

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 GEOCMON: OPERATOR LEARNING ON DEFORMABLE DOMAINS VIA DISENTANGLED GEOMETRIC CONDI- TIONING

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 Operator learning for PDEs on non-rigid, parametrically varying domains with
014 heterogeneous boundary conditions faces challenges from input modality entan-
015 glement, training instability, and generalization limitations. To address this, we
016 propose GeoCMON, a Geometric-Conditioned Multi-Branch Operator Network.
017 GeoCMON explicitly disentangles geometric and boundary features via special-
018 ized encoding branches, fused with a spatial trunk network using element-wise
019 multiplication and Einstein summation for expressive conditioning. Conditional
020 residual connections within branches enhance gradient flow and stability, while
021 a weighted MSE loss prioritizes physically significant solution magnitudes. Em-
022 pirical evaluations on 2D Laplace problems demonstrate GeoCMON’s superior
023 accuracy across varied difficulty, improved training dynamics (higher synchro-
024 nization, reduced activation variance), and enhanced feature orthogonality. Gra-
025 dient noise analyses confirm optimization stability. GeoCMON advances scalable
026 and interpretable operator learning for complex deformable domains, offering a
027 principled framework for scientific computing. We provide the detailed code in
028 [Supplementary Material](#).

029 1 INTRODUCTION

031 Neural operator learning is emerging as a cornerstone of Scientific Machine Learning (SciML),
032 promising to supplant traditional numerical solvers by directly learning mappings between infinite-
033 dimensional function spaces (Li et al., 2020c;b; Zhang et al., 2023). However, the paradigm’s suc-
034 cess has been largely confined to problems on domains with fixed geometries. When applied to the
035 more challenging setting of **deformable domains** parameterized by complex, evolving topologies,
036 the generalization capabilities of existing neural operators degrade sharply, as prior assumptions
037 of fixed meshes become inadequate (Hartman et al., 2023). This limitation constitutes a major
038 bottleneck to applying operator learning in critical scientific and engineering applications such as
039 fluid-structure interaction, structural optimization, and biomechanics, where the domain’s geometric
040 evolution is intrinsic to the problem.

041 At the heart of current failures lies a fundamental problem of **representational entanglement**. Ex-
042 isting architectures typically resort to simple concatenation or ad-hoc fusion strategies that mix fea-
043 tures describing the domain’s geometry with the physical conditions imposed on its boundary (e.g.,
044 Dirichlet or Neumann conditions), which is challenging due to their heterogeneous nature and dis-
045 tinct spatial structures (Ovsjanikov et al., 2016). This unprincipled approach forces a single network
046 to learn a brittle, entangled representation that risks a loss of representational fidelity and fails to
047 distinguish the solution field’s sensitivity to geometric deformation from its sensitivity to changes in
048 boundary conditions. The consequences are severe: unstable training dynamics, poor generalization
049 to unseen geometry-boundary combinations, and a lack of model interpretability (Tan & Bansal,
050 2019; Wandel et al., 2021; Kovachki et al., 2021; Bhattacharya et al., 2021). **Developing an archi-
051 tecture that can explicitly disentangle these heterogeneous input modalities and fuse them in a
052 principled manner is therefore paramount for robust operator learning.**

053 To address this challenge, we introduce the **Geometric-Conditioned Multi-branch Operator Network**
(**GeoCMON**), a novel, principled architecture for operator learning on deformable domains. Our

approach is founded on the core idea of *factorizing* the complex solution operator into three specialized, learnable components. We design separate *geometry* and *boundary condition branches* that allocate distinct subnetworks to encode their respective input modalities into independent, disentangled representations (Lu et al., 2021). To ensure stable gradient flow within these deep encoders, we augment our branches with *conditional residual connections*, an enhancement shown to improve stability on parametric PDEs (Jiang et al., 2023). Subsequently, a carefully designed *two-stage fusion mechanism* combines these representations: first, an interpretable conditioning is achieved via a Hadamard product, where the geometric representation modulates the boundary representation; second, a tensor contraction projects this fused conditional representation onto the query coordinates encoded by a *spatial trunk network*.

To further align the model’s learning with the underlying physics, we introduce a *magnitude-aware weighted loss function*. This loss disproportionately penalizes errors in high-magnitude regions of the solution field—which often correspond to the most physically critical phenomena—thereby directing the optimization focus toward the most challenging predictive regimes. Through extensive empirical evaluations, we demonstrate that GeoCMON significantly outperforms existing baselines in accuracy, training stability, and generalization.

