
Towards Universal Mesh Movement Networks

Mingrui Zhang¹ Chunyang Wang¹ Stephan Kramer¹ Joseph G. Wallwork²
Siyi Li¹ Jiancheng Liu Xiang Chen³ Matthew D. Piggott¹
¹Imperial College London ²University of Cambridge ³Noah's Ark Lab, Huawei
{mingrui.zhang18, chunyang.wang22, s.kramer}@imperial.ac.uk
{siyi.li20, m.d.piggott}@imperial.ac.uk
jw2423@cam.ac.uk xiangchen.ai@outlook.com ljcc0930@gmail.com

Abstract

Solving complex Partial Differential Equations (PDEs) accurately and efficiently is an essential and challenging problem in all scientific and engineering disciplines. Mesh movement methods provide the capability to improve the accuracy of the numerical solution without increasing the overall mesh degree of freedom count. Conventional sophisticated mesh movement methods are extremely expensive and struggle to handle scenarios with complex boundary geometries. However, existing learning-based methods require re-training from scratch given a different PDE type or boundary geometry, which limits their applicability, and also often suffer from robustness issues in the form of inverted elements. In this paper, we introduce the Universal Mesh Movement Network (UM2N), which – once trained – can be applied in a non-intrusive, zero-shot manner to move meshes with different size distributions and structures, for solvers applicable to different PDE types and boundary geometries. UM2N consists of a Graph Transformer (GT) encoder for extracting features and a Graph Attention Network (GAT) based decoder for moving the mesh. We evaluate our method on advection and Navier-Stokes based examples, as well as a real-world tsunami simulation case. Our method out-performs existing learning-based mesh movement methods in terms of the benchmarks described above. In comparison to the conventional sophisticated Monge-Ampère PDE-solver based method, our approach not only significantly accelerates mesh movement, but also proves effective in scenarios where the conventional method fails. Our project page can be found at <https://erizmr.github.io/UM2N/>.

1 Introduction

Various natural phenomena are modeled by Partial Differential Equations (PDE). The accurate and efficient approximation of the solutions to these complex and often nonlinear equations represents a fundamental challenge across all scientific and engineering disciplines. To solve real-world PDEs, many numerical methods require a computational mesh or grid to discretize the spatial domain. The quality of this mesh significantly impacts the balance between a numerical solution's accuracy and the computational cost required to obtain it. To maintain high-resolution everywhere in the domain is computationally expensive. Many systems modeled by PDEs however are multi-resolution in nature. For example, a small part of the system may be highly dynamic, while other regions are quasi-stationary; alternatively, some locations may be more important while others less important to the question being considered by the calculation [1–4]. While a uniform, high-resolution mesh can be utilized to capture the dynamics or what is most important accurately, this is often wasteful of finite, precious computational resources and comes with sustainability implications [5].

An active body of research is focused on developing deep learning based methods to accelerate the PDE-solving process by learning surrogate neural solvers [6–8]. A learned solver can output PDE

solutions directly given the physical parameters, and boundary conditions of a PDE. However, these methods encounter challenges such as not being able to guarantee physical plausibility of the PDE solution, weak generalization ability, and low data efficiency.

An alternative approach to improve the efficiency of a PDE solver is to utilize *mesh adaptation*, which is a technique for distributing the mesh according to numerical accuracy requirements. Two main categories of mesh adaptation techniques can be identified: *h*-adaptation and *r*-adaptation. *h*-adaptation refines or coarsens the mesh resolution dynamically through topological operations such as adding/deleting nodes and swapping element edges/faces. In contrast, *r*-adaptation (or *mesh movement*) relocates or moves mesh nodes, keeping the mesh topology and thus the number of elements and vertices in the mesh unchanged [9]. These traditional mesh adaptation techniques can help reduce PDE-solving costs, but the mesh adaptation process itself may come at the cost of significant computational overhead.

Deep learning based methods have been proposed to accelerate mesh adaptation. [10–16]. Most previous methods focus on *h*-adaptive refinement [11–13, 17]. Some existing works focus on learning error indicators or Riemannian metrics (which account for element shape and anisotropy, as well as size) [13–16] to guide the mesh adaptation process. The learned indicator or metric is then fed into a traditional remeshing for mesh generation, or a mesh optimizer, which overall limits the performance of these works to being no better than traditional methods, especially in terms of efficiency. Reinforcement learning based methods show potential to improve the mesh adaptation task, but are difficult to train with low data efficiency [11, 17, 18]. Only a small number of works focus on *r*-adaptation. [19, 20] investigated supervised and unsupervised learning based mesh movement methods. However, the proposed methods require re-training from scratch given a different PDE or geometry, which limits their applicability. In addition, the unsupervised learning method requires significant GPU memory and suffers from long training times, which makes re-training even more prohibitive for large scale problems.

Targeting the generalization ability and efficiency of mesh movement, here we introduce the Universal Mesh Movement Network (UM2N), a two stage, deep-learning based model that learns to move the computational mesh guided by an optimal-transport approach driven by solutions of equations of Monge-Ampère (MA) type. Given an underlying PDE to solve numerically, the process of MA based mesh movement includes choosing a suitable monitor function (*i.e.*, a measure for the desired mesh density) and solving an auxiliary PDE whose solution prescribes the movement of the mesh. The Monge-Ampère equation is an example of such an auxiliary PDE, which results in mesh movement with attractive theoretical properties [9], but whose solution comes at a significant computational cost. Therefore, we decouple the underlying PDE solve and the mesh movement process. The learning of the mesh movement process is essentially learning to solve the auxiliary Monge-Ampère equation.

The proposed UM2N consists of a Graph Transformer based encoder and a Graph Attention Network (GAT) [21] based decoder.

Element volume loss is selected within the training loss, instead of the coordinates loss used in previous works [19, 20]. This offers the advantage that element volume loss not only provides supervision signals but also penalizes negative volumes (*i.e.*, inverted elements), thereby reducing mesh tangling and enhancing robustness. We construct a PDE-independent dataset by generating random generic fields for model training. The trained model can be applied in a zero-shot manner to move meshes with different sizes and structure, for solvers applicable to different PDE types and boundary geometries. We comprehensively demonstrate the effectiveness of our methods on different cases including advection, Navier-Stokes examples, as well as a real-world tsunami simulations on meshes with different sizes and structures. Our UM2N outperforms existing learning-based mesh movement methods in terms of the benchmarks described above. In comparison to the sophisticated Monge-Ampère PDE-solver based method, our approach not only significantly accelerates mesh movement, but also proves effective in scenarios where the conventional method fails. The trained model can be directly integrated into any mesh based numerical solvers for error reduction and acceleration in a non-intrusive way, which can benefit various engineering applications suffering trade-offs between accuracy and computational cost.

Our main contributions are listed as follows:

- We present the UM2N framework, which once trained can be applied in a non-intrusive, zero-shot manner to move meshes with different sizes and structures, for solvers applicable to different PDE types and boundary geometries.
- We propose PDE-independent training strategies and mesh tangling aware loss functions for our Graph Transformer and GAT-based architecture.
- We demonstrate the universal ability of UM2N on different PDEs and boundary geometries: Advection, Navier-Stokes examples, and a real-world tsunami simulation case.

2 Related Works

Neural PDE solver. Neural PDE surrogate models can significantly accelerate the PDE solving process. There are three main groups of existing works. Physics-informed Neural Networks (PINNs) [6, 22, 23] are the first group, modeling the PDEs through implicit representations by neural networks. PINNs require knowledge of the governing equations and train neural networks with equation residuals from the PDEs with given boundary conditions and initial conditions. Neural operators are another group of methods for learning to solve PDEs which seek to learn a mapping from a function describing the problem to the corresponding solution function [7, 8, 24–26]. The third group is mesh-based PDE solvers with deep learning, including the utilization of CNNs and GNNs for processing structured and unstructured meshes [15, 27–32]. Our method can be considered to be related to the neural operator approach, although it learns a mapping from monitor function values to the potential field describing the mesh movement instead of the PDE solution directly. In addition, the networks in the proposed model share similar ideas to mesh-based PDE solvers for processing unstructured meshes.

Learning for mesh generation and adaptation. Deep learning methods have been advanced for applications in mesh generation and adaptation. Two tasks need to be differentiated here, the first involves modeling the ideal mesh resolution or density, *e.g.* through learning a metric or error indicator. The second involves the step that actually adapts the mesh in response, where we can distinguish between *h*-adaptivity, in which the connectivity structure of the mesh is changed to refine or coarsen the mesh, and *r*-adaptivity which maintains the connectivity and only moves the vertices.

An example of the first task is MeshingNet [10] which employs a neural network to establish the requisite local mesh density, subsequently utilized by a standard Delaunay triangulation-based mesh generator. Similarly, the determination of optimal local mesh density for the purpose of mesh refinement is described in [12]. In [13], the authors propose a learning-based model to determine optimal anisotropy to guide *h*-adaptation. [15] introduces a neural-network-based approach for predicting the sizing field for use with an external remesher. Instead of learning a metric, some works learn to manipulate the mesh directly. [11, 17, 18] formulate the *h*-adaptation process as a reinforcement learning paradigm, learning a policy that guides to adjust the mesh elements.

r-adaptation (or mesh movement) is generally formulated in a fundamentally different manner to *h*-adaptation. Leading *r*-adaptation methods are based on a mapping between a fixed reference computational domain and the physical domain encompassing the adapted mesh, which is established through the solution of a nonlinear auxiliary PDE. [19] first proposed the use of a GNN based mesh movement network. [20] investigated the unsupervised learning of mesh movement for improving performance of neural PDEs solvers. [33] introduced a learning based mesh movement method specially tailored for computation fluid dynamics simulation of airfoil design.

Our work focuses specifically on *r*-adaptation, *i.e.*, mesh movement methods. The approaches proposed in previous work require re-training the model when applied to different problems, or can only generalize within a limited range of PDE parameters. We propose a mesh movement network that aims to be universal across applications to different PDEs without retraining.

