

On the under-reaching phenomenon in message passing neural PDE solvers: Revisiting the CFL condition

Lucas Tesán^{a,1}, Mikel M. Iparraguirre^{a,1}, David González^a, Pedro Martins^{a,b},
Elías Cueto^{a,*}

^a ESI Group-UZ Chair of the National Strategy on Artificial Intelligence Aragon Institute of Engineering Research, Universidad de Zaragoza, Zaragoza, 50018, Spain

^b Aragonese Foundation for Research and Development (ARAID), Zaragoza, 50018, Spain

ARTICLE INFO

Keywords:

Neural PDE solvers
Message passing
Graph neural networks
Partial differential equations

ABSTRACT

This paper proposes sharp lower bounds for the number of message passing iterations required in graph neural networks (GNNs) when solving partial differential equations (PDE). This significantly reduces the need for exhaustive hyperparameter tuning. Bounds are derived for the three fundamental classes of PDEs (hyperbolic, parabolic and elliptic) by relating the physical characteristics of the problem in question to the message-passing requirement of GNNs. In particular, we investigate the relationship between the physical constants of the equations governing the problem, the spatial and temporal discretisation and the message passing mechanisms in GNNs.

When the number of message passing iterations is below these proposed limits, information does not propagate efficiently through the network, resulting in poor solutions, even for deep GNN architectures. In contrast, when the suggested lower bound is satisfied, the GNN parameterisation allows the model to accurately capture the underlying phenomenology, resulting in solvers of adequate accuracy.

Examples are provided for four different examples of equations that show the sharpness of the proposed lower bounds.

1. Introduction

Machine learning-based simulators have gained significant attention due to their ability to approximate complex physical phenomena at a fraction of the computational cost of traditional numerical solvers [1–3]. Once trained, these learned simulators can generate rapid predictions, making them particularly useful for tasks requiring real-time inference [4], inverse problem-solving [5] or chaotic-spatial dependent systems [6].

Traditional physics-based solvers rely on mechanistic models, typically expressed as partial or ordinary differential equations (PDE or ODE, respectively) and discretized into algebraic problems suitable for computation. Their accuracy hinges on the fidelity of the physical model and the discretization used as solver, such as finite differences [7], finite elements [8], or finite volumes [9]. These traditional approaches often demand substantial computational resources due to their limited ability to capture complex phenomena efficiently. To achieve accurate solutions, they require very fine spatial discretization and high temporal resolution, which significantly increases the number of iterations, and consequently their computational cost.

* Corresponding author.

E-mail address: ecueto@unizar.es (E. Cueto).

¹ Equal contributors

Neural Networks (NNs) have emerged as powerful simulation tools, enabling the development of solvers through two main paradigms. These are: physics-informed learning, where the residuals of the governing PDEs are minimized directly, possibly without the need for training data [10]; and data-driven learning, where models are trained on simulation or experimental data and subsequently used to predict solutions without explicitly solving the underlying equations at inference time [11]. In the data-driven approach, performance relies heavily on the quality and diversity of the training data, as well as the expressive capacity of the chosen neural architecture [2,12,13]. While training can be computationally demanding, inference is typically very efficient [14].

Physics-Informed Neural Networks (PINNs) emerge as the widely adopted alternative for scenarios where the governing equations are known [10,15,16]. This approach solves partial differential equations by embedding physical laws into the loss function, enabling solution approximation without labeled data or discretization. They are mesh-free and can directly optimize PDEs thanks to the differentiability of neural networks. However, they share their limited flexibility with traditional solvers—any modification to initial or boundary conditions requires re-optimizing the problem from scratch, and they generally lack the ability to generalize to unseen problems [17,18]. Despite the training challenges of this approach, PINNs can minimize the PDE at a finite number of collocation points, while enabling high-resolution simulations during inference by evaluating the model over a dense set of points. This represents a significant advantage over traditional solvers.

On the other hand, models based on dimensionality reduction techniques can condense the complex hidden dynamics within data into lower-dimensional representations. This facilitates the training of neural simulators and establishes a hybrid synergy between human expertise and artificial intelligence. However, they still suffer from the same fundamental limitations: an inability to locally interpret the underlying physics of the problem, and consequently, a lack of extrapolation capability to new geometries or domains [19,20]. Nevertheless, promising new research directions are emerging that enable the incorporation of domain-specific information, offering equally interesting alternatives for overcoming these challenges [21].

GNNs have emerged in the past years as one of the most powerful data-driven approaches for modeling physical phenomena. Similar to how convolutional networks [22] perform local optimization by acting as spatial operators capable of learning complex patterns [23,24], GNNs leverage their architecture to capture both local and global interactions within graph-based representations, such as mesh discretizations used in physical simulations. By establishing mathematical relationships between spatial discretization and the underlying physical processes, they can effectively represent nearly any type of phenomenology, regardless of complexity [25]. In pursuit of explainability and theoretical guarantees, an important extension of these models is Physics-Informed GNNs (PIGNNs), which enhance their capabilities by embedding physical laws such as conservation principles, symmetries, or governing partial differential equations directly into their architectures or loss functions [5,26,27]. This integration reflects the broader paradigm of hybrid artificial intelligence, where data-driven models are guided by domain knowledge. Alternatively, other hybrid approaches adopt a thermodynamic perspective, incorporating constraints inspired by thermodynamic principles [28]. By framing the phenomenon in a thermodynamic setting, these models offer a flexible framework for modeling open, non-equilibrium, or multiscale systems, while increasing interpretability [29–32,32–35].

Among the many factors that influence the success of deep learning models, hyperparameter tuning plays a particularly critical role. It directly affects the model's ability to learn effectively and generalize well. Poor choices in hyperparameter settings can lead to suboptimal training outcomes and degraded performance. This tuning process—an essential component of model selection—typically demands significant computational resources and time. Common hyperparameters include learning rate, number of hidden units, and batch size, which are often selected through trial-and-error strategies. Guided search methods [36] can help identify optimal configurations more efficiently. Nevertheless, hyperparameter tuning remains a challenge. In GNNs, a particularly critical hyperparameter is the number of message-passing iterations (MPI), which directly affects the model's ability to propagate information across the graph. While this has traditionally been treated as a standard tunable parameter, we propose establishing a lower bound based on the underlying system dynamics to reduce the hyperparameter search space and ensure sufficient information flow for accurate prediction [37].

To estimate the lower bound on MPI, we distinguish between two regimes. Hyperbolic systems (e.g., transport or wave equations [38]), are characterised by the presence of a wave which, in one form or another, propagates through the domain, so the proposed lower bound reflects the speed at which this wave travels. In parabolic and elliptic systems, in contrast, the described phenomena propagate with infinite speed, affecting the whole domain instantaneously—making the lower bound a function of the system's geometry and mesh size. In particular, for systems governed by elliptic equations, it is well known that they predict, contrary to experience, an infinite transmission rate. This phenomenon is called the infinite propagation speed paradox [39]. These principles apply to both standard message-passing networks [40] and more advanced models like MeshGraphNets [2].

By aligning hyperparameter selection with the underlying physics of the system, our approach ensures that learned graph simulations adhere to the principles of numerical stability while maintaining computational efficiency. This bridges a critical gap between data-driven modeling and the rigorous demands of scientific computing.

The main contribution of this article, therefore, is the proposition of a lower bound for the number of message steps as a function of the character of the equation governing the physical phenomenon. As will be seen below, the proposed bound for hyperbolic equations is strongly reminiscent of the famous Courant-Friedrichs-Lewy condition for finite differences applied to this type of equations. In problems governed by elliptic and parabolic equations, on the contrary, the proposed bound must ensure that the information about the phenomena taking place reaches all corners of the domain. It will be seen, therefore, how the bounds are very different in one case and the other, but how the predictions show great accuracy if they are fulfilled.

The layout of the paper is as follows. In Section 2 we review the basics of message-passing graph neural networks for the solution of PDEs. In Section 3 we review the problems arising from the insufficient number of message steps, a phenomenon known as under-reaching in the literature [41] and develop lower bounds to ensure that this problem does not occur. In Section 4 numerical tests




Fig. 1. Illustration of a GNN with an encoder-processor-decoder approach and $h = 1, \dots, M$ passes—we depict the case $M = 2$ —. We focus on the orange node of the graph, whose evolution from time step n to $n + 1$ is of primary interest. This evolution is driven by the perturbations originating from the red (located at a distance of two hops, $h = 2$) and blue nodes (located at a distance of six hops, $h = 6$). Therefore, it is crucial for the processor, through message passing iterations, to propagate the information from the red and blue nodes to the orange node. If the number of iterations is insufficient (less than the h -distance between nodes), the orange node's future state will not be estimated based on the actual nodes that influence it, but rather inferred from an incomplete or indirect representation. This can lead to suboptimal or incorrect predictions. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

are described which show the validity of the proposed limits. We follow in [Section 4.1](#) with the description of the examples used for analysis of the dynamics of message passing, while in [Section 4.2](#) we show the results obtained. The paper is closed with [Section 5](#) devoted to deriving the conclusions of this work.

2. Message-passing graph neural networks

Our model architecture, with which we have performed the experiments to be shown in the following sections, consists of a one-step predictor, except for elliptic problems, where the time variable is not present. The prediction at each time step becomes the current state used as input for the next time step. This setup makes the model easy to train, as it only requires training on one-step predictions, although more complex architectures exist, involving multi-step temporal integration, recurrent or memory-augmented schemes, etc. [\[42\]](#). Additionally, it is flexible during inference, allowing us to simulate long time periods (n -step rollouts) without explicitly requiring time as a parameter.

