}; \theta) = \mathcal{N}(\mathbf{c}_i(\mathbf{x}); \theta). \quad (20)$$

Here,  $P$  denotes a vector of polynomial basis functions,  $W$  denotes a radial kernel with support  $B_\varepsilon(x_i)$ , and  $\mathcal{N}$  denotes a multilayer perceptron with weights and biases  $\theta$  taking the optimal polynomial reconstruction as input. In Trask et al. (2019), this architecture is shown to be effective for both learning physical operators on unstructured data as well as offering an alternative to spectral GNNs for supervised learning on structured and unstructured data; see Figure 12. This framework builds upon the generalized moving least squares (GMLS) theory, which has traditionally been used to develop finite difference discretizations on unstructured points (Mirzaei et al., 2012). In relation to spectral GNNs, replacing a global spectral approximation with a localized polynomial approximation, which may be trivially accelerated with GPUs (Kuberry et al., 2019) is attractive from both performance and approximation perspectives.

### 3.3. Physics-informed graph networks

The stencil formulation of the previous section prompts a natural interpretation of GNNs in the context of finite difference methods. In staggered finite difference schemes for PDEs, one uses a stencil to compute a conservation balance (e.g., a divergence operator for conservation of fluxes in marker-and-cell schemes



**Figure 12.** Learning operators from stencils. From left to right, top to bottom: (a) On a Cartesian grid of data, CNNs may employ weight-sharing to fit finite-difference operators to data (Bar-Sinai et al., 2019). (b) On unstructured data, similar weight-sharing may be achieved by lifting data first to a space of polynomial coefficients and learning how an operator acts on polynomials. (c) Learning operators as difference stencils allows learning of higher-order corrections. Shown here at a CFL condition of 10, a stencil learned from an analytic solution to the advection–diffusion problem does not exhibit the numerical dissipation of a traditional finite difference/volume discretization of the PDE. (d) Beyond learning physics, these frameworks can be used for supervised learning tasks on unstructured scientific data. Shown here, the drag force acting on a cylinder is regressed from nodal velocities on an unstructured finite volume mesh. Figures adapted from Trask et al. (2019).

for transport (Harlow and Welch, 1965), or a curl operator for circulations in the Yee scheme for electromagnetism (Yee, 1966) from fluxes/circulations computed on a dual grid. The aggregation step of GNNs admits interpretation as a graph divergence, with message passing corresponding to learning fluxes/circulations. In Kumar and Chakraborty (2021), the authors apply this viewpoint, using an attention mechanism (Veličković et al., 2018) to learn messages from data. They demonstrate accurate prediction of solution operators for the viscous Burgers equations in one- and two-dimensions.

In Pfaff et al. (2020), the authors consider a similar message-passing GNN where the graph is taken to be a traditional simulation mesh. By augmenting the network with edge features encoding the

displacement of the domain from a reference configuration, they demonstrate how GNNs may be used to train surrogates of physics-based models from high-fidelity simulations on unstructured meshes in either an Eulerian or Lagrangian configuration. They demonstrate speed ups of  $10 - 100\times$  for cloth dynamics, elasticity, and compressible flows and attribute good generalization beyond parameters used in training to invariances preserved by the GNN architecture. Also, worth noting is favorable comparison against convolutional methods, where message-passing GNNs provide improved accuracy.

Whereas this is a useful perspective of learning operators to mimic integral balance equations corresponding to PDE models of physics, for design of “systems-of-systems” one may use a graph to represent individual subsystems as nodes and their interactions via edges. This often corresponds to characterizing a control volume consisting of an entire subsystem, rather than an infinitesimal volume. In Hall et al. (2021), the authors develop graph informed neural networks (GINNs) which employ a probabilistic graph model to perform system scale design of supercapacitor dynamics and a Langmuir adsorption model. While the Bayesian network architecture is distinct from the GNN architectures discussed in the previous section, their work provides an example of how graph architectures may be used to replace computationally intensive bottlenecks in system dynamics. In other frameworks aligning more closely with GNNs, similar strategies may be used to learn and design, for example, system-level circuit design with constituent data-driven device physics (Gao et al., 2020) or heat transfer of industrial buildings (Drgoña et al., 2021).

### 3.4. Graph exterior calculus

Exterior calculus has been used as a tool for designing mimetic discretizations of PDEs, which preserve structure and properties from continuum solutions such as conservation, maximum principles, nonoscillatory solutions near material discontinuities, and preservation of invariants (Arnold et al., 2007). These methods exploit connections to the generalized Stokes theorem to construct a data-driven *deRham complex*—a commuting diagram (see Figure 13) which allows a unified analysis of discretizations spanning mixed finite elements, staggered finite difference schemes, and mimetic difference methods (Bochev and Hyman, 2006). Graphs admit a similar *combinatorial de Rham complex* based upon graph extensions of the familiar vector calculus  $\text{div}/\text{grad}/\text{curl}$ , which allows rigorous analysis of Hodge Laplacians generalizing the previously discussed graph Laplacian. We discuss several recent works which use this as a basis for machine learning.

