GEMOVE: UNSUPERVISED GEOMETRIC MESH MOVEMENT BEYOND DENSITY

Anonymous authors
Paper under double-blind review

ABSTRACT

Mesh movement adaptively relocates mesh elements under a fixed computational budget to achieve an accuracy-efficiency tradeoff in scientific simulations. Existing methods largely rely on costly theoretical supervision, whereas unsupervised approaches offer a promising alternative with zero-shot generalizability across diverse scenarios. However, under limited unsupervised constraints, *local density equidistribution* is prone to physically inconsistent movements and poorly coordinated over-allocation. To address these issues, we propose GEMOVE, which enforces *global geometric equidistribution* in a unified mapped space, thereby deriving vectorized movement indicators for each mesh element. Furthermore, GEMOVE models the mesh topology as a spring system, enabling global coordination through Tutte equilibrium. Extensive experiments across four PDE families demonstrate that GEMOVE achieves an average error reduction of 24.62% , with strong zero-shot generalization to unseen equations and domain geometries. Beyond simulation benchmarks, GEMOVE effectively adapts to real-world ocean systems, driving mesh movements in faithful response to multiscale spatial structures and long-term temporal variability. The unsupervised model learned from ocean systems further generalizes across the global Earth system, demonstrating broad scientific applicability across 54 regions and 16 variables. Code is available at `https://anonymous.4open.science/r/GEMove/`.

1 INTRODUCTION

Meshes play a fundamental role in scientific computing by discretizing complex physical domains and fields (Pfaff et al., 2021; Yu et al., 2024; 2025; Garnier et al., 2025; Zamani et al., 2026; Feng et al., 2026). However, distributing a limited number of mesh vertices remains challenging because physical fields often exhibit highly nonuniform spatial structures (Le et al., 2024). Mesh movement (Budd et al., 2009) addresses this challenge under a fixed vertex budget by relocating vertices according to physical variations, concentrating resolution in regions with strong variations while allocating less resolution to smooth regions. This adaptive redistribution enables a mesh to better resolve heterogeneous physical structures under the same vertex budget (Kang et al., 2025).

Traditional mesh movement methods, such as the Monge–Ampère (MA) method (McRae et al., 2018), rely on theoretical calculations that require solving auxiliary PDEs at substantial computational cost. To reduce this burden, learning-based methods train neural networks to approximate the calculation process. M2N (Song et al., 2022) and UM2N (Zhang et al., 2024) learn from MA-generated pre-adapted meshes as theoretical supervision, enabling efficient mesh movement at inference. However, such theoretical supervision still requires costly target meshes in advance, limiting training scalability and generalization across diverse scientific scenarios.

Therefore, unsupervised mesh movement provides a promising alternative by removing the need for pre-generated target meshes. UGM2N (Wang et al., 2025) explores this direction through *local density equidistribution*, using a scalar mesh density ρ to allocate more resolution to regions with stronger physical variation. However, scalar density constrains only how much resolution is allocated, while leaving element shape and orientation undetermined (Huang, 2005). As illustrated in Fig. 1, geometrically different mesh configurations can satisfy the same density condition, i.e., $\rho_K|K| = 1$. This geometric indeterminacy shows that scalar density equidistribution alone provides insufficient constraints for determining mesh movement. See Appendix A for a detailed analysis.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

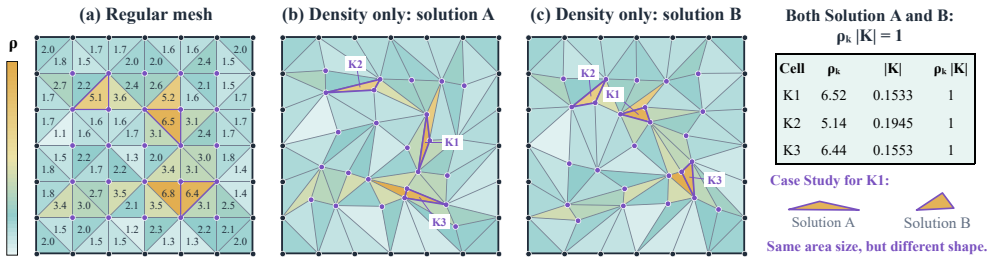

Figure 1: Under-constrained mesh geometry under scalar density equidistribution. Different mesh configurations satisfy the same density condition $\rho_K|K| = 1$, while their shapes and orientations remain undetermined.

This under-constrained formulation leads to two key limitations. **(i) Physically inconsistent movements** arise because scalar density provides insufficient geometric information to determine vertex relocation, especially for anisotropic physical structures. **(ii) Insufficient global coordination** arises because locally determined movements neglect interactions across the mesh, leading to conflicting movements and over-allocation of mesh resolution. Therefore, unsupervised mesh movement requires stronger geometric constraints for local movement and global coordination across the mesh.

To address these limitations, we propose GEMOVE, which moves beyond *local density equidistribution* through *global geometric equidistribution* in a unified mapped space. Instead of constraining only element density, GEMOVE drives mesh elements toward equal size and equilateral geometry in a new mapped space, providing vectorized movement indicators that characterize the magnitude and direction of local adaptation. To coordinate these local geometric demands, GEMOVE models the mesh topology as a spring system and jointly determines all interior vertices through Tutte equilibrium (Tutte, 1963). Under standard Tutte conditions, the resulting equilibrium guarantees inversion-free mesh movement while globally coordinating vertex relocation.

We evaluate GEMOVE from controlled PDE benchmarks to real-world Earth systems. Controlled experiments across Poisson, Helmholtz, Swirl, and Burgers equations evaluate error reduction and generalization across PDE families, while unseen polygonal domains further test adaptation across domain geometries. We then evaluate spatial adaptation to multiscale anisotropic structures in the Kuroshio Current system and temporal robustness under long-term variability in ENSO, followed by direct transfer from these local ocean systems to diverse global Earth-system regions and variables.

Our main contributions are summarized as follows:

1. We formulate unsupervised mesh movement as *global geometric equidistribution* without requiring pre-adapted target meshes, moving beyond scalar density allocation to jointly constrain element size and directional geometry in a unified mapped space.
2. We develop a global coordination mechanism based on Tutte equilibrium, modeling the mesh topology as a spring system to jointly determine all interior vertex relocations and coordinate local geometric demands across the mesh.
3. GEMOVE achieves a 24.62% average error reduction across four PDE families, generalizes zero-shot to unseen equations and domain geometries, adapts to multiscale spatial structures and long-term temporal variability in real-world ocean systems, and further transfers to 54 global Earth-system regions and 16 physical variables.

2 Related Work

Learning Mesh Movement with Theoretical Supervision. Learning-based mesh movement reduces the repeated cost of conventional mesh adaptation by learning mesh relocation. M2N (Song et al., 2022) introduces an end-to-end learning framework that approximates mesh relocation from target meshes generated by the Monge–Ampère (MA) method, substantially reducing the online cost of mesh adaptation. UM2N (Zhang et al., 2024) further improves generalizability, enabling mesh movement across different PDEs and mesh resolutions. However, both methods rely on computationally expensive MA-generated pre-adapted meshes for supervision. Other approaches embed mesh adaptation into PDE-solving pipelines, using numerical formulations or PDE-solution objec-

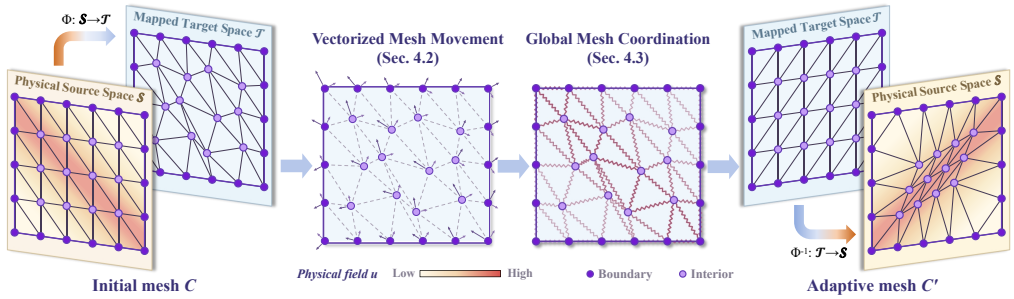

Figure 2: Overview of GEMOVE. Global geometric equidistribution determines vectorized mesh movement in the mapped space, while Tutte equilibrium globally coordinates vertex relocation to obtain the adaptive mesh in the physical space.

tives to guide mesh relocation (Hu et al., 2024; Omella & Pardo, 2024; Rowbottom et al., 2025). Overall, such theoretical supervision limits training scalability and broader generalization.

Unsupervised Mesh Movement. Unsupervised method removes the dependence on both pre-generated target meshes and solver-derived supervision, but remains unexplored. UGM2N (Wang et al., 2025) explores this direction through an M-uniform loss that enforces local density equidistribution. However, scalar density constrains only the amount of local resolution, leaving mesh geometry fundamentally under-constrained. This under-constrained formulation makes directional relocation ambiguous and provides limited coordination among local movements. In contrast, GEMOVE moves beyond local density equidistribution through global geometric equidistribution, combining geometric constraints for local movement with global coordination across the mesh.

3 PRELIMINARY

Notation. Let $\mathcal{S} \subset \mathbb{R}^2$ denote the physical space. An initial mesh in $\mathcal{S}$ is represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_i\}_{i=1}^N$ denotes the set of vertices and $\mathcal{E}$ denotes the fixed mesh connectivity. Each vertex v_i is associated with a coordinate $\mathbf{c}_i = (x_i, y_i)^\top \in \mathbb{R}^2$ and a physical feature $\mathbf{u}_i \in \mathbb{R}^{d_u}$. The vertex coordinates and physical features are collected as $\mathbf{C} \in \mathbb{R}^{N\times 2}$ and $\mathbf{U} \in \mathbb{R}^{N\times d_u}$, respectively. We denote the relocated vertex coordinates by $\mathbf{C}' \in \mathbb{R}^{N\times 2}$, while the mesh connectivity $\mathcal{E}$ remains unchanged during mesh movement.

Problem Formulation. In an unsupervised setting without pre-generated adaptive meshes as learning targets, we formulate mesh movement as an end-to-end relocation function $\mathcal{F}_\Theta$ that maps the initial mesh coordinates, physical features, and mesh connectivity to the relocated coordinates:

$$\mathbf{C}' = \mathcal{F}_\Theta(\mathbf{C}, \mathbf{U}, \mathcal{E}). \tag{1}$$

The relocated mesh aims to conform to the spatial structure of the physical field, allocating resolution according to both the magnitude and direction of local variations while maintaining global coordination across the mesh.

4 METHODOLOGY

4.1 GEMOVE OVERVIEW

As shown in Fig. 2, GEMOVE represents local adaptation demands through vectorized movement indicators in a unified mapped space (Sec. 4.2) and realizes these demands through globally coupled mesh relocation under Tutte equilibrium (Sec. 4.3). This formulation connects local geometric representation with mesh-wide coordination, translating heterogeneous physical-field variations into globally coordinated mesh movement.

For local geometric representation, we define $\mathcal{M}_\theta$ as the vectorized movement-indicator function. It predicts an indicator field $\mathbf{M} = \{\mathbf{M}_i\}_{i=1}^N$ from the initial mesh and physical features:

$$\mathbf{M} = \mathcal{M}_\theta(\mathbf{C}, \mathbf{U}, \mathcal{E}), \tag{2}$$

where $\mathbf{M}_i$ characterizes the direction and anisotropic strength of the desired local mesh geometry.

For global mesh coordination, we define f_ϕ as the relocation function that coordinates the local geometric demands to obtain the final coordinates. Combining the two components, the end-to-end relocation function $\mathcal{F}_\Theta$ is expressed as

$$\mathbf{C}' = \mathcal{F}_\Theta(\mathbf{C}, \mathbf{U}, \mathcal{E}) = f_\phi\left(\mathbf{C}, \mathbf{U}, \mathcal{M}_\theta(\mathbf{C}, \mathbf{U}, \mathcal{E}), \mathcal{E}\right), \qquad \Theta = \{\theta, \phi\}. \tag{3}$$

GEMOVE enables unsupervised training without pre-adapted target meshes by enforcing equal element size and equilateral geometry in the mapped space (Sec. 4.4).

4.2 VECTORIZED MESH MOVEMENT INDICATORS

To characterize how mesh resolution should be distributed locally, we seek a geometric representation that captures both the magnitude and direction of adaptation at each vertex. Accordingly, we propose a symmetric positive-definite vectorized movement indicator $\mathbf{M}_i \in \mathbb{S}^2_{++}$ for each vertex v_i, whose principal directions and values jointly encode the orientation and anisotropic strength of the desired mesh geometry. This matrix form provides a unified geometric representation of directional adaptation and naturally supports the mapped-space formulation introduced next.

A graph attention encoder (Brody et al., 2022) aggregates spatial and physical information over the mesh and produces a contextual representation $\mathbf{h}_i$ for each vertex. An indicator head predicts an orientation θ_i and anisotropic strength s_i, while $\rho_i > 0$ provides a graph-normalized scalar density factor derived from the local field variation. Further computational details of θ_i, s_i, and ρ_i are provided in Appendix B. We parameterize the indicator as

$$\mathbf{M}_i = \rho_i \mathbf{R}(\theta_i) \begin{bmatrix} e^{s_i} & 0 \\ 0 & e^{-s_i} \end{bmatrix} \mathbf{R}(\theta_i)^\top, \qquad \mathbf{R}(\theta_i) = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix}. \tag{4}$$

Here, ρ_i controls the isotropic scale, whereas θ_i and s_i determine the principal orientation and anisotropic intensity, respectively. The parameterization satisfies $\det(\mathbf{M}_i) = \rho_i^2$, while the ratio between its principal values is $e^{2|s_i|}$. It therefore separates density redistribution from directional deformation within a unified geometric representation.

The resulting field $\mathbf{M} = \{\mathbf{M}_i\}_{i=1}^N$ extends scalar resolution allocation to vectorized geometric adaptation, specifying not only where resolution should be concentrated but also how it should be distributed directionally.

4.3 GLOBAL MESH COORDINATION UNDER TUTTE EQUILIBRIUM

The vertex-wise movement indicators characterize local geometric demands, but independent vertex relocation provides only weak coordination across the mesh. We therefore model the mesh topology as a spring system, translating local geometric demands into edge-wise stiffnesses. The global force-balanced state of this multi-spring system is referred to as Tutte equilibrium, which jointly determines all interior coordinates.

Specifically, for each undirected edge $(v_i, v_j) \in \mathcal{E}$, let $\Delta\mathbf{c}_{ij} = \mathbf{c}_j - \mathbf{c}_i$ denote the coordinate difference vector and $\mathbf{M}_{ij} = (\mathbf{M}_i + \mathbf{M}_j)/2$ the edge-averaged movement indicator. The corresponding physical-space and mapped-space edge lengths are

$$\ell_{ij} = \|\Delta\mathbf{c}_{ij}\|_2, \qquad \ell_{ij}^M = \sqrt{\Delta\mathbf{c}_{ij}^\top \mathbf{M}_{ij} \Delta\mathbf{c}_{ij}}. \tag{5}$$

Using the vertex representations $\mathbf{h}_i$ from Sec. 4.2, we construct a symmetric edge descriptor from contextual and geometric information and predict a stiffness w_{ij} for each edge in the spring system:

$$\mathbf{z}_{ij} = \left[\mathbf{h}_i + \mathbf{h}_j \,\|\, |\mathbf{h}_i - \mathbf{h}_j| \,\|\, \ell_{ij} \,\|\, \ell_{ij}^M\right], \qquad w_{ij} = \mathrm{softplus}\big(\psi_\phi(\mathbf{z}_{ij})\big) + \varepsilon, \tag{6}$$

where $\|$ denotes concatenation, ψ_ϕ is a learnable edge function, and $\varepsilon > 0$ ensures $w_{ij} = w_{ji} > 0$ through the symmetric construction.

Under Tutte equilibrium, every interior vertex balances all incident interactions in the spring system:

$$\sum_{j \in \mathcal{N}(i)} w_{ij}\left(\mathbf{c}_i' - \mathbf{c}_j'\right) = \mathbf{0}, \qquad v_i \in \mathcal{V}_I. \tag{7}$$

Because every spring contributes to the equilibrium of the connected system, a change in any edge interaction affects the mesh-wide equilibrium. Local geometric demands are therefore coordinated jointly across the entire mesh rather than realized independently. With fixed boundary coordinates, all interior coordinates are therefore solved jointly as

$$\mathbf{L}_{II}\mathbf{C}'_I = -\mathbf{L}_{IB}\mathbf{C}_B, \qquad \mathbf{C}'_B = \mathbf{C}_B, \tag{8}$$

where $\mathbf{L}$ is the weighted graph Laplacian induced by $\{w_{ij}\}$. Under standard Tutte conditions, the resulting barycentric embedding is crossing-free and therefore avoids element inversion. Detailed derivations and theoretical conditions are provided in Appendix C.

4.4 Unsupervised Geometric Equidistribution Principle

To achieve unsupervised mesh movement, we define a mapped space under global geometric equidistribution, where heterogeneous physical-space demands share a common geometric basis. Given the vectorized movement indicators, each physical-space edge difference $\Delta\mathbf{c}_{ij}$ is transformed by $\mathbf{M}_{ij}^{1/2}\Delta\mathbf{c}_{ij}$, incorporating the resolution density and directional geometry into its mapped representation. Thus, differences in element size, orientation, and anisotropy are accounted for by $\mathbf{M}$. We therefore seek to drive each mesh element toward equal size and equilateral geometry in the mapped space. Equal size regulates resolution allocation, while equilateral geometry constrains directional deformation. Consequently, this unsupervised geometric constraint drives the physical-space mesh to redistribute its resolution and geometry according to heterogeneous local demands.

At the element level, for a triangular element $K = (i, j, k)$, let $\mathbf{J}_K = [\mathbf{c}'_j - \mathbf{c}'_i,\ \mathbf{c}'_k - \mathbf{c}'_i]$ denote its relocated Jacobian and $\mathbf{M}_K = (\mathbf{M}_i + \mathbf{M}_j + \mathbf{M}_k)/3$ its element-wise movement indicator. Using the edge matrix $\mathbf{E}_{\mathrm{ref}}$ of a canonical equilateral triangle as reference, we define

$$\widetilde{\mathbf{J}}_K = \mathbf{M}_K^{1/2}\mathbf{J}_K, \qquad \mathbf{G}_K = \left(\widetilde{\mathbf{J}}_K\mathbf{E}_{\mathrm{ref}}^{-1}\right)^{\top}\left(\widetilde{\mathbf{J}}_K\mathbf{E}_{\mathrm{ref}}^{-1}\right). \tag{9}$$

The eigenstructure of $\mathbf{G}_K$ characterizes the mapped-space scale and directional deformation relative to the canonical geometry, independent of rotation. See Appendix B.3 for further geometric details.

Mapped-space size equidistribution. We first require elements to approach a common mapped-space scale. Let $a_K^M = \sqrt{\det(\mathbf{G}_K)}$ denote the mapped-space element measure, and define its mesh-wise geometric mean α and normalized Gram matrix $\overline{\mathbf{G}}_K$ as

$$\alpha = \mathrm{sg}\left[\exp\left(\frac{1}{|\mathcal{K}|}\sum_{K\in\mathcal{K}}\log a_K^M\right)\right], \qquad \overline{\mathbf{G}}_K = \frac{\mathbf{G}_K}{\alpha}. \tag{10}$$

The corresponding size criterion is

$$\mathcal{L}_{\mathrm{size}} = \frac{2}{|\mathcal{K}|}\sum_{K\in\mathcal{K}}\left[\log\frac{a_K^M}{\alpha}\right]^2, \tag{11}$$

which drives all elements toward the same mapped-space measure while allowing their physical-space sizes to vary according to the local movement indicators.

Directional geometric alignment. Let $\lambda_{K,1}$ and $\lambda_{K,2}$ denote the positive eigenvalues of $\overline{\mathbf{G}}_K$. Their ratio measures the residual directional distortion relative to the canonical geometry, leading to

$$\mathcal{L}_{\mathrm{dir}} = \frac{1}{2|\mathcal{K}|}\sum_{K\in\mathcal{K}}\left[\log\frac{\lambda_{K,1}}{\lambda_{K,2}}\right]^2. \tag{12}$$

Minimizing this criterion drives the mapped elements toward the canonical equilateral geometry, while their heterogeneous orientations and anisotropy remain encoded by $\mathbf{M}_K$ in the physical space.

Unified geometric objective. We combine the scale and directional criteria to form a unified spectral geometric objective as follows:

$$\mathcal{L}_{\mathrm{geo}} = \mathcal{L}_{\mathrm{size}} + \mathcal{L}_{\mathrm{dir}} = \frac{1}{|\mathcal{K}|}\sum_{K\in\mathcal{K}}\left[(\log\lambda_{K,1})^2 + (\log\lambda_{K,2})^2\right]. \tag{13}$$

At the ideal geometric equilibrium, $\lambda_{K,1} = \lambda_{K,2} = 1$, equivalently $\overline{\mathbf{G}}_K = \mathbf{I}$ for every element. Thus, after accounting for the heterogeneous local demands encoded by $\mathbf{M}_K$, all relocated elements conform to the same canonical geometry in the mapped space, which defines the global geometric equidistribution principle of GEMove.

5 EXPERIMENTS

In this section, we empirically evaluate GEMOVE from controlled PDE benchmarks (Poisson, Helmholtz, Swirl, and Burgers) to real-world Earth-system fields (oceanographic, atmospheric, biogeochemical, and environmental) by addressing the following research questions (**RQs**):

- **RQ1:** How well does GEMOVE reduce the error and generalize across PDEs and domains?
- **RQ2:** How does each component contribute to GEMOVE's performance?
- **RQ3:** How well does GEMOVE adapt to multiscale spatial structures and long-term temporal variability in real-world ocean systems?
- **RQ4:** Can GEMOVE learned from local ocean systems transfer to global Earth-system fields across different geophysical spheres, regions, and variables?

5.1 EXPERIMENTAL SETUP

We summarize the experimental setup, including data, baselines, protocols, and evaluation metrics, with full details provided in Appendices D and F.

Simulation Data. Following UM2N (Zhang et al., 2024), we generate Poisson and Helmholtz training problems on rectangular triangular meshes. Swirl and Burgers test zero-shot cross-equation generalization, while polygonal Poisson and Helmholtz domains test geometric generalization.

Real-World Data. We use Kuroshio currents from Copernicus Marine (Lellouche et al., 2021; Copernicus Marine Service, 2026c) and ENSO SST anomalies from NOAA OISST v2.1 (Huang et al., 2021; 2020) as real-world ocean benchmarks. A broader global transfer benchmark further integrates multiple Earth-system datasets across 54 regions, 16 variables, and diverse fields.

Baselines and Protocol. We compare with M2N (Song et al., 2022), UM2N (Zhang et al., 2024), and UGM2N (Wang et al., 2025). M2N and UM2N require MA-generated target meshes; UGM2N and our GEMOVE are unsupervised. For real-world experiments, we retrain the two unsupervised methods on the ocean data and transfer them to global fields without further training. Comparisons use identical initial meshes, node budgets, and boundary constraints.

Evaluation. Following prior work (Zhang et al., 2024; Wang et al., 2025), we report error reduction (Error Red.) relative to the original mesh using high-resolution references and task-specific metrics. Values are reported as mean ± std over three independent training runs; higher is better.

5.2 GENERALIZATION ACROSS PDES AND DOMAIN SHAPES (RQ1)

We first examine whether GEMOVE generalizes beyond the PDE families and domain geometries seen during training. Table 1 reports the results under both equation and geometry shifts.

Table 1: Error reduction (%) across PDEs and unseen polygonal domains, averaged over resolutions and parameter settings. † marks training PDE families. * marks cases with unacceptable mesh failure due to element inversion. Best and second-best are bold and underlined; gray denotes accuracy decline after movement. Δ denotes the percentage-point gain of GEMOVE over strongest baseline.