Our primary contributions are:

- ❶ *A disentangled multi-branch operator learning architecture (GeoCMON)* that achieves robust, independent encoding of geometry and boundary conditions via specialized branches and conditional residual connections.
- ❷ *An expressive feature fusion strategy* that principledly combines the disentangled representations using a Hadamard product and tensor contraction for accurate, conditioned prediction at spatial points.
- ❸ *A physics-aware weighted loss function* that improves predictive fidelity by emphasizing physically critical regions without compromising optimization stability.
- ❹ *A comprehensive empirical analysis* validating our method’s superior performance on challenging families of parametric PDEs defined on non-rigid domains, setting a new state-of-the-art for surrogate modeling of complex physical systems.

2 RELATED WORK

Recent advances in physics-informed neural networks (PINNs) and operator learning have increasingly focused on developing frameworks capable of approximating solutions to parametric partial differential equations (PDEs) across complex and non-rigid domains. Traditional numerical solvers often suffer from prohibitive computational costs for parametric PDEs with varying boundary conditions and domain geometries (Liu et al., 2023). To address these challenges, operator learning approaches, such as the DIMON framework, utilize diffeomorphic mappings to transform functions from parameterized domains to reference domains, enhancing generalization over families of shapes and enabling efficient prediction on realistic 2D and 3D geometries (Yin et al., 2024). Concurrently, multi-branch neural network architectures have demonstrated effectiveness in separately encoding heterogeneous input modalities, including geometric features and boundary conditions, before fusing them for operator approximation (Kovachki et al., 2021; Wandel et al., 2021). This design facilitates improved representation capacity and generalization across varying and non-rigid domains, addressing limitations of single-branch architectures traditionally employed in operator learning (Bhattacharya et al., 2021).

Moreover, recent developments highlight adaptive weighting mechanisms within multi-branch frameworks to dynamically balance geometric and boundary condition influences, fostering greater robustness and accuracy in approximating operators over complex parametric spaces (Yin et al., 2022). Architectures such as the U-shaped Neural Operator (U-NO) further leverage domain contraction-expansion with skip connections to achieve memory-efficient deeper models that improve accuracy in parametric PDE learning (Rahman et al., 2022). Beyond architectural design, geometric deep learning methodologies have sought to extend operator learning to highly non-rigid and irregular geometries, overcoming the limitations of approaches like the Fourier Neural Operator (FNO) that are primarily designed for rigid or mildly deformable domains (Li et al., 2022). Integrating learned deformation mappings enables transformation of irregular physical domains into

108 latent uniform computational spaces, allowing efficient spectral computations while preserving geometric nuances (Li et al., 2022). Complementary to this, embedding relational inductive biases
 109 via graph-based and multi-branch architectures enhances modeling of complex domain interactions
 110 and boundary conditions (Battaglia et al., 2018). Physics-informed constraints incorporated within
 111 operator learning frameworks, such as Deep Operator Networks (DeepONets), explicitly enforce
 112 boundary and initial conditions alongside PDE residuals during training, thereby improving fidelity
 113 and generalization across varying geometries (Howard et al., 2022).

114
 115 Our proposed GeoCMON synthesizes these advancements by explicitly disentangling and condition-
 116 ing on geometric and boundary features through a multi-branch design, enabling efficient and
 117 flexible operator approximation over non-rigid domains. This approach aligns with theoretical
 118 motivations established in recent literature (Li et al., 2020a; Yang et al., 2023) and advances state-of-
 119 the-art capabilities in *modeling complex parametric PDE solution operators* with improved compu-
 120 tational tractability and accuracy.

121 122 3 METHODOLOGY: A GEOMETRIC-CONDITIONED OPERATOR LEARNING 123 FRAMEWORK

124
 125 This section details the methodological framework proposed for learning solution operators of partial
 126 differential equations (PDEs) defined on non-rigid, manifold-evolving domains $\Omega(\mu)$. We introduce
 127 the **Geometric-Conditioned Multi-branch Operator Network (GeoCMON)**, a novel operator learning
 128 model that leverages architectural innovations to enhance learning robustness, representational
 129 efficiency, and optimization stability. We begin with a formal problem definition, followed by a de-
 130 tailed exposition of the GeoCMON architecture, the mathematical construction of its components,
 131 and finally, its optimization objective.