This one-step predictor is a GNN that follows the encoder-processor-decoder architecture, which has demonstrated significant potential in modeling mesh-based and physical problems, see [Fig. 1](#). For the theoretical developments and the first numerical examples we wanted to keep the network architecture as simple as possible, in order to eliminate the possible influence of other factors in the analysis. In [Section 4.4](#), however, a more complex architecture is used, based on MeshGraphNets [\[2\]](#), the details of which will be explained in the section itself. For the time being, let us stick to the simplest possible architecture, which does not include any encoding on the edges of the graph and only stores and processes nodal variables:

- Every node in the graph shares a multi-layer perceptron that encodes the input $u_n = u(t = n\Delta t) \in \mathbb{R}^{n_{\text{dof}}}$ into a high-dimensional latent space $\xi_n \in \mathbb{R}^D$, with $D > n_{\text{dof}}$, in general. Here, n_{dof} represents the number of state variables at each node.
- At each time increment, the processor employs a message-passing algorithm to gather information from nodes within an h -hop neighborhood. This information is then integrated to update the node and possibly edge representations. Through this process, the Graph Neural Network (GNN) iteratively refines the latent representation of the graph G .
- The decoder transforms the nodal representations from the latent space ξ_{n+1} back into the physical space, as the prediction Δu_n per node.

The GNN is trained to predict one-step system evolutions, Δu_n , which depend solely on the input u_n and the learned network parameters θ ,

$$u_{n+1} = u_n + \Delta u_n, \quad (1)$$

where $\Delta u_n = \text{GNN}(u_n, \theta)$.

The processor is the core of this type of neural solver. In essence, at each time step n —which we omit for clarity—, it performs the learning of $h = 1, \dots, M$ message passing processes of the form

- Learn the message along edge ij :

$$m_{ij}^h = \phi(\xi_i^h, \xi_j^h, \theta), \quad (2)$$

- Update of node i :

$$\xi_i^{h+1} = \varphi\left(\xi_i^h, \sum_{j \in \mathcal{N}(i)} m_{ij}^h, \theta\right), \quad (3)$$

where $\mathcal{N}(i)$ represents the neighborhood of node i , i.e., the set of nodes j connected by the graph to node i .

The ability of GNNs to understand and process global information from the graph G , despite being a nodal level architecture, is entirely dependent on the transmission of the message passing algorithm. If the number of messages M is insufficient, the nodes are in practice disconnected, which prevents effective information propagation. This limitation can compromise the model's capacity

to capture regional or global dependencies necessary for accurately solving our problem. Even in a connected graph, inadequate message passing prevents effective communication between nodes, leading to suboptimal performance.

Previous work has shown that solvers based on message passing mechanisms have the ability to reproduce classical PDE solving schemes such as finite difference, finite volume or WENO-type schemes (schemes with $M = 1$, $M = 2$ and $M = 3$, respectively), see [42]. However, the authors of the mentioned paper use a higher number of passes, six hops. Additionally, at each hop they consider a larger number of neighbors for each node than would be indicated by the connectivity of the graph, using up to twenty neighbors per node. The result, if translated into a finite difference scheme, would result in a much higher order scheme and a highly populated stiffness matrix.

Almost all existing literature practices many rounds of message passing [6]. Pfaff et al state that increasing the number of graph net blocks (message passing steps) typically improves performance, but it incurs a higher computational cost, showing that additional passes reduce prediction error; they often employ fifteen hops as a compromise between computational cost and accuracy [2]. Yet, one cannot disguise that the observed improvement largely stems from the increased expressivity of deeper architectures rather than the message passing behavior itself. This phenomenon is now referred to as under-reaching and describes a situation where a node cannot detect or recognize nodes located beyond its specified number of layers [41]. It should be noted here that our architecture uses shared weights, so a higher number of passes does not imply greater expressiveness of the network in general.

If this were literally the case, the number of passes necessary to achieve the maximum accuracy of the method would be the number of passes necessary for the messages to get from one end of the mesh to the other (we do not consider the possibility of using extended neighborhoods as in [42], to alleviate this process, in order to simplify the hypotheses as much as possible). As will be seen in the following sections, this is not always the case. It is not always necessary for messages to reach from one end of the graph to the other, and in any case, this depends on the type of PDE that governs the physical phenomenon under study.

3. The phenomenon of under-reaching seen from a physical point of view

In this section, we investigate the impact of the number of message-passing iterations in GNNs across the three categories of partial differential equations: hyperbolic (Section 3.1), parabolic (Section 3.2), and elliptic (Section 3.3). Understanding the distinct characteristics of the physical phenomena described by each PDE family is crucial, as it reveals how information should be propagated across the graph. To this end, we analyze representative benchmark problems for each class: the wave equation (hyperbolic), the Fourier heat transfer equation (parabolic), and the Poisson equation for electrostatic potential (elliptic). Additionally, we include a more complex, real-world scenario involving incremental forming: a metal sheet is deformed by a tool, assumed to be rigid, into the required shape.

3.1. Phenomena governed by hyperbolic equations

A defining characteristic of hyperbolic problems is the presence of a wavefront—a localized disturbance that propagates through the domain over time—. This reflects the inherently local and directional nature of information flow in such systems, which poses specific challenges and requirements for message-passing in GNNs. One of the most well-known examples of a hyperbolic PDE is the unimodal wave equation.

The wave equation models phenomena such as sound waves, electromagnetic waves, and vibrations in a string. The general form of the wave equation in two spatial dimensions (2D) can be written as:

$$\frac{\partial^2 u}{\partial t^2} = c^2 \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right) = c^2 \nabla^2 u, \quad \text{in } \Omega \times (0, T], \quad (4)$$

where $u(x, y, t)$ represents the wave's height in a two-dimensional domain, and the wave propagates through both the x - and y -directions with speed c . In this scenario, the physical term that governs our problem is the wave propagation speed c , a constant that will be essential in defining the lower bound of the MPis.

Since, as we have seen, some authors have demonstrated the ability of GNNs using message passing to learn finite-difference type schemes from data, it is useful to study these message passing processes in the light of the physical constraints that apply on finite-difference schemes, particularly, the Courant-Friedrichs-Lewy condition [43].

To numerically approximate the solution of the described hyperbolic problem, we discretize the computational domain into a uniform spatial grid with spacing Δx and Δy , respectively, while time is discretized with a time step Δt . For the 2D finite difference scheme to be stable, the time step Δt the famous Courant-Friedrichs-Lewy (CFL) condition establishes that, for a given c_x , c_y , Δx , and Δy mesh resolution in space:

$$\frac{c_x \Delta t}{\Delta x} + \frac{c_y \Delta t}{\Delta y} \leq 1,$$

while for a square grid mesh ($\Delta y = \Delta x$) with constant radial speed,

$$\frac{c \Delta t}{\Delta x} \leq \frac{1}{\sqrt{2}}.$$

What the CFL condition is imposing on the discretization scheme is, plain and simple, that the wave described by Eq. (4) does not traverse more than the distance between two nodes in a single time increment. Otherwise, the scheme will be unstable and its accuracy will decline rapidly.

In other words, what the CFL condition is imposing is that, in each time increment, the physical wave does not go beyond the immediate neighborhood of each node. If we transfer this way of reasoning to the message passing scheme, what we would have is that the wave cannot reach, in each time increment, a greater distance than the message reaches. In other words, assuming for simplicity and without loss of generality a squared mesh,

$$M \Delta x > \sqrt{2} c \Delta t,$$

or

$$M > \sqrt{2} \frac{c \Delta t}{\Delta x}. \quad (5)$$

In general, this figure should be complied with, suitably rounded to a whole number of passes, so that

$$M > \left\lceil \sqrt{2} \frac{c \Delta t}{\Delta x} \right\rceil, \quad (6)$$

where $\lceil \cdot \rceil$ indicates the ceiling function.

3.2. Phenomena governed by parabolic equations

A parabolic PDE models diffusion processes, such as heat conduction and chemical diffusion, where the state of the system evolves over time. This group of problems is characterized by a diffusivity term, which causes perturbations to spread across the entire domain. The evolution of the system occurs at a rate determined by the flux field, which changes as the system approaches equilibrium. In contrast to hyperbolic problems, where a wavefront propagates the perturbation, parabolic problems involve a simultaneous, diffusive perturbation that affects the entire domain at each time step. This is the well-known paradox of infinite propagation velocity in diffusing systems [44].

The ubiquitous example of a parabolic problem is the heat transfer equation, which describes the distribution of temperature over time in a given domain. Consider the Fourier heat equation on a unit square domain Ω as a model for diffusive parabolic PDEs:

$$\frac{\partial u}{\partial t} = \alpha \nabla^2 u \quad \text{in } \Omega \times (0, T], \quad (7)$$

where $u(x, y, t)$ represents now the temperature distribution, $\alpha > 0$ is the thermal diffusivity constant, and $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ is the Laplacian operator. The problem is completed with appropriate initial and boundary conditions.

When numerically approximating parabolic partial differential equations, given the paradox of infinite propagation velocity, the entire domain must be informed about the origin of the perturbation at each temporal step [39]. This means that the state at every node in the domain is influenced by all other nodes. In other words, each point must “see” the global state of the system to correctly capture the underlying physical behavior. This observation motivates our design of the lower bound for message passing iterations in such systems. To ensure that the influence of boundary conditions and global geometric features is properly captured, the number of message passing iterations must be sufficient for information to propagate throughout the domain,

$$M = \frac{L}{\Delta x}, \quad (8)$$

where L denotes the size of the domain, Δx is the spatial discretization, and M represents the number of hops required to propagate information from one boundary of the domain to the opposite side. In this work, we assume, for simplicity, that all domains are square, with uniform spatial discretization such that $\Delta x = \Delta y$. This hypothesis, which might at first sight seem excessively rigid, is motivated by our desire not to mix another problem of graph networks, the so-called over-squashing [45]. The number of h -hop neighbors grows exponentially with h . If, in addition and due to the topology of the network, bottlenecks occur, the transmission of messages can become computationally very expensive. In order to isolate this effect conveniently, as mentioned above, we have preferred to analyze examples that are developed in square domains and that, therefore, do not present these difficulties.