Given a collection of  $k$ -cliques with associated scalar values, one may define a graph calculus. Denote the scalar value associated with a  $k$ -clique  $f_{i_1, \dots, i_k}$ , which is by convention skew-symmetric with respect to an odd permutation of subset indices. For example, a combinatorial gradient maps 1-cliques to 2-cliques

$$(\text{grad}\phi)_{ij} = \phi_j - \phi_i, \quad (21)$$

and a combinatorial curl maps 2-cliques to 3-cliques

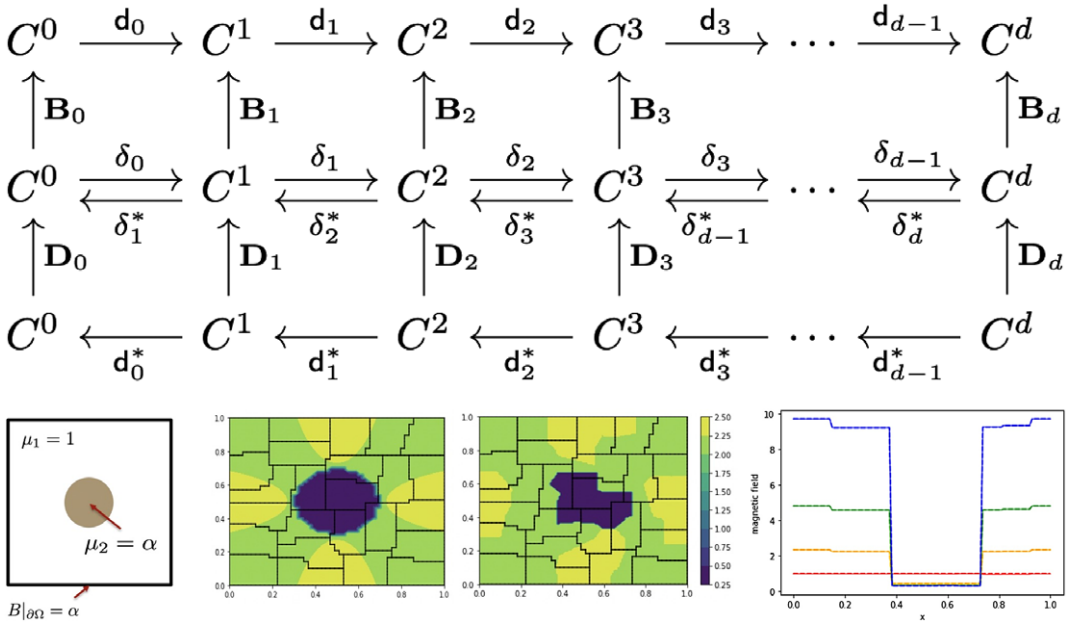
$$(\text{grad}\psi)_{ijk} = \psi_{ij} + \psi_{jk} + \psi_{ki}. \quad (22)$$

One may also introduce adjoint operators; for example the following combinatorial divergence is adjoint to the gradient

$$(\text{grad}^*\psi)_i = \sum_{j \sim i} \psi_{ij}. \quad (23)$$

Hence, we may easily interpret the aggregation step of a traditional GNN as a DNN composed with the combinatorial divergence.

These operators mimic algebraic structure from the familiar vector calculus. They form an *exact sequence* by guaranteeing  $\text{div}\text{curl} = \text{curl}\text{grad} = 0$ , and may be used to construct a graph Laplacian  $\delta = \text{grad}^* \cdot \text{grad}$  and higher-order Hodge Laplacians with well-characterized spectra, a Lax–Milgram theory, Poincare inequalities, and a Hodge decomposition, all of which function analogously to the familiar vector calculus setting.



**Figure 13.** Graph exterior calculus for physics-informed GNNs. Top: The graph calculus  $\text{div}/\text{grad}/\text{curl}$  ( $\delta_k$ ) and their adjoint ( $\delta_k^*$ ) may be augmented with machine learnable metric information ( $\mathbf{B}_k, \mathbf{D}_k$ ) to obtain a de Rham complex, which preserves algebraic structure related to stability and physical invariances. Bottom, left to right: To construct a surrogate for a heterogeneous magnetostatics inclusion problem parameterized by the jump in permittivity  $\alpha$  (a), we train a GNN to reproduce solution moments of a high-fidelity solution (b) on a coarse grid (c) while preserving conservation structure. A plot of the magnetic field across  $y = 0.5$  reveals the predicted solution (dashed) matches the data (solid) to machine precision for a range of jumps ( $\alpha = 1, 2, 4, 8$  in red/yellow/green/blue, respectively), demonstrating the scheme handles jumps in material properties for problems with nontrivial null-spaces, similar to mimetic PDE discretizations. Figure adapted from Trask et al. (2022).