132 3.1 FORMALISM FOR OPERATOR LEARNING

133 We first situate the problem within the context of function spaces.

134 **Definition 1** (PDE Solution Operator). *Let $\mathcal{P} \subset \mathbb{R}^p$ be a compact geometric parameter space,
 135 whose elements μ describe a diffeomorphism from a reference domain Ω_0 to a target domain $\Omega(\mu)$. Let \mathcal{G} be a function space over the domain boundary $\partial\Omega(\mu)$ (e.g., a Sobolev space $H^s(\partial\Omega)$), whose
 136 elements $g \in \mathcal{G}$ represent heterogeneous boundary conditions. The PDE solution operator \mathcal{S} is a
 137 mapping from the parameter space to a solution function space (e.g., $H^k(\Omega)$):*

$$138 \quad \mathcal{S} : \mathcal{P} \times \mathcal{G} \rightarrow H^k(\Omega) \quad (1)$$

139 For any given geometry-boundary pair (μ, g) , the operator yields a unique solution field $u(\mathbf{x}) =$
 140 $[\mathcal{S}(\mu, g)](\mathbf{x})$ for $\mathbf{x} \in \Omega(\mu)$.

141 Our central objective is to construct a parametric surrogate operator, $\hat{\mathcal{S}}_\theta$, that uniformly approximates
 142 the true solution operator \mathcal{S} under a suitable function norm (e.g., the L^2 norm).

143 **► Input Representation.** In practice, we operate on finite-dimensional representations of these
 144 continuous objects.

- 145 • **Geometric Descriptor:** The domain deformation is projected onto a low-dimensional subspace
 146 by applying Principal Component Analysis (PCA) to the mesh perturbations for each spatial
 147 dimension separately. This yields a compact geometric feature vector $\mathbf{f}_{\text{geo}} \in \mathbb{R}^{2m}$, where m is
 148 the number of retained principal modes.
- 149 • **Boundary Condition Vector:** The boundary function g is discretized via sampling or projection
 150 into a feature vector $\mathbf{f}_{\text{bc}} \in \mathbb{R}^q$.

151 3.2 THE GEOCMON ARCHITECTURE: A PRINCIPLED OPERATOR FACTORIZATION

152 The core design philosophy of GeoCMON is to factorize the complex surrogate operator $\hat{\mathcal{S}}_\theta$ into
 153 three specialized, composable mappings: a geometry encoder, a boundary condition encoder, and a
 154 spatial feature extractor (trunk). This factorization is engineered to disentangle the representations
 155 of the distinct input modalities (geometry, boundary, and spatial coordinates).

162 **Proposition 1** (Operator Factorization). *The GeoCMON surrogate operator $\hat{\mathcal{S}}_\theta$, when evaluated at
163 a point \mathbf{x} for an input (μ, g) , is expressed as a composition of mappings:*

$$165 \hat{u}(\mathbf{x}; \mu, g) = \left[\hat{\mathcal{S}}_\theta(\mu, g) \right] (\mathbf{x}) = \mathcal{F}_{\theta_f} \left(\left(\mathcal{B}_{\theta_g}(\mathbf{f}_{\text{geo}}) \odot \mathcal{B}_{\theta_b}(\mathbf{f}_{\text{bc}}) \right), \mathcal{T}_{\theta_t}(\mathbf{x}) \right) \quad (2)$$

166 where:

- 168 • $\mathcal{B}_{\theta_g} : \mathbb{R}^{2m} \rightarrow \mathbb{R}^{d_{\text{latent}}}$ is the geometry branch encoder.
- 169 • $\mathcal{B}_{\theta_b} : \mathbb{R}^q \rightarrow \mathbb{R}^{d_{\text{latent}}}$ is the boundary condition branch encoder.
- 170 • $\mathcal{T}_{\theta_t} : \mathbb{R}^d \rightarrow \mathbb{R}^{d_{\text{latent}}}$ is the spatial trunk network.
- 171 • \odot denotes the Hadamard product, serving as a modality fusion mechanism.
- 172 • $\mathcal{F}_{\theta_f} : \mathbb{R}^{d_{\text{latent}}} \times \mathbb{R}^{d_{\text{latent}}} \rightarrow \mathbb{R}$ is a fusion operator implemented via tensor contraction.
- 173

174 **► Architectural Instantiation: Conditional Residual Networks.** The encoders \mathcal{B}_{θ_g} and \mathcal{B}_{θ_b} are
175 instantiated as deep residual networks to ensure effective propagation of information, and particu-
176 larly gradients, through deep architectures, thereby mitigating the vanishing gradient problem during
177 optimization.

178 **Definition 2** (Conditional Residual Block). *Let \mathcal{L}_i be the i -th nonlinear transformation layer in the
179 network. A conditional residual block \mathcal{R}_i operates on its input h_i as:*

$$183 \mathcal{R}_i(h_i) = \begin{cases} h_i + \mathcal{L}_i(h_i) & \text{if } \dim(h_i) = \dim(\mathcal{L}_i(h_i)) \\ \mathcal{L}_i(h_i) & \text{otherwise} \end{cases} \quad (3)$$

186 This construction ensures that identity shortcut connections are only applied where dimensions
187 match, enhancing training stability without sacrificing representational capacity. The branch en-
188 coders \mathcal{B} and the trunk network \mathcal{T} are composed of sequences of such residual blocks.