3 Preliminaries and Problem Statement

Monge-Ampère mesh movement. A mesh movement problem can be defined as the search for the transformation between a fixed computational domain Ω_C , and a physical domain Ω_P . Coordinates in these spaces are denoted ξ and \mathbf{x} , respectively. The mesh \mathcal{H}_C defined in the computational domain is the original, often uniform mesh. The mesh \mathcal{H}_P defined in the physical domain is the

adapted (or moved) mesh.¹ The goal of mesh movement is to find a mapping $\mathbf{x} = f(\boldsymbol{\xi})$ between the continuous coordinate fields of Ω_C and Ω_P , although in practice we deduce a mapping of the discrete vertex sets of \mathcal{H}_C and \mathcal{H}_P . While the coordinate transformation can also be used to accommodate time-dependent moving boundaries [34] or the optimisation of the shape of the domain, here we focus on its application to achieve variable resolution such that the solution $u_{\mathbf{x}}$ of a PDE solved on the adapted mesh has higher accuracy than the solution $u_{\boldsymbol{\xi}}$ solved on the original mesh.

A monitor function m over the spatial domain is used to specify the desired mesh density, which can be based on various characteristics of the solution field u of the PDE. It prescribes where the mesh should be refined or coarsened, *i.e.* large monitor values indicate where high mesh resolution is required. Therefore, the goal of the mesh movement process can be rephrased as finding a mapping so that m is equidistributed over the adapted (*i.e.*, physical) mesh [35]:

$$m(\mathbf{x}) \det(\mathbf{J}) = \theta, \quad (1)$$

where \mathbf{J} is the Jacobian of the map $\mathbf{x} = f(\boldsymbol{\xi})$ with respect to the computational coordinates $\boldsymbol{\xi}$, $\det(\mathbf{J})$ denotes the determinant of the Jacobian which corresponds to the relative change in volume under the transformation, and θ is a normalization constant. By using concepts from optimal transport theory [36], the problem can be constrained to have a unique solution, with the deformation of the map expressed in terms of the gradient of a scalar potential ϕ such that:

$$\mathbf{x}(\boldsymbol{\xi}) = \boldsymbol{\xi} + \nabla_{\boldsymbol{\xi}}\phi(\boldsymbol{\xi}). \quad (2)$$

Substituting the additional constraint described by equation (2) into equation (1) gives a nonlinear PDE of Monge-Ampère type:

$$m(\mathbf{x}) \det(\mathbf{I} + \mathbf{H}(\phi(\boldsymbol{\xi}))) = \theta, \quad (3)$$

where \mathbf{I} is the identity matrix and $\mathbf{H}(\phi)$ is the Hessian of ϕ , with derivatives taken with respect to $\boldsymbol{\xi}$, *i.e.* $\mathbf{H}(\phi)_{ij} = \frac{\partial^2 \phi}{\partial \xi_i \partial \xi_j}$. It should be noted that guarantees for existence and uniqueness, and convergence of solution methods for the Monge-Ampère equation generally rely on the domain being diffeomorphic with a convex domain (*e.g.* [37]). In practice, the method is known to break down in some cases where the domain is not simply-connected or has non-smooth boundaries.

The MA equation (3) is an auxiliary PDE which is purely associated with the the mesh movement process, and is independent of the underlying PDE or physical problem we wish to solve, *i.e.*, the MA equation keeps the same form for different PDEs. This decoupling benefits learning-based methods in terms of generalization properties, *i.e.*, a well-learned MA neural solver has the potential to be applied to different problems with a small cost of fine-tuning or even without re-training.

4 UM2N: Universal Mesh Movement Network

4.1 Framework overview.

The proposed UM2N framework is shown in Fig. 1. Given an input mesh, vertex features and edge features are collected separately. The coordinates and monitor function values are gathered from the vertices and input into a graph transformer to extract embeddings. The embedding vector \mathbf{z} obtained from the graph transformer encoder is then concatenated with the extracted edge features \mathbf{e} to serve as the input for the Graph Attention Network (GAT) decoder. The decoder processes this combined input along with a mesh query to ultimately produce an adapted mesh.

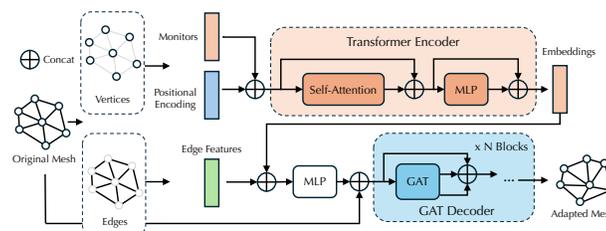


Figure 1: Overview of Universal Mesh Movement Network. The decoder processes this combined input along with a mesh query to ultimately produce an adapted mesh.

Graph Transformer encoder. Transformers [38] are a popular architecture which have been successfully applied for neural PDE solvers [25, 26, 39]. They have strong expressivity and can naturally handle irregular data structures such as unstructured meshes. Their attention mechanism can capture both the local and global information for a vertex of interest in a mesh. Here we use a

¹In what follows, we abuse notation slightly by also using $\boldsymbol{\xi}$ to refer the the (discrete) vertices of \mathcal{H}_C , rather than the (continuous) coordinate field. Similarly, we also use \mathbf{x} to denote the vertices of \mathcal{H}_P .

graph transformer encoder as a feature extractor. The input features include the coordinates of the original mesh ξ and monitor values m_ξ . The coordinates serve as the positional encoding. Features are concatenated and encoded into Q (query), K (key) and V (value) as inputs using MLPs for a self-attention module. The encoder output embedding z are fed into a downstream GAT based deformer for mesh movement.

Graph Attention Network (GAT) based decoder. We choose GATs to construct our decoder. Its attention mechanism can help constrain the mesh vertex movement to within one hop of its neighbors which assists in alleviation of mesh tangling issues. The decoder consists of N GAT blocks. The first GAT block takes a mesh query ξ , and the embedding z from the transformer encoder as well as the edge features e as inputs. For each following k^{th} block, the inputs consist of the coordinates of the initial mesh ξ , the coordinates of the intermediate moved mesh $\xi^{(k-1)}$ and extracted features from the previous layer $f^{(k-1)}$.

4.2 PDE-independent dataset and training.

In the existing work of neural mesh movement [19, 20], proposed models are trained on solutions of PDEs, with the training data generated by solving one specific type of PDE, limiting generalizability of the trained model. Aiming here to train universal mesh movement networks, we construct a PDE-independent training dataset $\mathcal{D} = \{d = (\xi, m_\xi; \mathbf{V}_r, \mathbf{T}_r)\}$, where ξ, m_ξ denotes the original mesh and monitor values as the model's input, and $\mathbf{V}_r, \mathbf{T}_r$ as the pre-calculated (reference) mesh's vertex coordinates and elements as ground truth. To build the dataset, we randomly generate generic solution fields, which are composed as the summation of a random number of Gaussian distribution functions centred in random locations, and with random widths in different directions to introduce anisotropy, samples are shown in Figure A1; refer to Appx. B for a detailed description. The generated fields can be interpreted as representing solutions of any PDE. For mesh movement, a widely used method to translate the PDE solution $u(\mathbf{x})$ into a monitor function, is the Hessian based formula [36]:

$$m(\mathbf{x}) = 1 + \alpha \frac{\|H(u)\|}{\max \|H(u)\|}, \quad (4)$$

where $\|H(u)\|$ indicates the Frobenius norm of the Hessian of u and α is a user chosen constant. This choice results in small cells where the curvature of the solution is high, with a ratio of $1 + \alpha$ between the largest and smallest cell volumes. Our network is trained on the values of the monitor function only, which can be computed relatively cheaply from the PDE solution. The expected movement of the mesh vertices is computed by solving the Monge-Ampère equation (3) for each of the monitor functions based on the randomly generated PDE solutions. After training, the model can then be applied using only the monitor values as input, independent of the PDE that is being solved. Other choices of formulae for the monitor function in terms of u may be appropriate in different cases, but again these can be applied without re-training the model, *e.g.*, the flow-past-a-cylinder and tsunami test cases in the following section use a monitor function based on the gradient of the solution.

4.3 Loss functions.

Given the dataset \mathcal{D} defined above, the final objective is to find model parameters θ , by minimizing the total loss L_{total} . Thus, we can formalize the process as

$$\arg \min_{\theta} L_{total}(\theta) := \lambda_{vol} L_{vol}(\theta) + \lambda_{cd} L_{cd}(\theta). \quad (5)$$

where L_{vol}, L_{cd} represents the element volume loss and Chamfer distance loss respectively, which are defined later. $\lambda_{cd}, \lambda_{vol} > 0$ represent hyper-parameters balancing these two effects.

We denote the modified (adapted) mesh with model parameters θ as $\mathcal{M}(\xi, m_\xi; \theta) = \{\mathbf{V}_a, \mathbf{T}_a\}$, where $\mathbf{V}_a = \{\mathbf{x}_i\}_{i=1}^{n_1}$ represents the coordinates of vertices and $\mathbf{T}_a = \{\mathbf{t}_i\}_{i=1}^{n_2}$ represents the elements with $\mathbf{x}_i, \mathbf{t}_i$ as the i^{th} vertex and element. For ease of presentation, we omit the θ here. Similarly, we define $\mathbf{y}_i, \mathbf{q}_i$ as $\mathbf{V}_r, \mathbf{T}_r$'s i^{th} vertex and element.

Element volume loss. In contrast to the coordinate loss used in previous work [19, 20], here we use the element volume loss for training. Element volume loss is computed as the averaged volume difference between each element in the adapted mesh (\mathbf{T}_a) and the reference mesh (\mathbf{T}_r) computed from the MA method. It is defined as

$$L_{vol}(\theta) = \mathbb{E}_{(\xi, m_\xi; V_r, T_r) \in \mathcal{D}} \left[\frac{1}{|T_a|} \sum_{i=1}^{n_2} |\text{Vol}(t_i) - \text{Vol}(q_i)| \right], \quad (6)$$

where Vol is the function computing the volume of a given element t .

The element volume loss is inspired by the equidistribution relation in equation (1). Intuitively, the equidistribution relation enforces deformed mesh elements such that their volumes correspond to rescaling by the monitor function value $m(\mathbf{x})$. Therefore, the element volume loss is utilised to encourage the model to learn to move the mesh so as to conform to the equidistribution relation, under the guidance of MA method, given that the MA method provides one good way to achieve this relation. In addition, the element volume loss penalizes negative Jacobian determinants, which helps prevent element inversion, *i.e.*, mesh tangling.