In Seo and Liu (2019), the authors use this vector calculus to fit Hodge-Laplacian operators to diffusion processes on scattered data, introducing a PINN-style regularizer to fit a GNN to data while penalizing deviations from an assumed diffusion; see also Figure 2. For example, for a diffusion problem, penalizing an explicit Euler update of the form  $\mathbf{x}^{n+1} = \mathbf{x}^n + \alpha \text{grad}^* \circ \text{grad} \mathbf{x}^n$  promotes diffusion, where  $n$  denotes a given timestep. This process is limited to describing diffusion processes only.

In Trask et al. (2022), the graph calculus is used to learn boundary value problems associated with conservation balances, such as conservation of flux or circulation. For example, a conservation law consisting of a diffusive plus nonlinear flux may be learned

$$\begin{aligned} \mathbf{F} + \kappa \nabla \phi + N(\nabla \phi) = 0 &\rightarrow \mathbf{w}_1 + \text{grad} \mathbf{u}_0 + \mathcal{N}[\text{grad} \mathbf{u}_0; \theta] = 0 \\ \nabla \cdot \mathbf{F} = f &\quad \text{grad}^* \mathbf{w}_1 = \mathbf{f}_0. \end{aligned}$$

Here, subscripts 0 and 1 denote scalar values associated with 1- and 2-cliques respectively,  $N$  is a nonlinearity, and  $\mathcal{N}$  denotes a trainable neural network with parameters  $\theta$ . An additional parameterization of a Hodge star operator provides a means of learning metric information. This generalizes the GNNs discussed previously: the first equation corresponds to an update function on edges encoding fluxes, while the second corresponds to an aggregate step enforcing conservation of flux. One arrives at an operator of the form

$$\text{grad}^* \text{grad} \mathbf{u}_0 + \text{grad}^* \mathcal{N}[\mathbf{u}_0; \theta] = \mathbf{f}_0, \tag{24}$$

so the operator is a nonlinear perturbation of a Hodge-Laplacian. This can equivalently be given in variational form, where the solution  $\mathbf{u}_0$  satisfies for all  $v$

$$a(\mathbf{u}_0, v) + N_v[\mathbf{u}_0] = (f, v), \quad (25)$$

with  $a$  denoting an elliptic bilinear form. This can rigorously be shown to have a unique solution, and therefore an equality constrained optimization strategy can be applied to fit such models to data while exactly enforcing physics, independent of available data.

A key feature of this framework is the ability to handle nontrivial null-spaces which arise from involution constraints in electromagnetism. This yields an attractive framework for AI/ML-enabled data-driven modeling in microelectronics and magneto-hydrodynamics which mandate structure preservation (Bochev et al., 2003; Gao et al., 2020). We comment, however, that this notion of compatibility is appropriate for nonlinear elliptic problems. Other notions of compatibility are important for other problems; for example, geometric mechanics is important for dynamical systems while for flow problems it is necessary to consider Lie derivatives for advection operators (Marsden et al., 1998; Lee et al., 2021; Hernández et al., 2022).

In nonphysics based applications, the graph exterior calculus plays an important role in topological data analysis. For example, statistical ranking may be used to indicate preferences as a flow on a graph of options (Jiang et al., 2011), and the Betti numbers of a graph corresponding to the null-space of combinatorial Hodge Laplacians may be used to identify connected components and other geometric features in datasets (Carlsson, 2009).

### 3.5. Scalable GNNs

There are several GNN frameworks available which utilize modern distributed and heterogeneous computing architectures to train over large graph-based data structures. The usage of these frameworks depends on (a) whether to employ full-graph or mini-batch training, and (b) whether they are optimized for CPU-only or hardware accelerated computing architectures. Frameworks for distributed and full graph training for CPU and GPUs were rigorously developed in Jia et al. (2020), Hall et al. (2021), and Md et al. (2021), whereas for mini-batch based graph training, packages were developed and optimized for CPU-based clusters (Tripathy et al., 2020; Zhang et al., 2020b). It is necessary to note that full-graph training requires many epochs in comparison to mini-batch training and also converges to lower accuracy (Wilson and Martinez, 2003; LeCun et al., 2012; Keskar et al., 2016; Zheng et al., 2021).

When considering hardware acceleration of GNNs, it is imperative to consider whether GPUs will be beneficial to the training of GNNs. For GPU acceleration of GNN training to be effective, we require a high on-GPU compute intensity relative to transfer of data between the CPU and GPU. In general, GNNs require irregular memory access and computation due to their unstructured nature. This is in contrast to training of dense networks which consist of matrix vector product kernels over layers of regular size. Therefore, for the training of large graphs, we require efficient strategies to move data from CPU to GPU and mitigate latency.

A second challenge for implementing GNNs on distributed frameworks is how to effectively load balance across mini-batches. Ideally, we aim to simultaneously balance both the vertices/edges across processors and the amount of data to be transferred from memory to register. Load balancing is also required to synchronize the parameter update across all processes during backprop.