Our principal hypothesis states that the optimal hyperparameter configuration for the GNN must guarantee that the information carried by the messages reaches the farthest side of the domain. This condition ensures that each inference step incorporates a global understanding of the system’s behavior.

3.3. Phenomena governed by elliptic equations

Elliptic phenomena are typically associated with steady-state processes, where the system has reached equilibrium, and there is no time dependence. These problems are characterized by partial differential equations that describe spatial relationships between variables, such as elasticity or potential in a given domain. Unlike hyperbolic and parabolic problems, elliptic equations do not involve wave propagation or diffusion over time but instead focus on spatial distributions that satisfy a given condition across the entire domain.

As in the elliptic case, problems governed by parabolic equations require messages to reach the far ends of the grid, which, in general, will add significant computational cost to the problem. Therefore, the stability condition will revert to that expressed by Eq. (8).

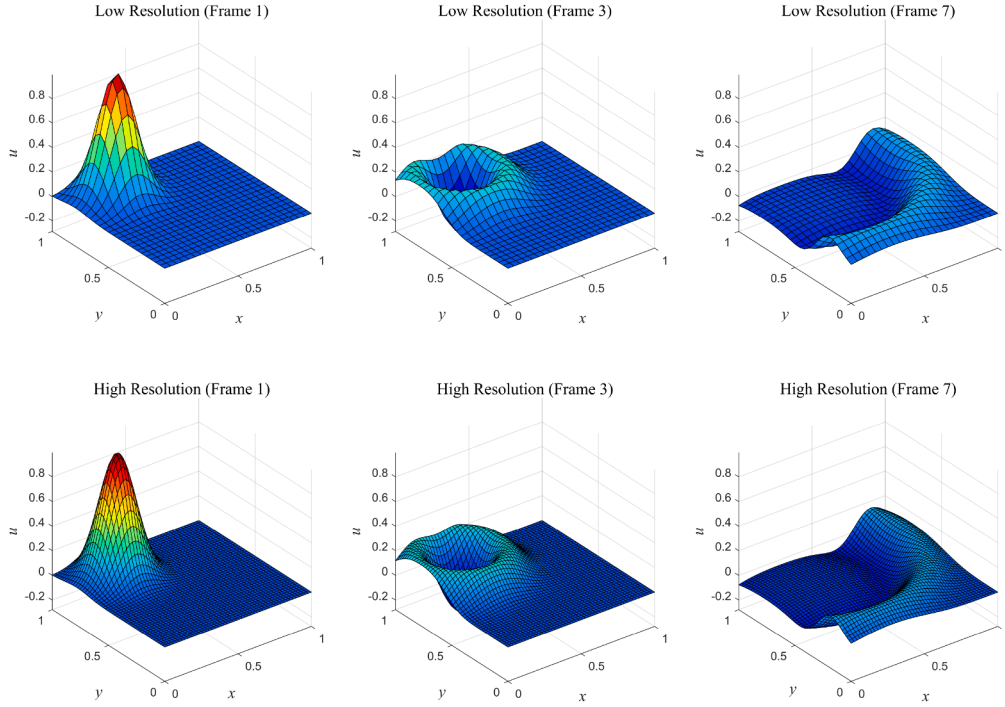


Fig. 2. Temporal evolution of one of the wave problems computed using finite difference methods with two distinct spatial resolutions (low and high—top and bottom rows, respectively). Each simulation uses a temporal sampling of $N = 200$ time steps, corresponding to a total duration of 2.0 s. These time steps are pruned and only ten of them kept for training. The wave speed is fixed at $c = 0.5$. Time steps 1, 3 and 5 are represented at the left, center and right columns, respectively.




Fig. 3. Evolution of the von Mises stress for the incremental forming problem simulated with finite element methods. The FEM simulations used 1000 time steps, of which only 10 are saved for training. The actuator impacts the plate from both upper and lower z -directions, producing plastic deformation. Clamped boundary conditions are applied at plate edges. Time steps 1, 7 and 11 of one of the data sets are shown, respectively, from left to right.

As an example, we consider, on one side, the electrostatic potential $u(x, y)$ for a two-dimensional domain $\Omega \subset \mathbb{R}^2$, governed by Poisson's equation in the presence of a charge distribution $\rho(x, y)$:

$$\nabla^2 u = -\frac{\rho}{\epsilon_0}, \quad (9)$$

where ϵ_0 is the permittivity of free space, $\rho(x, y)$ is the charge density distribution and $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ is the Laplacian operator. Again, the problem is completed with appropriate boundary conditions.

As a second elliptic example, we model quasi-static incremental forming process, that involves the collision between a rigid tool and a elasto-plastic square plate in 3D space, see Fig. 3. This setup introduces additional complexity, including the 3D nature of the problem, the use of multivariate fields—namely the displacement vector \mathbf{u} and stress tensor $\boldsymbol{\sigma}$ —and the presence of contact interactions between two bodies. In this context, the resulting deformations depend on the boundary conditions, the geometry of the system, and the magnitude and location of the imposed displacements from the actuator.

Table 1

Dataset characteristics and numerical methods. Hyperbolic and parabolic cases use finite differences (FD) on 2D rectangular meshes; elliptic cases use finite element methods (FEM) on 3D tetrahedral meshes. Columns show: dataset name, spatial step Δx , time step Δt , estimated lower-bound message passing iterations M , time sampling N , wave speed c (hyperbolic only), mesh nodes, discretization type, problem dimension, numerical solver, and domain dimensions. For the plastic plate problem, Δt represents the pseudo-time through which the loading process is discretized. The problem is actually quasi-static and the load is applied within a unit time period.

Dataset	Δx	Δt	M	$\times N$	c	Nodes	Discretiz.	Dim.	Solver	$L \times W(\times H)$
Hyperbolic, Fig. 2										
Wave-LowRes	0.04	0.001	4	200	0.5	625	Rect.	2D	FD	1×1
Wave-HighRes	0.02	0.001	8	200	0.5	2500	Rect.	2D	FD	1×1
Parabolic										
Fourier-LowRes	0.1	0.0004	10	50	–	100	Rect.	2D	FD	1×1
Fourier-LowRes2	0.1	0.0004	20	50	–	400	Rect.	2D	FD	2×2
Elliptic, Fig. 3										
Poisson-LowRes	0.1	–	10	500	–	100	Rect.	3D	Jacobi	1×1
Poisson-HighRes	0.05	–	20	500	–	400	Rect.	3D	Jacobi	1×1
Plastic Plate	0.065	0.001	15	100	–	625	Tethr.	3D	FEM	$1 \times 1 \times 0.1$

The governing partial differential equation for elasto-plastic phenomena in 3D is the equilibrium equation, for which we assume large strain settings,

$$\nabla_0 \mathbf{P} + \mathbf{f}_0 = \mathbf{0},$$

where \mathbf{P} represents the first Piola-Kirchhoff stress tensor, and $\nabla_0 = \frac{\partial \mathbf{P}}{\partial \mathbf{X}}$, with \mathbf{X} the reference configuration of the solid. \mathbf{f}_0 represents the body forces at the reference configuration. This equation must be supplemented with appropriate Dirichlet and Neumann boundary conditions. In addition, we assume that the material follows the classical multiplicative decomposition of the deformation gradient into elastic and plastic parts,

$$\mathbf{F} = \mathbf{F}^{el} \cdot \mathbf{F}^{pl},$$

with rate-independent plasticity, beginning when the yield function

$$f(\boldsymbol{\sigma}, H) < 0,$$

—considered here as J2 plasticity—is reached. H represents the (isotropic) hardening parameter of the constitutive model and $\boldsymbol{\sigma}$ the Cauchy stress tensor.

Since any of the above phenomena—and, in general, any other governed by elliptic equations—are essentially static, time plays no role in the problem and therefore information must be transmitted instantaneously throughout the domain. This results in a lower bound on the number of passes which is given, again, by Eq. (8).

4. Numerical results

4.1. Synthetic datasets

We have generated four distinct sets of data, one for each of the problems described in Section 3. Each dataset consists of 1000 simulations, which are split into 80 %, 10 %, and 10 % for training, validation, and testing, respectively.

As mentioned before, the hyperbolic and parabolic cases are generated with finite differences, while the elliptic case is solved with finite element methods (FEM). Both of these methods require fine spatial and temporal discretizations to ensure stability and convergence, leading to high computational costs due to the large number of iterations that the solvers must comply with. To address this cost, we pruned the spatial and time domains at a coarser ratio, where these methods would typically fail. As will be shown, message-passing GNNs show their potential in solving complex physical problems with coarser spatial and temporal discretizations.

To rigorously assess the impact of the number of message-passing iterations, we generate two datasets for each problem, intentionally varying the number of iterations required for convergence. This variation is achieved by altering the spatial resolution, thereby producing datasets of different granularities, as summarized in Table 1.

4.2. Experiments

The experiments presented in this work are designed to investigate how the number of MPI affects modeling physical phenomena using GNNs. Our focus is exclusively on assessing the effect of this hyperparameter, which we hypothesize to be critical to the performance and generalization of such architectures.