Chamfer distance loss. Chamfer distance finds for all vertices in the first mesh, the nearest vertex in a second mesh and sums the square of these distances. It encourages the model to output a mesh which has a similar spatial distribution of vertices to that of the reference mesh. To achieve vertex-to-vertex alignment, the bidirectional Chamfer distance is utilized without additional sampling. It is defined as

$$L_{cd}(\theta) = \mathbb{E}_{(\xi, m_\xi; V_r, T_r) \in \mathcal{D}} \left[\frac{1}{|V_a|} \sum_{\mathbf{x}_i \in V_a} \min_{\mathbf{y}_j \in V_r} \|\mathbf{x}_i - \mathbf{y}_j\|_2 + \frac{1}{|V_r|} \sum_{\mathbf{y}_j \in V_r} \min_{\mathbf{x}_i \in V_a} \|\mathbf{x}_i - \mathbf{y}_j\|_2 \right]. \quad (7)$$

5 Experiments

5.1 Experiment setups.

Unlike existing approaches to mesh movement, such as [19, 20], which re-train their model on a case-by-case basis, our study aims to explore the universal applicability of the proposed UM2N. Therefore, all examples shown in this section are tested in a zero-shot generalization manner, *i.e.*, without re-training. In addition to PDE type, our training meshes are of small size typically comprising 500 vertices, while our test meshes have far more vertices: Advection (2, 052 vertices), Cylinder (4, 993 vertices), Tsunami (8, 117 vertices), which also tests the model’s scalability. There are examples with more complicated settings, please refer to Appx. E.

Training. The training dataset consists of 600 randomly generated generic solution fields and original meshes with 463 and 513 vertices. The model is trained using the Adam optimizer. The training and experiments are performed on an Nvidia RTX 3090 GPU.

Metrics. The main metric for evaluating mesh movement quality is underlying PDE solution error reduction (ER) ratio. Here PDE error is approximated by the difference between the solutions obtained on a coarse mesh and accurate solutions which are obtained on a very high resolution mesh. The PDE error reduction ratio is the difference between the PDE error from an original mesh and that from an adapted mesh. Another important metric for mesh movement is mesh tangling. Once mesh tangling happens, the mesh is invalid for a numerical solver, which breaks the simulation.

5.2 Benchmarking mesh movement.

Table 1: Quantitative results across mesh movement methods. ER indicates PDE error reduction ratio. “Fail” indicates that mesh tangling happens during the simulation.

Methods	Swirl		Cylinder		Helmholtz	
	ER (%) ↑	Time (ms) ↓	ER (%) ↑	Time (ms) ↓	ER (%) ↑	Time (ms) ↓
MA [36]	37.43	5615.78	Fail	-	17.95	3265.23
M2N [19]	14.15	25.01	Fail	-	16.32	4.62
UM2N (Ours)	35.93	17.66	61.29	19.25	17.08	4.86

In the following section, we benchmark the mesh movement using both non-learned Monge-Ampère (MA) mesh movement implemented in *Movement* [40] and learning-based baseline M2N [19] over three different scenarios to compare with UM2N. Table 1 shows the quantitative results in three different scenarios, Swirl, Cylinder, and Helmholtz. An additional qualitative result of Tōhoku Tsunami simulation, is provided to show our model can handle highly complex boundaries in real-world scenarios. Moreover, to give a fair comparison, we re-trained the model in M2N on our training dataset before evaluation.

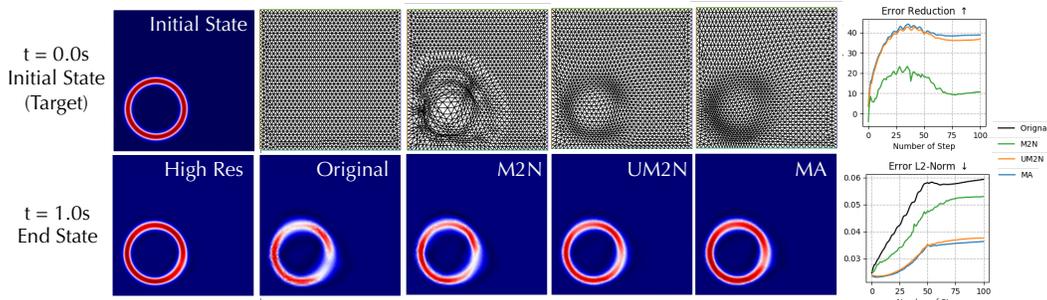


Figure 2: Results for swirl case. The plot shows the end state of the swirl simulation (bottom row) and corresponding meshes (top row). As the process is reversible, the end state should be consistent to the initial state. It can be observed that the solution obtained on high resolution almost recover the initial state. The proposed UM2N outperform M2N and achieve comparable performance to MA method. Quantitatively in the rightmost plot, the orange line (UM2N) and blue line (MA) best suppress error accumulation along timesteps.

Helmholtz. The Helmholtz case is a time-independent problem solving a 2D Helmholtz equation as in (A2). The source terms of the equation are generated use formula (A1). Our model achieves better error reduction than M2N and comparable to MA method in significantly less time.

Swirl. The swirl case considers a time-dependent pure advection equation. In this scenario, the initial tracer field takes a narrow ring shape and is advected within a swirling velocity field, see Eq. (A3) and (A4) in Appx. C. The velocity swirling direction is anti-clockwise in the first half of the simulation and clockwise for the second half. Given this setting, as the advection process is reversible, the final state should be consistent with the initial state. This constitutes a problem of significant multi-resolution complexity, as the extremely narrow ring requires a high resolution to accurately resolve the dynamics at the sharply defined interfaces of both inner and outer sides.

The results are shown in Fig. 2. As shown in Table 1, our UM2N shows an average 35.93% error reduction which clearly outperforms M2N's 14.15%. It is close to that of the full Monge-Ampère (MA) mesh movement approach which achieves an error reduction of 37.43%, but with significantly lower computational time.

Flow past cylinder. Simulation of flow past a cylinder in a channel is a classic and challenging multi-resolution case of computational fluid dynamics, solving the Navier-Stokes equations. The cylinder and the channel's top-bottom surfaces require a high resolution mesh to resolve the boundary layers [41]. Given appropriate Reynolds number, the well-known von Kármán vortex street phenomena can be observed within the wake flow which also benefits from higher resolution to resolve the time-varying evolution of vorticity structures. The drag C_D and lift C_L coefficients for the cylinder are other quantities of interest in this problem. Here we use a benchmark configuration from [42], please refer to Appx.C.3 for more details.

The simulation is run for 8,000 time steps of size 0.001 s, *i.e.*, 8 s physical time for the whole simulation. The original mesh is a uniform mesh with triangular elements and 4,993 vertices. The reference solutions are obtained on a high resolution mesh with 260,610 vertices.

We investigate the dynamics in the area around the cylinder and in the wake flow respectively. For the former, we use the measure of C_D as shown in Fig. 3. The results for C_L is shown in Fig. A10 in Appx. E. For the wake flow, we use cross-section probes at $x = [0.5, 1.0, 1.5, 2.0]$ to measure the intensity of vorticity spatially along the y -axis as shown in Fig. 4.

Qualitatively, it can be observed that the proposed method moves the mesh in a manner that captures the fluid dynamics for the whole domain in general (see upper part of Fig. 3). At the cylinder and top-bottom surface, the vertices are moved to increase the resolution in order to resolve the boundary layers as expected. Quantitatively, the drag C_D and lift C_L coefficients computed on our moved mesh show a 61.29% average error reduction compared to that of original mesh as shown in lower part of Fig. 3. The proposed UM2N approach not only improves the accuracy of these coefficients in magnitude but also prevents their periodic variation from phase shifting. For the wake flow, as shown in Fig. 4, we selected a single snapshot at $t = 4.0s$ for qualitative investigation. Our UM2N also shows an accuracy improvement for vorticity intensity on data sampled from all four cross-section probes comparing to the that on original mesh.

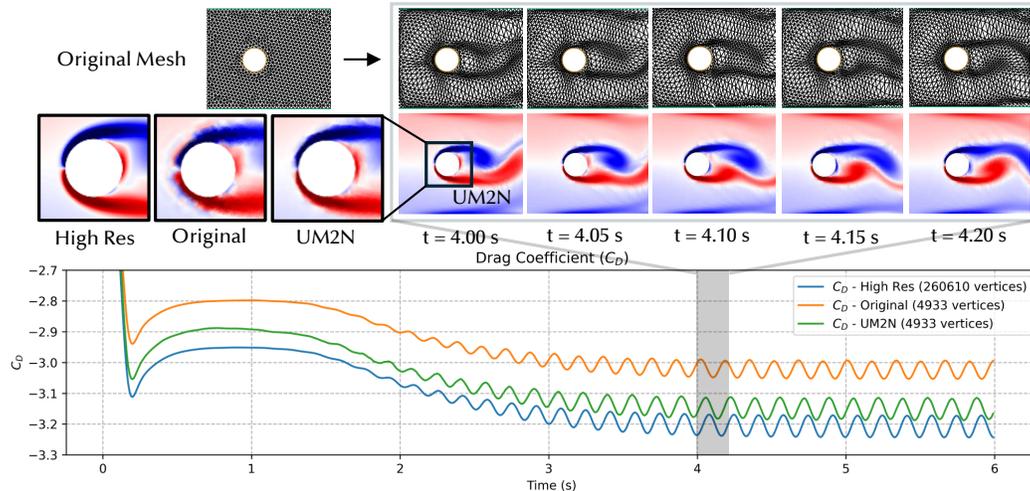


Figure 3: Results for the flow past cylinder case in time. In the upper part, we visualize 5 snapshots of adapted meshes and vorticity intensity from 4.00s to 4.20s. Please refer to the full video in the supplementary materials. It shows that our UM2N can output meshes which adapt to the dynamics around the cylinder. It can be observed from the zoom-in mini-plots at the top-left that UM2N reduces errors/noise in vorticity intensity compared to the original mesh. In the lower part, the blue, orange and green lines show the drag C_D coefficients obtained on a high resolution fixed mesh, UM2N adapted mesh and the original coarse mesh for the first 6 seconds (*i.e.*, 6000 steps). It can be observed that the UM2N output mesh improves the accuracy of C_D in magnitude and periodicity compared to the original mesh.