Natural graph networks (de Haan et al., 2020) treat complex and heterogeneous graphs with varying nodal feature sizes. To efficiently train heterogeneous graph architectures on hybrid CPU/GPUs based clusters, Zheng et al. (2021) proposed the DistDGLv2 framework, which is based on the deep graph library (DGL) (Wang et al., 2019a). In this framework, mini-batch training is performed on GPUs, whereas the graph structure along with the vertices and edges are stored in CPU memory. The authors have shown a  $20\times$  speedup for the GraphSage benchmark (Hamilton et al., 2017) and a  $36\times$  speedup on GAT benchmarks (Velickovic et al., 2018) using 64 GPUs on Amazon's EC2 cluster. Similar to Zeng et al. (2019), Bruss et al. (2019) propose a distributed framework for CPU-based clusters.

In a more general setting, scaling of unstructured graph algorithms is an important problem. Within the Department of Energy, the ExaGraph Co-design Center is a component of the Exascale Computing Project, which aims to develop broadly applicable graph kernels for exascale computation (1 exaflop =  $10^{18}$  flops) in applications spanning power grid, biology, chemistry, wind energy, and national security (Acer et al., 2021). The project aims to codesign algorithms, computing architecture, and hardware to build combinatorial kernels impacting general graph computation in addition to GNNs.

#### 4. Summary and Outlook

Unlike traditional machine learning, in scientific machine learning there are not big data available, and moreover, the few measurements available may be for auxiliary variables and not on the domain boundaries to provide proper boundary conditions. Physics-informed learning has emerged as a powerful approach in tackling such ill-defined problems and can use synergistically any available data from measurements together with data that are computed based on the physics of the problem and the corresponding governing equations. This can be done on random points as in PINNs or on the nodes of a graph as in GNNs.

PINNs and their extensions are relatively new methods that can tackle such problems, and they are especially effective for inverse problems, which may be challenging for traditional numerical methods. However, for realistic applications, new PINN implementations are needed that are amenable to multi-GPU computing. At the present time, a hybrid multi-GPU/multi-CPU paradigm based on domain decomposition seems to emerge as the most efficient solution but more work is required both on algorithms as well as on implementation, including mixed precision calculations that can significantly speed up PINNs.

GNNs offer effective alternative solutions, especially for unstructured data and very complex systems. Here, we have reviewed spectral methods, spatial methods and GATs that can all be used in the context of physics-informed learning, including discovering smart finite difference stencils and operators. We also reviewed PIGNs, where graph exterior calculus is used explicitly to construct the various differential operators on graphs, and preliminary results for simulation physical systems are very promising. In contrast to PINNs, PIGNs may impose inductive biases corresponding to exact enforcement of integral balances rather than differential operators. There are already several parallel frameworks for scaling up GNNs but the new effort should be directed toward scaling physics-informed GNNs.

Finally, we want to comment on the connection between graphs and PDEs. It seems that the scientific machine learning community is importing ideas from graph theory and algorithms to solve PDEs, which is largely true and could be exploited even further. For example, graph embedding methods offer drastic dimensionality reduction for complex data both for static and dynamic graphs, see Xu (2021), hence, graph embedding could also be exploited in the context of complex dynamical systems by the PDE community. However, recently there has been an effort to design new GNNs based on PDE theory, see, for example, the so-called “GRAND” scheme of Chamberlain et al. (2021), where standard PDE tools are used to understand existing GNN architectures and subsequently develop a broader class of GNNs that are stable even for very deep networks. Specifically, the authors of Chamberlain et al. (2021) have shown that many popular GNN architectures can be derived from a universal mathematical framework by choosing different forms of diffusion equations and time-discretization schemes. Hence, what we see emerging is a reciprocal “pull–push” relationship between graph theory and algorithms on one side and PDE theory and algorithms on the other side that will be beneficial in greatly accelerating the emerging field of scientific machine learning.

**Data Availability Statement.** Tensorflow based Python codes and data to reproduce Figure 12 can be obtained at <https://github.com/rgp62/gmls-nets>. However, the same implementation using PyTorch based python codes can be found at <https://github.com/atzberg/gmls-nets>. A detail description and implementation of domain decomposition algorithms shown in Figure 7 can be found in paper <https://www.sciencedirect.com/science/article/pii/S0021999121005787>. A detail description and implementation detail to reproduce results in Figure 5 can be found in <https://arxiv.org/pdf/2111.02801.pdf>.

**Author Contributions.** All authors wrote and edited the paper. All authors approved the final submitted draft.

**Funding Statement.** This research was supported by the DOE project PhILMs (no. DE-SC0019453) and the OSD/AFOSR MURI grant FA9550-20-1-0358. N.T.'s research was supported by the Department of Energy early career research program. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