Domain	PDE	M2N	UM2N	UGM2N	GEMOVE	Δ (pp) ↑
Rectangular	**Poisson**†	6.72 ± 0.85	<u>20.76 ± 0.61</u>	20.62 ± 0.65	**23.66 ± 0.55**	**+2.91**
	Helmholtz†	8.97 ± 0.80	<u>20.82 ± 0.59</u>	20.48 ± 0.64	**23.85 ± 0.53**	**+3.04**
	Swirl	-2.74 ± 0.80*	3.07 ± 0.68	<u>3.39 ± 0.62</u>	**17.55 ± 0.51**	**+14.16**
	Burgers	17.21 ± 0.95	13.14 ± 0.78	<u>21.26 ± 0.74</u>	**33.40 ± 0.62**	**+12.14**
	Overall	7.54 ± 0.65	14.45 ± 0.51	<u>16.44 ± 0.50</u>	**24.62 ± 0.42**	**+8.18**
Polygonal	**Poisson**	-6.33 ± 1.10*	-1.79 ± 0.89	<u>1.85 ± 0.75</u>	**9.88 ± 0.64**	**+8.03**
	Helmholtz	-3.27 ± 1.00*	2.53 ± 0.81	<u>3.09 ± 0.72</u>	**10.77 ± 0.61**	**+7.68**
	Overall	-4.80 ± 0.76	0.37 ± 0.61	<u>2.47 ± 0.55</u>	**10.33 ± 0.46**	**+7.86**

On the Poisson and Helmholtz families seen during training, GEMOVE achieves moderate gains of 2.91 and 3.04 percentage points over the strongest baselines, which also benefit from direct exposure to these PDE families during training. The advantage becomes substantially larger on the held-out

Swirl and Burgers problems, reaching 14.16 and 12.14 percentage points, respectively. Across all four PDE families, GEMOVE achieves a 24.62% average error reduction, compared with 16.44% for the strongest baseline.

A similar pattern appears under geometric shift. Although the polygonal tests retain the seen Poisson and Helmholtz equations, their unseen boundary geometries and mesh topologies substantially weaken the baselines, whereas GEMOVE maintains an overall 10.33% error reduction, compared with 2.47% for the strongest baseline. Together, these results show that the learned geometric adaptation transfers beyond both the PDE families and domain geometries encountered during training. Complete parameter-wise results and qualitative mesh comparisons are provided in Appendix E.

5.3 COMPONENT ABLATION AND MECHANISM ANALYSIS (RQ2)

To evaluate the contribution of each component, we construct three ablated variants of GEMOVE, with results reported in Table 2. *Only-Size* and *Only-Direction* retain only the mapped-space size and directional components of the geometric objective, respectively, while *w/o Coordination* preserves the complete geometric objective but removes global mesh coordination. All variants are evaluated on the same four PDE families, with results across independent runs reported in Appendix E.4.

Table 2: Ablation results (Error Red., %).

Model	Pois.	Helm.	Swirl	Burg.	Avg.
Only-Size	20.35	20.03	7.82	26.04	18.56
Only-Direction	19.51	18.94	13.39	24.71	19.14
w/o Coord.	21.82	22.16	15.70*	31.39	22.77
GEMOVE	**23.66**	**23.85**	**17.55**	**33.40**	**24.62**

GEMOVE benefits from complementary size and directional constraints. The size constraint contributes strongly on Poisson, Helmholtz, and Burgers, while the directional constraint is particularly important on Swirl, where pronounced anisotropic structures require directional adaptation. Combining both constraints further improves the average performance, showing that size and direction capture complementary local adaptation demands in GEMOVE.

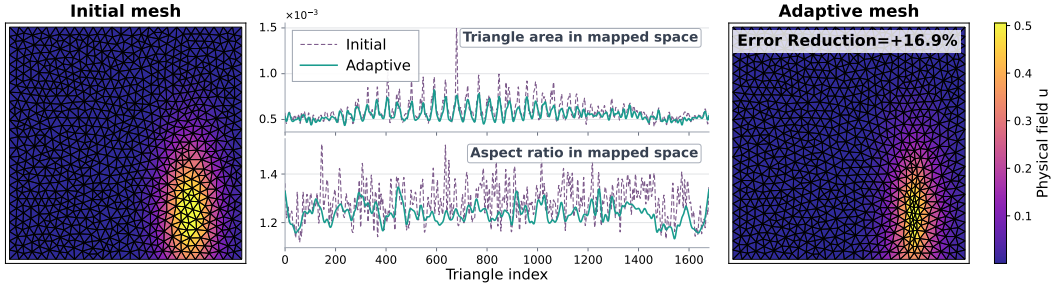

Figure 3: Geometric equidistribution in the mapped space. More regular element area and shape.

GEMOVE realizes geometric equidistribution through vectorized movement indicators. As shown in Fig. 3, although the adaptive mesh remains heterogeneous in physical space, GEMOVE transforms these spatially varying adaptation demands into uniform element areas and more regular aspect ratios in the mapped space, moving elements toward canonical geometry. This mapped-space redistribution shows that vectorized movement indicators encode amount and direction of adaptation, providing stronger geometric constraints for mesh relocation than scalar density alone.

Global coordination prevents pathological mesh movement. Removing global coordination reduces the average error reduction from 24.62% to 22.77%, while the complete GEMOVE remains consistently stronger across all four PDE families. More importantly, the numerical gap understates its geometric role. As shown in Fig. 4, locally reasonable adaptation demands can accumulate into excessive concentration or conflicting movements without mesh-wide coordination. Tutte equilibrium instead determines all interior vertices jointly,

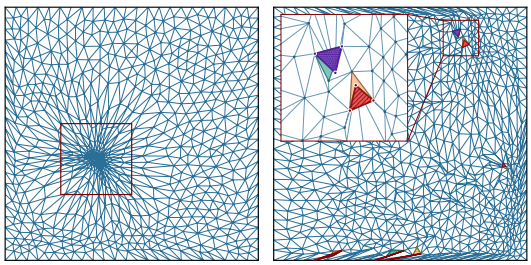

Figure 4: Failure modes without global coordination: over-allocation and element inversion.

coupling local geometric demands across the mesh and preventing these pathological movements under its theoretical conditions.

5.4 SPATIOTEMPORAL VARIABILITY IN REAL-WORLD OCEAN SYSTEMS (RQ3)

We next evaluate GEMOVE on two real-world ocean systems with complementary spatiotemporal variability. In this target-free setting, supervised M2N and UM2N would require pre-adapted MA meshes as supervision for each ocean dataset, whereas UGM2N and GEMOVE can be retrained directly from the physical fields without target meshes. We therefore compare the two unsupervised methods on Kuroshio and ENSO, which respectively stress spatial adaptation to competing anisotropic structures and temporal robustness under long-term climate variability. Detailed data settings, retraining protocols, and metric definitions are provided in Appendix F, with additional results in Appendix G.

GEMOVE better allocates resolution across competing Kuroshio structures. The left half of Table 3 shows that GEMOVE achieves a 14.4% error reduction and a 100% Win Rate, while increasing CoreIoU from 0.7759 to 0.8413 over UGM2N. Kuroshio contains a narrow energetic jet with surrounding fronts and mesoscale structures across spatial scales and orientations, creating strongly anisotropic and competing adaptation demands. As shown in Fig. 5, GEMOVE produces a clearer mesh response to physical structure, concentrating vertices along the rapidly varying jet core while preserving adaptation to adjacent regions. The error map shows stronger error reduction around these active structures, indicating that the movement better follows the spatial variation. For this snapshot, GEMOVE achieves a 13.3% error reduction, compared with 4.3% for UGM2N. These results indicate that the vectorized movement indicators better align mesh relocation with the underlying anisotropic physical structures, providing physically consistent adaptation in the Kuroshio system.

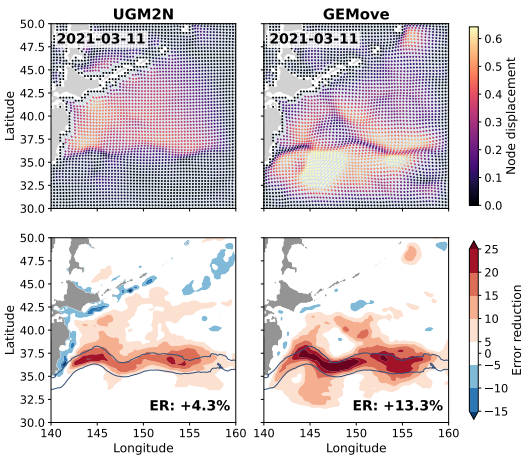

Figure 5: Spatial adaptation to the Kuroshio. Top: node displacement; bottom: spatial distribution of error reduction.

GEMOVE remains robust under long-term ENSO variability. The right half of Table 3 shows that the advantage of GEMOVE becomes more pronounced as the spatial resolution becomes coarser. Its gain over UGM2N increases from 18.3 percentage points at 2° to 28.3 at 3° and 35.9 at 4°, indicating that geometric adaptation becomes particularly beneficial when the available mesh resolution is more limited. Figure 6 further examines the 4° setting across El Niño, neutral, and La Niña conditions. Despite the substantial changes in the underlying ocean state, GEMOVE maintains similar error reductions of 69%, 66%, and 69% across the three regimes. This stability shows that the learned adaptation is not tied to a particular ENSO phase, but remains effective as the dominant spatial structure evolves over long-term climate variability.

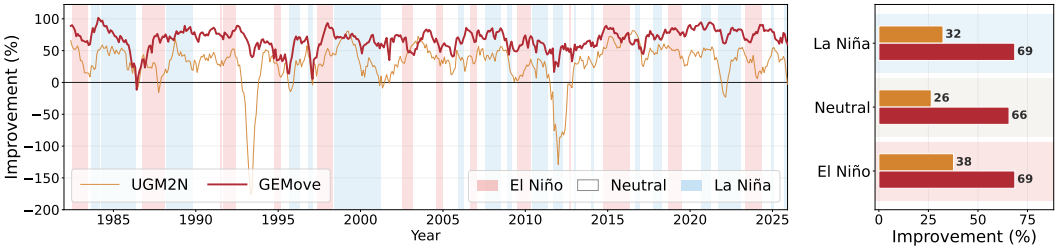

Figure 6: Long-term ENSO adaptation at 4° resolution. Left: Niño 3.4 error reduction over time, with colors indicating El Niño, neutral, and La Niña periods; right: mean error reduction for the three ENSO regimes over the held-out temporal evaluation period.

Table 3: Performance on the Kuroshio and ENSO benchmarks. Reported results are computed over the held-out temporal evaluation period. Best results are shown in bold.

Method	Kuroshio			ENSO: Error Red. ↑		
	Error Red. ↑	Win Rate ↑	CoreIoU ↑	2°	3°	4°
UGM2N	4.1 ± 1.1%	93.6 ± 1.5%	0.7759 ± 0.0062	27.0 ± 1.3%	29.2 ± 1.5%	32.2 ± 1.7%
GEMOVE	**14.4** ± 0.9%	**100.0** ± 0.0%	**0.8413** ± 0.0051	**45.3** ± 1.1%	**57.5** ± 1.3%	**68.1** ± 1.4%

5.5 GLOBAL TRANSFER ACROSS EARTH-SYSTEM FIELDS (RQ4)

Finally, we evaluate whether the adaptation rule learned from local ocean systems transfers to substantially different Earth-system fields. To avoid selecting favorable test cases, we randomly select 54 evaluation regions across four field groups (oceanographic, atmospheric, biogeochemical, and environmental fields), covering 16 predefined physical variables. We directly apply UGM2N and GEMOVE trained on the Kuroshio and ENSO data in Sec. 5.4, without retraining or fine-tuning on the target regions. The benchmark therefore introduces simultaneous shifts in geographic region, physical variable, and underlying spatial structure. Data sources, region definitions, and evaluation settings are provided in Appendix F.4, with additional region-level results in Appendix G.3.

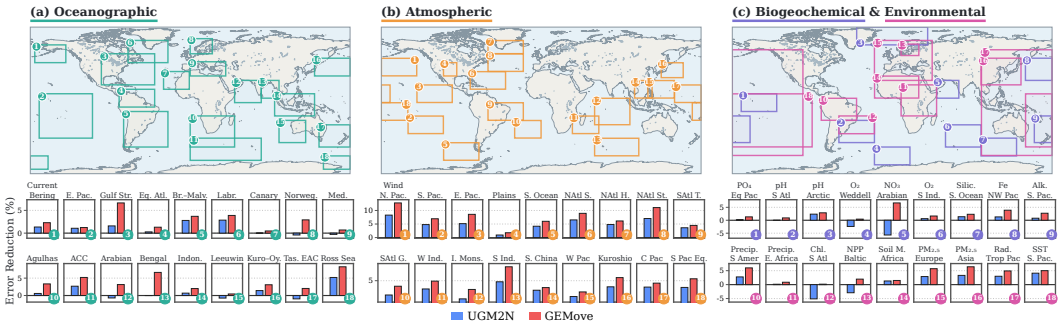

Figure 7: Zero-shot transfer from local ocean systems to global Earth-system fields. Error reduction (%) of UGM2N and GEMOVE across 54 randomly selected regions from oceanographic, atmospheric, biogeochemical, and environmental fields.

GEMOVE transfers consistently across diverse Earth-system fields. Figure 7 shows consistent improvements over UGM2N across all 54 evaluation regions. Across the complete benchmark, its average error reduction reaches 4.09%, compared with 1.80% for UGM2N, supporting the transferability of the geometric adaptation learned by GEMOVE. Moreover, GEMOVE maintains positive error reduction in all 54 regions, whereas UGM2N increases the error in 11 cases. The advantage remains consistent across the four field groups, with average gains of 2.22, 2.03, 2.70, and 2.56 percentage points for oceanographic, atmospheric, biogeochemical, and environmental fields, respectively. These results show that the learned geometric adaptation transfers across substantial geographic and physical shifts without parameter updates.

6 CONCLUSION

We present GEMOVE, an unsupervised mesh movement framework that advances adaptive discretization from scalar density redistribution toward globally coordinated geometric adaptation. By representing heterogeneous physical variations through vectorized movement indicators, GEMOVE adapts a fixed vertex budget to both the magnitude and direction of complex multiscale scientific structures without target meshes. Its effectiveness across PDEs and real-world Earth-system fields supports geometric adaptation as a general target-free approach for improving adaptive discretization across heterogeneous systems in scientific computing. A promising direction is to combine mesh movement with mesh refinement (Berger & Oliger, 1984; Wu et al., 2023): when additional computational budget allows refinement in under-resolved regions, applying GEMOVE afterward may further improve adaptation by redistributing the expanded mesh.

AI USE STATEMENT

In this work, we used generative AI tools to condense and polish the manuscript, improve readability, and correct language and grammatical errors. All AI-assisted text and references were reviewed and verified by the authors. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Robert Atlas, Ross N. Hoffman, Joseph Ardizzone, S. Mark Leidner, Juan Carlos Jusem, Deborah K. Smith, and Daniel Gombos. A cross-calibrated, multiplatform ocean surface wind velocity product for meteorological and oceanographic applications. *Bulletin of the American Meteorological Society*, 92(2):157–174, 2011. doi: 10.1175/2010BAMS2946.1.

Marsha J. Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53(3):484–512, 1984. doi: 10.1016/0021-9991(84)90073-1.

Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? In *International Conference on Learning Representations*, 2022.

Chris J. Budd, Weizhang Huang, and Robert D. Russell. Adaptivity with moving grids. *Acta Numerica*, 18:111–241, 2009. doi: 10.1017/S0962492906400015.

Eric P. Chassignet, Harley E. Hurlburt, Ole Martin Smedstad, George R. Halliwell, Patrick J. Hogan, Alan J. Wallcraft, Remy Baraille, and Rainer Bleck. The HYCOM (HYbrid Coordinate Ocean Model) data assimilative system. *Journal of Marine Systems*, 65(1–4):60–83, 2007. doi: 10.1016/j.jmarsys.2005.09.016.

Climate Hazards Center. Climate hazards center infrared precipitation with stations version 3 (CHIRPS3) data repository. University of California, Santa Barbara, 2025.

Copernicus Marine Service. Baltic sea biogeochemistry analysis and forecast. Swedish Meteorological and Hydrological Institute, 2026a. Product ID: BALTICSEA_ANALYSISFORECAST_BGC_003_007.

Copernicus Marine Service. Global ocean biogeochemistry analysis and forecast. Mercator Ocean International, 2026b. Product ID: GLOBAL_ANALYSISFORECAST_BGC_001_028.

Copernicus Marine Service. Global ocean physics reanalysis (GLORYS12V1). Mercator Ocean International, 2026c. Product ID: GLOBAL_MULTIYEAR_PHY_001_030.

David R. Doelling, Norman G. Loeb, Dennis F. Keyes, Michele L. Nordeen, Daniel Morstad, Cathy Nguyen, Bruce A. Wielicki, David F. Young, and Moguo Sun. Geostationary enhanced temporal interpolation for CERES flux products. *Journal of Atmospheric and Oceanic Technology*, 30(6):1072–1090, 2013. doi: 10.1175/JTECH-D-12-00136.1.

Wouter Dorigo, Wolfgang Wagner, Clement Albergel, Franziska Albrecht, Gianpaolo Balsamo, Luca Brocca, Daniel Chung, Martin Ertl, Matthias Forkel, Alexander Gruber, Eva Haas, Paul D. Hamer, Martin Hirschi, Jaakko Ikonen, Richard de Jeu, Richard Kidd, William Lahoz, Yi Y. Liu, Diego Miralles, Thomas Mistelbauer, Nadine Nicolai-Shaw, Robert Parinussa, Chiara Pratola, Christoph Reimer, Robin van der Schalie, Sonia I. Seneviratne, Tuomo Smolander, and Pascal Lecomte. ESA CCI soil moisture for improved earth system understanding: State-of-the art and future directions. *Remote Sensing of Environment*, 203:185–215, 2017. doi: 10.1016/j.rse.2017.07.001.

Wouter Dorigo, Wolfgang Preimesberger, Sebastian Hahn, Robin van der Schalie, Richard de Jeu, Richard Kidd, Nemesio Rodriguez-Fernandez, Martin Hirschi, Paolo Stradiotti, Thomas Frederikse, Alexander Gruber, and Dominique Duchemin. ESA soil moisture climate change initiative (Soil_Moisture_cci): COMBINED product, version 09.2. NERC EDS Centre for Environmental Data Analysis, 2026.

David C. Dowell, Curtis R. Alexander, Eric P. James, Stephen S. Weygandt, Stanley G. Benjamin, Geoffrey S. Manikin, Benjamin T. Blake, John M. Brown, Joseph B. Olson, Ming Hu, Tatiana G. Smirnova, Terra Ladwig, Jaymes S. Kenyon, Ravan Ahmadov, David D. Turner, Jeffrey D. Duda, and Trevor I. Alcott. The high-resolution rapid refresh (HRRR): An hourly updating convection-allowing forecast model. Part I: Motivation and system description. *Weather and Forecasting*, 37 (8):1371–1395, 2022. doi: 10.1175/WAF-D-21-0151.1.

Zhe Feng, Shilong Tao, Haonan Sun, Shaohan Chen, Zhanxing Zhu, and Yunhuai Liu. MAVEN: A mesh-aware volumetric encoding network for simulating 3d flexible deformation. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=XmULVr15E0.

Chris Funk, Pete Peterson, Laura Harrison, Robert Saldivar, Martin Landsfeld, Diego Pedreros, Shraddhanand Shukla, Andreas H. Fink, Frank Davenport, Seth H. Peterson, William Turner, Austin Sonnier, Michael Budde, Karyn Tabor, James Verdin, Disha Hauzaree, Mohamed Naim, Daniella Alaso, and Gregory Husak. The climate hazards center infrared precipitation with stations, version 3. *Scientific Data*, 13:718, 2026. doi: 10.1038/s41597-026-07096-4.

Paul Garnier, Vincent Lannelongue, Jonathan Viquerat, and Elie Hachem. MeshMask: Physics-based simulations with masked graph neural networks. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=bFHR8hNk4I.

Alexander Gruber, Tracy Scanlon, Robin van der Schalie, Wolfgang Wagner, and Wouter Dorigo. Evolution of the ESA CCI soil moisture climate data records and their underlying merging methodology. *Earth System Science Data*, 11:717–739, 2019. doi: 10.5194/essd-11-717-2019.

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. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/136302ea7874e2ff96d517f9a8eb0a35-Abstract-Conference.html.

Boyin Huang, Chunying Liu, Viva F. Banzon, Eric Freeman, Garrett Graham, William Hankins, Thomas M. Smith, and Huai-Min Zhang. NOAA 0.25-degree daily optimum interpolation sea surface temperature (OISST), version 2.1. NOAA National Centers for Environmental Information, 2020.

Boyin Huang, Chunying Liu, Viva Banzon, Eric Freeman, Garrett Graham, Bill Hankins, Tom Smith, and Huai-Min Zhang. Improvements of the daily optimum interpolation sea surface temperature (DOISST) version 2.1. *Journal of Climate*, 34(8):2923–2939, 2021. doi: 10.1175/JCLI-D-20-0166.1.

Weizhang Huang. Metric tensors for anisotropic mesh generation. *Journal of Computational Physics*, 204(2):633–665, 2005. doi: 10.1016/j.jcp.2004.10.024.

Namgyu Kang, Jaemin Oh, Youngjoon Hong, and Eunbyung Park. PIG: Physics-informed gaussians as adaptive parametric mesh representations. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=y5B0ca4mjt.

Christoph A. Keller, K. Emma Knowland, Bryan N. Duncan, Junhua Liu, Daniel C. Anderson, Sampa Das, Robert A. Lucchesi, Elizabeth W. Lundgren, Julie M. Nicely, Eric Nielsen, Lesley E. Ott, Emily Saunders, Sarah A. Strode, Pamela A. Wales, Daniel J. Jacob, and Steven Pawson. Description of the NASA GEOS composition forecast modeling system GEOS-CF v1.0. *Journal of Advances in Modeling Earth Systems*, 13(4):e2020MS002413, 2021. doi: 10.1029/2020MS002413.

Thanh-Tung Le, Khai Nguyen, Shanlin Sun, Kun Han, Nhat Ho, and Xiaohui Xie. Diffeomorphic mesh deformation via efficient optimal transport for cortical surface reconstruction. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=gxhRR8vUQb.

Jean-Michel Lellouche, Eric Greiner, Romain Bourdallé-Badie, Gilles Garric, Angélique Melet, Marie Drévillon, Clément Bricaud, Mathieu Hamon, Olivier Le Galloudec, Charly Regnier, Tony Candela, Charles-Emmanuel Testut, Florent Gasparin, Giovanni Ruggiero, Mounir Benkiran, Yann Drillet, and Pierre-Yves Le Traon. The copernicus global 1/12° oceanic and sea ice GLORYS12 reanalysis. *Frontiers in Earth Science*, 9:698876, 2021. doi: 10.3389/feart.2021.698876.

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.

NASA/LARC/SD/ASDC. CERES and GEO-enhanced TOA, within-atmosphere and surface fluxes, clouds and aerosols daily Terra-Aqua-NOAA20 Edition4B. NASA Langley Atmospheric Science Data Center DAAC, 2026.

Ángel J. Omella and David Pardo. r-adaptive deep learning method for solving partial differential equations. *Computers & Mathematics with Applications*, 153:33–42, 2024. doi: 10.1016/j.camwa.2023.11.005.

Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter W. Battaglia. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=roNqYL0_XP.

Wolfgang Preimesberger, Tracy Scanlon, Chun-Hsu Su, Alexander Gruber, and Wouter Dorigo. Homogenization of structural breaks in the global ESA CCI soil moisture multisatellite climate data record. *IEEE Transactions on Geoscience and Remote Sensing*, 59(4):2845–2862, 2021. doi: 10.1109/TGRS.2020.3012896.