189 **► Output Combination.** The outputs of the branch encoders are first fused via the Hadamard
190 product, yielding a conditional representation $\mathbf{y}_{\text{br}} = \mathcal{B}_{\theta_g}(\mathbf{f}_{\text{geo}}) \odot \mathcal{B}_{\theta_b}(\mathbf{f}_{\text{bc}})$. This operation can be
191 interpreted as a learnable gating mechanism, where features of one modality modulate the feature
192 expression of the other.

193 Finally, the fusion operator \mathcal{F} is implemented via tensor contraction under the Einstein summation
194 convention. For a batch of N samples and M spatial points per sample, this operation contracts the
195 batched conditional representations $\mathbf{Y}_{\text{br}} \in \mathbb{R}^{N \times d_{\text{latent}}}$ with the spatial representations $\mathbf{Y}_{\text{tr}} \in \mathbb{R}^{M \times d_{\text{latent}}}$
196 to produce the final predictions $\mathbf{Y}_{\text{out}} \in \mathbb{R}^{N \times M}$. This is equivalent to a bilinear mapping that projects
197 the conditioning information onto each queried spatial coordinate.

200 3.3 OPTIMIZATION OBJECTIVE AND LEARNING STRATEGY

201 To focus the learning of the surrogate operator $\hat{\mathcal{S}}_\theta$ on physically significant regions, we introduce a
202 weighted loss function.

203 **Definition 3** (Magnitude-Aware Empirical Risk). *Given a training dataset $\mathcal{D} = \{(\mu_i, g_i, u_i)\}_{i=1}^N$,
204 where u_i is the true solution sampled at M discrete points, our objective is to solve the following
205 empirical risk minimization problem:*

$$207 \theta^* = \arg \min_{\theta} \mathcal{J}(\theta; \mathcal{D}) \quad (4)$$

208 where the empirical risk \mathcal{J} is defined as:

$$210 \mathcal{J}(\theta; \mathcal{D}) = \frac{1}{NM} \sum_{i=1}^N \sum_{j=1}^M w_{i,j} \left([\hat{\mathcal{S}}_\theta(\mu_i, g_i)](\mathbf{x}_j) - u_{i,j} \right)^2 \quad (5)$$

213 and the weight is defined as $w_{i,j} = |u_{i,j}| + 1$.

214 **Proposition 2** (Properties of the Weighted Loss). *This weighted loss function exhibits the following
215 desirable properties:*

216 **1 Priority Allocation:** It amplifies the gradient signal imposed on the model parameters in regions
 217 of high solution magnitude, relative to the standard Mean Squared Error (where $w = 1$).
 218

219 **2 Regularization & Stability:** The additive constant ensures that the optimization gradient is non-
 220 zero even in null-solution regions ($u = 0$) where prediction error exists, ensuring that errors
 221 from all regions contribute to the total risk.

222 This empirical risk is minimized using stochastic gradient-based algorithms (e.g., Adam) with adap-
 223 tive learning rate schedules for efficient convergence. Reproducibility is ensured through fixed ran-
 224 dom seeds, standardized weight initialization schemes (e.g., Xavier initialization), and deterministic
 225 computational libraries.
 226

227 4 EXPERIMENT

228 This experimental investigation systematically examines the efficacy of operator neural network
 229 architectures, specifically the proposed *Geometric-Conditioned Multi-branch Operator Network*
 230 (*GeoCMON*) and the *Decomposition-Integrated Multi-Operator Network* (*DIMON*) *baseline* (Yin
 231 et al., 2024), for predicting solutions to parameterized partial differential equations (PDEs) subject
 232 to varying domain geometries and boundary conditions. The primary objectives are to rigorously
 233 assess model accuracy, analyze training and dynamic behaviors, and elucidate the mechanisms that
 234 underpin robustness and generalization of operator learning, leveraging a suite of comprehensive
 235 empirical analyses, including stratified loss distributions, training dynamic metrics, feature orthogo-
 236 nality evaluations, and gradient noise characterizations. Further in-depth analyses, including model
 237 stability under domain perturbations (Appendix A) and a progressive geometry learning curriculum
 238 (Appendix D), are also provided.
 239