**Competing Interests.** The authors declare no competing interests exist.

## References

- Acer S, Azad A, Boman EG, Buluç A, Devine KD, Ferdous S, Gawande N, Ghosh S, Halappanavar M, Kalyanaraman A, Khan A, Minutoli M, Pothen A, Rajamanickam S, Selvitopi O, Tallent NR and Tumeo A (2021) Exagraph: Graph and combinatorial methods for enabling exascale applications. *The International Journal of High Performance Computing Applications* 35(6), 553–571.
- Alnæs M, Blechta J, Hake J, Johansson A, Kehlet B, Logg A, Richardson C, Ring J, Rognes ME and Wells GN (2015) The fenics project version 1.5. *Archive of Numerical Software* 3(100), 9–23.
- Arnold DN, Bochev PB, Lehoucq RB, Nicolaides RA and Shashkov M (2007) *Compatible Spatial Discretizations*, Vol. 142. New York: Springer Science & Business Media.
- Atwood J and Towsley D (2016) Diffusion-convolutional neural networks. In *Advances in Neural Information Processing Systems*, Vol. 29. New York: Curran Associates.
- Bar-Sinai Y, Hoyer S, Hickey J and Brenner MP (2019) Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences* 116(31), 15344–15349.
- Bochev PB, Hu JJ, Robinson AC and Tuminaro RS (2003) Towards robust 3d z-pinch simulations: Discretization and fast solvers for magnetic diffusion in heterogeneous conductors. *Electronic Transactions on Numerical Analysis* 15, 186–210.
- Bochev PB and Hyman JM (2006) Principles of mimetic discretizations of differential operators. In *Compatible Spatial Discretizations*. New York: Springer, pp. 89–119.
- Bruna J, Zaremba W, Szlam A and LeCun Y (2013) Spectral networks and locally connected networks on graphs. arXiv preprint arXiv:1312.6203.
- Bruss CB, Khazane A, Rider J, Serpe R, Nagrecha S and Hines KE (2019) Graph embeddings at scale. arXiv preprint arXiv:1907.01705.
- Bubenik P (2015) Statistical topological data analysis using persistence landscapes. *Journal of Machine Learning Research* 16(1), 77–102.
- Carlsson G (2009) Topology and data. *Bulletin of the American Mathematical Society* 46(2), 255–308.
- Chamberlain B, Rowbottom J, Gorinova MI, Bronstein M, Webb S and Rossi E (2021) Grand: Graph neural diffusion. In *International Conference on Machine Learning*. PMLR, pp. 1407–1418. (Online)
- Chen Y, Lu L, Karniadakis GE and Dal Negro L (2020) Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Optics Express* 28(8), 11618–11633.
- Chen RTQ, Rubanova Y, Bettencourt J and Duvenaud DK (2018) Neural ordinary differential equations. In Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N and Garnett R (eds), *Advances in Neural Information Processing Systems*, Vol. 31. New York: Curran Associates.
- Chung FR and Graham FC (1997) *Spectral Graph Theory*, Vol. 92. Providence, RI: American Mathematical Society.
- Daghghi S, Meisburger N, Zhao M and Shrivastava A (2021) Accelerating slide deep learning on modern cpus: Vectorization, quantizations, memory optimizations, and more. *Proceedings of Machine Learning and Systems* 3, 156–166.
- Davies A, Veličković P, Buesing L, Blackwell S, Zheng D, Tomašev N, Tanburn R, Battaglia P, Blundell C, Juhász A, Lackenby M, Williamson G, Hassabis D and Kohli P (2021) Advancing mathematics by guiding human intuition with AI. *Nature* 600(7887), 70–74.
- de Haan P, Cohen TS and Welling M (2020) Natural graph networks. *Advances in Neural Information Processing Systems* 33, 3636–3646.
- Defferrard M, Bresson X and Vandergheynst P (2016) Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, Vol. 29. New York: ACM.
- Drogoňa J, Tuor AR, Chandan V and Vrabie DL (2021) Physics-constrained deep learning of multi-zone building thermal dynamics. *Energy and Buildings* 243, 110992.
- Gao X, Huang A, Trask N and Reza S (2020) Physics-informed graph neural network for circuit compact model development. In *2020 International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*. Kobe: IEEE, pp. 359–362.
- Gao H, Zahr MJ and Wang J-X (2022) Physics-informed graph neural galerkin networks: A unified framework for solving pde-governed forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering* 390, 114502.