Bo Qiu and Shuiming Chen. Variability of the kuroshio extension jet, recirculation gyre, and mesoscale eddies on decadal time scales. *Journal of Physical Oceanography*, 35(11):2090–2103, 2005.

James Rowbottom, Georg Maierhofer, Teo Deveney, Eike Hermann Müller, Alberto Paganini, Katharina Schratz, Pietro Lio, Carola-Bibiane Schönlieb, and Chris Budd. G-Adaptivity: Optimised graph-based mesh relocation for finite element methods. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 52173–52197. PMLR, 2025.

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. In *Advances in Neural Information Processing Systems*, volume 35, pp. 7199–7210, 2022.

Kevin E. Trenberth. The definition of El Niño. *Bulletin of the American Meteorological Society*, 78 (12):2771–2777, 1997.

W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, s3-13(1): 743–767, 1963. doi: 10.1112/plms/s3-13.1.743.

Zhichao Wang, Xinhai Chen, Qinglin Wang, Xiang Gao, Qingyang Zhang, Menghan Jia, Xiang Zhang, and Jie Liu. UGM2N: An unsupervised and generalizable mesh movement network via M-Uniform loss. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Frank J. Wentz, Joel Scott, Ross N. Hoffman, S. Mark Leidner, Robert Atlas, and Joseph Ardizzone. Remote sensing systems cross-calibrated multi-platform (CCMP) 6-hourly ocean vector wind analysis product on 0.25 degree grid, version 2.0. Remote Sensing Systems, Santa Rosa, California, 2015. Dataset.

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.

Pingping Xie, Robert Joyce, Shaorong Wu, Soo-Hyun Yoo, Yelena Yarosh, Fengying Sun, and Roger Lin. Reprocessed, bias-corrected CMORPH global high-resolution precipitation estimates from 1998. *Journal of Hydrometeorology*, 18(6):1617–1641, 2017. doi: 10.1175/JHM-D-16-0168.1.

Pingping Xie, Robert Joyce, Shaorong Wu, Soo-Hyun Yoo, Yelena Yarosh, Fengying Sun, and Roger Lin. NOAA climate data record of CPC morphing technique (CMORPH) high resolution global precipitation estimates, version 1. NOAA National Centers for Environmental Information, 2019.

Youn-Yeol Yu, Jeongwhan Choi, Woojin Cho, Kookjin Lee, Nayong Kim, Kiseok Chang, ChangSeung Woo, Ilho Kim, SeokWoo Lee, Joon Young Yang, Sooyoung Yoon, and Noseong Park. Learning flexible body collision dynamics with hierarchical contact mesh transformer. In *The Twelfth International Conference on Learning Representations*, 2024.

Youn-Yeol Yu, Jeongwhan Choi, Jaehyeon Park, Kookjin Lee, and Noseong Park. PIORF: Physics-informed ollivier–ricci flow for long-range interactions in mesh graph neural networks. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=qkBBHixPow.

Amirhossein Zamani, Bruno Roy, and Arianna Rampini. Unsupervised representation learning for 3d mesh parameterization with semantic and visibility objectives. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=9LYsvna4Sk.

Mingrui Zhang, Chunyang Wang, Stephan Kramer, Joseph G. Wallwork, Siyi Li, Jiancheng Liu, Xiang Chen, and Matthew D. Piggott. Towards universal mesh movement networks. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Appendix Contents

A Under-Constrained Geometry of Scalar Equidistribution **15**
A.1 Under-Constrained Element Geometry . . . 15
A.2 Multi-Solution Ambiguity . . . 15

B Detailed Construction of the Vectorized Movement Indicator **17**
B.1 Contextual Encoding and Geometric Parameters . . . 17
B.2 Indicator Construction . . . 17
B.3 Connection to the Geometric Objective . . . 18

C Tutte Equilibrium and Inversion-Free Guarantee **19**
C.1 Tutte Equilibrium . . . 19
C.2 Inversion-Free Guarantee . . . 20

D Synthetic PDE Experimental Details **20**
D.1 Ablation Settings . . . 22

E Additional Synthetic PDE Results **22**
E.1 Qualitative Mesh Adaptation . . . 23
E.2 Parameter-wise PDE Results . . . 23
E.3 Generalization across Polygonal Domains . . . 25
E.4 Ablation Results Across Runs . . . 25
E.5 Geometric Redistribution in the Mapped Space . . . 25
E.6 Inference Efficiency . . . 27

F Real-World Earth-System Data and Experimental Details **27**
F.1 Real-World Ocean Data . . . 27
F.2 Kuroshio Current Data and Evaluation . . . 28
F.3 ENSO Data and Evaluation . . . 30
F.4 Global Earth-System Transfer Data (RQ4) . . . 32
F.5 Temporal Coverage of the Global Transfer Benchmark . . . 37

G Additional Real-World Results **39**
G.1 Additional Kuroshio Results . . . 39
G.2 Additional ENSO Temporal Results . . . 41
G.3 Additional Global Transfer Visualization . . . 42

A UNDER-CONSTRAINED GEOMETRY OF SCALAR EQUIDISTRIBUTION

This appendix develops the geometric analysis underlying GEMOVE, specifies the experimental settings, and provides additional results. Appendices A–C connect scalar ambiguity, vectorized movement indicators, and global coordination. Appendices D and E cover the synthetic PDE benchmarks, while Appendices F and G cover the real-world experiments.

We first characterize the multi-solution ambiguity introduced in Sec. 1. For a prescribed positive scalar density, an area constraint leaves two intrinsic degrees of freedom in a triangular element and does not select its orientation relative to the physical field.

A.1 UNDER-CONSTRAINED ELEMENT GEOMETRY

Modulo rigid translation and rotation, a nondegenerate triangular element is completely characterized by its three side lengths (a, b, c). Hence, its intrinsic geometry lies in a three-dimensional shape space. For a prescribed density $\rho_K > 0$ and a common allocation constant $q > 0$, scalar equidistribution imposes only

$$\rho_K |K| = q, \qquad |K| = A_K^* = \frac{q}{\rho_K}. \tag{14}$$

Using the polynomial form of Heron's relation, the admissible geometries are therefore constrained by a single scalar equation

$$F(a, b, c) = 2a^2b^2 + 2b^2c^2 + 2c^2a^2 - a^4 - b^4 - c^4 - 16(A_K^*)^2 = 0. \tag{15}$$

Proposition 1 (Under-constrained geometry). For a positive prescribed area A_K^*, the set of nondegenerate triangles satisfying Eq. 15 is locally a two-dimensional manifold in the three-dimensional intrinsic shape space.

Proof. The gradient of F cannot vanish simultaneously for a nondegenerate triangle: $\nabla F = \mathbf{0}$ would require

$$a^2 = b^2 + c^2, \qquad b^2 = a^2 + c^2, \qquad c^2 = a^2 + b^2, \tag{16}$$

which is incompatible with $a, b, c > 0$. Thus, 0 is a regular value of F on the positive-area triangle space. By the regular level-set theorem,

$$\dim \{(a, b, c) : F(a, b, c) = 0\} = 3 - 1 = 2. \tag{17}$$

Therefore, scalar equidistribution removes only one intrinsic degree of freedom and leaves two geometric degrees of freedom undetermined.

Accordingly, fixing the element area determines how much local resolution is allocated, but does not uniquely determine edge-length ratios, angles, or directional geometry.

A.2 MULTI-SOLUTION AMBIGUITY

The under-constrained geometry above can be characterized more explicitly through the family of area-preserving affine transformations. Let $\widehat{K}$ be any nondegenerate reference triangle satisfying

$$|\widehat{K}| = A_K^* = \frac{q}{\rho_K}. \tag{18}$$

Consider the two-parameter family

$$\mathbf{A}(s, \tau) = \begin{bmatrix} e^s & \tau e^s \\ 0 & e^{-s} \end{bmatrix}, \qquad (s, \tau) \in \mathbb{R}^2. \tag{19}$$

Since

$$\det \mathbf{A}(s,\tau) = e^{s}e^{-s} = 1, \tag{20}$$

we have $\mathbf{A}(s,\tau) \in \mathrm{SL}(2,\mathbb{R})$. The transformed element

$$K_{s,\tau} = \mathbf{A}(s,\tau)\widehat{K} \tag{21}$$

therefore satisfies

$$|K_{s,\tau}| = |\det \mathbf{A}(s,\tau)|\,|\widehat{K}| = A_K^*, \tag{22}$$

and consequently

$$\rho_K |K_{s,\tau}| = q, \qquad \forall (s,\tau) \in \mathbb{R}^2. \tag{23}$$

Thus, scalar equidistribution is invariant over the complete area-preserving orbit

$$\mathcal{O}(\widehat{K}) = \left\{ \mathbf{A}\widehat{K} \;:\; \mathbf{A} \in \mathrm{SL}(2,\mathbb{R}) \right\}. \tag{24}$$

Proposition 2 (Infinitely many distinct realizations). The orbit in Eq. 24 contains infinitely many non-congruent triangular elements satisfying exactly the same scalar equidistribution condition.

Proof. The intrinsic distortion induced by $\mathbf{A}(s,\tau)$ is characterized by

$$\mathbf{C}(s,\tau) = \mathbf{A}(s,\tau)^{\top}\mathbf{A}(s,\tau) = \begin{bmatrix} e^{2s} & \tau e^{2s} \\ \tau e^{2s} & \tau^2 e^{2s} + e^{-2s} \end{bmatrix}. \tag{25}$$

Its determinant remains fixed,

$$\det \mathbf{C}(s,\tau) = 1, \tag{26}$$

while its trace is

$$\operatorname{tr} \mathbf{C}(s,\tau) = e^{2s}(1+\tau^2) + e^{-2s}. \tag{27}$$

For a labeled reference triangle, a rigid transformation of the resulting physical element preserves its edge Gram matrix. Since two independent reference edges determine the affine deformation, different tensors $\mathbf{C}(s,\tau)$ correspond to different labeled edge geometries. For example, along $\tau = 0$,

$$\operatorname{tr} \mathbf{C}(s,0) = e^{2s} + e^{-2s} = 2\cosh(2s), \tag{28}$$

which is strictly increasing for $s > 0$. Relabeling the three vertices introduces only finitely many equivalent descriptions and cannot collapse this continuum to finitely many congruence classes. More directly, at least one reference edge has a nonzero horizontal component, so its length under $\mathrm{diag}(e^{s}, e^{-s})$ grows without bound as $s \to \infty$. The orbit therefore contains infinitely many non-congruent triangles of the same area. The shear parameter τ provides an additional way to change edge lengths and angles without changing scalar allocation.

Scalar equidistribution is also invariant to rigid orientation. For any $\phi \in [0, 2\pi)$,

$$K_{s,\tau,\phi} = \mathbf{Q}(\phi)\mathbf{A}(s,\tau)\widehat{K}, \qquad \mathbf{Q}(\phi) \in \mathrm{SO}(2), \tag{29}$$

has the same element area, while its alignment relative to the physical field changes with ϕ. Therefore, scalar density neither uniquely determines intrinsic element shape nor selects a preferred physical orientation.

Together with Proposition 1, this establishes the multi-solution ambiguity: fixing scalar density restricts the element to an equal-area manifold, but leaves a continuum of anisotropic, sheared, and directionally distinct geometries admissible. This motivates the vectorized movement indicator in Sec. 4.2, which supplements scalar allocation with explicit geometric information.

Local ambiguity and mesh compatibility. The construction above concerns an individual element with prescribed ρ_K. Adjacent triangles share vertices, so their affine transformations cannot generally be chosen independently. Likewise, a spatially varying density must be evaluated consistently when vertices move. The local construction therefore identifies information absent from a scalar criterion; it does not assert that every equal-area element deformation extends to a valid mesh with fixed boundaries. This distinction motivates the two components of GEMOVE: the indicator supplies directional geometry, while Tutte equilibrium coordinates compatible vertex positions through the shared connectivity.

B DETAILED CONSTRUCTION OF THE VECTORIZED MOVEMENT INDICATOR

We provide the implementation details of the vectorized movement indicator introduced in Sec. 4.2. The construction preserves the scalar density information for local resolution allocation and supplements it with learned orientation and anisotropy.

B.1 CONTEXTUAL ENCODING AND GEOMETRIC PARAMETERS

For each vertex v_i, the model combines node-wise physical features with local graph features and a graph-level field representation. The local and global feature extractors each produce a 16-dimensional representation, which is concatenated with the five non-coordinate node features and projected to a 512-dimensional hidden state. Three GATv2(Brody et al., 2022) layers with six attention heads then aggregate information over the fixed mesh connectivity, yielding the contextual representation $\mathbf{h}_i$. The same encoder architecture and input dimensionality are used across all experiments, including the frozen-checkpoint transfer in RQ4.

The scalar density component is obtained from the node-wise Hessian magnitude and normalized independently within each graph. For a graph g with vertex set $\mathcal{V}_g$, we use

$$\rho_i = \mathrm{clip}\left(\frac{\left|h_i^{\mathrm{Hess}}\right|}{\frac{1}{|\mathcal{V}_g|}\sum_{j\in\mathcal{V}_g}\left|h_j^{\mathrm{Hess}}\right| + \varepsilon}, 10^{-3}, 10^{3} \right), \qquad i \in \mathcal{V}_g, \tag{30}$$

where $\varepsilon = 10^{-8}$. Before clipping, graph-wise normalization preserves relative density demands and reduces sensitivity to the overall Hessian scale. The positive lower bound keeps the indicator well-defined even in regions with vanishing Hessian magnitude; clipping limits extreme density contrasts.

The contextual representation $\mathbf{h}_i$ is further mapped by a two-layer head with a 128-dimensional hidden layer to two unconstrained outputs $(\widehat{\theta}_i, \widehat{s}_i)$. They are transformed as

$$\theta_i = \pi \tanh(\widehat{\theta}_i), \qquad s_i = \frac{1}{2}\log \kappa_{\max} \tanh(\widehat{s}_i), \qquad \kappa_{\max} = 20. \tag{31}$$

Here, θ_i determines the principal orientation, while s_i controls the anisotropic strength.

B.2 INDICATOR CONSTRUCTION

The three components are assembled into the symmetric positive-definite movement indicator

$$\boxed{\mathbf{M}_i = \rho_i \mathbf{R}(\theta_i) \begin{bmatrix} e^{s_i} & 0 \\ 0 & e^{-s_i} \end{bmatrix} \mathbf{R}(\theta_i)^\top}, \tag{32}$$

where

$$\mathbf{R}(\theta_i) = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix}. \tag{33}$$

Its two principal values are $\rho_i e^{s_i}$ and $\rho_i e^{-s_i}$, which directly give

$$\det(\mathbf{M}_i) = \rho_i^2, \qquad \kappa(\mathbf{M}_i) = e^{2|s_i|} \le \kappa_{\max}. \tag{34}$$

Thus, ρ_i determines the nodal determinant scale, whereas θ_i and s_i determine orientation and anisotropy at that fixed scale. This determinant identity is nodal: averaging differently oriented indicators into $\mathbf{M}_K$ does not, in general, preserve the square of the average nodal density. The implementation finally applies $\mathbf{M}_i \leftarrow (\mathbf{M}_i + \mathbf{M}_i^\top)/2$ to remove numerical asymmetry.

This construction complements the analysis in Appendix A: scalar density determines how much resolution is allocated locally, while the vectorized indicator additionally specifies how that resolution is distributed geometrically.

Directional interpretation. For a displacement $\mathbf{d}$ along a principal direction with indicator value μ, the squared mapped length is $\mu\|\mathbf{d}\|_2^2$. Consequently, a common target mapped length corresponds to a physical length proportional to $\mu^{-1/2}$. Larger principal values therefore request finer spacing in the corresponding direction. When $s_i = 0$, the indicator reduces to $\rho_i \mathbf{I}$ and orientation has no effect. In general, $\mathbf{M}_i$ specifies a local geometric demand rather than a displacement vector; actual movement follows from the coupled system in Sec. 4.3.

B.3 Connection to the Geometric Objective

Throughout this paper, the term *unified mapped space* refers to the common geometric representation induced by the movement indicators, rather than to a single explicitly integrated Euclidean coordinate map. The mapped space in Sec. 4.4 is implemented through local metric measurements. Each element uses the positive-definite average $\mathbf{M}_K$ to transform its Jacobian. A spatially varying indicator need not integrate to one global Euclidean coordinate map; the objective requires only the local Gram matrices and their common normalization.

Let $\sigma_{K,1}, \sigma_{K,2} > 0$ be the singular values of $\mathbf{M}_K^{1/2} \mathbf{J}_K \mathbf{E}_{\text{ref}}^{-1}$. The normalized Gram eigenvalues satisfy $\lambda_{K,r} = \sigma_{K,r}^2/\alpha$. In two dimensions,

$$a_K^M = \frac{|\det \mathbf{J}_K| \sqrt{\det \mathbf{M}_K}}{\det \mathbf{E}_{\text{ref}}} = \frac{4}{\sqrt{3}} |K| \sqrt{\det \mathbf{M}_K}. \tag{35}$$

Thus, a_K^M is proportional to metric-weighted physical area, while the ratio of the two eigenvalues measures deformation relative to an equilateral reference. Setting $x = \log \lambda_{K,1}$ and $y = \log \lambda_{K,2}$ gives $x^2 + y^2 = \frac{1}{2}(x+y)^2 + \frac{1}{2}(x-y)^2$. Equal mapped areas alone constrain $x + y$; the complete objective also constrains $x - y$.

The stop-gradient in Eq. 10 treats the mesh-wise reference scale as fixed during each backward pass. The condition $\overline{\mathbf{G}}_K = \mathbf{I}$ describes the ideal canonical target when the numerical stabilizers are neglected. Fixed boundaries, shared vertices, and spatially varying indicators can prevent simultaneous attainment. Moreover, a Gram matrix is insensitive to the sign of $\det \mathbf{J}_K$, so the geometric objective alone does not certify orientation preservation. That role belongs to the equilibrium construction under the conditions stated next.

Interpretation of the learned metric. Because the movement indicator and relocated mesh are optimized jointly, mapped-space regularity measures consistency between the resulting mesh and the learned geometric representation; it does not by itself establish that the learned principal directions coincide with a unique physical anisotropy. This co-adaptation is constrained by the prescribed density scale, bounded anisotropy, shared mesh connectivity, fixed boundaries, and global equilibrium. We therefore use mapped-space statistics to characterize the geometric mechanism, while the independent reconstruction errors and ablation results evaluate whether the learned directional representation improves adaptation to the underlying physical field.

C Tutte Equilibrium and Inversion-Free Guarantee

We provide the derivation of the Tutte equilibrium used in Sec. 4.3 and state the geometric condition under which the resulting mesh is inversion-free.

C.1 Tutte Equilibrium

Consider the mesh graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with positive symmetric edge weights $w_{ij} = w_{ji} > 0$. For a fixed forward pass, the predicted weights are held fixed during the coordinate solve. Interpreting each undirected edge, counted once, as a zero-rest-length linear spring with stiffness w_{ij}, the elastic energy is

$$\mathcal{E}_{\mathrm{T}}(\mathbf{C}') = \frac{1}{2} \sum_{(i,j)\in\mathcal{E}} w_{ij} \left\|\mathbf{c}'_i - \mathbf{c}'_j\right\|_2^2. \tag{36}$$

The Tutte equilibrium is obtained by minimizing Eq. 36 with the boundary vertices fixed. For each interior vertex $v_i \in \mathcal{V}_I$, the first-order stationary condition gives

$$\frac{\partial \mathcal{E}_{\mathrm{T}}}{\partial \mathbf{c}'_i} = \sum_{j\in\mathcal{N}(i)} w_{ij} \left(\mathbf{c}'_i - \mathbf{c}'_j\right) = \mathbf{0}. \tag{37}$$

Equation 37 is exactly the Hooke's-law force balance at vertex v_i. Equivalently,

$$\mathbf{c}'_i = \sum_{j\in\mathcal{N}(i)} \lambda_{ij} \mathbf{c}'_j, \qquad \lambda_{ij} = \frac{w_{ij}}{\sum_{k\in\mathcal{N}(i)} w_{ik}}, \tag{38}$$

where

$$\lambda_{ij} > 0, \qquad \sum_{j\in\mathcal{N}(i)} \lambda_{ij} = 1. \tag{39}$$

Thus, every interior coordinate is a convex combination of its neighbors with strictly positive coefficients. These coefficients couple neighboring positions; they do not specify independent vertex displacements.

Let $\mathbf{L} = \mathbf{D} - \mathbf{W}$ denote the weighted graph Laplacian. After partitioning the vertices into interior and boundary sets,

$$\mathbf{L} = \begin{bmatrix} \mathbf{L}_{II} & \mathbf{L}_{IB} \\ \mathbf{L}_{BI} & \mathbf{L}_{BB} \end{bmatrix}, \tag{40}$$

the equilibrium conditions for all interior vertices are compactly written as

$$\boxed{\mathbf{L}_{II}\mathbf{C}'_I = -\mathbf{L}_{IB}\mathbf{C}_B}, \qquad \mathbf{C}'_B = \mathbf{C}_B. \tag{41}$$

Therefore, the relocated interior coordinates are obtained simultaneously from a single equilibrium system rather than from independent vertex displacements.

Existence and uniqueness. Suppose every connected component containing an interior vertex also contains a fixed boundary vertex. Extend an interior scalar vector $\mathbf{x}_I$ by $\mathbf{x}_B = \mathbf{0}$. Then

$$\mathbf{x}_I^\top \mathbf{L}_{II} \mathbf{x}_I = \sum_{(i,j)\in\mathcal{E}} w_{ij}(x_i - x_j)^2 > 0 \quad \text{for } \mathbf{x}_I \neq \mathbf{0}. \tag{42}$$

Indeed, equality would force $\mathbf{x}$ to be constant along every edge and hence zero in each boundary-anchored component. Therefore, $\mathbf{L}_{II}$ is positive definite and the two coordinate columns have a unique solution. This algebraic uniqueness holds for fixed weights and boundary positions; it does not imply a unique learned indicator or a unique optimum of the training objective.

Differentiating the equilibrium. With the boundary coordinates fixed, differentiation of Eq. 41 gives

$$\mathbf{L}_{II}\, d\mathbf{C}'_I = -(d\mathbf{L}_{II})\mathbf{C}'_I - (d\mathbf{L}_{IB})\mathbf{C}_B. \tag{43}$$

This relation explains how the geometric loss can update the edge predictor and indicator through the coordinate solve. A local change in conductance can affect distant interior vertices through the same connected system, providing the global coordination used by GEMOVE.

C.2 INVERSION-FREE GUARANTEE

The convex-combination form in Eq. 38 provides the geometric guarantee of the Tutte construction.

Proposition. Consider a planar triangular mesh satisfying the standard Tutte conditions: the mesh graph is a planar 3-connected graph, the outer boundary cycle is fixed at the vertices of a strictly convex polygon in its cyclic order, and all edge weights are strictly positive. Then the equilibrium defined by Eq. 41 is a crossing-free straight-line embedding. Consequently, every triangular element has positive oriented area and no element inversion occurs.

Proof. Because $w_{ij} > 0$, Eq. 38 expresses every interior vertex as a strict convex combination of its neighbors. The resulting coordinates therefore constitute a barycentric embedding of the planar graph. Under these boundary and connectivity assumptions, Tutte's embedding theorem guarantees that this realization preserves planarity and maps every face to a nondegenerate convex polygon.

Since the mesh faces are triangles, each face remains a nondegenerate triangle. The fixed orientation of the outer boundary determines a consistent orientation for all faces in the planar embedding. Hence, for every element $K = (i, j, k)$ ordered consistently with the positive orientation of the domain,

$$\mathcal{A}'_K = \frac{1}{2} \det \begin{bmatrix} \mathbf{c}'_j - \mathbf{c}'_i & \mathbf{c}'_k - \mathbf{c}'_i \end{bmatrix} > 0, \tag{44}$$

which excludes element inversion.