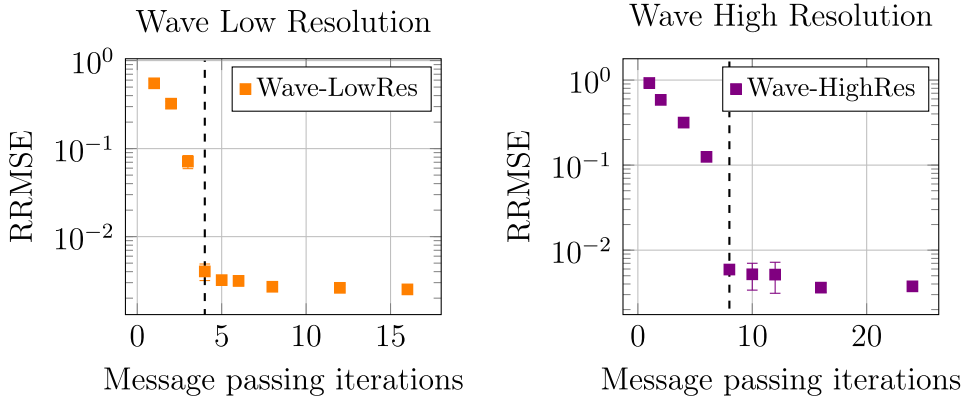


Fig. 4. Relative Root Mean Square Error (RRMSE) over three random seeds for 10-step hyperbolic rollouts, capturing cumulative prediction error. Error bars indicate standard deviation across different model initializations. The dashed vertical lines are the theoretical lower bounds predicted by this work. Left: for the low-resolution dataset, the predicted lower bound is 4 MPI. Right: for the high-resolution dataset, the predicted lower bound is 8 message passes per time step. Beyond the threshold, errors converge. All models share the same number of parameters (shared weights), with a latent space size of 64 and 128, respectively.

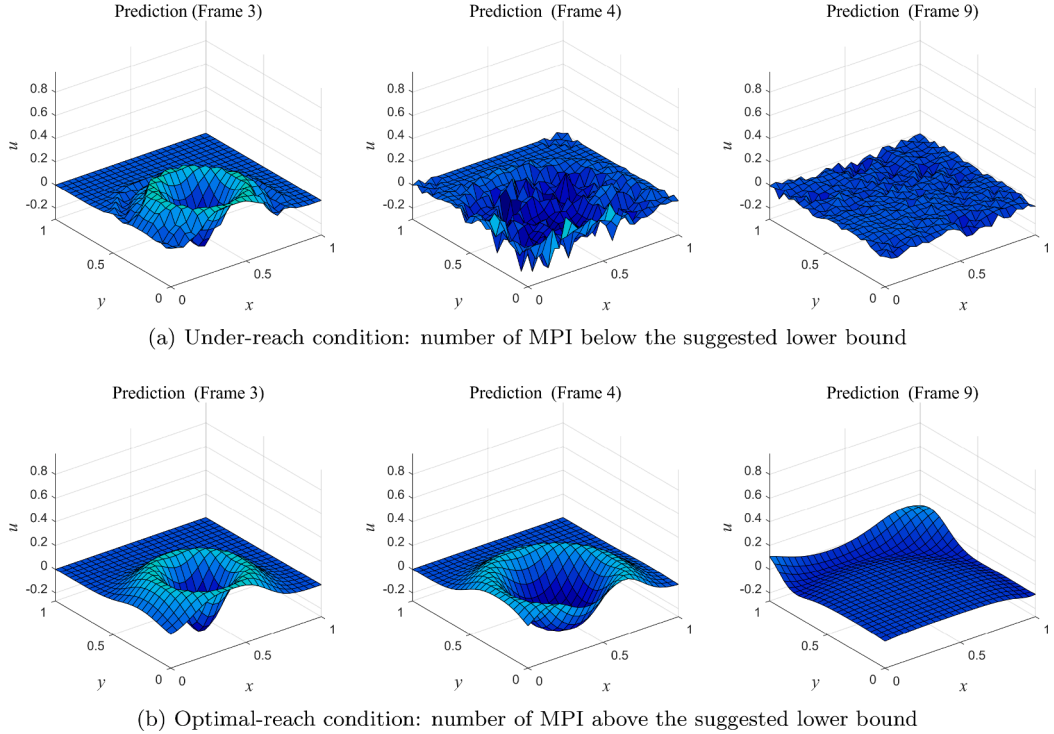


Fig. 5. Comparison of GNN prediction performance across different time steps and message passing iterations for the wave problem. Left: early time step ($n = 3$), middle: intermediate ($n = 4$), right: later stage ($n = 9$). (a) With $M = 2$ MPI (under-reaching regime, below optimal threshold); (b) $M = 8$ MPI (optimal regime, demonstrating prediction accuracy and stable rollout performance).

To enable a fair comparison across different models, we fix the number of trainable parameters, ensuring that each model maintains sufficient learning capacity for the given tasks. This is achieved through the shared-weight design of the processor module of the model, where the only varying factor among different architectures is the number of MPI. This design choice prevents performance differences from being attributed to the model's expressiveness, thereby isolating the effect of the number of MPI. To validate our hypothesis, we conduct a sensitivity analysis on model performance with respect to the number of MPI. Each configuration is trained using multiple random initialization seeds, allowing us to derive statistically robust insights into the performance trends.

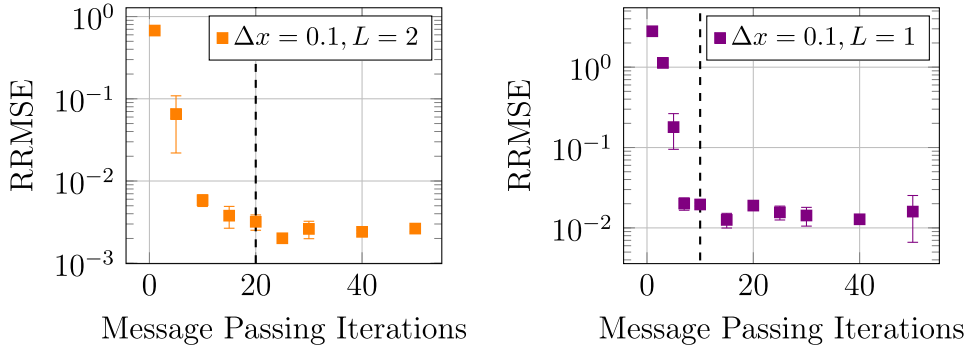


Fig. 6. Error over five random seeds for n -step parabolic rollouts. Dashed lines represent the suggested lower bound. Left: 2×2 domain, giving an optimal number of message passing iterations of $M = 20$. Right: 1×1 domain, thus giving $M = 10$. Errors remain virtually stable from this point onwards. In any case, for a very large number of passes it seems that the error increases slightly, which may be due to the phenomenon of over-smoothing. All models use a fixed 128-dimensional latent space and share parameters.

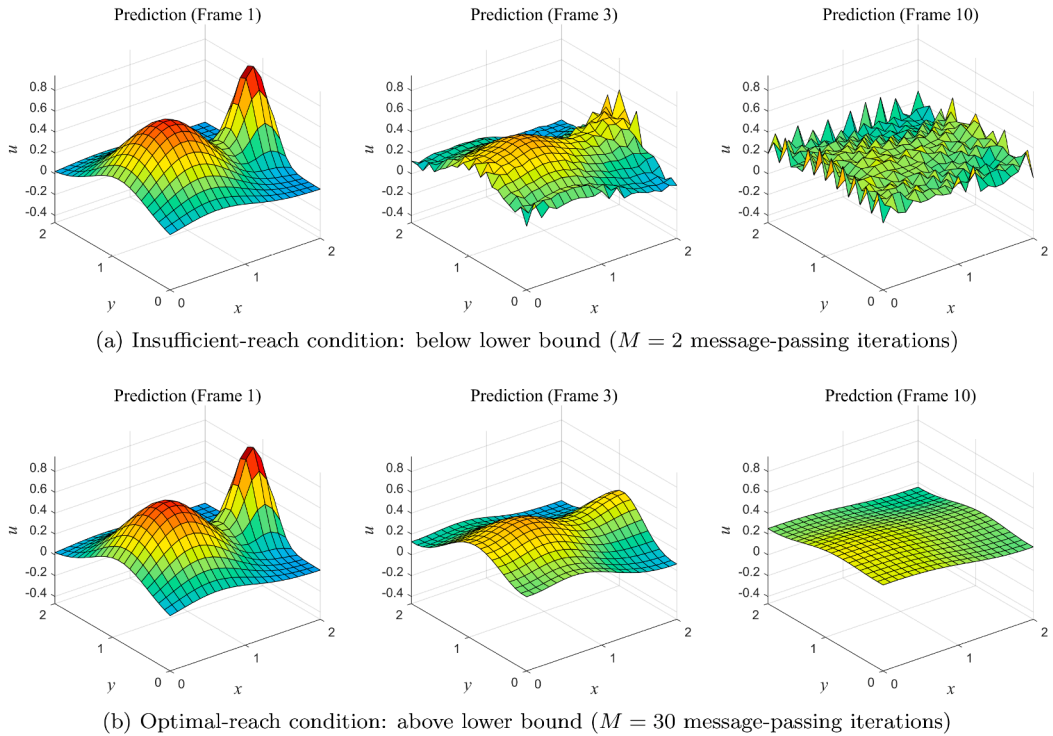


Fig. 7. Comparison of GNN performance for the Fourier problem. (a) Insufficient message passing ($M = 2$ MPI) fails to capture the solution dynamics. (b) Optimal design with $M = 30$ MPI achieves peak accuracy, demonstrating robust spatial interpretability and stable predictions.

Section 4.3 examines the proposed lower bound for hyperbolic equations, studying the relationship between the distance travelled by the message as it passes through the graph, and the distance physically travelled by the wave whose transmission is modeled by the equation in question. We demonstrate, in the context of the hyperbolic problem, that this relationship can effectively guide the optimal selection of MPI as a hyperparameter in GNNs.

Section 4.4 investigates the proposed lower bound for parabolic and elliptic equations by analyzing the influence of spatial discretization for a domain geometry. In this context, we focus on parabolic and elliptic problems, where the system's behavior is not governed by wavefront propagation, but instead emerges from the global structure of the geometry itself.