- Gilmer J, Schoenholz SS, Riley PF, Vinyals O and Dahl GE** (2017) Neural message passing for quantum chemistry. In *International Conference on Machine Learning*. PMLR, pp. 1263–1272. (Sydney, Australia)
- Goswami K, Sharma A, Pruthi M and Gupta R** (2021) Study of drug assimilation in human system using physics informed neural networks. arXiv preprint [arXiv:2110.05531](https://arxiv.org/abs/2110.05531).
- Goyal P, Dollár P, Girshick R, Noordhuis P, Wesolowski L, Kyrola A, Tulloch A, Jia Y and He K** (2017) Accurate, large minibatch sgd: Training imagenet in 1 hour. arXiv preprint [arXiv:1706.02677](https://arxiv.org/abs/1706.02677).
- Gustafsson B, Kreiss H-O and Olinger J** (1995) *Time Dependent Problems and Difference Methods*, Vol. 24. New York: John Wiley & Sons.
- Hall EJ, Taverniers S, Katsoulakis MA and Tartakovsky DM** (2021) GINNs: Graph-informed neural networks for multiscale physics. *Journal of Computational Physics* 433, 110192.
- Hamilton W, Ying Z and Leskovec J** (2017) Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems* 30, 1025–1035.
- Harlow FH and Welch JE** (1965) Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface. *The Physics of Fluids* 8(12), 2182–2189.
- Hennigh O, Narasimhan S, Nabian MA, Subramaniam A, Tangsali K, Rietmann M, Ferrandis JA, Byeon W, Fang Z and Choudhry S** (2020) Nvidia simnet™: An ai-accelerated multi-physics simulation framework. arXiv preprint [arXiv:2012.07938](https://arxiv.org/abs/2012.07938).
- Hernández Q, Badas A, Chinesta F and Cueto E** (2022) Thermodynamics-informed graph neural networks. arXiv preprint [arXiv:2203.01874](https://arxiv.org/abs/2203.01874).
- Iakovlev V, Heinonen M and Lähdesmäki H** (2020) Learning continuous-time pdes from sparse data with graph neural networks. arXiv preprint [arXiv:2006.08956](https://arxiv.org/abs/2006.08956).
- Jagtap AD and Karniadakis GE** (2020) Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations. *Communications in Computational Physics* 28(5), 2002–2041.
- Jagtap AD, Kharazmi E and Karniadakis GE** (2020) Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering* 365, 113028.
- Jia Z, Lin S, Gao M, Zaharia M and Aiken A** (2020) Improving the accuracy, scalability, and performance of graph neural networks with ROC. *Proceedings of Machine Learning and Systems* 2, 187–198.
- Jiang X, Lim L-H, Yao Y and Ye Y** (2011) Statistical ranking and combinatorial Hodge theory. *Mathematical Programming* 127(1), 203–244.
- Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, Bates R, Ždek A, Potapenko A, Bridgland A, Meyer C, Kohl SAA, Ballard AJ, Cowie A, Romera-Paredes B, Nikolov S, Jain R, Adler J, Back T, Petersen S, Reiman D, Clancy E, Zielinski M, Steinegger M, Pacholska M, Berghammer T, Bodenstein S, Silver D, Vinyals O, Senior AW, Kavukcuoglu K, Kohli P and Hassabis D** (2021) Highly accurate protein structure prediction with alphafold. *Nature* 596(7873), 583–589.
- Karniadakis GE, Karniadakis G and Sherwin S** (2005) *Spectral/HP Element Methods for Computational Fluid Dynamics*. Oxford: Oxford University Press on Demand.
- Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S and Yang L** (2021) Physics-informed machine learning. *Nature Reviews Physics* 3(6), 422–440.
- Karypis G and Kumar V** (2009) MeTis: Unstructured graph partitioning and sparse matrix ordering system, Version 4.0.
- Kashyap S, Kumar S, Agarwal V, Misra DP, Phadke SR and Kapoor A** (2018) Protein protein interaction network analysis of differentially expressed genes to understand involved biological processes in coronary artery disease and its different severity. *Gene Reports* 12, 50–60.
- Keskar NS, Mudigere D, Nocedal J, Smelyanskiy M and Tang PTP** (2016) On large-batch training for deep learning: Generalization gap and sharp minima. arXiv preprint [arXiv:1609.04836](https://arxiv.org/abs/1609.04836).
- Kharazmi E, Zhang Z and Karniadakis GE** (2021) Hp-VPINNs: Variational physics-informed neural networks with domain decomposition. *Computer Methods in Applied Mechanics and Engineering* 374, 113547.
- Kipf TN and Welling M** (2016) Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907).
- Kuberry P, Bosler P and Trask N** (2019) Compadre toolkit.
- Kumar Y and Chakraborty S** (2021) Grade: A graph based data-driven solver for time-dependent nonlinear partial differential equations. arXiv preprint [arXiv:2108.10639](https://arxiv.org/abs/2108.10639).
- Lagari PL, Tsoukalas LH, Safarkhani S and Lagaris IE** (2020) Systematic construction of neural forms for solving partial differential equations inside rectangular domains, subject to initial, boundary and interface conditions. *International Journal on Artificial Intelligence Tools* 29(5), 2050009.
- Lagaris IE, Likas A and Fotiadis DI** (1998) Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks* 9(5), 987–1000.
- LeCun YA, Bottou L, Orr GB and Müller K-R** (2012) Efficient backprop. In *Neural Networks: Tricks of the Trade*. Berlin: Springer, pp. 9–48.
- Lee K, Trask N and Stinis P** (2021) Machine learning structure preserving brackets for forecasting irreversible processes. *Advances in Neural Information Processing Systems* 34, 5696–5707.