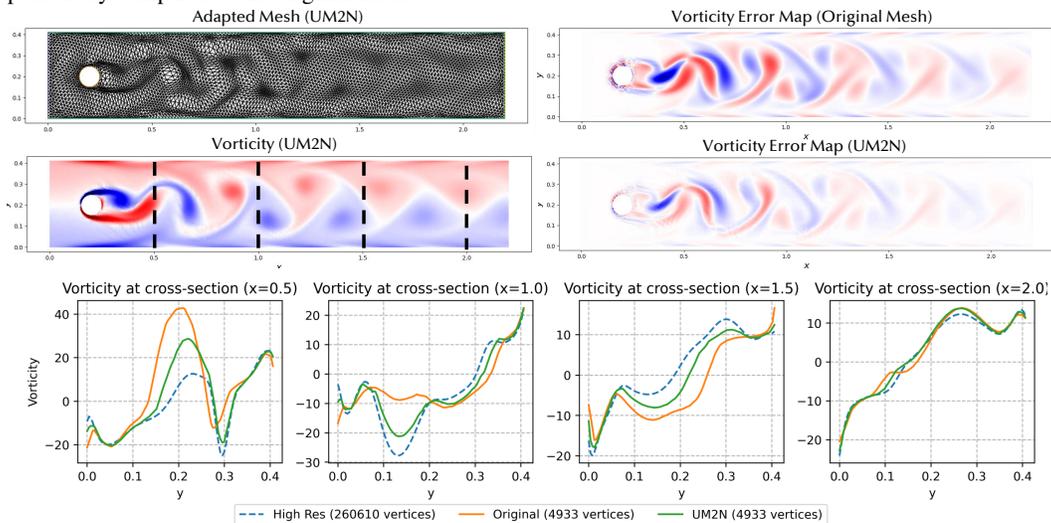


Figure 4: Analysis on wake flow vorticity for the cylinder case. The adapted mesh output by UM2N (see top-left part of the plot). Qualitatively, the error maps at the top-right part indicate that our UM2N reduce the vorticity difference to the high resolution mesh result in general. The quantitative comparison results along the probes are shown in the lower part of the figure. There are four cross-section probes place at $x = [0.5, 1.0, 1.5, 2.0]$ marked by the black dash line in the middle left plot. The green line (our UM2N) is much more consistent to the blue dash line (high resolution) compared to the orange line (original mesh), which further verifies that our method can reduce the PDE errors.

Note that we intended to compare our UM2N to M2N and the conventional full MA method. However, both MA and M2N fail *i.e.*, the solver diverges, within 30 time steps due to mesh tangling and thus provide insufficient data points for comparison. For comparisons with a non-tangled M2N case, please refer to the simplified laminar flow-past-cylinder example presented in Appendix D.1.

Tōhoku tsunami simulation. Here we present the case of a simulation of the 2011 Tōhoku Tsunami, to show that our methods can be applied to real world scenarios with boundaries describing highly complex geometries. The tsunami is simulated by solving the shallow water equations using the Thetis framework [43]. For details on the datasets and software used to set up the Tōhoku case study, see Appx. B.2. Note that the mesh (8, 117 vertices) has an arbitrarily irregular boundary. It can be

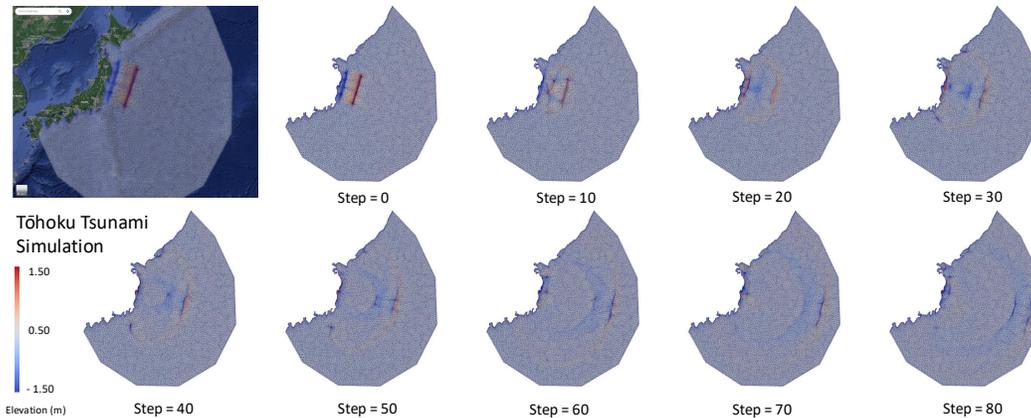


Figure 5: Qualitative results for the Tōhoku tsunami simulation. The boundary of the simulation is generated based on the real coastline data (see the mesh overlaid on the satellite map). We show 9 snapshots of the wave elevation from 80 steps of the tsunami simulation enhanced by our UM2N. In these visualizations, red and blue hues indicate wave elevations that are, respectively, above or below mean sea level. Regions exhibiting significant elevation magnitudes necessitate increased resolution to accurately resolve the underlying dynamics. Our method dynamically and robustly adjusts the mesh to increase resolution at the wave front, thereby effectively tracking its propagation.

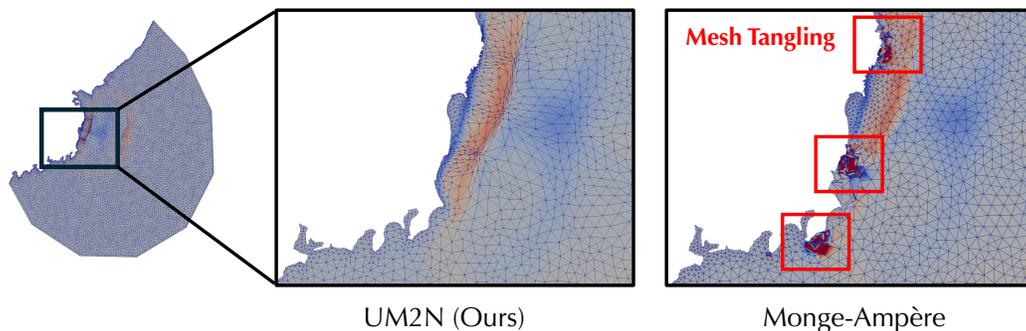


Figure 6: Handling complex boundaries. The Monge-Ampère method generates inverted elements, resulting in mesh tangling near the coastline as shown in the areas delineated by red rectangles. Our UM2N can smoothly handle the elements near the coastline without mesh tangling.

observed that the mesh moves in a manner that helps track the wave propagation with enhanced resolution, as shown in Fig. 5. Our model can also robustly handle the highly complex boundary without mesh tangling near the coastline. As a comparison, the conventional Monge-Ampère (MA) mesh movement approach struggles to handle the complex geometries, *i.e.*, fails to converge when solving the non-linear MA equation, resulting in highly tangled elements near the coastline as shown in Fig. 6. Therefore, our method not only provides a more efficient approach but also show advantages dealing with complex boundaries compared to the full MA mesh movement approach. Please see the full video of this case in the supplementary materials.

5.3 Tangled results of the UM2N model

There are in fact still cases where UM2N might fail. To investigate these, we perform additional experiments as shown in Figure 7. Starting from a rectangle, we gradually distort the geometry to a more and more non-convex shape, *i.e.*, the case becomes more challenging from top to bottom. The UM2N finally fails in the extreme case at sudden jumps in the boundary accompanied by a large variation in the required resolution as shown in the bottom row of Figure R1: a sudden constriction of the flow leads to tangling with UM2N in the two left corners of the channel. In less extreme cases, UM2N does produce a valid mesh, whereas the original MA method still fails (as seen in the second to bottom rows in Figure 7).

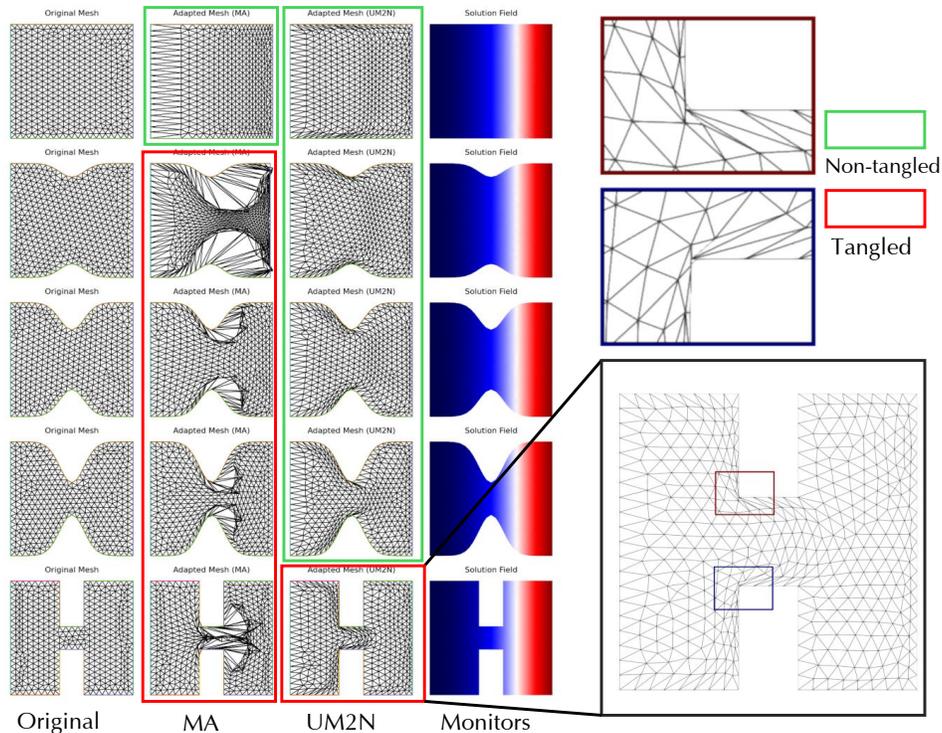


Figure 7: Tangled Case of UM2N and MA. Red/blue indicate high/low monitor values, *i.e.*, the mesh is expected to stretch towards the right direction. The case is more and more challenging from top to bottom.