Scope of the guarantee. The proposition is a sufficient-condition result. A unique equilibrium with positive weights does not by itself establish an inversion-free embedding for a nonconvex boundary or a domain with holes. Collinear boundary vertices also require care beyond the strictly convex statement above. This distinction matters for polygonal domains and coastline-following wet meshes: their empirical reconstruction results should not be read as an unconditional extension of the theorem. The guarantee concerns element orientation, not a uniform lower bound on element angles or a guarantee of reduced PDE error. The theoretical guarantee concerns planarity and element orientation under the stated Tutte conditions; it does not by itself guarantee a particular resolution distribution or exclude excessive local concentration. The reduction of over-allocation observed in our experiments is therefore an empirical effect of global coordination rather than a consequence of the Tutte embedding theorem.

D SYNTHETIC PDE EXPERIMENTAL DETAILS

The experimental details referenced in the main text are organized here and in the following appendices. This section specifies the synthetic setup; Appendix E gives the parameter-wise results, and Appendix F defines the ocean and global-transfer data and metrics.

PDEs and data. We consider four PDE families with complementary solution structures: Poisson and Helmholtz are included during training, while Swirl and Burgers are held out for cross-equation evaluation. Their governing equations are

$$\begin{aligned} -\Delta u &= f, && \text{Poisson,} \\ -\Delta u + u &= f, && \text{Helmholtz,} \\ \frac{\partial u}{\partial t} + \mathbf{c} \cdot \nabla u &= 0, && \text{Swirl,} \\ \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} - \nu \Delta \mathbf{u} &= \mathbf{0}, && \text{Burgers.} \end{aligned} \tag{45}$$

Poisson and Helmholtz use exact solutions constructed from random mixtures of one to five anisotropic Gaussian components with corresponding Dirichlet boundary conditions. The Swirl velocity-field parameter takes values $\kappa \in \{1.0, 1.5\}$, matching the notation in Table 4. This notation distinguishes it from the normalization α in Eq. 10. Burgers uses one- and two-Gaussian initial conditions with $\nu \in [10^{-4}, 2 \times 10^{-3}]$. The static problems are generated on unstructured triangular meshes over the unit square with

$$l_c \in \{0.040,\ 0.045,\ 0.050,\ 0.055\}, \tag{46}$$

and fine reference solutions use $l_c = 0.01$. Each static PDE-resolution setting contains 400 samples, split into $300/40/60$ training, validation, and test samples. Independently generated polygonal Poisson and Helmholtz meshes are further used to evaluate generalization to unseen boundary geometries and mesh topologies.

Generalization protocol. The two transfer tests isolate different changes. Swirl and Burgers change the governing equation while retaining the rectangular-domain evaluation setting. Polygonal Poisson and Helmholtz retain the training PDE families while changing the domain geometry and connectivity. Neither test introduces additional vertices during movement. Within each comparison, the initial mesh, boundary constraints, and node budget are shared across methods.

Features and training. Each vertex is represented by coordinates, boundary indicators, the solution value, its gradient, and Hessian magnitude, together with regular-grid representations of the solution and Hessian fields. The movement-indicator construction is detailed in Appendix B. GEMOVE is trained with Adam for 1000 epochs using a learning rate of 5×10^{-5}, weight decay of 5×10^{-4}, training/validation batch sizes of $20/10$, and gradient clipping at unit norm. The checkpoint with the lowest validation geometric loss is retained. M2N and UM2N use MA-generated adaptive meshes as supervision, whereas UGM2N and GEMOVE require no target meshes.

Reference solutions. The synthetic PDE experiments use high-resolution numerical solutions as the reference for evaluating mesh-movement accuracy. For all four PDE families, the reference solution is obtained by solving the corresponding PDE with Firedrake on a fine triangular mesh with $l_c = 0.01$. The solution reconstructed from each original or moved coarse mesh is compared against this fine-mesh solution when computing the Fine-L_2 error and the corresponding Error Red.

For Poisson and Helmholtz, the prescribed solutions are additionally constructed analytically from random mixtures of anisotropic Gaussian components. These analytical solutions provide an independent reference for verifying the numerical solutions produced by Firedrake. We emphasize that the Monge–Ampère-generated adaptive meshes used by M2N and UM2N serve only as supervision targets for training these baselines. They are not used as reference PDE solutions or ground truth when evaluating numerical accuracy.

Evaluation. Numerical accuracy is measured against the corresponding fine-grid reference solution. Let E_{orig} and E_{method} denote the Fine-L_2 errors before and after mesh movement, respectively. We report

$$\text{Error Red.} = 100 \left(1 - \frac{E_{\text{method}}}{E_{\text{orig}}}\right) \%. \tag{47}$$

Positive values indicate reduced numerical error. Results are averaged over the samples and parameter configurations associated with each evaluation setting. The overall rectangular-domain score in the main text assigns equal weight to the four PDE families. In particular, the two Swirl settings are first averaged within the Swirl family, rather than counted as two separate PDE families. Percentage-point differences compare methods at the same aggregation level; a difference of two average scores is not generally the average gain over a baseline selected separately in every row.

D.1 Ablation Settings

We construct three ablated variants to isolate the contributions of mapped-space size, directional geometry, and global mesh coordination. All variants use the same training data, input features, model configuration, and optimization protocol as the complete GEMove, unless otherwise specified.

Only-Size. The complete geometric objective contains two complementary components: mapped-space element-measure equidistribution and directional regularization. The *Only-Size* variant retains only the size component,

$$\mathcal{L}_{\text{size}} = \frac{1}{|\mathcal{K}|} \sum_{K \in \mathcal{K}} 2 \left[\log \frac{a_K^M}{\alpha} \right]^2, \tag{48}$$

which encourages different elements to have similar measures in the mapped space while removing the constraint on directional deformation.

Only-Direction. The *Only-Direction* variant retains only the directional component,

$$\mathcal{L}_{\text{direction}} = \frac{1}{|\mathcal{K}|} \sum_{K \in \mathcal{K}} \frac{1}{2} \left[\log \frac{\lambda_{K,1}}{\lambda_{K,2}} \right]^2, \tag{49}$$

which penalizes anisotropic distortion relative to the canonical mapped-space geometry while removing explicit element-measure equidistribution. These two variants therefore separate the size and directional components that are jointly optimized by the complete geometric objective.

Without global coordination. To evaluate the contribution of the Tutte equilibrium decoder, the *w/o Coordination* variant replaces the globally coupled relocation module in Sec. 4.3 with a direct coordinate decoder. The contextual node representation $\mathbf{h}_i$ produced by the graph encoder is passed to an MLP that predicts the relocated coordinate of each interior vertex:

$$\mathbf{c}_i' = g_\phi(\mathbf{h}_i), \qquad v_i \in \mathcal{V}_I. \tag{50}$$

Here, each coordinate is decoded independently from its contextual node representation, without the mesh-wide equilibrium solve or explicit coupling between the final relocation decisions. Boundary vertices remain fixed as in the complete model.

The *w/o Coordination* variant preserves the same movement-indicator construction and the complete geometric objective $\mathcal{L}_{\text{geo}}$ in Eq. 13. The comparison therefore evaluates the effect of replacing the Tutte-based equilibrium decoder with direct node-wise coordinate prediction while retaining the same encoder, movement-indicator construction, geometric objective, and boundary treatment. Because the two decoders realize geometric information through different computational pathways, this ablation should be interpreted as evaluating the contribution of equilibrium-based relocation rather than isolating global coordination as a single independent factor.

E Additional Synthetic PDE Results

We provide additional quantitative and qualitative results for the synthetic PDE benchmarks. The detailed parameter-wise results complement the averages reported in the main text, while the mesh visualizations further illustrate the adaptation patterns of different methods across both seen and unseen PDE families.

E.1 Qualitative Mesh Adaptation

Figure 8 compares the moved meshes produced by all methods on representative Poisson, Helmholtz, Burgers, and Swirl fields. Poisson and Helmholtz are included during training, whereas Burgers and Swirl are unseen during training. Across these different solution structures, GEMove consistently adapts the mesh to regions with strong spatial variation while retaining coordinated mesh geometry.

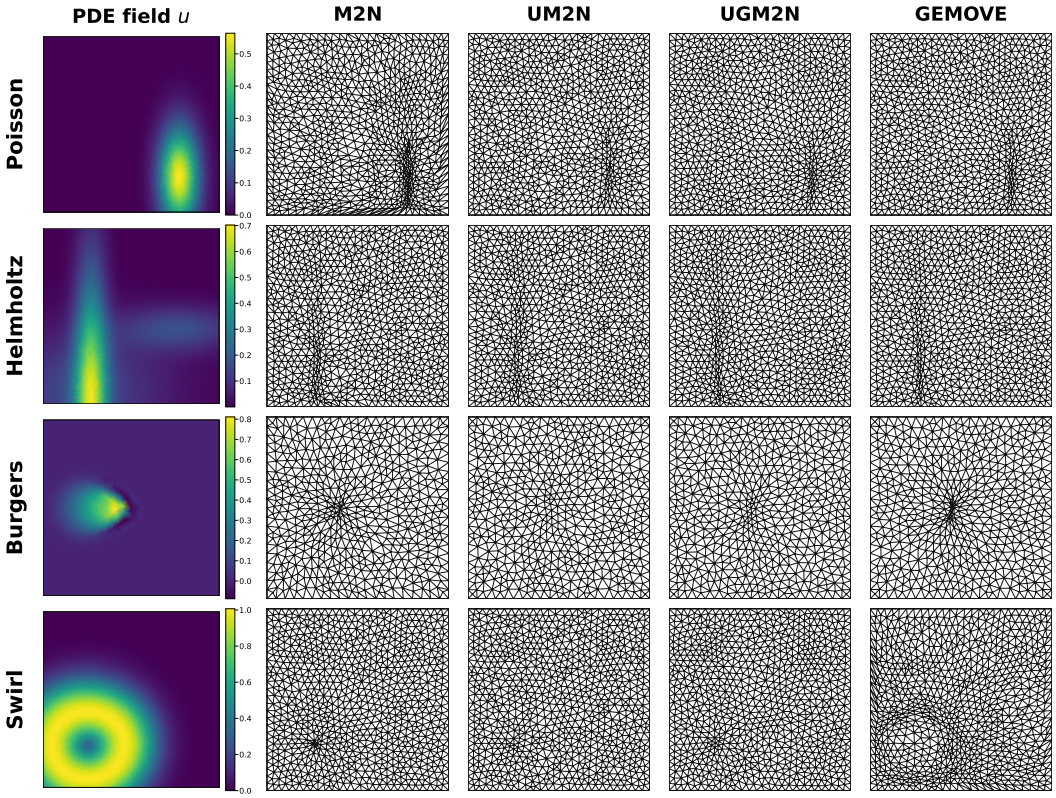

Figure 8: Representative mesh adaptation across the four PDE families. The first column shows the corresponding PDE field, followed by the moved meshes produced by M2N, UM2N, UGM2N, and GEMove. Poisson and Helmholtz are included during training, while Burgers and Swirl are evaluated zero-shot.

E.2 Parameter-wise PDE Results

Table 4 reports the complete results across mesh resolutions and PDE configurations. The advantage of GEMove is particularly pronounced on the unseen Swirl and Burgers problems, where changes in solution structure introduce substantially stronger distribution shifts than those in the training PDEs.

Dependence on the PDE and mesh scale. GEMove gives positive error reduction in every reported rectangular setting. The two Swirl configurations show that the aggregate advantage is not confined to one velocity-field parameter: the reported gains over the strongest baseline are 15.37 and 10.53 percentage points, respectively. For Burgers, the error reduction remains above 32% at all four mesh scales. These results support transfer across the tested solution structures, while the parameter-wise table also shows that gains need not vary monotonically with l_c.

For comparison under a setting-wise weighting scheme, averaging Poisson, Helmholtz, both Swirl settings, and Burgers as five equally weighted entries gives 5.49%, 12.17%, 13.83%, and 23.20% for M2N, UM2N, UGM2N, and GEMove, respectively, with a reported gain of 9.38 percentage points over the strongest baseline. That convention assigns twice as much weight to Swirl as to

Table 4: Fine-L_2 error reduction (%) across four PDE families, with the two Swirl parameter settings shown separately. The overall row averages the four families after averaging the two Swirl settings. † denotes PDE families included during training. Best and second-best results are shown in bold and underlined, respectively; negative improvements are shown in gray. Δ is the percentage-point gain of GEMOVE over the strongest baseline.

PDE	Setting	M2N	UM2N	UGM2N	GEMOVE	Δ (pp)
Poisson†	$l_c = 0.040$	10.45 ± 1.08	21.80 ± 0.79	20.85 ± 0.85	**24.65 ± 0.71**	**+2.85**
	$l_c = 0.045$	0.19 ± 1.24	20.08 ± 0.85	20.04 ± 0.91	**23.43 ± 0.78**	**+3.35**
	$l_c = 0.050$	12.33 ± 1.10	20.83 ± 0.81	21.01 ± 0.87	**23.44 ± 0.72**	**+2.43**
	$l_c = 0.055$	3.93 ± 1.18	20.32 ± 0.83	20.60 ± 0.89	**23.13 ± 0.75**	**+2.53**
	Avg.	6.72 ± 0.85	20.76 ± 0.61	20.62 ± 0.65	**23.66 ± 0.55**	**+2.91**
Helmholtz†	$l_c = 0.040$	10.43 ± 1.03	21.78 ± 0.77	20.86 ± 0.83	**24.65 ± 0.70**	**+2.87**
	$l_c = 0.045$	9.21 ± 1.09	20.37 ± 0.84	19.41 ± 0.90	**24.16 ± 0.75**	**+3.79**
	$l_c = 0.050$	12.32 ± 1.07	20.81 ± 0.78	21.02 ± 0.84	**23.45 ± 0.71**	**+2.43**
	$l_c = 0.055$	3.91 ± 1.13	20.30 ± 0.81	20.61 ± 0.87	**23.14 ± 0.72**	**+2.53**
	Avg.	8.97 ± 0.80	20.82 ± 0.59	20.48 ± 0.64	**23.85 ± 0.53**	**+3.04**
Swirl ($\kappa = 1.0$)	$l_c = 0.040$	-4.65 ± 1.50	-4.94 ± 1.29	3.82 ± 1.14	**20.41 ± 0.97**	**+16.59**
	$l_c = 0.045$	-3.58 ± 1.44	0.84 ± 1.26	4.28 ± 1.11	**21.61 ± 0.99**	**+17.33**
	$l_c = 0.050$	-6.20 ± 1.52	5.16 ± 1.22	3.38 ± 1.13	**17.16 ± 0.93**	**+12.00**
	$l_c = 0.055$	0.19 ± 1.34	1.40 ± 1.19	3.28 ± 1.10	**17.05 ± 0.91**	**+13.77**
	Avg.	-3.56 ± 1.08	0.62 ± 0.92	3.69 ± 0.83	**19.06 ± 0.70**	**+15.37**
Swirl ($\kappa = 1.5$)	$l_c = 0.040$	-1.79 ± 1.44	5.60 ± 1.19	-1.04 ± 1.10	**18.57 ± 0.91**	**+12.97**
	$l_c = 0.045$	-1.86 ± 1.39	4.89 ± 1.15	1.10 ± 1.07	**16.51 ± 0.89**	**+11.62**
	$l_c = 0.050$	-3.02 ± 1.41	5.65 ± 1.17	6.02 ± 1.08	**14.05 ± 0.87**	**+8.03**
	$l_c = 0.055$	-0.95 ± 1.32	5.91 ± 1.13	6.22 ± 1.07	**15.03 ± 0.89**	**+8.81**
	Avg.	-1.91 ± 1.03	5.51 ± 0.86	3.08 ± 0.80	**16.04 ± 0.66**	**+10.53**
Burgers	$l_c = 0.040$	26.11 ± 1.32	13.68 ± 1.08	22.88 ± 1.02	**34.78 ± 0.86**	**+8.67**
	$l_c = 0.045$	17.54 ± 1.29	12.86 ± 1.04	23.38 ± 1.00	**33.02 ± 0.83**	**+9.64**
	$l_c = 0.050$	13.78 ± 1.27	13.30 ± 1.05	19.26 ± 0.99	**33.19 ± 0.84**	**+13.93**
	$l_c = 0.055$	11.40 ± 1.24	12.72 ± 1.03	19.53 ± 0.99	**32.62 ± 0.83**	**+13.09**
	Avg.	17.21 ± 0.95	13.14 ± 0.78	21.26 ± 0.74	**33.40 ± 0.62**	**+12.14**
Overall	*Family avg.*	7.54 ± 0.65	14.45 ± 0.51	16.44 ± 0.50	**24.62 ± 0.42**	**+8.18**

each other family. We use the four-family value 24.62% in the overall row to match the main-text comparison.

E.3 Generalization across Polygonal Domains

Table 5 further reports the complete resolution-wise results on unseen polygonal domains. While the baselines provide only limited or negative improvements in several settings, GEMove maintains positive error reduction across all tested resolutions.

Table 5: Fine-L_2 error reduction (%) on unseen polygonal domains. Best and second-best results are shown in bold and underlined, respectively; negative improvements are shown in gray. Δ is the percentage-point gain of GEMove over the strongest baseline.

PDE	Setting	M2N	UM2N	UGM2N	GEMove	Δ (pp)
Poisson	$l_c = 0.040$	-9.53 ± 1.55	-3.06 ± 1.24	2.24 ± 1.04	**9.16 ± 0.88**	**+6.92**
	$l_c = 0.045$	-3.95 ± 1.46	-2.81 ± 1.20	1.79 ± 1.00	**8.22 ± 0.84**	**+6.43**
	$l_c = 0.050$	-5.12 ± 1.49	-1.45 ± 1.18	1.31 ± 0.99	**11.65 ± 0.87**	**+10.34**
	$l_c = 0.055$	-6.73 ± 1.42	0.18 ± 1.14	2.06 ± 1.01	**10.49 ± 0.85**	**+8.43**
	Avg.	-6.33 ± 1.10	-1.79 ± 0.89	1.85 ± 0.75	**9.88 ± 0.64**	**+8.03**
Helmholtz	$l_c = 0.040$	-4.05 ± 1.38	1.37 ± 1.10	2.86 ± 0.98	**11.34 ± 0.83**	**+8.48**
	$l_c = 0.045$	-3.62 ± 1.36	2.75 ± 1.08	3.07 ± 0.96	**10.13 ± 0.82**	**+7.06**
	$l_c = 0.050$	-1.68 ± 1.31	2.38 ± 1.05	3.49 ± 0.94	**9.35 ± 0.80**	**+5.86**
	$l_c = 0.055$	-3.74 ± 1.35	3.61 ± 1.09	2.94 ± 0.96	**12.27 ± 0.83**	**+8.66**
	Avg.	-3.27 ± 1.00	2.53 ± 0.81	3.09 ± 0.72	**10.77 ± 0.61**	**+7.68**
Overall	*Avg.*	-4.80 ± 0.76	0.37 ± 0.61	2.47 ± 0.55	**10.33 ± 0.46**	**+7.86**

The polygonal results provide a distinct test of geometric transfer. GEMove retains a positive gain for both equations at every reported resolution, whereas some baseline scores become negative. The overall 10.33% reduction uses equal weight for Poisson and Helmholtz and matches Table 1. These accuracy comparisons concern the evaluated meshes; the theoretical validity conditions remain those in Appendix C.2.

E.4 Ablation Results Across Runs

Table 6 reports the ablation results across independent runs, including the corresponding standard deviations. The complete GEMove model achieves the strongest average performance across all four PDE families with consistently low variability.

Table 6: Ablation results across independent runs (Error Red., %, mean ± std).

Model	Pois.	Helm.	Swirl	Burg.	Avg.
Only-Size	20.35 ± 1.08	20.03 ± 1.05	7.82 ± 1.43	26.04 ± 1.19	18.56 ± 0.93
Only-Direction	19.51 ± 0.97	18.94 ± 0.95	13.39 ± 1.18	24.71 ± 1.07	19.14 ± 0.84
w/o Coord.	21.82 ± 0.91	22.16 ± 0.89	15.70 ± 1.12 *	31.39 ± 0.96	22.77 ± 0.78
GEMove	**23.66 ± 0.55**	**23.85 ± 0.53**	**17.55 ± 0.51**	**33.40 ± 0.62**	**24.62 ± 0.42**

E.5 Geometric Redistribution in the Mapped Space

We further examine how mesh movement changes element geometry in the mapped space across representative PDE examples. In addition to the qualitative mesh adaptation shown in Fig. 8, we compare the distributions of mapped-space element area and edge aspect ratio before and after movement. According to the geometric equidistribution principle, successful adaptation should reduce the dispersion of element areas while driving element shapes toward the canonical geometry. Figures 9 and 10 show that this geometric redistribution consistently emerges across different PDE fields.

Reading area and shape together. Figure 9 should be interpreted through both geometric statistics. A concentrated area distribution alone is compatible with the equal-area distortions constructed in Appendix A. More regular edge ratios provide complementary evidence that the movement addresses shape as well as allocation. Together, these diagnostics characterize mapped-space size redistribution and directional regularity, although edge aspect ratio is a visual shape statistic rather than the training objective itself.

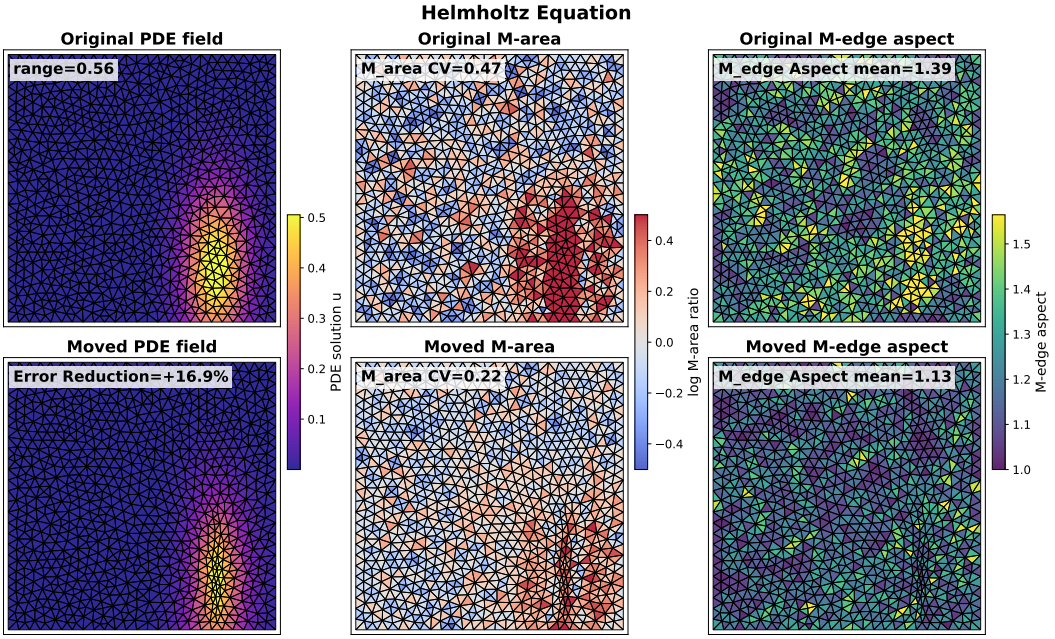

Figure 9: Mapped-space geometric redistribution for a representative PDE example. After mesh movement, the mapped-space element-area distribution becomes more concentrated, while edge aspect ratios move toward more regular element geometry.

Figure 10 examines the same mechanism on a Poisson example. Physical-space elements need not become equal-sized or equilateral: an anisotropic indicator can map a stretched physical triangle to a regular metric-space triangle. The relevant comparison is therefore geometric consistency after accounting for $\mathbf{M}_K$, rather than uniformity in physical coordinates.