4.3. Wave equation

For systems governed by hyperbolic equations, as mentioned above, our hypothesis is that the network must send messages that are always ahead of the physical wave modeled by the equation. The sending of messages can never lag behind the

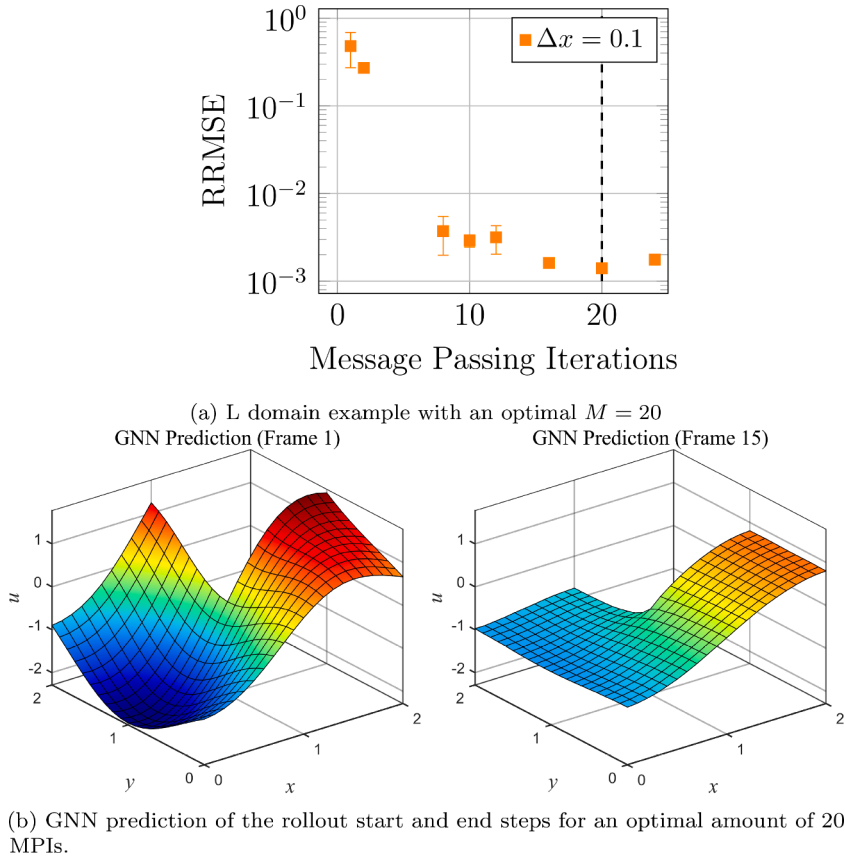


Fig. 8. Error convergence for n -step parabolic rollouts computed over an L-shaped domain, with results averaged across five random seeds. Panel (a) depicts the asymptotic evolution of the numerical error for the resolved geometric limit, where the dashed lines indicate the theoretical lower bounds on convergence. Panel (b) presents the resulting prediction error from the policy that satisfies the optimal condition of 20 MPIs.

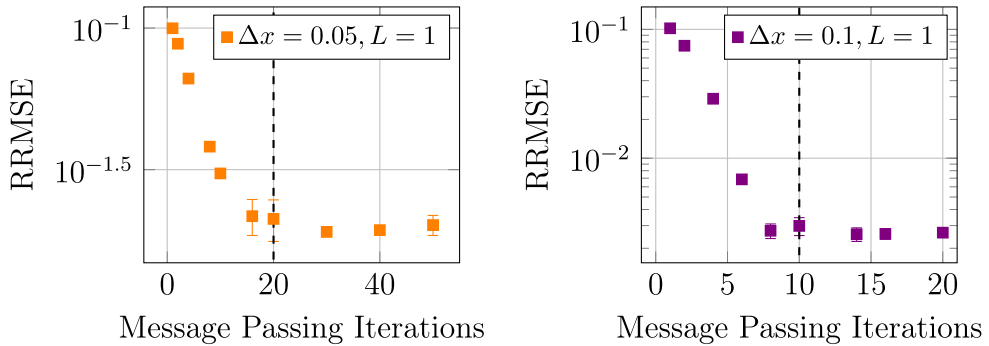


Fig. 9. Relative root mean square error over four random seeds for n -step parabolic rollouts for the electrostatic potential problem. Dashed lines represent the predicted lower bound. Left: High-resolution dataset with $M = 20$. Right: Low resolution dataset, with $M = 10$. Errors converge asymptotically beyond these thresholds. All models use a fixed 64-dimensional latent space and share parameters across number of MPI.

physics, or else the affected nodes will not receive the necessary information to be able to reproduce the physical phenomenon, the wave.

This relationship establishes a physics-guided lower bound for the hyperparameter settings, ensuring model convergence-defined here as the model's ability to integrate multiple rollout steps without accumulating significant drift or deviating into physically implausible scenarios. This reduces uncertainty during hyperparameter tuning and, consequently, lowers computational cost. The experimental results shown in Fig. 4 validate the proposed hypothesis in both scenarios. The plots display the n -step rollout error on

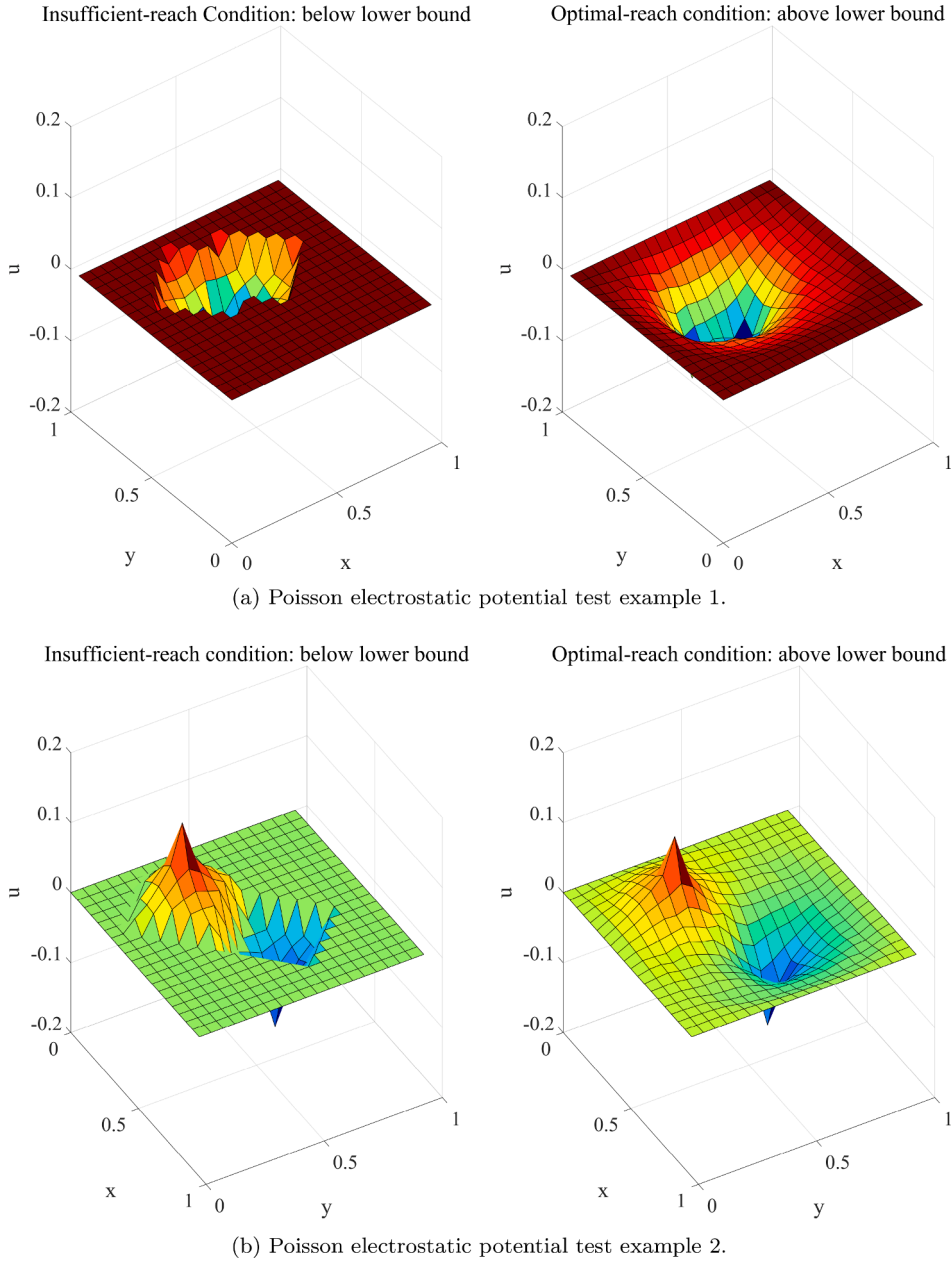


Fig. 10. Comparison of GNN performance for elliptic problems for two different examples. The left column shows Insufficient-reach design below the lower bound with only 4 MPI. The right column shows optimal-reach design above the lower bound with 30 MPI and achieve peak accuracy, demonstrating robust spatial interpretability and stable rollout predictions.

the test set for varying numbers of MPI across two different lower-bound configurations. As illustrated, when the number of MPI is set below the lower bound (indicated by the dashed line), the models fail to capture the underlying physical behavior, resulting in poor performance. It is only once the lower bound is exceeded (optimal setting) that the models begin to effectively learn the physics, ultimately producing accurate predictions.

Moreover, an asymptotic behavior can clearly be observed: as the number of MPI enters the optimal regime, the performance tends to saturate, flattening the relationship between prediction error and the number of iterations. This proves that the number of MPI is a crucial hyperparameter in the model's design, and that can be bounded using expert knowledge. Although all models have the same number of trainable parameters, increasing the number of MPI significantly improves problem understanding and convergence-reducing relative errors by up to three orders of magnitude. This highlights that performance gains come not from model size, but from more effective information propagation.