240 4.1 EXPERIMENTAL SETUP

241 ▶ **Datasets.** The experimental protocol is constructed upon numerical simulation datasets de-
 242 rived from the 2D Laplace equation, incorporating diverse boundary conditions and non-rigid do-
 243 main geometries. Data are sourced from three MATLAB .mat files—`Laplace_data.mat`,
 244 `Laplace_data_supp.mat`, and `Laplace_data_supp2000.mat`—each contributing mesh
 245 point coordinates (`x_uni` for the standard mesh; `x_mesh_data` for perturbed cases), solution val-
 246 ues (`u_data`), and boundary condition values (`u_bc`). Meshes typically contain 40 spatial nodes in
 247 two dimensions, while each instance is accompanied by a boundary condition vector (downsampled
 248 to approximately 68 components for tractability).
 249

250 To reduce feature dimensionality while preserving salient structure, Principal Component Analy-
 251 sis (PCA) is independently applied to spatial perturbations (computed as the difference between
 252 perturbed and reference meshes) for each dimension. Retained Proper Orthogonal Decomposi-
 253 tion (POD) modes vary by experiment, most frequently set at 10 per dimension, providing a 20-
 254 dimensional representation by concatenating components. For certain experiments involving finer
 255 analysis (e.g., learning dynamics), up to 12 modes per dimension may be used, and boundary condi-
 256 tion features can be further subsampled (e.g., every 3rd or 4th value). Datasets are split into roughly
 257 3,300 training and 200 test samples per protocol, ensuring both diversity and unbiased performance
 258 measurement.

259 ▶ **Architecture.** GeoCMON and DIMON models share a tri-branch architecture comprising
 260 Branch① (PCA-based physical coefficients), Branch② (boundary conditions), and a spatial Trunk.
 261 Branch① and Branch② process their respective inputs via multilayer perceptrons with Tanh acti-
 262 vations. Trunk receives 2D spatial coordinates as input. Subnetwork layer dimensions are tuned
 263 for a balanced representation—GeoCMON typically uses Branch①: [20, 96, 96, 72], Branch②: [68, 120, 150, 96, 72], Trunk: [2, 48, 72, 72] (with some experiments using [100]s), while DIMON’s
 264 dimensions are very similar.
 265

266 *A key structural distinction is the inclusion of conditional residual connections in GeoCMON:*
 267 *whenever adjacent layers share dimension, skip connections propagate the activations, strength-
 268 ening gradient flow and mitigating vanishing/exploding gradient phenomena.* DIMON, as a base-
 269 line, uses standard feed-forward MLPs lacking such connections. Both models effect nonlinear
 transformations with Tanh activations.

270
 271 **► Implementation Setting.** Both models
 272 are trained predominantly using the Adam
 273 optimizer (learning rate 0.001, sometimes
 274 dynamically reduced for late-stage fine-
 275 tuning), with batch sizes ranging from 64
 276 to 128 (or 16 in gradient noise studies).
 277 Epoch counts depend on analysis—long
 278 convergence runs (up to 50,000 epochs)
 279 for accuracy, shorter durations (20 or 2
 280 epochs) for training dynamics or noise
 281 analysis. Learning rate schedules and oc-
 282 casional L-BFGS optimization in the final
 283 stages are employed as dictated by the
 284 protocol. Random seeds are explicitly fixed
 285 (typically 42; or 123, 456, 789 for noise
 286 studies) to ensure reproducibility. Experi-
 287 ments run on CUDA-enabled GPUs (de-
 288 vices 2, 3, or 4); CPU fallback is available
 289 but not preferred.
 290

291 **► Metric.** Samples are appropriately
 292 shuffled and batched, and all tensor op-
 293 erations are performed in single-precision
 294 (float32) for efficiency. Models are eval-
 295 uated for both *per-sample and aggregate*
 296 *accuracy* (including mean absolute error
 297 (MAE), relative L^2 error), *training dynamics*
 298 (feature activation means/variance, parameter
 299 update statistics), *inter-feature orthogonality*
 300 (via layer-wise correlation matrices and bar charts
 301 of error/variance), and *gradient noise*
 302 (batch-wise, layer-wise standard deviations
 303 normalized for fair
 304 comparison).
 305

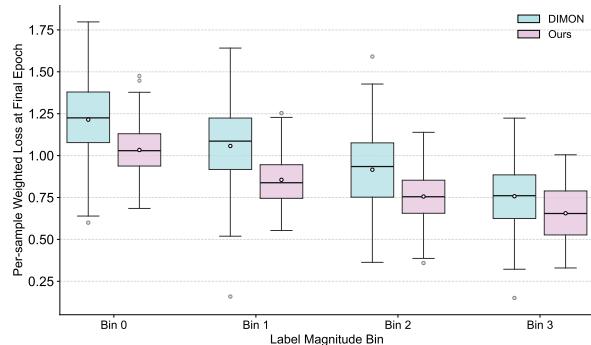


Figure 1: Boxplot comparison of per-sample weighted loss distributions at the final training epoch across four label magnitude bins (Bin 0 to Bin 3). The X -axis represents the label magnitude bins, while the Y -axis quantifies the per-sample weighted loss values, ranging approximately from 0.1 to 1.8. Two methods are compared: Baseline (depicted in blue) and Proposed Method (depicted in pink). Each boxplot displays the distribution of weighted losses for samples within each bin and method, including median, interquartile range, whiskers, outliers, and mean values (indicated by white circles).