- Liu K, Sun X, Jia L, Ma J, Xing H, Wu J, Gao H, Sun Y, Boulnois F and Fan J (2019) Chemi-net: A molecular graph convolutional network for accurate drug property prediction. *International Journal of Molecular Sciences* 20(14), 3389.
- Lu L, Meng X, Mao Z and Karniadakis GE (2021a) DeepXDE: A deep learning library for solving differential equations. *SIAM Review* 63(1), 208–228.
- Lu L, Pestourie R, Yao W, Wang Z, Verdugo F and Johnson SG (2021b). Physics-informed neural networks with hard constraints for inverse design. arXiv preprint arXiv:2102.04626.
- Mao Z, Jagtap AD and Karniadakis GE (2020) Physics-informed neural networks for high-speed flows. *Computer Methods in Applied Mechanics and Engineering* 360, 112789.
- Marsden JE, Patrick GW and Shkoller S (1998) Multisymplectic geometry, variational integrators, and nonlinear pdes. *Communications in Mathematical Physics* 199(2), 351–395.
- McClenny L and Braga-Neto U (2020) Self-adaptive physics-informed neural networks using a soft attention mechanism. arXiv preprint arXiv:2009.04544.
- Mad V, Misra S, Ma G, Mohanty R, Georganas E, Heinecke A, Kalamkar D, Ahmed NK and Avancha S (2021) DistGNN: Scalable distributed training for large-scale graph neural networks. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. New York: ACM, pp. 1–14.
- Mirzaei D, Schaback R and Dehghan M (2012) On generalized moving least squares and diffuse derivatives. *IMA Journal of Numerical Analysis* 32(3), 983–1000.
- Monti F, Boscaini D, Masci J, Rodola E, Svoboda J and Bronstein MM (2017) Geometric deep learning on graphs and manifolds using mixture model CNNs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. San Juan, PR: IEEE, pp. 5115–5124.
- Niepert M, Ahmed M and Kutzkov K (2016) Learning convolutional neural networks for graphs. In *International Conference on Machine Learning*. PMLR, pp. 2014–2023. (New York, NY, USA)
- Pang G, Lu L and Karniadakis GE (2019) fPINNs: Fractional physics-informed neural networks. *SIAM Journal on Scientific Computing* 41(4), A2603–A2626.
- Pareja A, Domeniconi G, Chen J, Ma T, Suzumura T, Kanezashi H, Kaler T, Schardl T and Leiserson C (2020) EvolveGCN: Evolving graph convolutional networks for dynamic graphs. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 5363–5370.
- Pfaff T, Fortunato M, Sanchez-Gonzalez A and Battaglia PW (2020) Learning mesh-based simulation with graph networks. arXiv preprint arXiv:2010.03409.
- Poli M, Massaroli S, Park J, Yamashita A, Asama H and Park J (2019) Graph neural ordinary differential equations. arXiv preprint arXiv:1911.07532.
- Raissi M, Perdikaris P and Karniadakis GE (2018) Numerical Gaussian processes for time-dependent and nonlinear partial differential equations. *SIAM Journal on Scientific Computing* 40(1), A172–A198.
- Raissi M, Perdikaris P and Karniadakis GE (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378, 686–707.
- Raissi M, Yazdani A and Karniadakis GE (2020) Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science* 367(6481), 1026–1030.
- Rajbhandari S, Rasley J, Ruwase O and He Y (2020) Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. Atlanta, GA: IEEE, pp. 1–16.
- Rasley J, Rajbhandari S, Ruwase O and He Y (2020) Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. New York: ACM, pp. 3505–3506.
- Rosenthal G, Váša F, Griffa A, Hagmann P, Amico E, Goñi J, Avidan G and Sporns O (2018) Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes. *Nature Communications* 9(1), 1–12.
- Rossi E, Chamberlain B, Frasca F, Eynard D, Monti F and Bronstein M (2020) Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637.
- Rusu RB and Cousins S (2011) 3D is here: Point cloud library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*. Shanghai: IEEE.
- Sahli Costabal F, Yang Y, Perdikaris P, Hurtado DE and Kuhl E (2020) Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics* 8, 42.
- Seo S and Liu Y (2019) Differentiable physics-informed graph networks. arXiv preprint arXiv:1902.02950.
- Sergeev A and Del Balso M (2018) Horovod: fast and easy distributed deep learning in tensorflow. arXiv preprint arXiv:1802.05799.
- Shukla K, Di Leoni PC, Blackshire J, Sparkman D and Karniadakis GE (2020) Physics-informed neural network for ultrasound nondestructive quantification of surface breaking cracks. *Journal of Nondestructive Evaluation* 39(3), 1–20.
- Shukla K, Jagtap AD, Blackshire JL, Sparkman D and Karniadakis GE (2021a) A physics-informed neural network for quantifying the microstructural properties of polycrystalline nickel using ultrasound data: A promising approach for solving inverse problems. *IEEE Signal Processing Magazine* 39(1), 68–77.