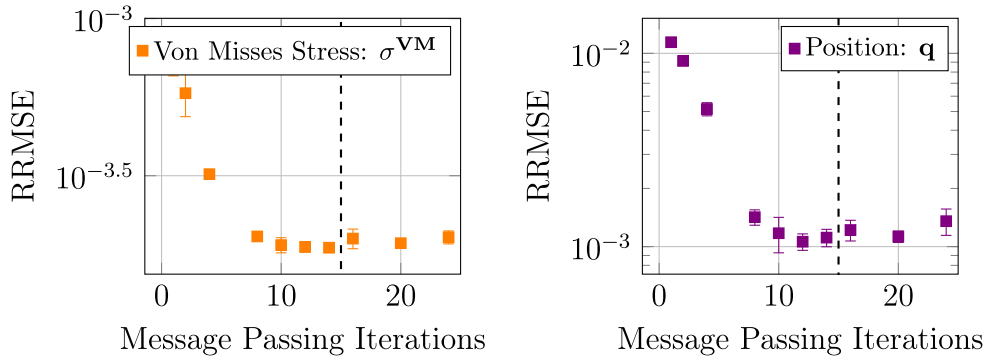


Fig. 11. Error over five random seeds for n -step predictions in the incremental forming problem, showing cumulative error across rollout. The dashed vertical line represents the suggested lower bound for the number of message passing iterations, M . When the model's design is below the lower bound the results reveal higher errors. Both von Mises stress (left) and position (right) errors are shown, each showing a minimum for the predicted optimal MPI value. Models use a 64-dimensional latent space with shared weights across MPI to maintain constant model size.

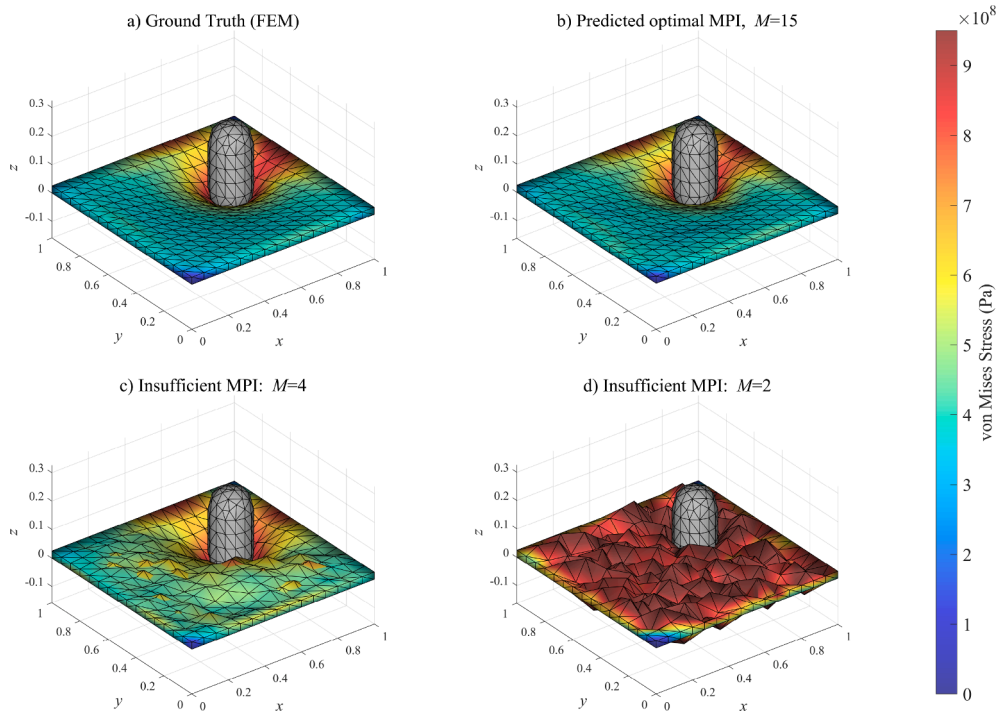


Fig. 12. Comparison of GNN performance as a function of the number of passes for the incremental forming problem. a) The ground truth FEM solution; b) Prediction produced by the model with 15 message-passing iterations, above the lower bound; c) and d) display the divergence observed when the message-passing iterations does not satisfy the lower bound, with 4 and 2 no. of message passing iterations respectively.

The primary reason for poor performance when the model is insufficiently or sub-optimally hyperparameterized is that it is forced to learn the physical phenomenon from incomplete information. This limitation truncates the model's ability to capture the underlying physics, even when the latent space has sufficient expressiveness. As a result, instead of producing meaningful predictions, this can lead to divergence, see Fig. 5.

4.4. Phenomena governed by parabolic and elliptic PDEs

When simulating phenomena governed by parabolic or elliptic PDEs, we hypothesize that message-passing GNN convergence requires each node to eventually acquire a global view of the system state. Since the GNN updates node states solely through the message passing algorithm, as described in Sections 3.2–3.3, this motivates defining the lower bound from a geometric perspective.

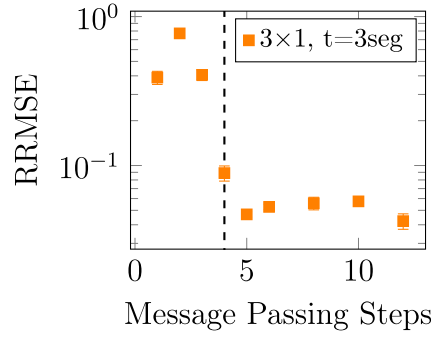


Fig. 13. Relative Root Mean Square Error (RRMSE) over a test extrapolation out of the training domain for the wave problem. Error bars indicate standard deviation across different model initializations. Errors are higher in under-reaching regimes-below the message aggregation limit (dotted line: 4 passes for low-res). Beyond this threshold, errors converge, validating the under-reach theory also for extrapolation scenarios. All models share the same number of parameters (shared weights), with a fixed latent space of 64 dimensions.

4.4.1. Fourier equation

We validate the hypothesis using two different cases, with low and high resolution, respectively. As shown in Fig. 6, an asymptotic behavior is clearly observed: as the number of MPI enters the optimal regime, performance tends to saturate, flattening the relationship between prediction error and the number of iterations. Moreover, increasing the number of iterations beyond this point does not lead to significant reductions in error, indicating saturation and identifying the minimal number of MPI. A high number of MPI brings higher computational cost, which scales as $\mathcal{O}(M)$, where M is the number of MPI.

Fig. 7 compares the model's performance between optimal and insufficient message-passing configurations. The results demonstrate that avoiding under-reaching significantly enhances convergence efficiency and mitigates biased or noisy predictions. Fig. 8 illustrates the same experiment within an L-shaped domain, thereby validating the hypothesis in non-convex geometrical settings.

4.4.2. Electrostatic potential problem

For the electrostatic potential problem, we validate the hypothesis using two different lower-bounded cases, with low- and high-resolution discretisations, that have lower bounds of $M = 10$ and $M = 20$, respectively. Here, the input to the GNN is a distribution of charge densities ρ_k and the output the resulting potential field u . This confirms that the lower-bound behavior is dictated purely by the spatial discretization size and the geometry of the domain. This outcome is expected in elliptic systems, where the solution influences the entire domain instantaneously (i.e., the whole state is effectively updated in a single inference step) and time is not a variable of the problem.

As shown in Fig. 9, both datasets exhibit an asymptotic performance trend as the number of MPI increases. Additionally, Fig. 10 illustrates the impact of meeting or violating the geometric lower bound. When the model is under-reached, it fails to capture the global behavior of the physical phenomenon, as its predictions are limited to the region within the restricted message-passing range. In contrast, models configured within the optimal regime effectively propagate information across the entire domain, enabling accurate and physically consistent predictions.

4.4.3. Incremental forming

For the incremental forming problem, we employ the MeshGraphNet architecture [2]. This setup differs from the previously discussed problems in that it involves a multi-graph configuration, where a rigid actuator contacts a deformable plate, and the network is tasked with predicting both displacement and von Mises stress fields. More information about this architecture can be found at [14]. Despite the increased complexity of the scenario, the core hypothesis regarding MPI remains valid, as the underlying PDE remains elliptic (the problem is assumed quasi-static).

As shown in Fig. 11, the multi-graph configuration also exhibits an asymptotic performance trend as the number of MPI approaches the predicted lower bound. It is noted, however, that an increase in the number of passes beyond the predicted figure provides some increase in error, albeit slight. This increase may be due to another phenomenon, known as over-smoothing [46,47]. Over-smoothing is a phenomenon whereby, if the number of message passes grows too large, all nodes in the graph end up carrying very similar information, losing their uniqueness and causing a drop in prediction accuracy.

Fig. 12 illustrates the effect of correctly configuring the number of MPI. As can be seen, the accuracy of the simulation can be seriously affected if the number of passes is much lower than necessary. As in the preceding cases, setting the number of MPIs to the number predicted by our work ensures excellent accuracy for all the simulations studied. We hypothesize that, in some cases, convergence may even be reached before the predicted MPI limit due to the effect of the multi-graph actuator, which can act as an information bridge and reduce the need for several intermediate message iterations.