5.4 Ablation study

Volume loss vs coordinate loss. We denote the UM2N variant trained with coordinate loss [19] as UM2N-coord. As shown in Table 2, UM2N-coord achieves comparable error reduction (31.21%) to the proposed UM2N (35.93%) on the Swirl case. However, it fails (*i.e.*, encounters mesh tangling) on the Cylinder case which has a more complex boundary geometry. This indicates that the element volume loss is mesh-tangling aware, *i.e.*, helps to prevent producing inverted elements.

Table 2: Ablation Study for design choices of learning monitors and element volume loss.

Method	Helmholtz ER (%) \uparrow	Swirl ER \uparrow (%)	Cylinder ER (%) \uparrow
UM2N-coord	16.64	31.21	Fail
UM2N-sol	13.52	-0.35	39.94
UM2N (Ours)	17.08	35.93	61.29

Monitors vs PDE solutions as inputs. We denote the UM2N variant using the PDE solution as input, instead of monitor function values, as UM2N-sol. The UM2N-sol is trained on solutions of the Helmholtz PDE. The UM2N-sol shows a much inferior performance (-0.35%) compared to UM2N (35.93%) on the Swirl case as well as on the Cylinder case. This indicates that learning mesh movement from monitors directly improves the generalisability of the model.

6 Conclusions

We introduce the Universal Mesh Movement Network (UM2N), which once trained on our generated PDE-independent dataset, can be applied in a non-intrusive, zero-shot manner to move meshes with different sizes and structures, for solvers applicable to different PDE types and boundary geometries. Our method demonstrates superior mesh movement results on Advection and Navier-Stokes PDE examples as well as a real-world tsunami simulation case. Our UM2N outperforms the existing learning based method and achieves comparable PDE error reduction to the far more costly Monge-Ampère (MA) PDE based approach. Our UM2N demonstrates effectiveness in cases with complex boundary geometries where existing learning and conventional MA based methods fail. Our method also shows significant acceleration compared to MA based mesh adaptation. See Appx. G for discussions on limitations and broader impacts.

Acknowledgment

Thanks to Jinghan Jia for testing UM2N performance on Nvidia A100.

References

- [1] Joseph G. Wallwork, Nicolas Barral, Stephan C. Kramer, David A. Ham, and Matthew D. Piggott. Goal-oriented error estimation and mesh adaptation for shallow water modelling. *SN Applied Sciences*, 2(6):1053, 2020. ISSN 2523-3971. doi: 10.1007/s42452-020-2745-9.
- [2] Pantelis R Vlachas, Georgios Arampatzis, Caroline Uhler, and Petros Koumoutsakos. Multi-scale simulations of complex systems by learning their effective dynamics. *Nature Machine Intelligence*, 4(4):359–366, 2022.
- [3] Chi Zhang, Massoud Rezavand, and Xiangyu Hu. A multi-resolution SPH method for fluid-structure interactions. *Journal of Computational Physics*, 429:110028, 2021. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2020.110028>.
- [4] Michael Athanasopoulos, Hassan Ugail, and Gabriela González Castro. Parametric design of aircraft geometry using partial differential equations. *Advances in Engineering Software*, 40(7):479–486, 2009.
- [5] The carbon footprint of computational research. *Nature Computational Science*, 3(8):659–659, 2023.
- [6] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [7] Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*, 2019.
- [8] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *International Conference on Learning Representations (ICLR)*, 2021.
- [9] Weizhang Huang and Robert D. Russell. *Adaptive Moving Mesh Methods*. Applied Mathematical Sciences. Springer Science+Business Media, LLC, 2 edition, 2011.
- [10] Zheyang Zhang, Yongxing Wang, Peter K. Jimack, and He Wang. MeshingNet: A new mesh generation method based on deep learning, 2020.
- [11] Jiachen Yang, Tarik Dzanic, Brenden K. Petersen, Jun Kudo, Ketan Mittal, Vladimir Z. Tomov, Jean-Sylvain Camier, Tuo Zhao, Hongyuan Zha, Tzanio V. Kolev, Robert W. Anderson, and Daniel Faissol. Reinforcement learning for adaptive mesh refinement. *CoRR*, abs/2103.01342, 2021. URL <https://arxiv.org/abs/2103.01342>.
- [12] Keefe Huang, Moritz Krügener, Alistair Brown, Friedrich Menhorn, Hans-Joachim Bungartz, and Dirk Hartmann. Machine learning-based optimal mesh generation in computational fluid dynamics. *arXiv preprint arXiv:2102.12923*, 2021.
- [13] Krzysztof J. Fidkowski and Guodong Chen. Metric-based, goal-oriented mesh adaptation using machine learning. *Journal of Computational Physics*, 426:109957, 2021. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2020.109957>. URL <https://www.sciencedirect.com/science/article/pii/S0021999120307312>.
- [14] Wu Tingfan, Liu Xuejun, An Wei, Zenghui Huang, and Lyu Hongqiang. A mesh optimization method using machine learning technique and variational mesh adaptation. *Chinese Journal of Aeronautics*, 2021.
- [15] Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W Battaglia. Learning mesh-based simulation with graph networks. *International Conference on Machine Learning (ICLR)*, 2021.

- [16] Joseph G Wallwork, Jingyi Lu, Mingrui Zhang, and Matthew D Piggott. E2N: error estimation networks for goal-oriented mesh adaptation. *arXiv preprint arXiv:2207.11233*, 2022.
- [17] Tailin Wu, Takashi Maruyama, Qingqing Zhao, Gordon Wetzstein, and Jure Leskovec. Learning controllable adaptive simulation for multi-resolution physics. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=PbfgkZ2HdbE>.
- [18] Corbin Foucart, Aaron Charous, and Pierre F.J. Lermusiaux. Deep reinforcement learning for adaptive mesh refinement. *Journal of Computational Physics*, 491:112381, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112381>. URL <https://www.sciencedirect.com/science/article/pii/S002199912300476X>.
- [19] Wenbin Song, Mingrui Zhang, Joseph G Wallwork, Junpeng Gao, Zheng Tian, Fanglei Sun, Matthew D Piggott, Junqing Chen, Zuoqiang Shi, Xiang Chen, and Jun Wang. M2N: Mesh movement networks for PDE solvers. *Advances in Neural Information Processing Systems*, 35, 2022.
- [20] Peiyan Hu, Yue Wang, and Zhi-Ming Ma. Better neural PDE solvers through data-free mesh movers. In *The Twelfth International Conference on Learning Representations*, 2024.
- [21] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [22] Shengze Cai, Zhiping Mao, Zhicheng Wang, Minglang Yin, and George Em Karniadakis. Physics-informed neural networks (pinns) for fluid mechanics: A review. *Acta Mechanica*, 37(12):1727–1738, 2021.
- [23] Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [24] Zijie Li, Kazem Meidani, and Amir Barati Farimani. Transformer for partial differential equations’ operator learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=EPPqt3uERT>.
- [25] Zijie Li, Dule Shu, and Amir Barati Farimani. Scalable transformer for PDE surrogate modeling. *Advances in neural information processing systems (NeurIPS)*, 36, 2023.
- [26] Zhongkai Hao, Chengyang Ying, Zhengyi Wang, Hang Su, Yinpeng Dong, Songming Liu, Ze Cheng, Jun Zhu, and Jian Song. Gnot: A general neural operator transformer for operator learning. *International Conference on Machine Learning (ICML)*, 2023.
- [27] Nils Thuerey, Konstantin Weissenow, Lukas Prantl, and Xiangyu Y. Hu. Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows. *AIAA Journal*, 2018. URL <https://api.semanticscholar.org/CorpusID:67855877>.
- [28] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. *International Conference on Machine Learning (ICML)*, 2020.
- [29] Johannes Brandstetter, Daniel Worrall, and Max Welling. Message passing neural PDE solvers. *International Conference on Learning Representations (ICLR)*, 2022.
- [30] Steeven JANNY, Aurélien Bènèteau, Madiha Nadri, Julie Digne, and Nicolas THOME. Eagle: Large-scale learning of turbulent fluid dynamics with mesh transformers. *International Conference on Learning Representations*, 2023.
- [31] Lukas Prantl, Benjamin Ummenhofer, Vladlen Koltun, and Nils Thuerey. Guaranteed conservation of momentum for learning particle-based fluid dynamics. In *Conference on Neural Information Processing Systems*, 2022.
- [32] Xu Han, Han Gao, Tobias Pfaff, Jian-Xun Wang, and Li-Ping Liu. Predicting physics in mesh-reduced space with temporal attention. *International Conference on Learning Representations*, 2022.