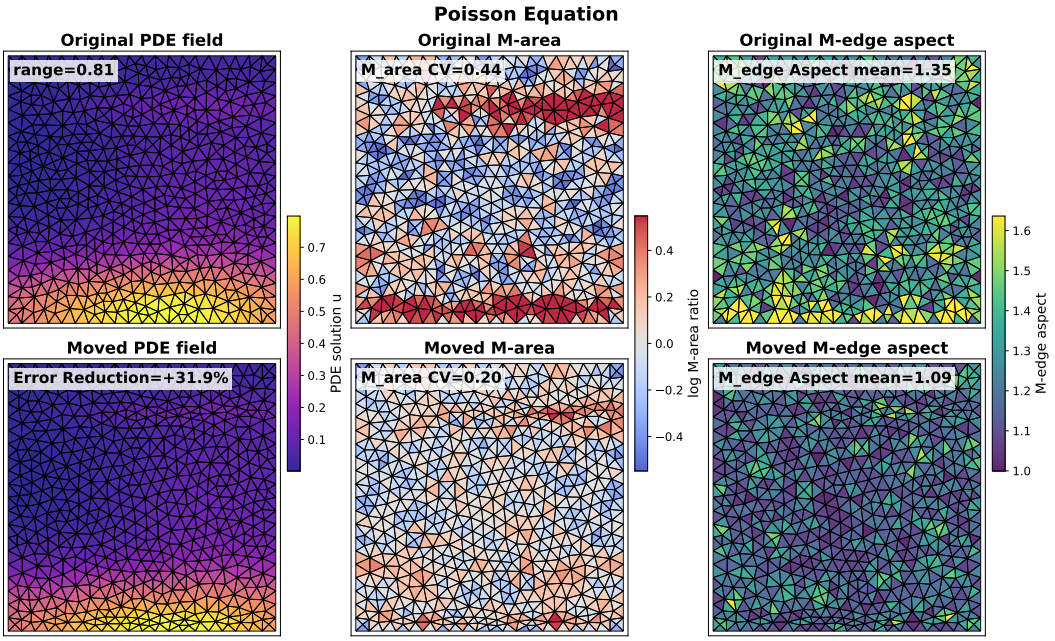

Figure 10: Mapped-space geometric redistribution for a representative Poisson example. The moved mesh exhibits a narrower mapped-space element-area distribution and more regular edge aspect ratios than the original mesh.

Together, the two examples illustrate the mapped-space geometric mechanism on selected fields, while Tables 4 and 5 provide the broader accuracy comparisons. Regularity in the mapped space

and numerical error reduction are complementary observations; neither statistic alone establishes the other. In particular, mapped-space regularity should be interpreted as consistency with the learned movement indicator rather than as an independent measurement of physical directional alignment.

E.6 Inference Efficiency

We further evaluate the inference efficiency of the learning-based mesh movement methods on the four PDE benchmarks. The timing records cover Helmholtz, Poisson, Swirl, and Burgers across 17 data configurations. After removing the first CUDA warm-up sample from each configuration, the reported statistics contain 238 evaluated samples under the same runtime environment.

Table 7 reports the wall-clock time for one complete mesh-movement step. The measured time includes the full forward procedure of each method. For GEMove, this includes graph encoding, movement-indicator and edge conductance prediction, and the Tutte equilibrium solve.

Table 7: Single-step inference time on the four PDE benchmarks. Results are averaged over 238 samples from 17 data configurations after removing the first CUDA warm-up sample of each configuration. All methods are measured under the same runtime environment.

Method	**Inference Time (ms)**
M2N	23.46
UM2N	22.72
UGM2N	24.97
GEMove	25.11

All four methods require approximately 23–25 ms for one complete mesh-movement step under the same runtime environment. The small differences among M2N, UM2N, UGM2N, and GEMove indicate comparable inference costs across the evaluated PDE settings. Thus, the additional geometric representation and global Tutte coordination in GEMove introduce only limited runtime overhead.

The timings are obtained from existing wall-clock measurements using `time.time()` without explicit CUDA synchronization. They are therefore intended for relative runtime comparison under the same experimental environment rather than as a strict GPU-latency benchmark.

F Real-World Earth-System Data and Experimental Details

We provide additional details for the real-world experiments in Sec. 5.4 and Sec. 5.5. The experiments are organized in two stages. First, UGM2N and GEMove are retrained in an unsupervised manner using temporally selected subsets of the Kuroshio and ENSO records. Kuroshio tests spatial allocation around a narrow energetic jet, while ENSO tests reconstruction across long-term changes in SST anomalies. Second, the resulting ocean-trained models are kept fixed and directly transferred to the global Earth-system datasets in RQ4, without target-region or target-variable fine-tuning.

The following subsections define the data sources, field representations, spatial domains, and evaluation metrics. RQ3 uses ocean-specific unsupervised retraining, whereas RQ4 evaluates transfer of those fixed checkpoints. These protocols are distinct from the synthetic cross-equation test in RQ1.

F.1 Real-World Ocean Data

We use daily Kuroshio ocean-current fields from the Copernicus Marine global ocean reanalysis (Lellouche et al., 2021; Copernicus Marine Service, 2026c) and monthly ENSO SST anomalies derived from NOAA OISST v2.1 (Huang et al., 2021; 2020). Figure 11 shows the spatial domains used for the two ocean experiments. UGM2N and GEMove are retrained without target meshes on a temporal subset of each record. The spatial and temporal analyses retain the natural chronological order of the available records. The resulting checkpoints are then used without parameter updates in RQ4 (Appendix F.4).

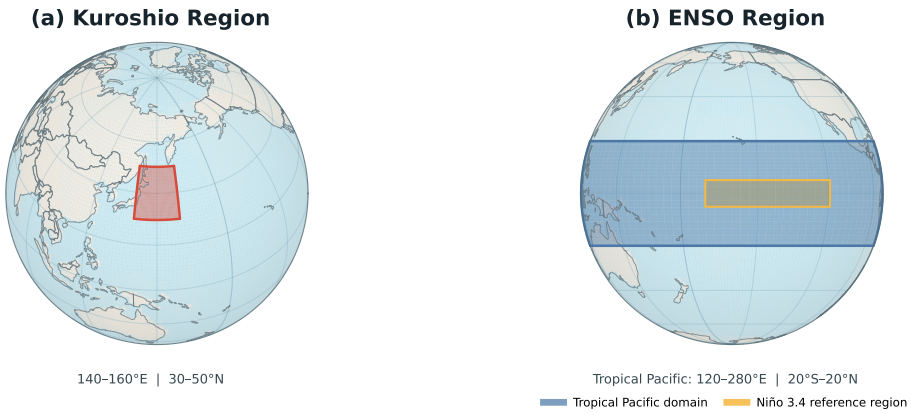

Figure 11: Spatial domains of the two real-world ocean experiments. (**a**) Kuroshio region. (**b**) Tropical Pacific domain used for ENSO, with the Niño 3.4 reference region highlighted.

F.2 Kuroshio Current Data and Evaluation

Data and spatial domain. The Kuroshio experiment uses daily ocean-current fields from the Copernicus Marine global ocean reanalysis (Lellouche et al., 2021; Copernicus Marine Service, 2026c). We extract a western North Pacific domain covering $140°$–$160°$E and $30°$–$50°$N. The held-out evaluation record spans January 1, 2020 to June 30, 2021 and contains 547 daily snapshots.

The domain contains the Kuroshio Extension and its surrounding subtropical–subpolar transition. A narrow high-speed jet, meanders, fronts, and mesoscale structures (Qiu & Chen, 2005) create competing adaptation demands at different scales and orientations.

Physical variables and mesh features. The ocean velocity field is represented by its zonal and meridional components,

$$\mathbf{v} = (u, v)^\top, \tag{51}$$

where u and v denote the eastward and northward velocity components, respectively. We further define the current speed as

$$s = \sqrt{u^2 + v^2}, \tag{52}$$

which provides a scalar measure of local current intensity.

The speed field is used to construct the physical features supplied to the mesh-movement model. In particular, its local gradient describes the strength and direction of first-order spatial variation, while Hessian-derived information captures higher-order variation associated with fronts, curved jets, and localized mesoscale features. These quantities provide the physical information from which the movement indicator is constructed.

Mesh construction and spatial resolution. For the Kuroshio experiments, the computational mesh is constructed from a coarse ocean field with an effective spatial resolution of approximately $1/3°$. The higher-resolution ocean field is retained only as the reference field for reconstruction-based evaluation.

Only wet ocean points are retained when constructing the triangular mesh; land and missing-value locations are excluded. This produces an irregular wet-domain mesh whose boundary follows both the outer geographic window and coastlines represented by the land–sea mask. The geographic coordinates of the retained coarse-grid points are mapped to a normalized computational domain before mesh movement. After relocation, the ocean field is reconstructed on the high-resolution geographic grid for evaluation.

Kuroshio training protocol. The Kuroshio portion of the joint ocean training consists of 634 ocean-current snapshots sampled every 15 days from 1993–2018, while 30 samples from 2019 are

reserved for validation. Each sample is constructed from the ocean-current components (u, v) by forming the current-speed field, the wet-node triangular mesh, and the corresponding gradient and Hessian features. The model inputs include coordinates, boundary indicators, physical-field features, gradients, Hessian norms, and convolutional features. Training uses only the geometric objective in Eq. 13, without pre-adapted target meshes, Monge–Ampère meshes, high-resolution reference fields, reconstruction-error supervision, movement penalties, or inversion penalties. We optimize the model with Adam using a learning rate of 5×10^{-5}, weight decay of 5×10^{-4}, batch size 2, gradient clipping at unit norm, and a cosine-annealing learning-rate schedule. Training runs for at most 1000 epochs, and applies early stopping after 20 epochs without validation improvement.

Reference field. The Kuroshio experiment uses the original high-resolution ocean-current field as the reference for reconstruction-based evaluation. The reference field is defined on the fine grid and is represented by the zonal and meridional velocity components (u, v). After mesh movement, the velocity field reconstructed from the moved mesh is interpolated back onto the same fine grid and compared with this reference field. Therefore, the evaluation ground truth is the uncoarsened high-resolution ocean-current field rather than any pre-adapted mesh.

Vector-field reconstruction error. Let $\mathbf{v}_{p,t} = (u_{p,t}, v_{p,t})$ and $\widehat{\mathbf{v}}_{p,t}$ denote the reference and reconstructed velocity vectors at fine-grid location p and time t. We compute an area-weighted fine-grid RMSE as

$$E_t = \sqrt{\frac{\sum_{p \in \Omega_{\text{wet}}} A_p \left\| \widehat{\mathbf{v}}_{p,t} - \mathbf{v}_{p,t} \right\|_2^2}{\sum_{p \in \Omega_{\text{wet}}} A_p}}, \tag{53}$$

where A_p denotes the geographic cell-area weight and Ω_{wet} denotes the valid ocean region. Area weighting accounts for the latitude-dependent physical area represented by geographic grid cells.

The per-snapshot error reduction is

$$\text{ErrorRed}_t = 100 \left(1 - \frac{E_t^{\text{method}}}{E_t^{\text{orig}}} \right) \%. \tag{54}$$

The aggregate Error Red. reported in the main text summarizes these snapshot-wise improvements over the analysis period. Each snapshot-wise percentage uses the original-mesh error for the same date.

Win Rate. Average error reduction measures overall reconstruction improvement, but does not directly describe temporal consistency. We therefore additionally report Win Rate,

$$\text{WinRate} = \frac{1}{T} \sum_{t=1}^{T} \mathbb{I} \left[E_t^{\text{method}} < E_t^{\text{orig}} \right], \tag{55}$$

which measures the fraction of evaluated snapshots for which mesh movement reduces reconstruction error relative to the original mesh. The main table expresses this fraction as a percentage. A 100% Win Rate means every evaluated snapshot improves in this aggregate error; it does not imply that every location within each snapshot has lower error.

Kuroshio-core overlap. A domain-averaged reconstruction metric alone does not indicate whether the high-speed current core is correctly represented. We therefore additionally measure the spatial overlap of the reconstructed Kuroshio core.

For each snapshot, the core threshold τ_t is defined as the 85th percentile of the reference fine-grid speed over wet ocean points. The reference and reconstructed core masks are

$$\mathcal{C}_t = \{ p : s_{p,t} \geq \tau_t \}, \qquad \widehat{\mathcal{C}}_t = \{ p : \widehat{s}_{p,t} \geq \tau_t \}. \tag{56}$$

The threshold is determined from the reference field and then applied to both fields on the same valid wet-grid support. Using a common threshold avoids redefining the target core separately for each reconstruction.

CoreIoU is defined as

$$\mathrm{CoreIoU}_t = \frac{|\mathcal{C}_t \cap \widehat{\mathcal{C}}_t|}{|\mathcal{C}_t \cup \widehat{\mathcal{C}}_t|}. \tag{57}$$

This metric complements the reconstruction error by directly evaluating whether mesh adaptation preserves the spatial organization of the energetic current core.

F.3 ENSO Data and Evaluation

SST anomaly data. The ENSO experiment uses NOAA OISST v2.1 (Huang et al., 2021; 2020). OISST provides a long-term, spatially complete sea-surface temperature record by combining satellite observations with in-situ measurements from multiple observing platforms on a regular 0.25° global grid.

We use monthly SST fields beginning in January 1982. For each calendar month, the SST anomaly is computed relative to the corresponding 1991–2020 monthly climatology. The resulting anomaly field is denoted by $T'(\lambda, \varphi, t)$, where λ, φ, and t represent longitude, latitude, and month, respectively.

The chronological monthly sequence contains warm, neutral, and cold states with different spatial anomaly patterns. A temporal subset is used for unsupervised retraining, and the continuous record supports the long-term reconstruction analysis in RQ3.

Spatial resolutions. We evaluate mesh adaptation at coarse resolutions of 2°, 3°, and 4°. Each coarse representation is constructed directly from the native 0.25° SST anomaly field through area-aware aggregation rather than by successively downsampling one coarse representation into another. The native 0.25° field is retained as the reference for reconstruction and diagnostic calculation.

These resolutions progressively reduce the available spatial degrees of freedom. At each resolution, the moved and original meshes use the same node budget. Error reductions are computed against that resolution's own original mesh, so a larger percentage reduction on a coarser mesh does not by itself imply lower absolute error than on a finer mesh.

Niño 3.4 index. The principal climate diagnostic used in the ENSO analysis is the Niño 3.4 SST anomaly. The Niño 3.4 region is defined as

$$\mathcal{R}_{3.4} = [5^\circ\mathrm{S}, 5^\circ\mathrm{N}] \times [170^\circ\mathrm{W}, 120^\circ\mathrm{W}]\,. \tag{58}$$

Under the 0°–360° longitude convention used by the data pipeline, this corresponds to 190°–240°E (Trenberth, 1997).

The reference Niño 3.4 index is computed as an area-weighted average of the SST anomaly over the region:

$$I_t = \frac{\sum_{p \in \mathcal{R}_{3.4}} A_p T'_{p,t}}{\sum_{p \in \mathcal{R}_{3.4}} A_p}. \tag{59}$$

The same spatial weighting is applied to the reconstructed SST anomaly field to obtain the predicted index $\widehat{I}_t$. Using the same regional averaging operator for the reference and reconstructed fields isolates differences caused by spatial reconstruction rather than by the index definition itself.

ENSO phase definition. To characterize temporal performance under different background climate states, we form a centered three-month mean of the reference Niño 3.4 index:

$$\overline{I}_t = \frac{I_{t-1} + I_t + I_{t+1}}{3}. \tag{60}$$

For months with the three-month window available, the reference mean, measured in degrees Celsius, determines one of three states:

$$\begin{aligned} \overline{I}_t \geq 0.5 \quad &: \text{El Niño}, \\ -0.5 < \overline{I}_t < 0.5 \quad &: \text{Neutral}, \\ \overline{I}_t \leq -0.5 \quad &: \text{La Niña}. \end{aligned} \tag{61}$$

The phase label is always determined from the reference SST record rather than from a model reconstruction. This definition is used only to stratify the mesh-movement evaluation across warm, neutral, and cold ENSO conditions. It does not impose the additional persistence criterion used by the official Oceanic Niño Index event definition and should therefore not be interpreted as an independent event-classification procedure.

ENSO training protocol. The ENSO portion of the joint ocean training consists of monthly SST-anomaly fields from 1982–1984 at 2° and 3° coarse resolutions. The monthly fields from 1985 at the same two resolutions are used for validation, resulting in 24 validation samples. All months from January 1986 onward are held out from training and validation for temporal evaluation. The SST anomaly is computed by subtracting the corresponding 1991–2020 monthly climatology from the monthly mean SST, after which the native 0.25° field is conservatively aggregated directly to each coarse resolution and converted into an ocean-domain triangular mesh according to ocean coverage. The model uses the same input features and geometric objective as in the Kuroshio experiment and does not use target meshes, high-resolution reference fields, or reconstruction errors as supervision. Optimization uses Adam with a learning rate of 5×10^{-5}, weight decay of 5×10^{-4}, batch size 2, unit-norm gradient clipping, a cosine-annealing and learning-rate schedule. Training runs for at most 1000 epochs with early stopping after 20 epochs without validation improvement.

Reference field. The ENSO experiment uses the native 0.25° NOAA OISST v2.1 SST anomaly field as the reference for reconstruction-based evaluation. For each month, the anomaly is computed by subtracting the corresponding 1991–2020 monthly climatology from the monthly mean SST. The reference Ni no 3.4 index is then obtained by taking an ocean-area-weighted average of this high-resolution anomaly field over the Ni no 3.4 region. The SST anomaly reconstructed from the moved mesh is evaluated against the same native-resolution reference field, and its reconstructed Ni no 3.4 index is computed using the same area-weighted regional averaging operator.

Temporal evaluation protocol. All quantitative ENSO results are computed on the held-out record from January 1986 onward. For temporal context, the complete record beginning in 1982 is shown in the long-term visualizations, including the short training and validation periods, which are excluded from the reported held-out summary metrics. The 2° and 3° settings therefore evaluate temporal generalization at seen spatial resolutions, while the 4° setting additionally evaluates generalization to an unseen resolution.

Niño 3.4 reconstruction error. For an analysis interval containing T months, we compute

$$\mathrm{RMSE}_{3.4} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\widehat{I}_t - I_t\right)^2}. \tag{62}$$

The resolution-wise Error Red. is defined relative to the corresponding regular/original mesh:

$$\mathrm{ErrorRed}_{3.4} = 100\left(1 - \frac{\mathrm{RMSE}_{3.4}^{\mathrm{method}}}{\mathrm{RMSE}_{3.4}^{\mathrm{orig}}}\right)\%. \tag{63}$$

For the phase-conditioned analysis in Fig. 6, months are grouped according to Eq. 61, allowing the same model to be examined under El Niño, neutral, and La Niña conditions while preserving the continuous chronological structure of the underlying SST record.

Interpretation of the ENSO metric. The index RMSE evaluates preservation of a regional SST diagnostic. Spatial errors can partially cancel in an area average, so reduced Niño 3.4 error is not equivalent to reduced error at every SST grid point. Likewise, a ratio of interval-level RMSE values differs from an average of month-wise relative errors. The temporal curves complement the resolution-wise scores by showing when gains occur; the phase labels stratify those reconstructions and do not turn the experiment into a forecast of future ENSO states.

F.4 GLOBAL EARTH-SYSTEM TRANSFER DATA (RQ4)

Transfer setting. RQ4 evaluates whether the mesh-adaptation rule learned from the local ocean systems transfers across changes in geographic location, characteristic spatial structure, and physical variable. The UGM2N and GEMOVE checkpoints obtained from the Kuroshio and ENSO training data are kept fixed and directly applied to the target Earth-system fields without target-region retraining, fine-tuning, or parameter updates.

The benchmark contains 54 geographic regions organized into four groups: 18 ocean-current regions, 18 wind-field regions, 9 sea-water chemistry regions, and 9 regions containing other environmental variables. Together, the benchmark spans 16 physical variable categories. Current and wind each count as one vector-valued category rather than counting their horizontal components as separate variables.

The four groups cover vector flows, scalar tracers, and environmental fields with substantially different spatial structures. The benchmark therefore extends the real-world evaluation from local ocean systems to geographically and physically heterogeneous Earth-system fields while retaining a common mesh-movement protocol.

Unified field representation. The source datasets are provided by different geophysical products with different variables, units, resolutions, and missing-data conventions. Before mesh movement, each field is converted to a common two-component spatial representation used by the mesh-movement pipeline.

For a native vector field, such as an ocean current or atmospheric wind, the two components correspond directly to the zonal and meridional physical components,

$$\mathbf{v} = (u, v)^\top. \tag{64}$$

We define its scalar magnitude as

$$s = \sqrt{u^2 + v^2}. \tag{65}$$

The magnitude field provides the primary scalar description from which the geometric node features used for mesh adaptation are constructed, while the two physical components are retained for field sampling, reconstruction, and evaluation.

For a native scalar variable q, such as pH, nutrient concentration, chlorophyll, precipitation, $PM_{2.5}$, soil moisture, radiation, or SST, we construct the two-component spatial-gradient representation

$$\mathbf{g}_q = \begin{bmatrix} \partial q/\partial x \\ \partial q/\partial y \end{bmatrix}, \tag{66}$$

where the spatial derivatives are computed on the corresponding coarse-grid representation used for mesh movement. Its magnitude is

$$s_q = \|\mathbf{g}_q\|_2 = \sqrt{\left(\frac{\partial q}{\partial x}\right)^2 + \left(\frac{\partial q}{\partial y}\right)^2}. \tag{67}$$

This construction provides a common direction-aware representation for scalar and vector Earth-system fields while retaining the spatial structures that drive mesh adaptation. Accordingly, native vector variables are reconstructed through their physical components, whereas scalar variables are reconstructed through their two-component spatial-gradient representations under the same experimental pipeline.

Mesh features. The same geometric feature construction is applied to all target fields. For native vector data, features are derived from the magnitude s in Eq. 65; for scalar data, they are derived from the gradient magnitude s_q in Eq. 67. At each valid coarse-grid node, the mesh representation contains the normalized spatial coordinates, boundary indicators, field magnitude, its discrete gradient, the Frobenius norm of its Hessian, and the Hessian components. These quantities are assembled into the same fixed feature representation used during ocean training, so RQ4 preprocessing does not change the encoder input dimensionality or network architecture.

In addition to the node-wise graph features, two regular 20×20 feature maps encode the field magnitude and Hessian norm for the convolutional feature extractor. Missing entries required only

for constructing these feature maps are filled from the nearest available values before interpolation to the common 20×20 representation. This feature-level filling is used only to construct model inputs and does not modify the high-resolution reference field or the evaluation mask.

The discrete gradient and Hessian used by the mesh-movement network are computed on the corresponding coarse-grid representation. They therefore serve as local geometric descriptors of spatial variation rather than as independently evaluated physical quantities. This common construction allows the same learned mesh-movement architecture to process fields with substantially different physical meanings and units.

Input normalization across physical variables. The physical variables in RQ4 have heterogeneous units and numerical scales, ranging from velocity and chemical concentration to precipitation, atmospheric composition, soil moisture, and radiative quantities. We do not convert these quantities into a common physical unit. Instead, model inputs are normalized independently for every field snapshot.

For each non-coordinate node-feature channel c, we use

$$f'_{c,i} = \frac{f_{c,i}}{\max\left(\max_j |f_{c,j}|, \varepsilon_{\mathrm{norm}}\right)}, \tag{68}$$

where the maximum is taken over the valid nodes of the current graph and $\varepsilon_{\mathrm{norm}} > 0$ prevents division by zero for vanishing feature channels. The two convolutional feature maps use the same normalization.