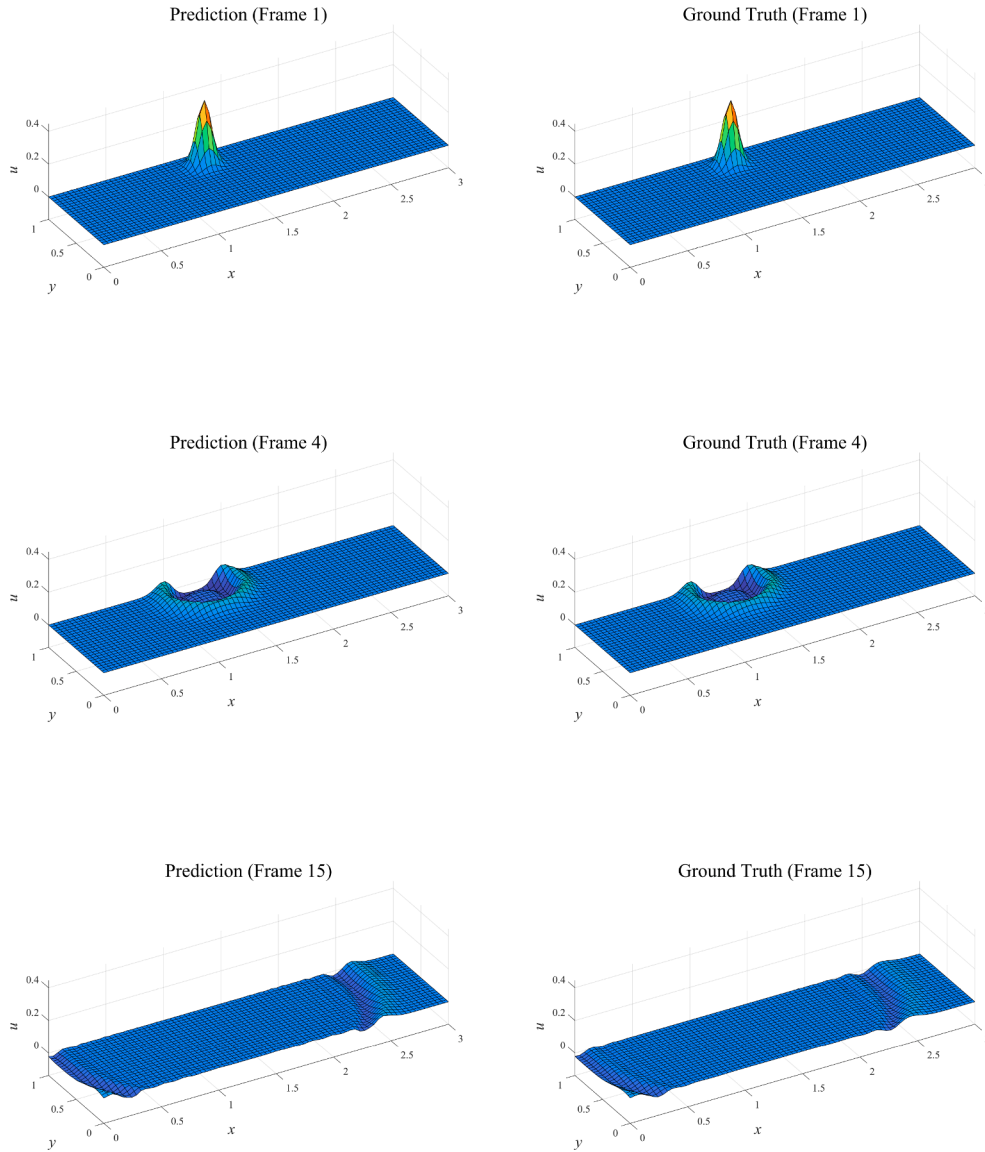


Fig. 14. A comparison between the dynamic rollout predicted by the Graph Neural Network (GNN) and the numerical solution in an extrapolative inference scenario for the low-resolution wave propagation problem. The model was trained on a 1×1 plate domain and evaluated on an extrapolated 3×1 plate, highlighting its ability in generalizing to larger spatial domains. 5 message passes were employed.

4.5. Extrapolation performance

GNN are a powerful tool for handling different geometries as well as varying initial and boundary conditions during inference, enabling the modeling and extrapolation to unseen scenarios. However, like all deep learning methods, this is usually achieved not without significant effort. It is crucial to understand both the underlying physics of each problem and the technology being used. Only then can we achieve meaningful extrapolation, grounded in the idea that, in many cases, we are essentially performing interpolation within a learned latent space, something at what NN are excellent at.

In this work, we use the proposed lower-bound principle to identify scenarios where models are likely to succeed or fail. For hyperbolic problems, geometric variations do not pose a significant challenge due to the local nature of the modeled phenomena—a travelling wave—. As shown in Fig. 13, the model generalizes well to unseen geometries and even to longer rollout horizons, demonstrating its ability to extrapolate in both space and time.

Conversely, for elliptic and parabolic problems, successful extrapolation to unseen scenarios depends on whether the MPI span the entire domain. Fig. 14 shows that a model trained on the Poisson high-resolution mesh dataset with a maximum of 3 charges can successfully solve an unseen case with 5 charges. However, when the domain itself changes, from 1×1 to 3×3 as in Fig. 15, the GNN




Fig. 15. A comparison between the dynamic rollout predicted by the Graph Neural Network (GNN) and the numerical solution in an extrapolative inference setting for the high-resolution Poisson problem. In this case, the trained domain was extended from three initial source points to five, demonstrating the model's capacity to generalize beyond its training configuration.

fails to cover the entire domain. Despite the simpler load configuration (only 3 charges), the model's MPI design no longer satisfies the new lower-bound condition and becomes insufficient.

5. Conclusions

This study systematically investigates the influence of message-passing iterations in graph neural networks when modeling physical systems governed by partial differential equations. Our findings demonstrate that message passing is not merely a computational mechanism but a physics-aligned process that directly impacts model accuracy, generalization, and stability. Specifically, we identify and validate physics-guided lower bounds for message propagation that can be used to inform optimal model design.

For hyperbolic systems, we propose a lower bound based on the relationship between the distance travelled by the message sent and by the physical wave whose evolution is to be resolved. Empirical results from wave propagation datasets confirm that convergence is only achieved when message passing travels ahead of the physical wave. This validates message passing as a critical hyperparameter and establishes a novel, principled approach for hyperparameter tuning in GNNs.

For parabolic and elliptic systems, we demonstrate that convergence requires global information propagation across the domain. This leads us to define a lower bound on message-passing steps, dependent on mesh topology and the geometry of the domain. Across heat diffusion and electrostatic potential datasets, we observe that performance saturates once this bound is reached. This behavior persists even in complex scenarios such as plastic deformation in incremental forming. Notably, even under varying geometries and boundary conditions, the optimal number of iterations remains tightly linked to the spatial discretization, confirming the universality of the proposed criterion.

These experiments demonstrate that the extrapolation capability of GNNs strongly depends on the type of lower-bound condition governing the problem. For hyperbolic problems, the locality of the solution allows the model to generalize well to unseen geometric domains. For elliptic and parabolic problems, success requires that MPIs span the entire domain at each time step. When this is not the case—as in larger, unseen domains—the model fails due to limited information propagation. However, it performs well on unseen scenarios as long as the domain remains within the reach of the MPIs. To address these challenges, prior research has suggested incorporating multi-graph or multi-resolution message-passing algorithms into graph processors, which can further enhance the efficiency of information propagation [48,49]. Other studies on graph pooling indicate that similar improvements can be achieved for the issues discussed in this work, while also preserving interpretability of the structured data [50]. This is particularly relevant for mitigating the mentioned domain under-reach in phenomena governed by parabolic and elliptic partial differential equations (PDEs).

Overall, our work offers a rigorous framework for physics-aware GNN design, where the number of message-passing iterations is explicitly tied to the physical and geometric properties of the modeled system. These findings provide guidelines for architecture selection and hyperparameter tuning, advancing the use of GNNs in scientific machine learning and data-driven physics simulation.

CRedit authorship contribution statement

Lucas Tesán: Writing – original draft, Visualization, Investigation, Data curation; **Mikel M. Iparraguirre:** Writing – original draft, Investigation, Data curation; **David González:** Writing – review & editing, Formal analysis, Conceptualization; **Pedro Martins:** Writing – review & editing, Investigation, Conceptualization; **Elías Cueto:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition.

Data availability

Data will be made available on request.

Declaration of competing interest

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

Acknowledgement

This work was supported by the Spanish Ministry of Science and Innovation, AEI/10.13039/501100011033, through Grant number PID2023-147373OB-I00, and by the Ministry for Digital Transformation and the Civil Service, through the ENIA 2022 Chairs for the creation of university-industry chairs in AI, through Grant TSI-100930-2023-1.

This material is also based upon work supported in part by the [Army Research Laboratory](#) and the [Army Research Office](#) under contract/grant number [W911NF2210271](#).

The authors also acknowledge the support of ESI Group through the chair at the University of Zaragoza.