- [33] Jian Yu, Hongqiang Lyu, Ran Xu, Wenxuan Ouyang, and Xuejun Liu. Flow2Mesh: A flow-guided data-driven mesh adaptation framework. *Physics of Fluids*, 36(3):037124, 03 2024.
- [34] J.M. Nunez Rattia, J.R. Percival, S.J. Neethling, and M.D. Piggott. Modelling local scour near structures with combined mesh movement and mesh optimisation. *Journal of Computational Physics*, 375:1220–1237, 2018.
- [35] Chris J Budd, Weizhang Huang, and Robert D Russell. Adaptivity with moving grids. *Acta Numerica*, 18:111–241, 2009.
- [36] Andrew T. T. McRae, Colin J. Cotter, and Chris J. Budd. Optimal-transport-based mesh adaptivity on the plane and sphere using finite elements. *SIAM Journal on Scientific Computing*, 40(2):A1121–A1148, 2018. doi: 10.1137/16M1109515.
- [37] Mohamed M. Sulman, J.F. Williams, and Robert D. Russell. An efficient approach for the numerical solution of the Monge–Ampère equation. *Applied Numerical Mathematics*, 61(3): 298–307, 2011.
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems (NeurIPS)*, 30, 2017.
- [39] Shuhao Cao. Choose a transformer: Fourier or Galerkin. *Advances in neural information processing systems (NeurIPS)*, 34, 2021.
- [40] Joseph G. Wallwork, Stephan C. Kramer, Mingrui Zhang, and Davor Dundovic. Movement, May 2024. URL <https://doi.org/10.5281/zenodo.11239647>.
- [41] Hermann Schlichting and Klaus Gersten. *Boundary-Layer Theory*. Springer, Berlin Heidelberg, 9 edition, 2017. ISBN 978-3-662-52917-1. doi: 10.1007/978-3-662-52919-5_2.
- [42] M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher. *Benchmark Computations of Laminar Flow Around a Cylinder*, pages 547–566. Vieweg+Teubner Verlag, Wiesbaden, 1996. ISBN 978-3-322-89849-4. doi: 10.1007/978-3-322-89849-4_39. URL https://doi.org/10.1007/978-3-322-89849-4_39.
- [43] T. Kärnä, S. C. Kramer, L. Mitchell, D. A. Ham, M. D. Piggott, and A. M. Baptista. Thetis coastal ocean model: discontinuous Galerkin discretization for the three-dimensional hydrostatic equations. *Geoscientific Model Development*, 11(11):4359–4382, 2018. doi: 10.5194/gmd-11-4359-2018. URL <https://gmd.copernicus.org/articles/11/4359/2018/>.
- [44] Florian Rathgeber, David A Ham, Lawrence Mitchell, Michael Lange, Fabio Luporini, Andrew TT McRae, Gheorghe-Teodor Bercea, Graham R Markall, and Paul HJ Kelly. Firedrake: automating the finite element method by composing abstractions. *ACM Transactions on Mathematical Software (TOMS)*, 43(3):1–27, 2016.
- [45] Martin S Alnæs, Anders Logg, Kristian B Ølgaard, Marie E Rognes, and Garth N Wells. Unified form language: A domain-specific language for weak formulations of partial differential equations. *ACM Transactions on Mathematical Software (TOMS)*, 40(2):1–37, 2014.
- [46] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Steven Benson, Jed Brown, Peter Brune, Kris Buschelman, Emil Constantinescu, Lisandro Dalcin, Alp Dener, Victor Eijkhout, Jacob Faibussowitsch, William D. Gropp, Václav Hapla, Tobin Isaac, Pierre Jolivet, Dmitry Karpeev, Dinesh Kaushik, Matthew G. Knepley, Fande Kong, Scott Kruger, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Lawrence Mitchell, Todd Munson, Jose E. Roman, Karl Rupp, Patrick Sanan, Jason Sarich, Barry F. Smith, Stefano Zampini, Hong Zhang, Hong Zhang, and Junchao Zhang. PETSc/TAO users manual. Technical Report ANL-21/39 - Revision 3.21, Argonne National Laboratory, 2024.
- [47] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009. doi: <https://doi.org/10.1002/nme.2579>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.2579>.

- [48] Alexandros Avdis, Adam S. Candy, Jon Hill, Stephan C. Kramer, and Matthew D. Piggott. Efficient unstructured mesh generation for marine renewable energy applications. *Renewable Energy*, 116:842–856, 2018. ISSN 0960-1481. doi: <https://doi.org/10.1016/j.renene.2017.09.058>. URL <https://www.sciencedirect.com/science/article/pii/S0960148117309205>.
- [49] Mingrui Zhang, Chunyang Wang, Stephan C. Kramer, and Joseph G. Wallwork. Um2n, October 2024. URL <https://doi.org/10.5281/zenodo.14009265>.
- [50] Pål Wessel and Walter H. F. Smith. A global, self-consistent, hierarchical, high-resolution shoreline database. *Journal of Geophysical Research: Solid Earth*, 101(B4):8741–8743, 1996. doi: <https://doi.org/10.1029/96JB00104>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/96JB00104>.
- [51] Christopher Amante and Barry W Eakins. ETOPO1 arc-minute global relief model: procedures, data sources and analysis, 2022. Accessed: 2022-03-22.
- [52] Yoshimitsu Okada. Surface deformation due to shear and tensile faults in a half-space. *Bulletin of the seismological society of America*, 75(4):1135–1154, 1985.

Appendix

We include additional experimental details and results in the appendix. Detailed information on the software used is provided in Appendix A. A description of the dataset generation process is presented in Appendix B. We provide the details of experimental settings of cases in the paper in Appendix C. To provide a more holistic analysis for the proposed UM2N, we include evaluations across various test cases, performance benchmarks, more ablation studies, and assessments on low-quality initial meshes in Appendix D. To demonstrate the the universal ability of the UM2N, additional experiments conducted under diverse settings are detailed in Appendix E. We introduce the mesh movement-enhanced simulation, illustrating the non-intrusive integration of UM2N into a numerical simulation pipeline in Appendix F. A discussion about the limitations and broader impacts of our approach is provided in Appendix G.

A Software used in the paper.

In this paper, PDEs are solved using finite element methods with *Firedrake* [44]. *Firedrake* is written in Python, uses the *Unified Form Language (UFL)* [45] domain-specific language to represent finite element forms, and automatically generates C kernels for computational efficiency. *Firedrake* uses *PETSc* [46] for solving linear and nonlinear equations that result from discretisation, as well as for its underlying unstructured mesh representation.

Original meshes (*i.e.*, meshes before adaptation) are generated by either *Firedrake* or *Gmsh* [47]. *Gmsh* is a meshing software especially designed for finite element methods, providing access to several meshing algorithms. In the case of the Tōhoku tsunami simulation, the mesh of the Japan Sea is generated using *qmesh* [48] – a meshing tool specifically designed for coastal ocean modelling applications, which manages GIS (Geoscientific Information System)-based datasets to construct a valid input geometry used by *Gmsh* to generate the mesh.

The conventional mesh movement strategy used to generate training data is implemented in *Movement* [40], which is itself written in UFL and *Firedrake*. *Movement* implements several mesh movement approaches, including two based on solutions of Monge-Ampère type equations. In this paper, we use the ‘relaxation’ approach, which solves the nonlinear MA equation in an iterative fashion by introducing a pseudo-timestep (see [36] for details).

The code used to produce the results presented in this paper is publicly available at <https://github.com/mesh-adaptation/UM2N>. (Release [49].) Code documentation is hosted at <https://mesh-adaptation.github.io>.

B Dataset description.

B.1 Dataset generation pipeline.

We developed a pipeline for generating PDE-independent training data and further research. Through our pipeline, generic fields are generated and can be utilized as solutions of any PDE. The monitor values computed from the generic field are used to guide the MA-based method to perform mesh adaption on the original mesh.

We stack N randomly generated Gaussian distributions to form the generic field u . Concretely, they can be generated by the formula:

$$u = \sum_{k=1}^N \exp\left(\frac{(x - \mu_x)^2}{\sigma_x^2} + \frac{(y - \mu_y)^2}{\sigma_y^2}\right), \quad (\text{A1})$$

where μ_x and μ_y denote the means of the distribution along two orthogonal directions in a 2D domain, while σ_x and σ_y represent the respective standard deviations.

All training data are generated within a 2-D unit square domain. Equivalent size $n \times n$ meshes ($n \in \{18, 20\}$) are generated for model training. Here we show samples of the generated fields in Fig. A1.

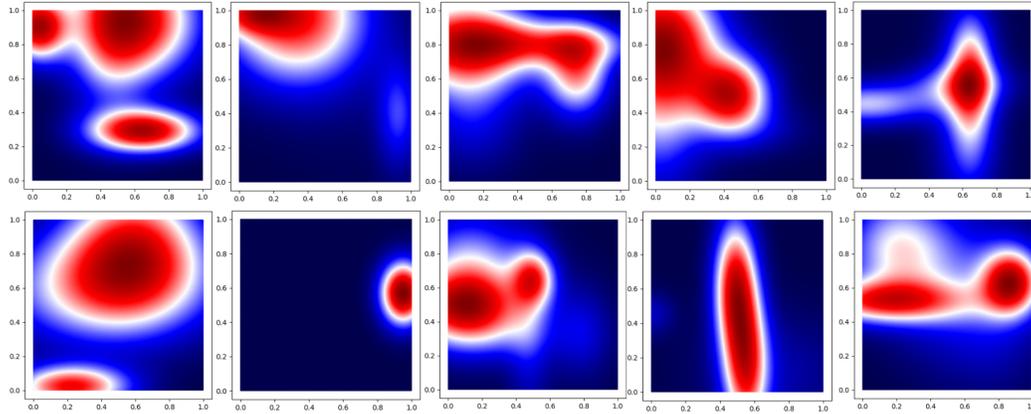


Figure A1: Samples of generated generic fields.

B.2 Evaluation data for the Tōhoku tsunami simulation.

As mentioned in Sec. A, the mesh of the Japan Sea is generated using qmesh [48]. For this, we provide qmesh with coastline data from [50]. Bathymetry (sea bed topography) data is taken from the ETOPO1 dataset [51]. A highly idealised tsunami source condition is used, comprised of a sum of Okada functions [52].

C Details of experimental settings.

C.1 Helmholtz.

The Helmholtz equation used in the paper is defined as:

$$-\nabla^2 u + u = f, \quad \nabla u \cdot \vec{n} = 0 \quad \text{on } \Gamma, \quad (\text{A2})$$

where f is the source term and Γ denotes the boundary.

C.2 Swirl.

We use a similar setting to the unsteady 2D linear advection case in [18].

Initial Condition.

$$u_0 = \exp\left(-\frac{1}{2\sigma^2}(\sqrt{(x-x_0)^2 + (y-y_0)^2} - r_0)^2\right). \quad (\text{A3})$$

The initial condition defines a ring centered at (x_0, y_0) with an inner radius r_0 . The thickness of the ring is 3σ . In this paper, $r_0 = 0.2$, $(x_0, y_0) = (0.30, 0.30)$, and $\sigma = 0.05/3$.

Velocity Field.

$$\mathbf{c}(x, y, t) = \left(\frac{9}{10}a(t)\sin^2(\pi x)\sin(2\pi y), -\frac{9}{10}a(t)\sin^2(\pi y)\sin(2\pi x)\right). \quad (\text{A4})$$

In this paper, $0 < t < 1$, if $t < 0.5$ then $a(t) = 1$, if $0.5 \leq t < 1$, $a(t) = -1$, *i.e.*, the velocity direction is reversed at $t = 0.5$

Monitor Function.

$$m(\mathbf{x}) = 1 + \max\left(\alpha \frac{\|H(u)\|}{\max\|H(u)\|}, \beta \frac{\|G(u)\|}{\max\|G(u)\|}\right), \quad (\text{A5})$$

where we set $\alpha = 5$ and $\beta = 10$. The $\|H(u)\|$ and $\|G(u)\|$ are Hessian norm and gradient norm respectively.