This snapshot-wise normalization places heterogeneous model inputs on comparable numerical scales without requiring dataset-specific parameter updates. Importantly, the normalization is applied only to the features supplied to the mesh-movement model. The final reconstruction errors are computed from the corresponding unnormalized physical vector or scalar-gradient fields.

Fine and coarse spatial representations. Each RQ4 dataset is represented at two spatial resolutions with distinct roles: a native Fine grid used exclusively as the high-resolution evaluation reference and a Coarse regular grid used to construct the model input and initial mesh.

The *Fine* grid is obtained directly from the native coordinate grid of the corresponding data product after extracting the specified geographic window. We do not resample the Fine field to a common resolution across datasets. Consequently, the reference resolution remains dataset-dependent. For example, the evaluated products include HYCOM fields with approximately 0.08° longitudinal and 0.04° latitudinal spacing, CCMP fields near 0.25°, high-resolution HRRR fields, Copernicus biogeochemical products near 0.25°, and CERES fields near 1°. The native Fine field is retained only for final reconstruction-based evaluation and is not used to assign field values to relocated mesh nodes.

The *Coarse* regular field is constructed over the same geographic window by block-mean resampling of the Fine field. For each output location, finite Fine-grid values within the corresponding spatial block are averaged, while land and missing values are excluded. The valid Coarse grid points define the initial node set from which the triangular mesh and its physical features are constructed.

The target coarse resolution is selected to obtain approximately 3000 valid nodes while maintaining a substantial reduction from the native Fine resolution. The exact number of valid nodes varies across regions because the geographic aspect ratio, coastline, land mask, missing-data pattern, and native grid resolution differ between datasets. Within each individual region, however, all compared methods use exactly the same Coarse field, initial node set, node budget, mesh connectivity, and boundary constraints. Thus, the high-resolution Fine field provides a common evaluation reference, while all mesh movement and field reconstruction originate from the same lower-resolution information.

Missing values, valid nodes, and mesh boundaries. Source fill values and product-specific invalid values are decoded before spatial processing and converted to missing values. A grid point is treated as valid only when all components required by the corresponding two-component field representation are finite. For native vector fields, this requires both physical vector components to be available; for scalar-derived fields, both spatial-gradient components must be defined.

Block-mean resampling ignores missing entries instead of treating them as physical zeros. Fine and Coarse validity masks are retained separately throughout the pipeline.

The valid Coarse nodes are triangulated to form the graph used by the mesh-movement model. A node is marked as a boundary node when it lies near the exterior geographic boundary, is adjacent to a land or missing-data location, or belongs to the topological boundary of the triangular mesh. These nodes remain fixed during mesh movement. The boundary treatment therefore preserves the valid geographic support of each target region while allowing interior vertices to adapt to the spatial field structure.

Field sampling and reconstruction. All methods use the same Coarse regular field as the source of information for field reconstruction. The Original baseline retains the initial Coarse node coordinates, whereas each mesh-movement method relocates the same set of interior nodes under the same fixed node budget.

After mesh movement, the field components at the relocated node coordinates are obtained by bilinear interpolation from the same Coarse regular field used to construct the initial mesh. The resulting nodal values are then placed on the moved triangular mesh and reconstructed onto the native Fine-grid locations using piecewise-linear interpolation.

For a Fine-grid point $\mathbf{x}$ contained in a moved triangle with nodal values f_1, f_2, and f_3, the reconstructed value is

$$\widehat{f}(\mathbf{x}) = \lambda_1 f_1 + \lambda_2 f_2 + \lambda_3 f_3, \qquad \lambda_1 + \lambda_2 + \lambda_3 = 1, \tag{69}$$

where λ_1, λ_2, and λ_3 are the barycentric coordinates of $\mathbf{x}$ within the triangle. The same reconstruction procedure is applied independently to the two field components and consistently across all compared methods.

The native Fine field is used only after reconstruction to evaluate the resulting field and therefore does not provide high-resolution information to the mesh-movement or reconstruction process. The mesh-movement model itself preserves the original mesh connectivity throughout relocation.

Fine-grid reference and reconstruction error. All methods within the same target region use the same Fine reference field, Fine validity mask, Coarse input, and evaluation time steps. For a two-component field

$$\mathbf{V}_{p,t} = (u_{p,t}, v_{p,t})^\top \tag{70}$$

at Fine-grid location p and time t, let $\widehat{\mathbf{V}}_{p,t}$ denote the reconstructed field. We use latitude-dependent area weights

$$w_p = \cos(\varphi_p), \tag{71}$$

where φ_p denotes latitude. The Fine-grid reconstruction error is defined as

$$E_t = \sqrt{\frac{\sum_{p\in\Omega_t} w_p \left[(\widehat{u}_{p,t} - u_{p,t})^2 + (\widehat{v}_{p,t} - v_{p,t})^2\right]}{\sum_{p\in\Omega_t} w_p}}, \tag{72}$$

where Ω_t denotes the common valid Fine-grid evaluation support shared by all compared methods at time t. It is determined from the source validity mask and is fixed before evaluating any reconstructed field. Land and source missing values are excluded consistently for all methods.

For native vector variables, $\mathbf{V}$ corresponds directly to the physical vector field. For a native scalar variable q, the same reconstruction and error definition is applied to the two-component field $\mathbf{V} = \nabla q$ from Eq. 66. Thus, all RQ4 cases use the same component-wise reconstruction and area-weighted Fine-grid evaluation procedure under their corresponding field representation.

Error reduction. Because the target variables have different physical units and native magnitudes, absolute FineRMSE values are not compared directly across datasets. Instead, the primary RQ4 metric measures the relative decrease in reconstruction error with respect to the unchanged Original mesh.

For every time step,

$$\text{Error Red.}_t = 100 \left(1 - \frac{E_t^{\text{method}}}{E_t^{\text{orig}}}\right) \%. \tag{73}$$

Positive values indicate that mesh movement reduces reconstruction error, whereas negative values indicate that the moved mesh performs worse than the Original mesh for that snapshot.

The region-level score reported in RQ4 is obtained by first computing Eq. 73 independently for each valid time step and then averaging the resulting relative improvements over the evaluated record. This normalization makes the reported scores dimensionless and permits comparison of relative mesh-adaptation gains across fields with different physical units. It does not imply that their underlying absolute reconstruction errors are physically interchangeable.

Temporal evaluation protocol. The quantitative RQ4 results use all valid snapshots in the configured evaluation record rather than only the beginning, middle, or end snapshots used for qualitative visualization. The available record length depends on the native data source and therefore varies across regions.

For the 18 ocean-current regions, the audited records contain between 122 and 547 evaluated snapshots depending on the regional product and available period. The ocean-surface wind cases contain approximately 1461 six-hourly snapshots for each of the 17 CCMP-based regions, while the HRRR Great Plains case uses 124 available three-hourly snapshots. Each of the 9 sea-water chemistry cases contains 366 configured daily snapshots. The remaining environmental datasets contain between 31 and 366 snapshots, reflecting their different native records and sampling frequencies.

These differences in temporal coverage are retained rather than artificially equalized. Each target region is evaluated over its own available record, and comparisons between methods are always performed on the same snapshots within that region.

Variables and data sources. The ocean-current group is derived from HYCOM global ocean circulation products (Chassignet et al., 2007), which provide data-assimilative ocean velocity fields and allow transfer to circulation systems far outside the Kuroshio training region.

Most wind regions use the Cross-Calibrated Multi-Platform (CCMP) ocean-surface vector wind product (Atlas et al., 2011; Wentz et al., 2015), while one continental case uses NOAA HRRR (Dowell et al., 2022). Together, these fields extend the evaluation from ocean circulation to atmospheric vector dynamics.

The marine biogeochemical variables are obtained from Copernicus Marine biogeochemical products (Copernicus Marine Service, 2026b;a). They include macronutrients, dissolved oxygen, pH, iron, alkalinity, chlorophyll, and net primary production, providing scalar tracer fields whose spatial organization differs fundamentally from velocity and SST.

The remaining environmental variables further broaden the physical scope. Precipitation is drawn from CMORPH CDR (Xie et al., 2017; 2019) and CHIRPS v3 (Funk et al., 2026; Climate Hazards Center, 2025); soil moisture is taken from the ESA CCI combined soil-moisture product (Dorigo et al., 2026; Gruber et al., 2019; Dorigo et al., 2017; Preimesberger et al., 2021); $PM_{2.5}$ is obtained from NASA GEOS-CF (Keller et al., 2021); radiative fields are obtained from CERES SYN1deg (NASA/LARC/SD/ASDC, 2026; Doelling et al., 2013); and the SST transfer case again uses NOAA OISST v2.1 (Huang et al., 2021; 2020).

Table 8 summarizes the variables and the number of geographic regions associated with each source.

Ocean-current regions. The 18 current regions are designed to span geographically and dynamically distinct circulation regimes rather than repeatedly sampling one basin. They include the Bering transition, eastern Pacific, Gulf Stream, equatorial Atlantic, Brazil–Malvinas confluence, Labrador Sea, Canary Current, Norwegian Sea, Mediterranean Sea, Agulhas system, Antarctic Circumpolar Current, Arabian Sea, Bay of Bengal, Indonesian Throughflow, Leeuwin Current, Kuroshio–Oyashio transition, Tasman/East Australian Current, and Ross Sea.

The collection therefore covers several characteristic forms of ocean circulation, including energetic western-boundary-current systems, eastern-boundary currents, current confluences, equatorial circulation, marginal seas, inter-basin throughflows, and high-latitude circulation. The associated HYCOM fields (Chassignet et al., 2007) provide a geographically broad test of whether a movement rule learned from one western North Pacific current system remains useful for unrelated ocean circulation structures.

Table 8: Physical variables and data sources used in the global Earth-system transfer benchmark. The final column gives the number of geographic regions associated with each variable in the 54-region evaluation.

Group	Variable	Physical quantity	Data source	# Regions
Ocean currents	Current	Horizontal ocean velocity	HYCOM	18
Wind	Wind	Ocean-surface wind	CCMP	17
	Wind	Near-surface atmospheric wind	NOAA HRRR	1
Ocean chemistry	PO_4	Phosphate	Copernicus Marine BGC	1
	pH	Seawater pH	Copernicus Marine BGC	2
	O_2	Dissolved oxygen	Copernicus Marine BGC	2
	NO_3	Nitrate	Copernicus Marine BGC	1
	Silicate	Silicate	Copernicus Marine BGC	1
	Fe	Dissolved iron	Copernicus Marine BGC	1
	Alkalinity	Total alkalinity	Copernicus Marine BGC	1
Environment	Precipitation	Precipitation	CMORPH / CHIRPS	2
	Chlorophyll	Chlorophyll	Copernicus Marine BGC	1
	NPP	Net primary production	Copernicus Baltic BGC	1
	Soil moisture	Surface soil moisture	ESA CCI	1
	$PM_{2.5}$	Fine particulate matter	NASA GEOS-CF	2
	Radiation	Radiative field	CERES	1
	SST	Sea-surface temperature	NOAA OISST v2.1	1

Atmospheric wind regions. The 18 wind regions similarly span diverse atmospheric flow regimes. They include North and South Pacific storm, subtropical, trade-wind, and equatorial regions; the eastern-Pacific ITCZ; a Great Plains continental case; Southern Ocean winds; North and South Atlantic subtropical, storm-track, high-latitude, trade-wind, and gyre regions; western and southern Indian Ocean winds; the Indian monsoon; the South China Sea; and western and central Pacific wind regimes.

CCMP provides spatially complete ocean-surface vector wind analyses by cross-calibrating information from multiple observing platforms (Atlas et al., 2011; Wentz et al., 2015). The HRRR case introduces a contrasting high-resolution continental near-surface wind field (Dowell et al., 2022). These cases test transfer from ocean-current vectors to atmospheric vector fields whose coherent structures arise from substantially different dynamics.

Marine biogeochemical regions. The 9 sea-water chemistry regions contain phosphate, pH, dissolved oxygen, nitrate, silicate, dissolved iron, and alkalinity fields from Copernicus Marine biogeochemical products (Copernicus Marine Service, 2026b). The selected regions span the Pacific, Atlantic, Arctic, Southern Ocean, and Indian Ocean.

Unlike velocity fields, these variables are transported and transformed as biogeochemical tracers and often exhibit broad background gradients together with localized transitions associated with circulation, biological activity, and water-mass structure. They therefore provide a qualitatively different class of scalar fields for testing whether geometric mesh adaptation transfers beyond the physical variables observed during ocean-system training.

Other environmental fields. The final 9 regions extend the transfer benchmark beyond ocean circulation and marine chemistry. They contain precipitation over South America and East Africa from CMORPH and CHIRPS (Xie et al., 2017; 2019; Funk et al., 2026; Climate Hazards Center, 2025), South Atlantic chlorophyll and Baltic net primary production from Copernicus Marine products (Copernicus Marine Service, 2026b;a), African soil moisture from ESA CCI (Dorigo et al., 2026; Gruber et al., 2019), European and Asian $PM_{2.5}$ from NASA GEOS-CF (Keller et al., 2021), a tropical-Pacific radiative field from CERES (NASA/LARC/SD/ASDC, 2026; Doelling et al., 2013), and South-Pacific SST from NOAA OISST v2.1 (Huang et al., 2021; 2020).

These cases extend the evaluation to the hydrological cycle, land surface, atmospheric composition, radiation, and marine ecosystems.

Spatial-region selection. Each RQ4 case is defined over a fixed latitude–longitude window. We first randomly sample candidate regions without reference to model performance, and then remove geographically or physically redundant cases while retaining regions that provide broad geographic

and physical coverage. The final benchmark therefore includes current and frontal systems, transition zones, throughflows, equatorial and high-latitude oceans, marginal seas, atmospheric storm tracks, trade-wind and monsoon regions, the ITCZ, and geographically separated biogeochemical and environmental fields.

The benchmark varies geographic location, physical variable, native resolution, characteristic spatial structure, and temporal sampling jointly. It therefore evaluates broad transfer across Earth-system fields rather than isolating a causal contribution from any single source of distribution shift.

Evaluation interpretation. For each target region, UGM2N and GEMOVE receive the corresponding field and initial mesh while their learned parameters remain fixed. Target fields provide the spatial information needed to determine mesh movement, but no target-region samples are used for model optimization or parameter updates.

All reported methods are compared using the same Coarse input mesh and the same Fine reference field within a region. The Fine grid remains native to the corresponding data product, so absolute reconstruction errors are meaningful within a dataset but should not be interpreted as directly comparable physical quantities across different variables.

The dimensionless Error Red. metric instead evaluates how much each method improves reconstruction relative to the same unchanged mesh. The RQ4 results therefore measure whether an ocean-trained mesh-adaptation rule continues to allocate a fixed spatial resolution budget effectively when transferred across new geographic regions, physical systems, and field structures.

F.5 TEMPORAL COVERAGE OF THE GLOBAL TRANSFER BENCHMARK

The RQ4 benchmark retains the native temporal coverage and sampling frequency available for each target dataset rather than artificially equalizing the number of snapshots across regions. Table 9 reports the geographic window, evaluation period, nominal temporal sampling, and number of selected snapshots for all 54 target regions. The index is defined independently within each of the four field groups and follows the region ordering used in the region-level transfer visualizations.

The temporal coverage reflects the characteristics and availability of the underlying data products. Ocean-current cases are evaluated on daily records containing between 122 and 547 snapshots. Most wind cases use one year of 6-hourly CCMP data, resulting in 1461 selected snapshots per region, while the HRRR Great Plains case uses the available 3-hourly record. The sea-water chemistry cases use daily records over 2024, and the remaining environmental fields retain their product-specific evaluation periods. Within each target region, UGM2N and GEMOVE are always evaluated on the same spatial domain and temporal samples.

Table 9: Temporal and spatial coverage of the 54 target regions in the global Earth-system transfer benchmark. Spatial windows are consistently reported as longitude followed by latitude using E/W and N/S geographic coordinates. The evaluation period denotes the configured temporal interval, Sampling denotes the nominal temporal resolution, and # Snapshots is the number of selected time steps. Indices are defined independently within each field group.

Idx.	Region / Dataset	Spatial Window (Lon; Lat)	Evaluation Period	Sampling	# Snap.
Ocean currents (18 regions)					
1	Bering transition	174.96°W–140.00°W; 45.00°N–65.00°N	2020-03-01–2021-02-28	Daily	365
2	Eastern Pacific	170.00°W–110.00°W; 40.00°S–10.00°N	2020-01-01–2020-12-31	Daily	366
3	Gulf Stream	100.00°W–40.00°W; 20.00°N–55.00°N	2019-08-01–2020-01-31	Daily	184
4	Equatorial Atlantic current	80.00°W–40.00°W; 5.00°S–15.00°N	2020-01-01–2020-12-31	Daily	366
5	Brazil–Malvinas	74.96°W–25.04°W; 50.00°S–10.00°S	2019-06-01–2020-05-31	Daily	366
6	Labrador Sea	70.00°W–25.04°W; 45.00°N–70.00°N	2019-11-01–2020-04-30	Daily	182
7	Canary Current	30.00°W–1.00°W; 15.00°N–35.00°N	2020-01-01–2020-12-31	Daily	366
8	Norwegian Sea	0.00°E–24.96°E; 55.00°N–72.00°N	2019-10-01–2020-09-30	Daily	366
9	Mediterranean	0.00°E–40.00°E; 30.00°N–46.00°N	2019-01-01–2020-06-30	Daily	547

Table 9 continued

Idx.	Region / Dataset	Spatial Window (Lon; Lat)	Evaluation Period	Sampling	# Snap.
10	Agulhas	0.00°E–50.00°E; 50.00°S–15.00°S	2019-01-01–2020-06-30	Daily	547
11	ACC	0.00°E–80.00°E; 65.00°S–40.00°S	2021-01-01–2022-06-30	Daily	546
12	Arabian Sea	50.00°E–80.00°E; 0.00°N–25.00°N	2019-06-01–2019-09-30	Daily	122
13	Bay of Bengal	80.00°E–100.00°E; 5.00°N–25.00°N	2019-06-01–2020-05-31	Daily	366
14	Indonesian Throughflow	95.04°E–140.00°E; 15.00°S–10.00°N	2020-01-01–2020-12-31	Daily	366
15	Leeuwin Current	100.00°E–130.00°E; 45.00°S–20.00°S	2020-01-01–2020-12-31	Daily	366
16	Kuroshio–Oyashio transition	140.00°E–180.00°E; 30.00°N–50.00°N	2020-01-01–2020-12-31	Daily	366
17	Tasman/EAC	145.04°E–180.00°E; 50.00°S–25.00°S	2019-01-01–2020-06-30	Daily	547
18	Ross Sea	150.00°E–160.00°W; 75.00°S–60.00°S	2020-01-01–2020-12-31	Daily	366
Wind fields (18 regions)					
1	N. Pacific storm	179.88°W–140.12°W; 30.12°N–49.88°N	2020-01-01–2020-12-31	6-hourly	1461
2	CCMP S. Pacific subtropical	149.88°W–110.12°W; 34.88°S–15.12°S	2020-01-01–2020-12-31	6-hourly	1461
3	CCMP E. Pacific ITCZ	139.88°W–100.12°W; 0.12°N–19.88°N	2020-01-01–2020-12-31	6-hourly	1461
4	HRRR Great Plains	110.00°W–95.05°W; 30.00°N–45.00°N	2026-08-01–2026-08-31	3-hourly	124
5	Southern Ocean	109.88°W–70.12°W; 64.88°S–45.12°S	2020-01-01–2020-12-31	6-hourly	1461
6	CCMP N. Atlantic subtropical	79.88°W–40.12°W; 15.12°N–34.88°N	2020-01-01–2020-12-31	6-hourly	1461
7	CCMP N. Atlantic high-lat.	59.88°W–20.12°W; 55.12°N–69.88°N	2020-01-01–2020-12-31	6-hourly	1461
8	N. Atlantic storm	59.88°W–20.12°W; 35.12°N–54.88°N	2020-01-01–2020-12-31	6-hourly	1461
9	CCMP S. Atlantic trade	59.88°W–20.12°W; 19.88°S–0.12°S	2020-01-01–2020-12-31	6-hourly	1461
10	S. Atlantic gyre	29.88°W–0.12°W; 39.88°S–20.12°S	2020-01-01–2020-12-31	6-hourly	1461
11	CCMP W. Indian monsoon	35.12°E–64.88°E; 34.88°S–15.12°S	2020-01-01–2020-12-31	6-hourly	1461
12	Indian monsoon	60.12°E–99.88°E; 24.88°S–4.88°N	2020-01-01–2020-12-31	6-hourly	1461
13	CCMP S. Indian westerlies	60.12°E–109.88°E; 59.88°S–40.12°S	2020-01-01–2020-12-31	6-hourly	1461
14	CCMP South China Sea	105.12°E–124.88°E; 5.12°N–24.88°N	2020-01-01–2020-12-31	6-hourly	1461
15	W. Pacific trade/typhoon	120.12°E–149.88°E; 5.12°N–24.88°N	2020-01-01–2020-12-31	6-hourly	1461
16	CCMP Kuroshio	135.12°E–159.88°E; 28.12°N–44.88°N	2020-01-01–2020-12-31	6-hourly	1461
17	CCMP Central Pacific trade	150.12°E–170.12°W; 0.12°N–19.88°N	2020-01-01–2020-12-31	6-hourly	1461
18	CCMP S. Pacific equatorial	170.12°E–150.12°W; 19.88°S–0.12°S	2020-01-01–2020-12-31	6-hourly	1461
Ocean chemistry (9 regions)					
1	CMEMS phosphate	170.00°W–130.00°W; 10.00°S–10.00°N	2024-01-01–2024-12-31	Daily	366
2	South Atlantic pH	60.00°W–20.00°W; 45.00°S–20.00°S	2024-01-01–2024-12-31	Daily	366
3	Arctic pH	40.00°W–39.92°E; 65.00°N–85.00°N	2024-01-01–2024-12-31	Daily	366
4	Weddell oxygen	20.00°W–20.00°E; 70.00°S–50.00°S	2024-01-01–2024-12-31	Daily	366
5	Arabian nitrate	50.00°E–75.00°E; 5.00°N–25.00°N	2024-01-01–2024-12-31	Daily	366
6	South Indian dissolved oxygen	60.00°E–100.00°E; 50.00°S–25.00°S	2024-01-01–2024-12-31	Daily	366
7	CMEMS silicate	100.00°E–140.00°E; 60.00°S–40.00°S	2024-01-01–2024-12-31	Daily	366
8	CMEMS iron	150.00°E–180.00°E; 25.00°N–50.00°N	2024-01-01–2024-12-31	Daily	366
9	CMEMS alkalinity SP	160.00°E–160.00°W; 45.00°S–15.00°S	2024-01-01–2024-12-31	Daily	366
Environment (9 regions)					
1	CMORPH South America	79.88°W–40.12°W; 19.88°S–4.88°N	2026-01-01–2026-03-31	Daily	90
2	CHIRPS East Africa	10.02°E–49.97°E; 14.98°S–19.98°N	2020-01-01–2020-01-31	Daily	31

Table 9 continued

Idx.	Region / Dataset	Spatial Window (Lon; Lat)	Evaluation Period	Sampling	# Snap.
3	South Atlantic chlorophyll	60.00°W–20.00°W; 35.00°S–15.00°S	2024-01-01–2024-12-31	Daily	366
4	Baltic Sea NPP	10.01°E–29.99°E; 53.01°N–65.89°N	2024-01-01–2024-12-31	Daily	366
5	ESA soil Africa	19.88°W–29.88°E; 5.12°N–29.88°N	2020-01-01–2020-12-31	Daily	366
6	GEOS-CF $PM_{2.5}$ Europe	20.00°W–45.00°E; 25.00°N–70.00°N	2025-08-04–2025-11-02	Daily	90
7	GEOS-CF $PM_{2.5}$ Asia	100.00°E–145.00°E; 20.00°N–50.00°N	2025-08-04–2025-11-02	Daily	90
8	CERES tropical Pacific	100.50°E–100.50°W; 59.50°S–59.50°N	2020-01-01–2020-12-31	Daily	366
9	OISST South Pacific	150.12°E–90.12°W; 49.88°S–9.88°N	2020-01-01–2020-12-31	Daily	366

The different snapshot counts arise from the native temporal resolution and available evaluation period of each data product rather than from differences in the spatial node budget. Region-level Error Red. is computed independently within each target region using its own Original-mesh reference and is not obtained by pooling absolute errors across datasets. Accordingly, datasets with different temporal sampling frequencies remain comparable through their within-region relative reconstruction improvements.