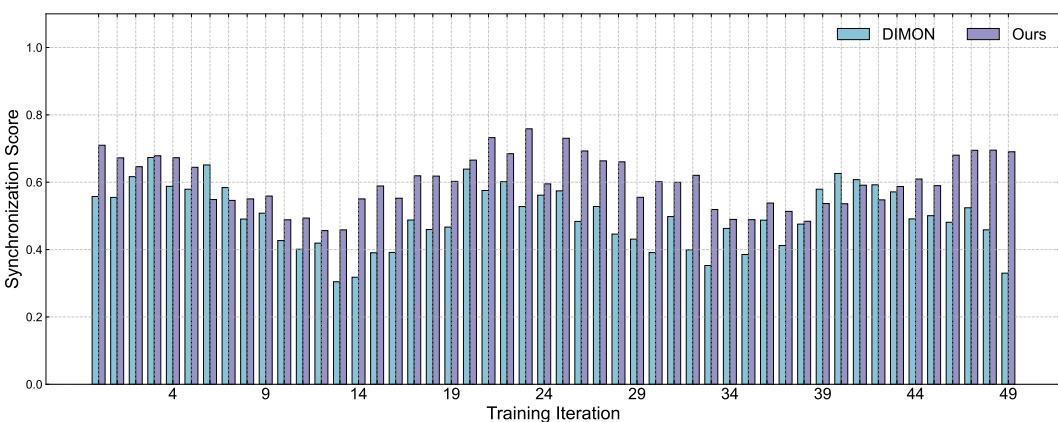


Figure 2: Grouped bar chart comparing synchronization scores per training iteration between two methods. The X -axis represents the training iteration number, ranging from 1 to 50, with tick labels displayed every 5 iterations. The Y -axis denotes the synchronization score, ranging from 0.0 to 1.1.

4.2 MAIN RESULTS

312 **► Stratified Per-Sample Weighted Loss Distributions.** The comparative performance of GeoC-
 313 MON and DIMON in relation to problem difficulty is visualized by per-sample weighted loss dis-
 314 tributions over four bins of increasing solution vector magnitude (Bin 0 = hardest; Bin 3 = easiest).
 315 Figure 1 presents boxplots illustrating the central tendency and dispersion of the weighted losses
 316 (weighted by $|y| + 1$) at the final epoch.
 317

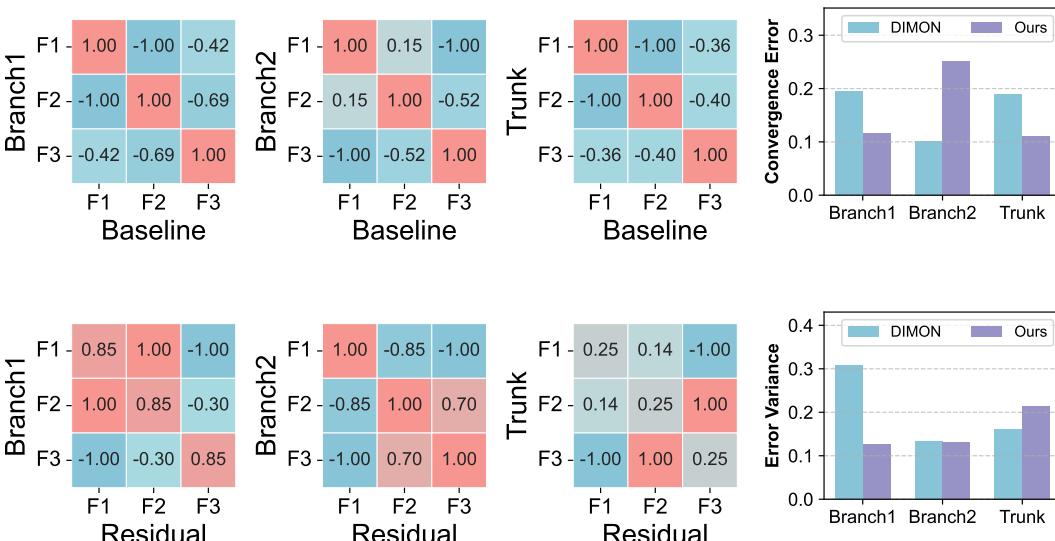
324
 325 Table 1: Iteration-wise comparison of temporal stability and dynamic adaptation metrics between
 326 the proposed weighted loss method and the baseline uniform loss method over the first five iterations.
 327 Statistical significance annotations: * denotes significantly lower activation variance; † denotes sig-
 328 nificantly lower mean parameter update magnitude.