Table A2: Performance benchmark on test case shown in Figure. A3.

Hardware	MA CPU	UM2N CPU	UM2N GPU (Nvida RTX 3090)	UM2N GPU (Nvida A100)
Avg. inference time (s)	114.3	1.113	0.0487	0.0342
Speedup	-	~ 102×	~ 2347×	~ 3342×

C.3 Flow past a cylinder.

Setting The Reynolds number is 100, a quadratic velocity profile is imposed at the left boundary as inflow condition. At the right boundary a zero pressure outflow condition is applied. The cylinder and top-bottom surface are imposed with non-slip boundaries.

D More Evaluations.

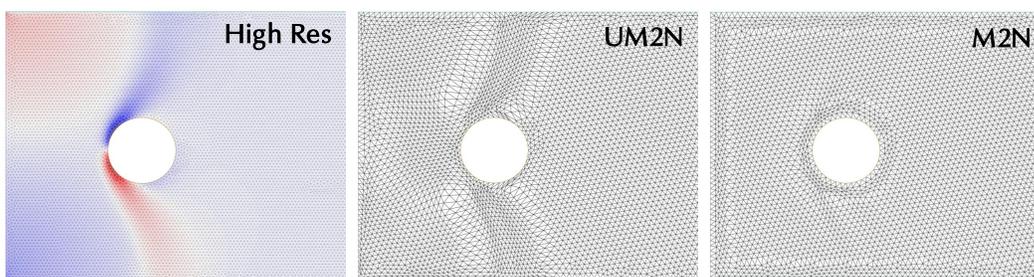


Figure A2: Non-tangled results of M2N on flow-past-a-cylinder case.

D.1 Non-tangled results of other methods on the flow-past-cylinder.

We additionally perform a non-tangled flow-past-cylinder shown in Fig. A2. We simplified the setting: lower the Reynolds number by reducing the inflow velocity and set the top and bottom slip boundary, which finally makes it a laminar flow. As shown in Fig A2, both the UM2N and M2N give non-tangled meshes and it can be observed that UM2N better captures the dynamics *i.e.*, adapts to the PDE solution of stable state compared to the High Res solution. The quantitative results shown in the Table A1 also indicate that UM2N perform better than M2N regarding to error reduction.

Table A1: Measure between UM2N and M2N model output mesh on the flow-past-a-cylinder in Fig. A2.

Property	UM2N	M2N
	ER ↑ (%)	ER ↑ (%)
Vorticity	27.43 %	10.22 %

D.2 Performance benchmark.

To evaluate the runtime performance of UM2N, we constructed a benchmark case based on the Swirl scenario, consisting of 11,833 vertices and 23,668 triangular elements. We extended the one-ring setup to a nine-ring configuration. In this benchmark, we compare the UM2N output mesh and the reference mesh generated by the MA method, as well as the solutions obtained on these meshes, against those computed on high-resolution and original meshes as shown in Fig. A3.

We tested the inference time of UM2N on two GPUs (Nvidia GeForce RTX 3090 and Nvidia A100) and compared its performance to the MA method. For a fair comparison, we also included CPU benchmarks, as the current state-of-the-art MA implementation is limited to CPU execution. The CPU used for testing was an 11th Gen Intel(R) Core(TM) i9-11900K @ 3.50GHz. The performance results are presented in Table A2.

The results indicate that UM2N achieved more than 2000x speedup on the Nvidia RTX 3090 and over 3000x on the Nvidia A100, compared to the MA method. Even on the CPU, UM2N demonstrated an approximate 100x speedup. In summary, UM2N provides comparable solution accuracy while delivering a significant performance improvement.

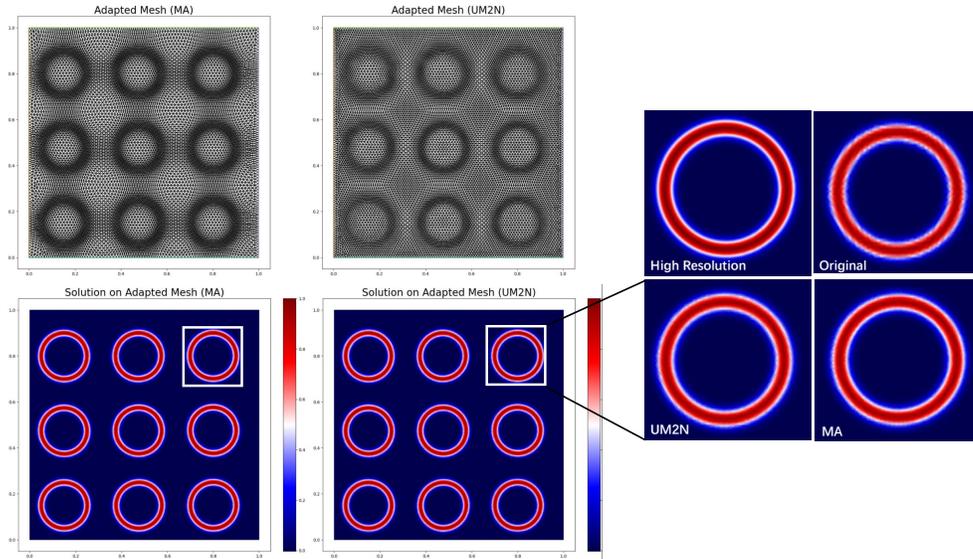


Figure A3: Visualization of UM2N output mesh and reference mesh from MA. Comparison between solutions obtained on high resolution, original, UM2N and MA meshes.

D.3 Boundary layer case.

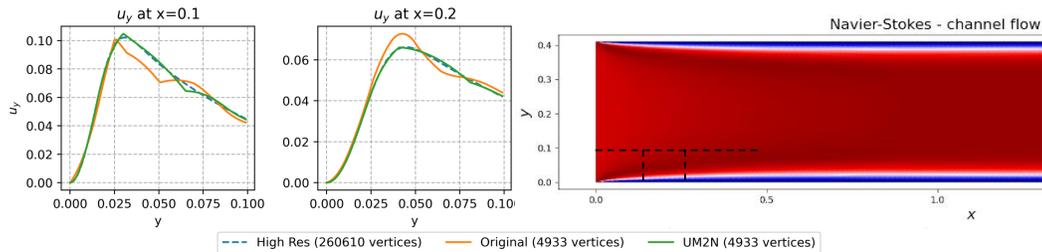


Figure A4: Boundary Layer. The two plots at left visualize the u_y from boundary to $y = 0.1$. It can be observed that UM2N improves the accuracy of the velocity obtained on uniform mesh compared to High Res solution.

In our flow past the cylinder case, at the top and bottom we imposed non-slip boundary conditions which already promotes a boundary layer similar to the suggested flow over a flat plate. For a clearer illustration, we have now additionally conducted a flow past two parallel plates experiment. We observe that the velocity parallel to the boundary (u_y) obtained from UM2N moved mesh better aligns with the high resolution results compared to the results on uniform mesh as shown in the Fig. A4.

D.4 Poor quality initial mesh.

We include an additional experiment starting from a poor quality initial mesh as shown in Fig. A5. The input mesh has highly anisotropic elements and several vertices with valence 6. This can certainly be considered a “low quality” mesh. Consider the monitor function to be the L^2 -norm of the recovered gradient of the solution of an anisotropic Helmholtz problem. With this monitor function, we tried applying conventional MA solvers and found that both the quasi-Newton and relaxation approaches failed to converge and/or resulted in tangled meshes, even though the elements are already aligned with the anisotropy of the Helmholtz solution. The UM2N approach, however, was able to successfully apply mesh movement without tangling. We would also like to clarify that our work targeting mesh movement that dynamically adapts to a PDE solution, which is not a replacement for mesh generation in general, so we would always expect a reasonable input mesh when applying our model.

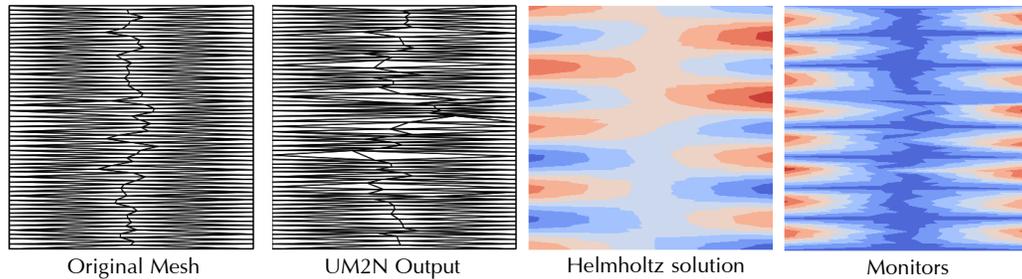
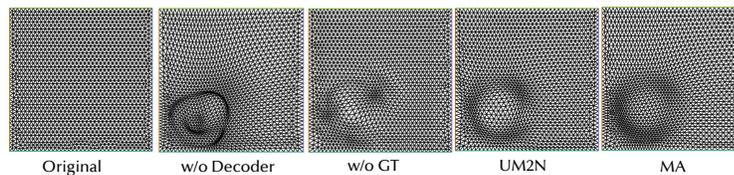


Figure A5: A very poor initial mesh case. UM2N perform mesh movement for an anisotropic Helmholtz problem.

Method	Swirl ER ↑ (%)
UM2N-w/o-Decoder	10.68
UM2N-w/o-GT	12.35
UM2N (Ours)	35.93
MA (reference)	37.43

(a) Error reduction of UM2N variants on Swirl case.



(b) Moved from UM2N variants and reference mesh from MA on Swirl case.

Figure A6: Ablation of network components. Full UM2N gives the best error reduction.

D.5 Further Ablation studies.

We perform an additional ablation study of our network architecture on the Swirl case. We construct two models, which replace the GAT Decoder with one layer GAT denoted as UM2N-w/o-Decoder and remove the graph transformer denoted as UM2N-w/o-GT. The results are shown in the Table A6a. It can be observed that both variants give worse results compared to the full UM2N. A visualization is also shown in Fig. A6b. It can be observed that without the Decoder, the model distorts the shape of the ring, i.e., missing relative information between vertices; without the graph transformer, the model fails to capture the details of the ring shape.