References

- [1] F.J. Montáns, E. Cueto, K.-J. Bathe, Machine learning in computer aided engineering, in: *Machine Learning in Modeling and Simulation: Methods and Applications*, Springer, 2023, pp. 1–83.
- [2] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, P. Battaglia, Learning mesh-based simulation with graph networks, in: *International Conference on Learning Representations*, 2021. https://openreview.net/forum?id=roNqYLO_XP.
- [3] F. Aldakheel, E.S. Elsayed, Y. Heider, O. Weeger, Physics-based machine learning for computational fracture mechanics, *Mach. Learn. Comput. Sci. Eng.* 1 (1) (2025) 18.
- [4] R. Sarkar, S. Abi-Karam, Y. He, L. Sathidevi, C. Hao, FlowGNN: a dataflow architecture for real-time workload-agnostic graph neural network inference, in: 2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA), 2023, pp. 1099–1112. <https://doi.org/10.1109/HPCA56546.2023.10071015>
- [5] H. Gao, M.J. Zahr, J.-X. Wang, Physics-informed graph neural Galerkin networks: a unified framework for solving PDE-governed forward and inverse problems, *Comput. Methods Appl. Mech. Eng.* 390 (2022) 114502. <https://doi.org/10.1016/j.cma.2021.114502>
- [6] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirmsberger, M. Fortunato, F. Alet, P. Battaglia, GraphCast: learning skillful medium-range global weather forecasting, *arXiv preprint arXiv:2212.12794* (2022).
- [7] K. Jordán, *Calculus of Finite Differences*, 33, American Mathematical Society, 1965.
- [8] K.J. Bathe, Finite element method, *Wiley Encycl. Comput. Sci. Eng.* (2007) 1–12. <https://doi.org/10.1002/9780470050118.ecse159>
- [9] R. Eymard, T. Gallouët, R. Herbin, Finite volume methods, *Handb. Numer. Anal.* 7 (2000) 713–1018.
- [10] G.E. Karniadakis, I.G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, L. Yang, Physics-informed machine learning, *Nat. Rev. Phys.* 3 (6) (2021) 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
- [11] Y. Zhao, H. Li, H. Zhou, H.R. Attar, T. Pfaff, N. Li, A review of graph neural network applications in mechanics-related domains, *Artif. Intell. Rev.* 57 (11) (2024) 315. <https://doi.org/10.1007/s10462-024-10931-y>
- [12] K.R. Allen, T. Lopez-Guevara, K. Stachenfeld, A. Sanchez-Gonzalez, P. Battaglia, J. Hamrick, T. Pfaff, Physical design using differentiable learned simulators, 2022. <https://arxiv.org/abs/2202.00728>. 2202.00728
- [13] K.R. Allen, Y. Rubanova, T. Lopez-Guevara, W. Whitney, A. Sanchez-Gonzalez, P. Battaglia, T. Pfaff, Learning rigid dynamics with face interaction graph networks, 2022. <https://arxiv.org/abs/2212.03574>. 2212.03574
- [14] A. Tierz, M.M. Iparraguirre, I. Alfaro, D. González, F. Chinesta, E. Cueto, On the feasibility of foundational models for the simulation of physical phenomena, *Int. J. Numer. Methods Eng.* 126 (6) (2025) e70027. <https://doi.org/10.1002/nme.70027>
- [15] S. Cuomo, V.S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics-informed neural networks: where we are and what's next, *J. Sci. Comput.* 92 (3) (2022) 88. <https://doi.org/10.1007/s10915-022-01939-z>
- [16] S. Cai, Z. Mao, Z. Wang, M. Yin, G.E. Karniadakis, Physics-informed neural networks (PINNs) for fluid mechanics: a review, *Acta Mech. Sin.* 37 (12) (2021) 1727–1738. <https://doi.org/10.1007/s10409-021-01148-1>
- [17] J. Kim, K. Lee, D. Lee, S.Y. Jhin, N. Park, DPM: A novel training method for physics-informed neural networks in extrapolation 35 (2021) 8146–8154. <https://doi.org/10.1609/aaai.v35i9.16992>
- [18] L. Fesser, L. D'Amico-Wong, R. Qiu, Understanding and mitigating extrapolation failures in physics-informed neural networks, 2023. <https://arxiv.org/abs/2306.09478>. 2306.09478
- [19] S. Fresca, L. Dede', A. Manzoni, A comprehensive deep learning-based approach to reduced order modeling of nonlinear time-dependent parametrized PDEs, *J. Sci. Comput.* 87 (2) (2021) 61. <https://doi.org/10.1007/s10915-021-01462-7>
- [20] S. Fresca, A. Manzoni, POD-DL-ROM: Enhancing deep learning-based reduced order models for nonlinear parametrized PDEs by proper orthogonal decomposition, *Comput. Methods Appl. Mech. Eng.* 388 (2022) 114181. <https://doi.org/10.1016/j.cma.2021.114181>
- [21] S. Brivio, S. Fresca, A. Manzoni, Handling geometrical variability in nonlinear reduced order modeling through continuous geometry-aware DL-ROMs, *Comput. Methods Appl. Mech. Eng.* 442 (2025) 117989. <https://doi.org/10.1016/j.cma.2025.117989>
- [22] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (11) (1998) 2278–2324.
- [23] M.M. Bronstein, J. Bruna, T. Cohen, P. Veličković, Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges, *arXiv*, 2021, p. 50. <https://doi.org/10.48550/ARXIV.2104.13478>
- [24] M.M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, P. Vandergheynst, Geometric deep learning: going beyond euclidean data, *IEEE Signal Process. Mag.* 34 (4) (2017) 18–42.
- [25] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, P.S. Yu, A comprehensive survey on graph neural networks, *IEEE Trans. Neural Netw. Learn. Syst.* 32 (1) (2021) 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- [26] A. Thangamuthu, G. Kumar, S. Bishnoi, R. Bhattoo, N.M.A. Krishnan, S. Ranu, Unravelling the performance of physics-informed graph neural networks for dynamical systems, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), *Advances in Neural Information Processing Systems*, 35, Curran Associates, Inc., 2022, pp. 3691–3702. https://proceedings.neurips.cc/paper_files/paper/2022/file/17b598fda495256bef6785c2b76c3217-Paper-Datasets_and_Benchmarks.pdf
- [27] W. Liu, M.J. Pyrcz, Physics-informed graph neural network for spatial-temporal production forecasting, *Geoenery Sci. Eng.* 223 (2023) 211486. <https://doi.org/10.1016/j.geoen.2023.211486>
- [28] Q. Hernandez, A. Badias, F. Chinesta, E. Cueto, Thermodynamics-informed graph neural networks, *EEE Trans. Artif. Intell.* (2022).

- [29] Q. Hernández, A. Badías, D. González, F. Chinesta, E. Cueto, Structure-preserving neural networks, *J. Comput. Phys.* 426 (2021) 109950.
- [30] Q. Hernández, A. Badías, F. Chinesta, E. Cueto, Port-metaplectic neural networks: thermodynamics-informed machine learning of complex physical systems, *Comput. Mech.* 72 (3) (2023) 553–561.
- [31] Q. Hernández, A. Badías, F. Chinesta, E. Cueto, Thermodynamics-informed neural networks for physically realistic mixed reality, *Comput. Methods Appl. Mech. Eng.* 407 (2023) 115912.
- [32] A. Tierz, I. Alfaro, D. González, F. Chinesta, E. Cueto, Graph neural networks informed locally by thermodynamics, *Eng. Appl. Artif. Intell.* 144 (2025) 110108.
- [33] A. Gruber, K. Lee, N. Trask, Reversible and irreversible bracket-based dynamics for deep graph neural networks, *Adv. Neural Inf. Process. Syst.* 36 (2023) 38454–38484.
- [34] B. Moya, A. Badías, D. González, F. Chinesta, E. Cueto, A thermodynamics-informed active learning approach to perception and reasoning about fluids, *Comput. Mech.* 72 (3) (2023) 577–591.
- [35] C. Bermejo-Barbanoj, B. Moya, A. Badías, F. Chinesta, E. Cueto, Thermodynamics-informed super-resolution of scarce temporal dynamics data, *Comput. Methods Appl. Mech. Eng.* 430 (2024) 117210.
- [36] M. Feurer, F. Hutter, Hyperparameter Optimization, 2019, pp. 3–33. https://doi.org/10.1007/978-3-030-05318-5_1
- [37] M. Balcilar, P. Heroux, B. Gauzere, P. Vasseur, S. Adam, P. Honeine, Breaking the limits of message passing graph neural networks, in: M. Meila, T. Zhang (Eds.), *Proceedings of the 38th International Conference on Machine Learning*, 139 of *Proceedings of Machine Learning Research*, PMLR, 2021, pp. 599–608. <https://proceedings.mlr.press/v139/balcilar21a.html>.
- [38] P.D. Lax, *Hyperbolic Systems of Conservation Laws and the Mathematical Theory of Shock Waves*, number 11 in *Conference Board of the Mathematical Sciences Regional Conference Series in Applied Mathematics*, Society for Industrial and Applied Mathematics, Philadelphia, PA, Philadelphia, PA, 1973.
- [39] J.-L. Auriault, The paradox of fourier heat equation: a theoretical refutation, *Int. J. Eng. Sci.* 118 (2017) 82–88.
- [40] J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals, G.E. Dahl, *Message Passing Neural Networks*, Springer International Publishing, Cham, 2020, pp. 199–214. https://doi.org/10.1007/978-3-030-40245-7_10
- [41] W. Lu, Z. Guan, W. Zhao, Y. Yang, L. Jin, NodeMixup: tackling under-reaching for graph neural networks, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 38, 2024, pp. 14175–14183.
- [42] J. Brandstetter, D. Worrall, M. Welling, Message passing neural PDE solvers, *arXiv preprint arXiv:2202.03376* (2022).
- [43] C.A. De Moura, C.S. Kubrusly, The courant–Friedrichs–Lewy (CFL) condition, *AMC* 10 (12) (2013) 45–90.
- [44] V. Bertola, E. Cafaro, On the speed of heat, *Phys. Lett. A* 372 (1) (2007) 1–4.
- [45] J. Topping, F. Di Giovanni, B.P. Chamberlain, X. Dong, M.M. Bronstein, Understanding over-squashing and bottlenecks on graphs via curvature, in: *International Conference on Learning Representations*, 2021.
- [46] S. Qureshi, et al., Limits of depth: over-smoothing and over-squashing in GNNs, *Big Data Min. Anal.* 7 (1) (2023) 205–216.
- [47] Y. Jiang, D. Zou, X. Zhang, Y.G. Wang, Limiting over-smoothing and over-squashing of graph message passing by deep scattering transforms, *arXiv preprint arXiv:2407.06988* (2024).
- [48] R. Gladstone, H. Rahmani, V. Suryakumar, H. Meidani, M. D’Elia, A. Zareei, Mesh-based GNN surrogates for time-independent PDEs, *Sci. Rep.* 14 (1) (2024). Publisher Copyright: © The Author(s) 2024. <https://doi.org/10.1038/s41598-024-53185-y>
- [49] S. Mousavi, S. Wen, L. Lingsch, M. Herde, B. Raonić, S. Mishra, RIGNO: a graph-based framework for robust and accurate operator learning for PDEs on arbitrary domains, 2025, [arXiv:2501.19205](https://arxiv.org/abs/2501.19205)
- [50] S. Barwey, H. Kim, R. Maulik, Interpretable A-posteriori error indication for graph neural network surrogate models, 2024, [arXiv:2311.07548](https://arxiv.org/abs/2311.07548)