G Additional Real-World Results

The following visualizations complement the quantitative results in Secs. 5.4 and 5.5. The Kuroshio gallery shows spatial allocation across dates, the flow-conditioned diagnostics examine where reconstruction gains occur, the ENSO curves show temporal variation at additional resolutions, and the radar plots expose region-level variation within the four transfer groups.

G.1 Additional Kuroshio Results

The Kuroshio case study is designed to test whether the learned movement rule can consistently redistribute a limited mesh budget toward the narrow, energetic current core and surrounding frontal structures. Beyond the aggregate metrics reported in the main text, it is therefore useful to inspect representative moved meshes at multiple times as well as diagnostics conditioned on physically meaningful flow characteristics.

Figure 12 shows representative moved meshes for UGM2N and GEMove at eight dates spanning the analysis period. Across seasons, both methods move mesh nodes toward the dynamically active Kuroshio jet and adjacent transition zones rather than preserving a nearly uniform spatial allocation. The fixed boundary constraints retain the prescribed coastal and wet-domain boundary vertices. The selected examples show a more focused allocation by GEMove along the jet and nearby frontal regions, consistent with the aggregate reconstruction comparison in the main text.

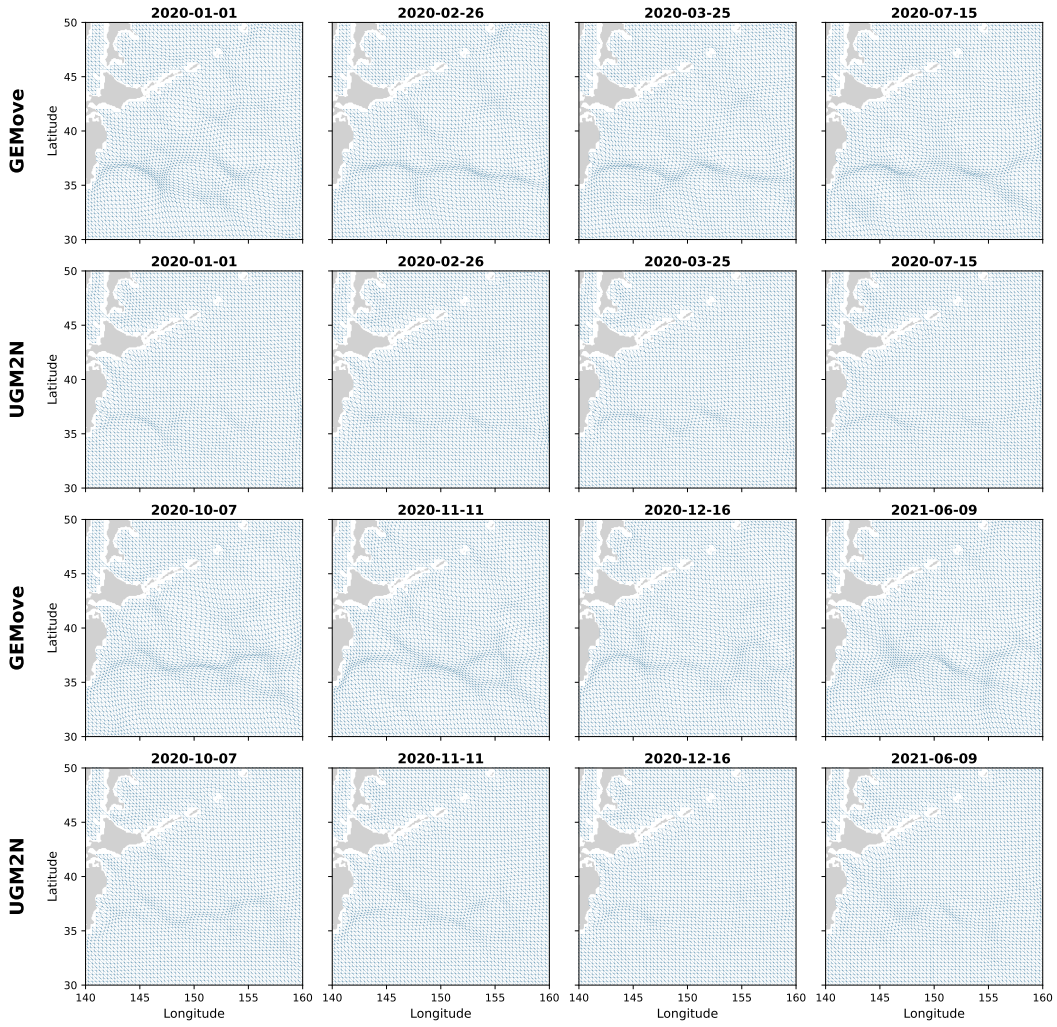

Figure 12: Representative moved meshes for the Kuroshio experiment at eight dates spanning the analysis period. Each date is shown for both GEMOVE and UGM2N. Mesh movement reallocates the fixed node budget toward the jet and adjacent fronts, with a more focused concentration for GEMOVE in the selected examples.

Allocation around competing structures. The gallery is useful for assessing how the same node budget is shared between the current core and neighboring transitions. Concentrating vertices along the strongest jet can leave adjacent structures poorly resolved; conversely, spreading them uniformly can miss the narrow core. The intended behavior is a coordinated redistribution across these demands, with smooth connections between densely and sparsely sampled regions. The selected dates illustrate this spatial behavior, while the full-period Error Red., Win Rate, and CoreIoU summarize its accuracy, temporal consistency, and core-region overlap, respectively.

Flow-conditioned diagnostics. To further explain where the performance gains arise, we provide the diagnostics in Fig. 13. The left panel evaluates cross-stream reconstruction skill as a function of distance from the reference jet axis, thereby revealing how the improvement is distributed across the anticyclonic flank, the jet core, and the cyclonic flank. The center and right panels condition the improvement on the local magnitude of the speed gradient and vorticity, respectively. These quantities serve as compact dynamical proxies for fronts, shear, and rotational flow structure.

In the reported diagnostics, larger gains occur around the jet core and in regions with stronger speed gradients or vorticity. The advantage of GEMOVE over UGM2N also tends to increase in the more

active regimes. This pattern is consistent with improved allocation around sharp fronts and shear, complementing the mesh gallery with error-based evidence.

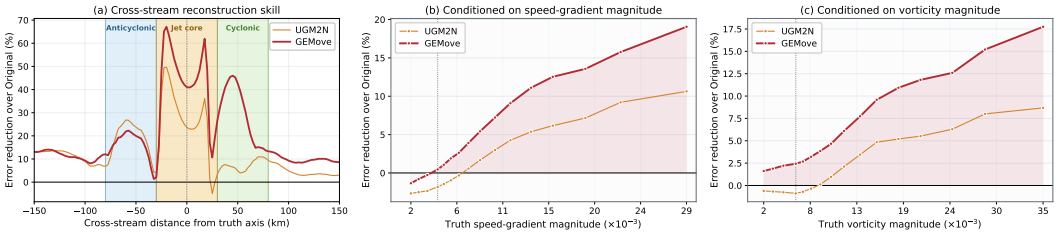

Figure 13: Additional Kuroshio diagnostics. (**a**) Cross-stream reconstruction skill as a function of distance from the reference jet axis. Positive values indicate error reduction relative to the original mesh. The shaded background partitions the section into the anticyclonic flank, jet core, and cyclonic flank. (**b**) Error reduction conditioned on the magnitude of the reference speed gradient. (**c**) Error reduction conditioned on the magnitude of the reference vorticity. In the conditioned diagnostics, GEMOVE shows a systematically larger gain, particularly in dynamically more complex regions.

The conditioned plots describe associations between flow structure and reconstruction gain. They do not measure a change in the underlying ocean dynamics. In particular, vorticity magnitude and speed-gradient magnitude identify different aspects of the reference flow and can highlight overlapping locations. Reading them alongside the cross-stream profile helps distinguish where the gain occurs without treating the three panels as independent experiments.

G.2 ADDITIONAL ENSO TEMPORAL RESULTS

The ENSO temporal plots complement aggregate RMSE reduction by showing how reconstruction gains vary across the chronological climate record.

Figure 14 extends the temporal visualization to the additional 2° and 3° coarse resolutions, using the same phase-colored background as in the main text. Each curve shows the time-resolved improvement over the original mesh, while the shaded intervals indicate the climate phase determined from the reference Niño 3.4 index. Presenting all resolutions in the same visual format makes it easier to assess whether the performance advantage persists as the mesh budget changes.

The reported curves show predominantly positive improvement over the multidecadal record, with a stronger and more sustained gain for GEMOVE at both resolutions. Its advantage remains visible across the reference-defined warm, neutral, and cold intervals. These temporal patterns complement the 4° comparison in the main text and show that the gain is not confined to one selected climate episode.

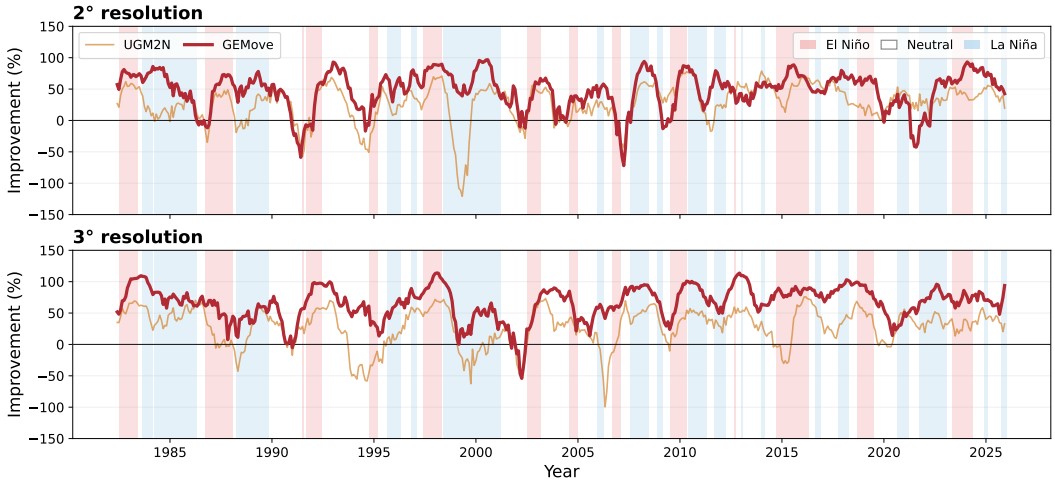

Figure 14: Additional ENSO temporal results at 2° and 3° coarse resolution. The curves show time-resolved improvement relative to the original mesh, and the background shading indicates the climate phase determined from the reference Niño 3.4 index. The reference-defined phase shading is shared across methods, allowing comparison across warm, neutral, and cold intervals.

Comparing temporal and aggregate behavior. Temporal variation matters even when two methods have similar aggregate scores: a method can concentrate its gains in a few intervals or maintain them across changing conditions. The additional curves make this distinction visible at 2° and 3°. Comparisons should be made within each resolution, using its original-mesh baseline and the same reference phase labels. The phase shading provides physical context, while the quantitative resolution-wise metric remains the Niño 3.4 RMSE reduction defined in Appendix F.3.

G.3 Additional Global Transfer Visualization

Figure 15 presents the region-level RQ4 results as radar plots, complementing the bar-chart view in the main text.

Each radar panel corresponds to one of the four field groups: oceanographic, atmospheric, biogeochemical, and environmental fields. The numbered axes indicate the region index within each group, and the radial value records the corresponding error reduction relative to the original mesh. This visualization makes it possible to compare both group-level behavior and the consistency of the advantage across individual regions.

Across all four groups, GEMove shows larger error reduction than UGM2N across the evaluated regions, consistent with the bar-chart comparison in Fig. 7. Variation between spokes further shows that the magnitude of improvement depends on the target region and physical field. The radar panels therefore provide a region-level view of the same transfer results rather than an additional independent evaluation.

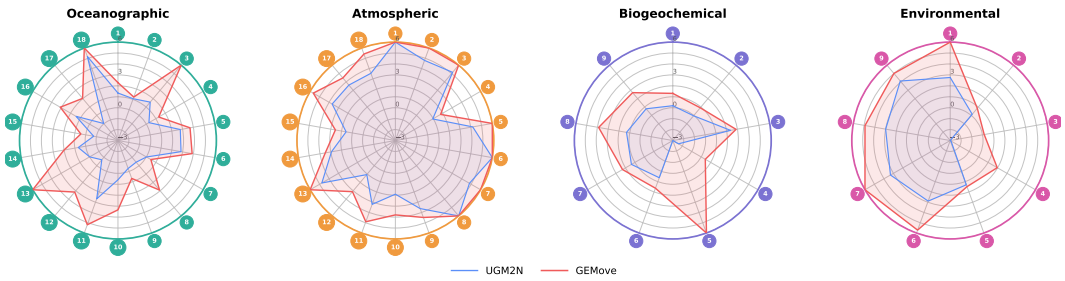

Figure 15: Radar-plot view of the RQ4 transfer results across oceanographic, atmospheric, biogeochemical, and environmental fields. Each spoke corresponds to one geographic region within the group, and the radial value denotes error reduction relative to the original mesh. Larger radial values indicate greater reconstruction improvement.

Reading region-level transfer. Each spoke should be compared directly between methods and against zero error reduction. A larger polygon area is not itself an aggregate accuracy metric because it depends on the ordering and number of spokes, which differ between the 18-region and 9-region groups. The plots therefore support a region-by-region reading of transfer consistency under the fixed-checkpoint protocol.

Complete global mesh-movement visualization. To complement the quantitative transfer results, we further visualize GEMOVE's mesh movement for one representative snapshot from every RQ4 target region. The following six pages cover all 54 regions across the four field groups: oceanographic, atmospheric, biogeochemical, and environmental fields. For each case, the globe indicates the evaluated geographic domain, while the corresponding panel shows the moved mesh and node-displacement magnitude in geographic degrees. Together, these visualizations provide a complete qualitative view of how the same ocean-trained movement rule adapts across diverse target fields without parameter updates.

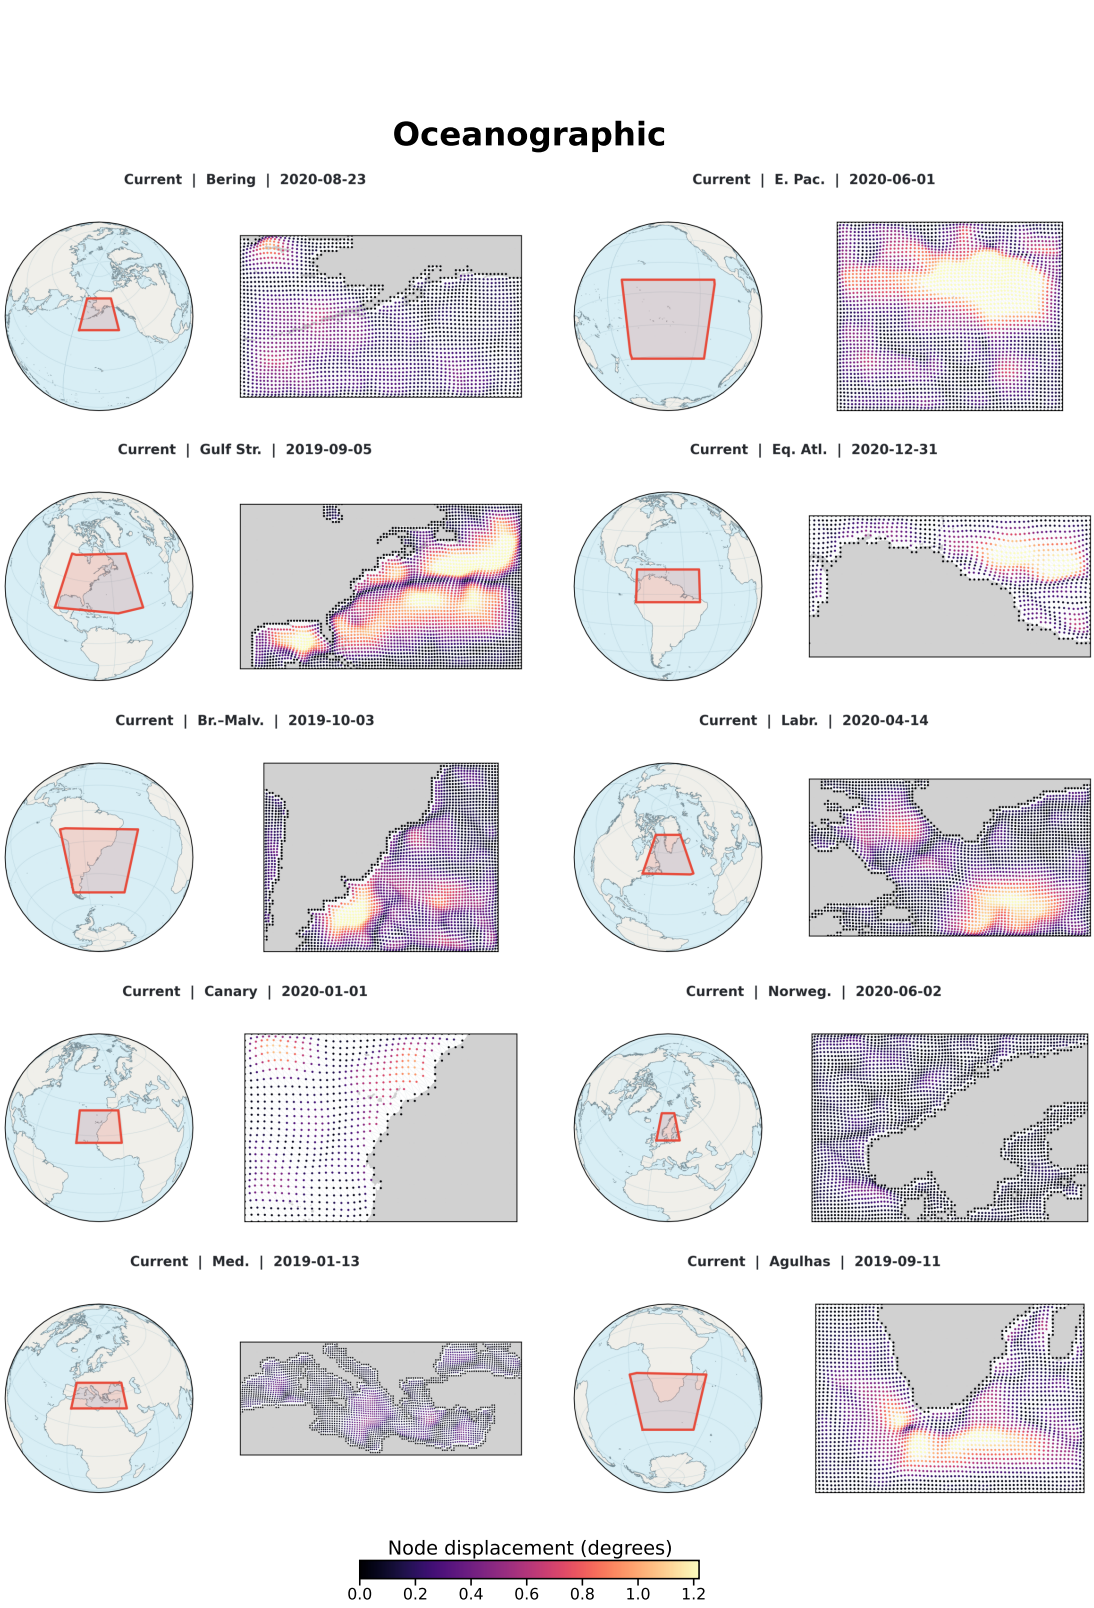
Oceanographic
Current | Bering | 2020-08-23
Current | E. Pac. | 2020-06-01
Current | Gulf Str. | 2019-09-05
Current | Eq. Atl. | 2020-12-31
Current | Br.-Malv. | 2019-10-03
Current | Labr. | 2020-04-14
Current | Canary | 2020-01-01
Current | Norweg. | 2020-06-02
Current | Med. | 2019-01-13
Current | Agulhas | 2019-09-11
Node displacement (degrees)
0.0
0.2
0.4
0.6
0.8
1.0
1.2

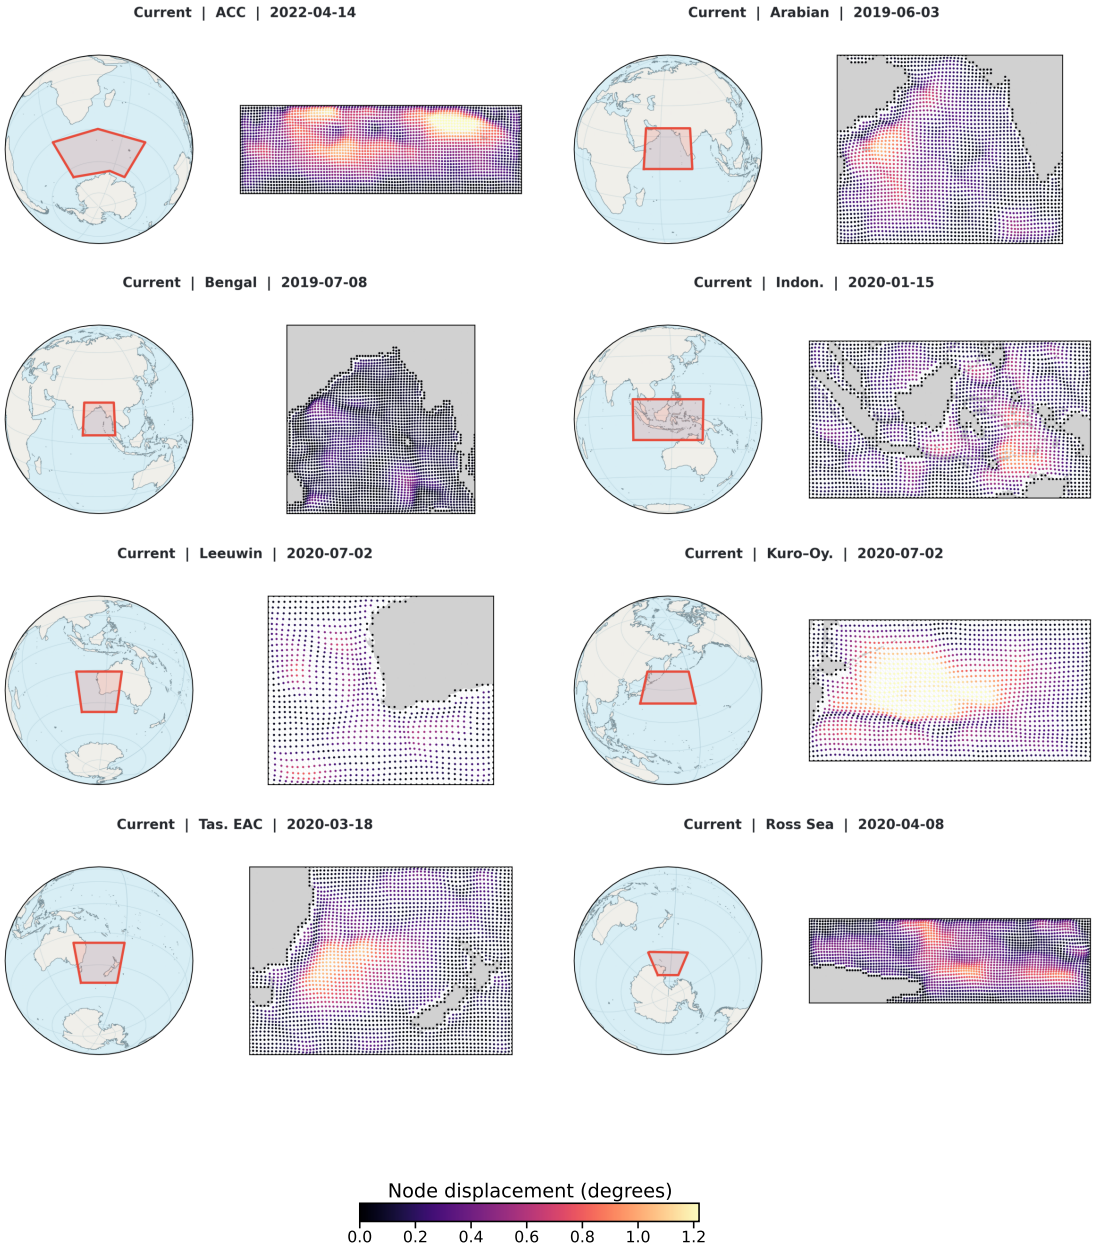
Oceanographic
Current | ACC | 2022-04-14
Current | Arabian | 2019-06-03
Current | Bengal | 2019-07-08
Current | Indon. | 2020-01-15
Current | Leeuwin | 2020-07-02
Current | Kuro-Oy. | 2020-07-02
Current | Tas. EAC | 2020-03-18
Current | Ross Sea | 2020-04-08
Node displacement (degrees)
0.0
0.2
0.4
0.6
0.8
1.0
1.2