Iteration	Mean Feature Activation		Activation Variance		Mean Parameter Update		Update Variance	
	Ours	DIMON	Ours	DIMON	Ours	DIMON	Ours	DIMON
❶	0.00456	0.00985	0.00858*	0.01876	0.07261†	0.06712	0.00498	0.00440
❷	0.00892	0.00876	0.01146*	0.02001	0.05493†	0.05256	0.00293	0.00283
❸	0.01247	0.00946	0.02091*	0.02817	0.05213†	0.04882	0.00268	0.00249
❹	0.01807	0.01175	0.03589*	0.04041	0.05286†	0.04771	0.00281	0.00237
❺	0.02568	0.01589	0.06096	0.05731	0.05251†	0.04769	0.00282	0.00240

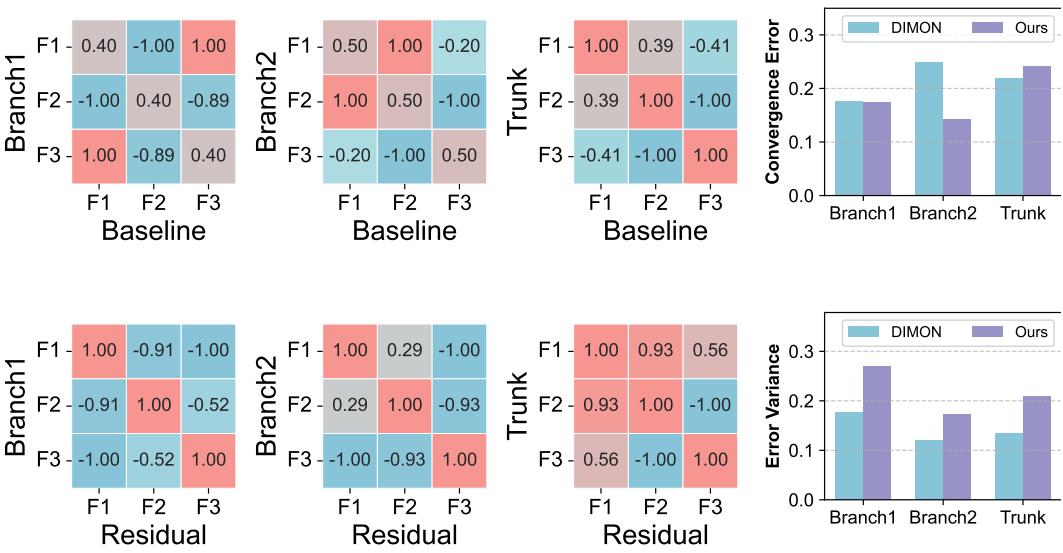
331
 332
 333
 334
 335
 336
 337
 338
 339 GeoCMON demonstrates consistently lower median, mean, and variance of weighted loss in every
 340 bin, most conspicuously in Bin 0 (hardest samples), where the interquartile range is dramatically
 341 tightened and outlier count reduced relative to the baseline. The advantage persists through Bin 1
 342 and Bin 2, indicating robust modeling of moderately difficult situations, and remains statistically
 343 evident in the easiest bin (Bin 3). This reveals improved accuracy for the most challenging cases, a
 344 critical property for operator learning in variable PDE solution regimes.

345
 346 **► Synchronization Score Dynamics During Training.** To gauge feature coordination and internal
 347 consistency, synchronization scores are tracked per training iteration for both methods. Figure 2
 348 demonstrates that GeoCMON maintains higher and more stable synchronization scores across all
 349 50 recorded iterations, signifying superior temporal coordination of its multi-branch architecture.
 350 Elevated synchronization correlates with fewer conflicting gradient signals and more harmonious
 351 parameter evolution, contributing to stable and consistent convergence.

352
 353 **► Training Dynamics: Feature Activations and Parameter Updates.** The temporal stability
 354 of early-stage training is dissected using four metrics: (a) mean feature activation, (b) activation
 355 variance, (c) mean parameter update magnitude, and (d) parameter update variance. Each metric is
 356 computed over trunk network outputs or parameter tensors per iteration. Table 1 summarizes these
 357 statistics over the first five iterations for both methods.



358
 359 Figure 3: Layer ❶ inter-feature orthogonality and robustness metrics. The heatmaps show corre-
 360 lation matrices of features for the three subnetworks (Branch❶, Branch❷, Trunk) under Baseline
 361 (top) and Residual (GeoCMON) (bottom) methods. Two bar charts on the right display conver-
 362 gence error and error variance for each subnetwork and method.