E Additional experimental results.

Flow past multiple cylinders. The Fig. A7 shows a challenge scenario with 5 cylinders in the domain. The proposed UM2N can move mesh capturing the complex dynamics without mesh tangling.

Flow past V-formation cylinders. The Fig. A8 shows a bio-inspired V-formation setting. The proposed UM2N can move mesh capturing the wake flows.

Flow past cylinders and squares. The Fig. A9 shows the proposed UM2N can handle another challenging scenario with a mixture of cylinders and squares of different sizes.

Lift coefficient (C_L) of flow past cylinder. In Fig. A10, we visualized the lift coefficient over timestep across different methods.

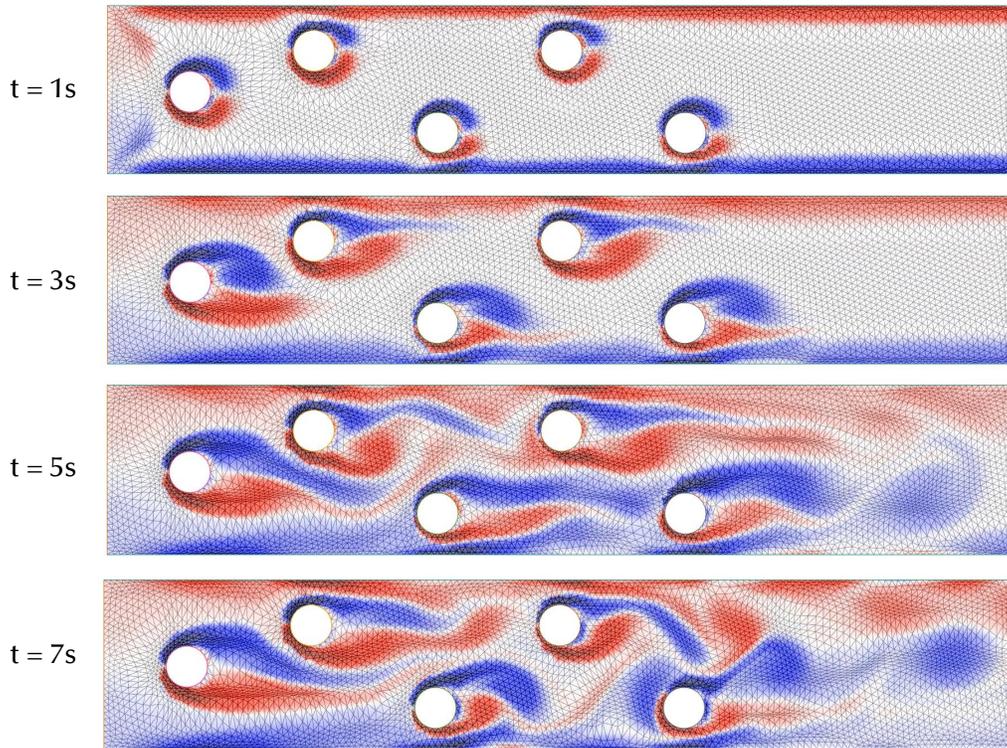


Figure A7: Multiple Cylinders

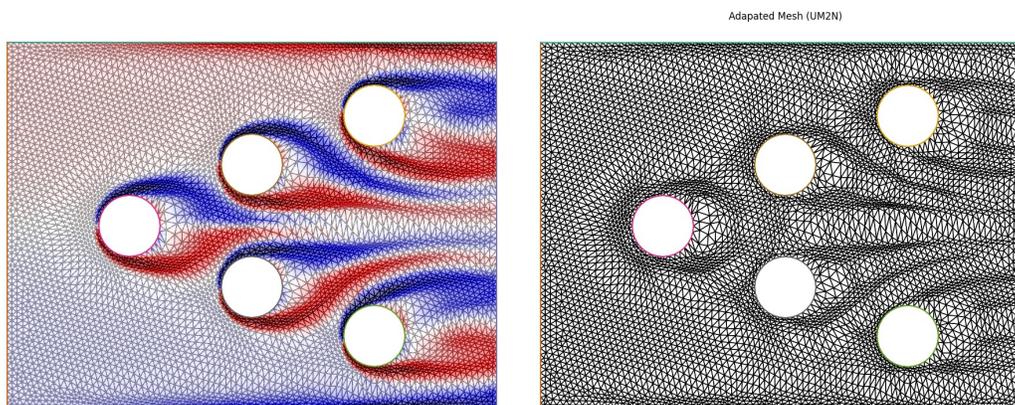


Figure A8: V-formation

F Mesh movement enhanced simulation.

The Mesh Movement Enhanced Simulation Algorithm, as outlined in Alg. 1, integrates our mesh adaption strategies into the simulation process to enhance the quality and accuracy of computational solutions. The proposed MMES procedure can be adapted into almost any time-dependent numerical simulation processes non-intrusively.

For a given time-dependent PDE \mathcal{P} and an initial mesh ξ_{init} , the algorithm evaluates the monitor function m based on the solution field of the PDE at each time step t . The monitor value m is then used by our proposed model to guide mesh movement. For robustness, a mesh integrity check is performed before time stepping (a tangled mesh will interrupt the procedure).

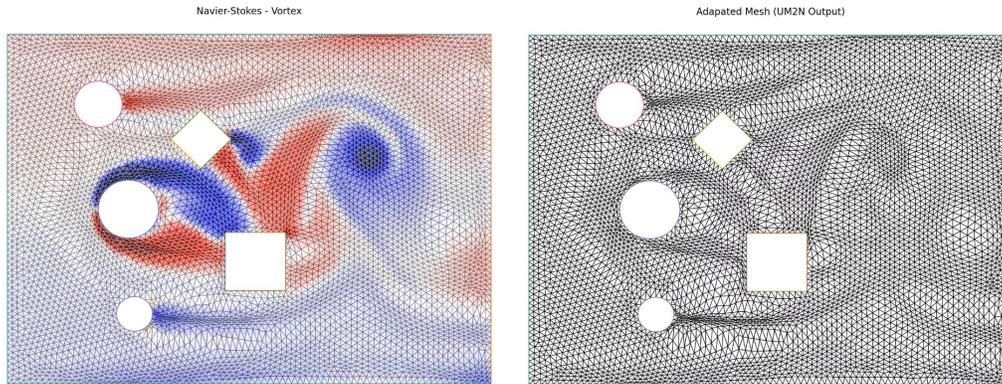


Figure A9: Cylinder and Square

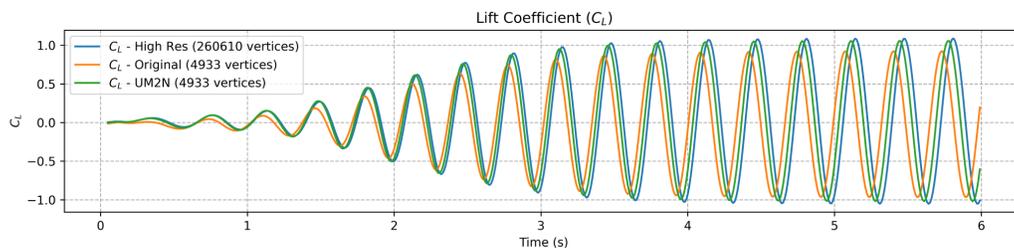


Figure A10: Results of C_L coefficients. The blue, orange and green lines show the drag C_L coefficients obtained on a high resolution fixed mesh, UM2N adapted mesh and the original coarse mesh for the first 6 seconds (i.e., 6000 steps). It can be observed that the UM2N output mesh improves the accuracy of C_L in magnitude and periodicity compared to the original mesh.

G Limitations and broader impacts.

Limitations. 1) The GAT to an extent constrains the mesh movement to lie in a restricted range, which may be less effective in scenarios requiring substantial large deformation within a single iteration of the mesh mover. 2) There is no theoretical guarantee that mesh tangling will be completely prevented. 3) The model's best performance relies on a manually chosen monitor function, though a gradient based monitor function can yield reasonable results. 3) Given a low-quality or extremely coarse original mesh, the efficacy of the proposed method may be constrained by the inherent

Algorithm 1 Mesh Movement Enhanced Simulation

```

1: procedure SIMULATION( $\mathcal{P}$ ,  $\xi_{\text{init}}$ ,  $T$ )
2:    $\xi_t \leftarrow \xi_{\text{init}}$ 
3:   while  $t < T$  do
4:      $u \leftarrow \text{PDESolver}(\mathcal{P}, \xi_t)$ 
5:      $m_{\xi} \leftarrow \text{MonitorFunc}(u, \xi_t)$ 
6:      $\xi_{t+1} \leftarrow \text{UM2N}(\xi_t, m_{\xi})$ 
7:      $\text{MeshIsValid} \leftarrow \text{CheckMeshNotTangle}(\xi_{t+1})$ 
8:     if  $\text{MeshIsValid}$  then
9:        $\xi_t \leftarrow \xi_{t+1}$ 
10:    else
11:      Break
12:    end if
13:     $\text{TimeStepUpdate}()$ 
14:  end while
15: end procedure

```

limitation on the degree of freedom count, as it redistributes rather than increases the number of mesh vertices. A promising future direction could be exploring the integration of mesh movement with *h*-adaptation, potentially developing a hybrid *hr*-adaptation approach that combines the advantages of both techniques.

Broader Impacts. This research contributes to the field of AI for physical simulation by introducing an advanced learning-based mesh movement method that enhances the accuracy and efficiency of numerical simulations. The primary application of this method lies in areas requiring high-fidelity simulations of physical phenomena, such as geophysics, renewable energy, climate modeling, and aerodynamics. This method can aid in better weather or hazards forecasting, more efficient energy harvesting, fast prototyping of aircraft design etc.

NeurIPS 2024 Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Please see abstract and introduction

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Please see section G

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Please see the section 3

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We would like to release the code for reproducibility. Please refer to Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We would like to release the code for reproducibility. Please refer to Appendix A.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please see section 5.1 and appendix C

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We perform the experiments multiple times and compute the average.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Please see section 5.1

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research in the paper conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Please see section G

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Please see the appendix A

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We would provide a README alongside the assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.