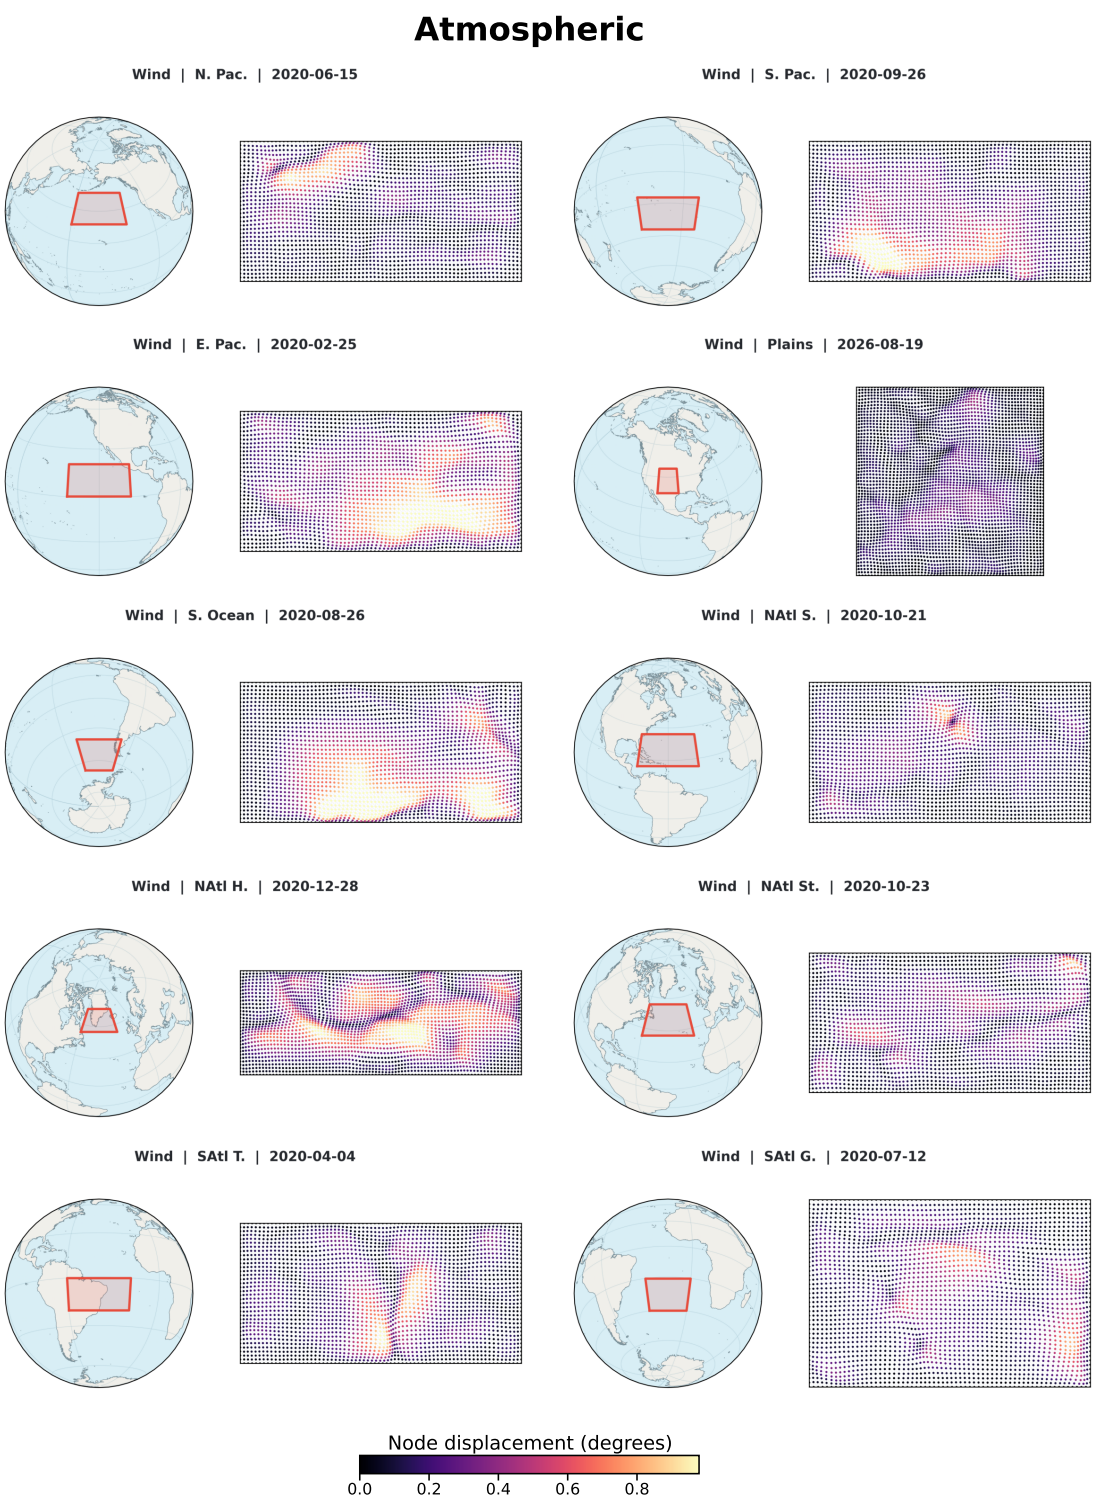
Atmospheric
Wind | N. Pac. | 2020-06-15
Wind | S. Pac. | 2020-09-26
Wind | E. Pac. | 2020-02-25
Wind | Plains | 2026-08-19
Wind | S. Ocean | 2020-08-26
Wind | NAtl S. | 2020-10-21
Wind | NAtl H. | 2020-12-28
Wind | NAtl St. | 2020-10-23
Wind | SAtl T. | 2020-04-04
Wind | SAtl G. | 2020-07-12
Node displacement (degrees)
0.0
0.2
0.4
0.6
0.8

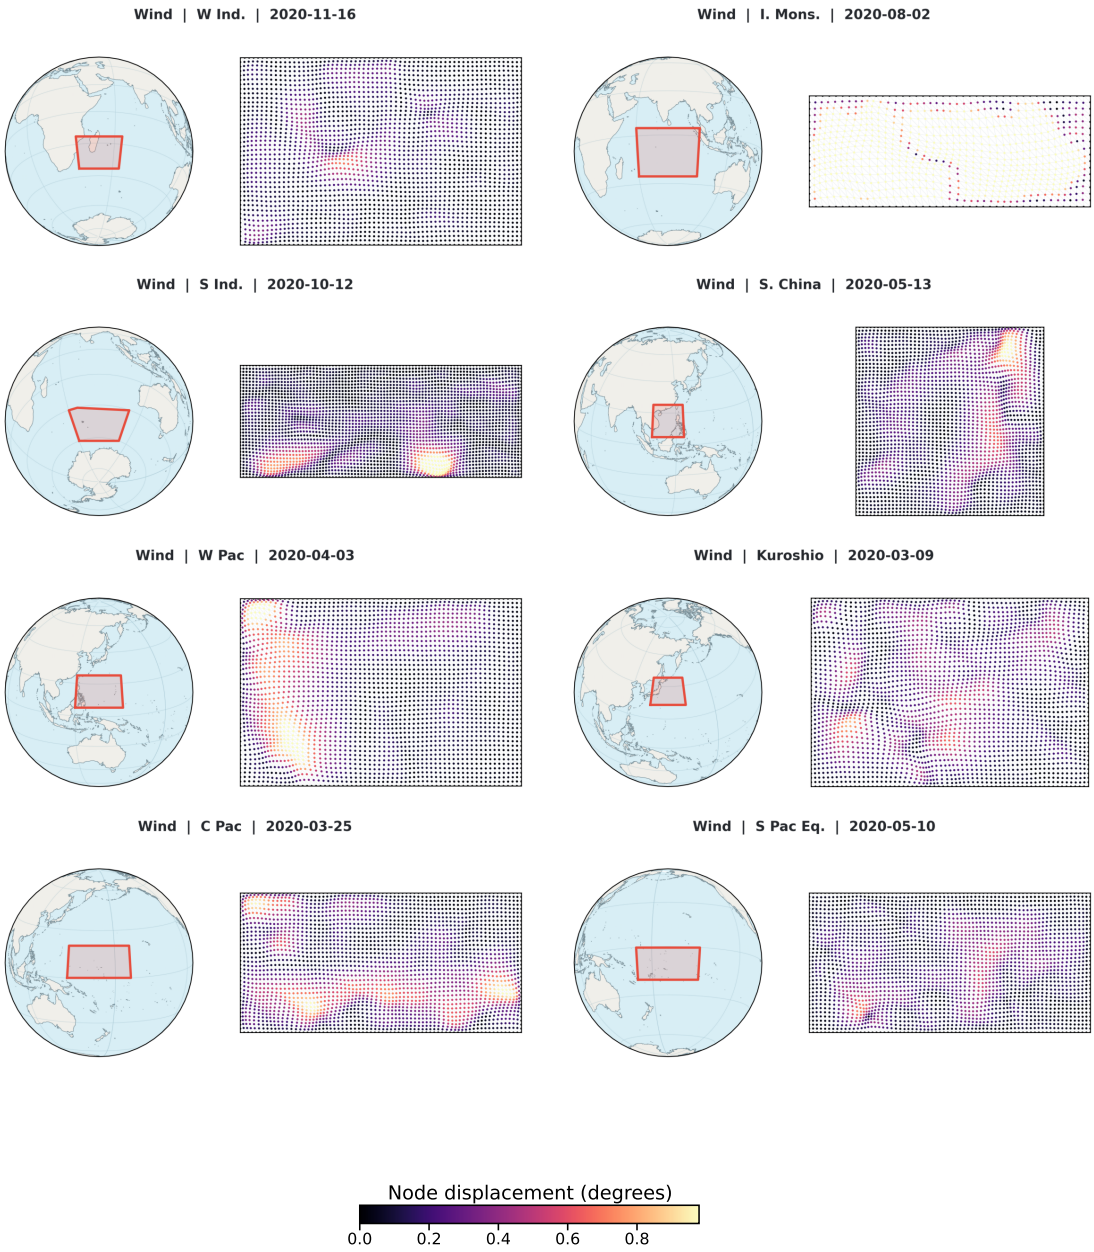
Atmospheric
Wind | W Ind. | 2020-11-16
Wind | I. Mons. | 2020-08-02
Wind | S Ind. | 2020-10-12
Wind | S. China | 2020-05-13
Wind | W Pac | 2020-04-03
Wind | Kuroshio | 2020-03-09
Wind | C Pac | 2020-03-25
Wind | S Pac Eq. | 2020-05-10
Node displacement (degrees)
0.0
0.2
0.4
0.6
0.8

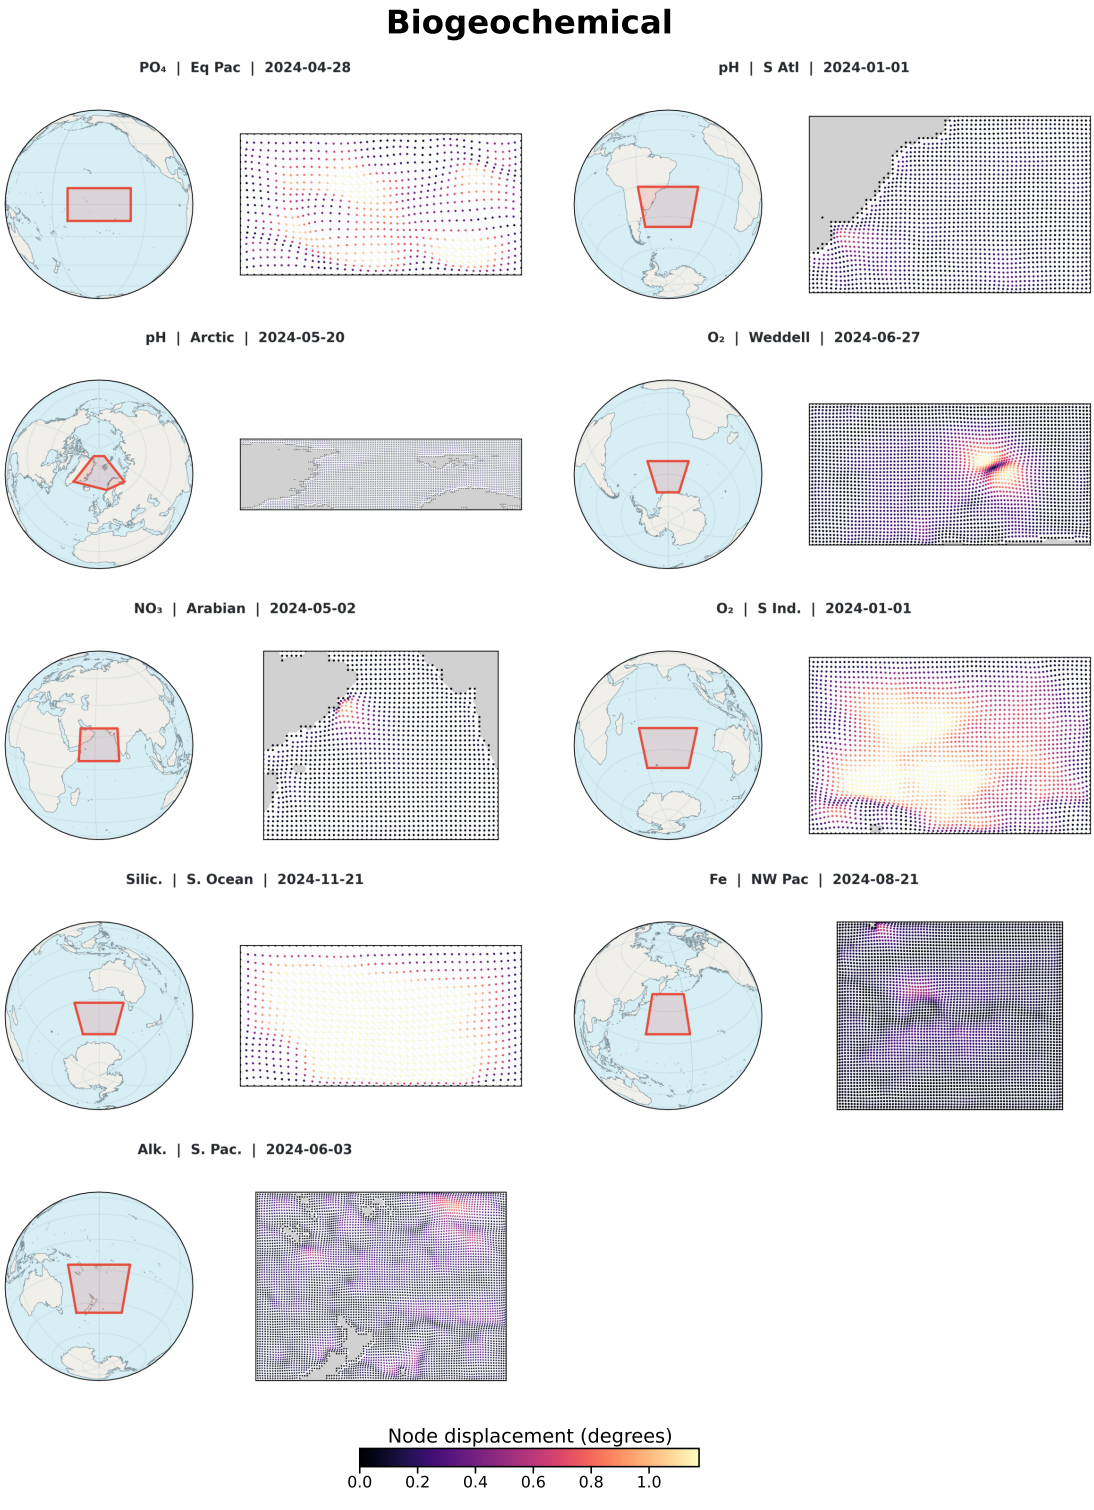
Biogeochemical
PO₄ | Eq Pac | 2024-04-28
pH | S Atl | 2024-01-01
pH | Arctic | 2024-05-20
O₂ | Weddell | 2024-06-27
NO₃ | Arabian | 2024-05-02
O₂ | S Ind. | 2024-01-01
Silic. | S. Ocean | 2024-11-21
Fe | NW Pac | 2024-08-21
Alk. | S. Pac. | 2024-06-03
Node displacement (degrees)
0.0
0.2
0.4
0.6
0.8
1.0

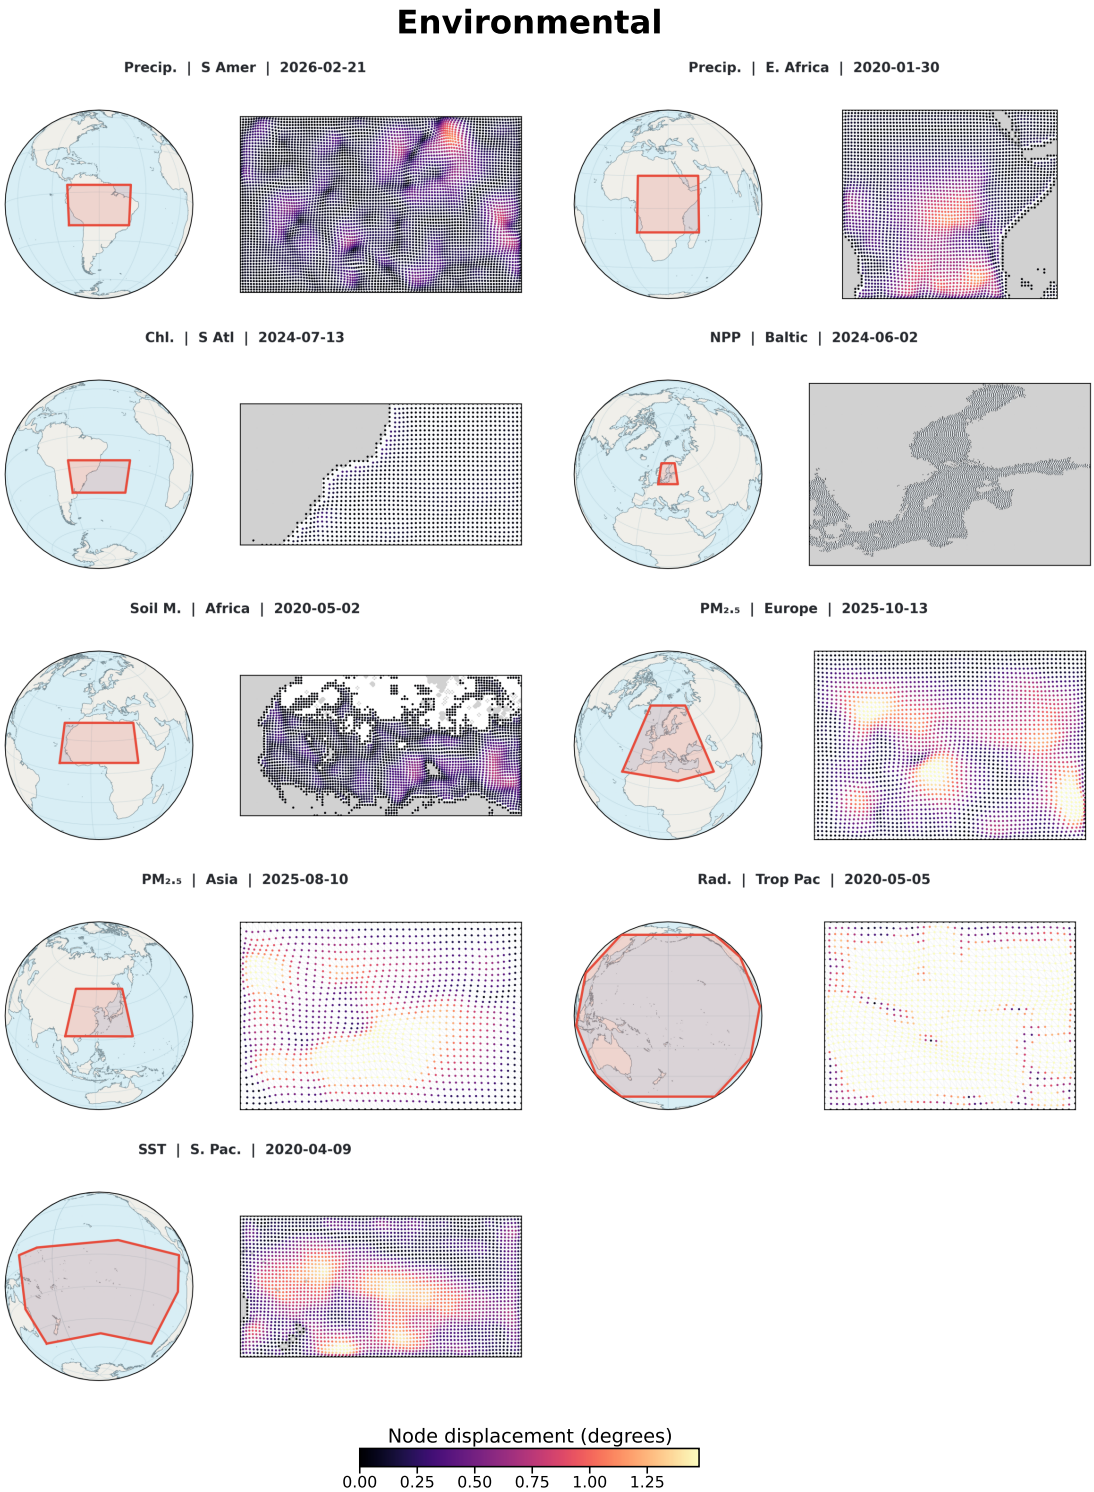
Environmental
Precip. | S Amer | 2026-02-21
Precip. | E. Africa | 2020-01-30
Chl. | S Atl | 2024-07-13
NPP | Baltic | 2024-06-02
Soil M. | Africa | 2020-05-02
PM2.5 | Europe | 2025-10-13
PM2.5 | Asia | 2025-08-10
Rad. | Trop Pac | 2020-05-05
SST | S. Pac. | 2020-04-09
Node displacement (degrees)
0.00 0.25 0.50 0.75 1.00 1.25