- Shukla K, Jagtap AD and Karniadakis GE** (2021b) Parallel physics-informed neural networks via domain decomposition. *Journal of Computational Physics* 447, 110683.
- Trask N, Huang A and Hu X** (2022) Enforcing exact physics in scientific machine learning: A data-driven exterior calculus on graphs. *Journal of Computational Physics* 456, 110969.
- Trask N, Patel RG, Gross BJ and Atzberger PJ** (2019) Grmls-nets: A framework for learning from unstructured data. arXiv preprint arXiv:1909.05371.
- Tripathy A, Yelick K and Buluç A** (2020) Reducing communication in graph neural network training. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. Atlanta, GA: IEEE, pp. 1–14.
- Trivedi R, Farajtabar M, Biswal P and Zha H** (2019) Dyrep: Learning representations over dynamic graphs. In *International Conference on Learning Representations*, pp.1–25. (New Orleans, Louisiana, USA)
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł and Polosukhin I** (2017) Attention is all you need. In *Advances in Neural Information Processing Systems*, Vol. 30. New York: Curran Associates.
- Velickovic P, Cucurull G, Casanova A and Romero A** (2018) Pietro lio, and yoshua bengio. *Graph attention networks*.
- Velicković P, Cucurull G, Casanova A, Romero A, Liò P and Bengio Y** (2017) Graph attention networks. arXiv preprint arXiv:1710.10903.
- Velicković P, Cucurull G, Casanova A, Romero A, Liò P and Bengio Y** (2018) Graph attention networks.
- Waheed U, Haghghat E, Alkhalifah T, Song C and Hao Q** (2020) Eikonal solution using physics-informed neural networks. In *82nd EAGE Annual Conference & Exhibition*, Vol. 2020. European Association of Geoscientists & Engineers, pp. 1–5. (Online)
- Wang Y, Feng C, Chen L, Yin H, Guo C and Chu Y** (2019b) User identity linkage across social networks via linked heterogeneous network embedding. *World Wide Web* 22(6), 2611–2632.
- Wang M, Zheng D, Ye Z, Gan Q, Li M, Song X, Zhou J, Ma C, Yu L, Gai Y, Xiao T, He T, Karypis G, Li J and Zhang Z** (2019a). Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv preprint arXiv:1909.01315.
- Wasserman L** (2018) Topological data analysis. *Annual Review of Statistics and Its Application* 5, 501–532.
- West JD, Wesley-Smith I and Bergstrom CT** (2016) A recommendation system based on hierarchical clustering of an article-level citation network. *IEEE Transactions on Big Data* 2(2), 113–123.
- Wilson DR and Martinez TR** (2003) The general inefficiency of batch training for gradient descent learning. *Neural Networks* 16 (10), 1429–1451.
- Xu M** (2021) Understanding graph embedding methods and their applications. *SIAM Review* 63(4), 825–853.
- Xu M, Sanz DL, Garcés P, Maestu F, Li Q and Pantazis D** (2021a) A graph Gaussian embedding method for predicting Alzheimer’s disease progression with meg brain networks. *IEEE Transactions on Biomedical Engineering* 68(5), 1579–1588.
- Xu M, Singh AV and Karniadakis GE** (2021b). Dyng2g: An efficient stochastic graph embedding method for temporal graphs. arXiv preprint arXiv:2109.13441.
- Xu M, Wang Z, Zhang H, Pantazis D, Wang H and Li Q** (2020b) A new graph Gaussian embedding method for analyzing the effects of cognitive training. *PLoS Computational Biology* 16(9), e1008186.
- Xu K, Zhu W and Darve E** (2020a). Distributed machine learning for computational engineering using MPI. arXiv preprint arXiv:2011.01349.
- Yazdani A, Lu L, Raissi M and Karniadakis GE** (2020) Systems biology informed deep learning for inferring parameters and hidden dynamics. *PLoS Computational Biology* 16(11), e1007575.
- Yee K** (1966) Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation* 14(3), 302–307.
- Yin M, Zheng X, Humphrey JD and Karniadakis GE** (2021) Non-invasive inference of thrombus material properties with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering* 375, 113603.
- Yu J, Lu L, Meng X and Karniadakis GE** (2021) Gradient-enhanced physics-informed neural networks for forward and inverse pde problems. arXiv preprint arXiv:2111.02801.
- Zeng H, Zhou H, Srivastava A, Kannan R and Prasanna V** (2019) Accurate, efficient and scalable graph embedding. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. Rio de Janeiro: IEEE, pp. 462–471.
- Zhang D, Guo L and Karniadakis GE** (2020a) Learning in modal space: Solving time-dependent stochastic PDEs using physics-informed neural networks. *SIAM Journal on Scientific Computing* 42(2), A639–A665.
- Zhang D, Huang X, Liu Z, Hu Z, Song X, Ge Z, Zhang Z, Wang L, Zhou J, Shuang Y and Qi Y** (2020b). Agl: A scalable system for industrial-purpose graph machine learning. arXiv preprint arXiv:2003.02454.
- Zhang D, Lu L, Guo L and Karniadakis GE** (2019) Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems. *Journal of Computational Physics* 397, 108850.
- Zheng D, Song X, Yang C, LaSalle D, Su Q, Wang M, Ma C and Karypis G** (2021) Distributed hybrid CPU and GPU training for graph neural networks on billion-scale graphs. arXiv preprint arXiv:2112.15345.
- Zhou Y and Li H** (2019) Asset diversification and systemic risk in the financial system. *Journal of Economic Interaction and Coordination* 14(2), 247–272.

**Cite this article:** Shukla K, Xu M, Trask N and Karniadakis GE (2022). Scalable algorithms for physics-informed neural and graph networks. *Data-Centric Engineering*, 3, e24. doi:10.1017/dce.2022.24