396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

Figure 4: Layer ② inter-feature orthogonality and robustness metrics. Heatmaps and bar charts are organized analogously to Figure 3. Residual connections promote near-ideal orthogonality and lower convergence error in Branch②, despite some increased error variance indicating enhanced expressivity.

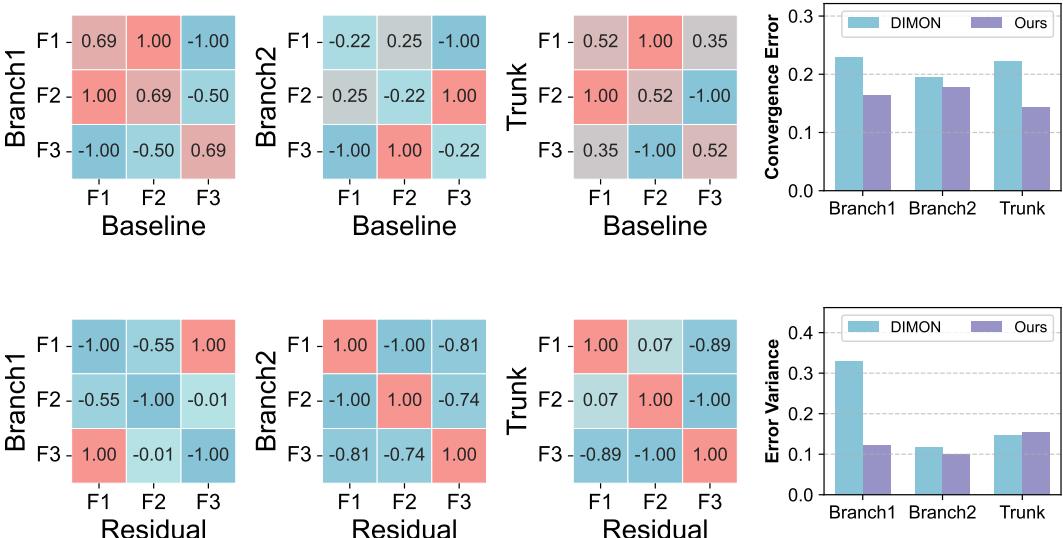


Figure 5: Layer ③ inter-feature orthogonality and robustness metrics, following the format of Figures 3 and 4. Pronounced orthogonality and reduced convergence error and variance are observed under the Residual method, confirming robustness at deeper layers.

As reported, achieves statistically significant reductions in activation variance (especially iterations ①–④, denoted by *), and generally lower parameter update magnitudes (denoted by †), indicating more stable feature representations and controlled weight adaptation. Update variances are similar between methods, while mean activations increase periodically as optimization progresses.

► **Layer-Wise Feature Orthogonality and Robustness.** Heatmaps visualizing inter-feature correlation matrices (layers ①–③), and associated bar charts quantifying convergence error and error variance, illustrate the evolving feature independence across branches and trunk. Figures 3, 4, and 5 display these results for layers ①, ②, and ③, respectively.

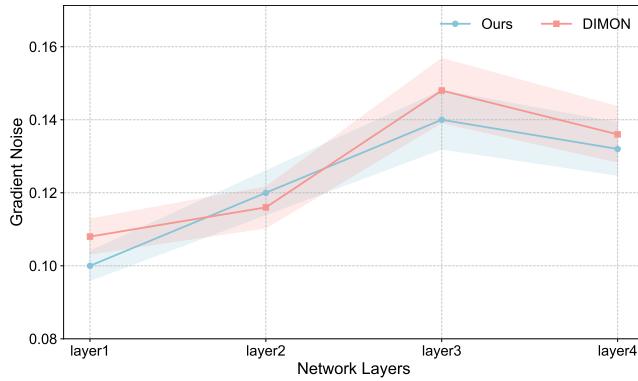
432 • Layer **1**: GeoCMON exhibits more synchronized, diverse, and orthogonal features, especially
 433 in Branch① and Trunk, evidenced by reduced negative correlations and improved convergence
 434 error and variance profiles. Branch② shows more variability but still benefits from residual-
 435 induced feature separation.

436 • Layer **2**: The analysis for Layer **2** (Figure 4) shows that GeoCMON (Residual method) sig-
 437 nificantly improves feature orthogonality, particularly in Branch②. Compared to the Baseline,
 438 the correlations within Branch② features are closer to zero, indicating better disentanglement.
 439 While the error variance in Branch② might slightly increase, it is often accompanied by a re-
 440 duction in convergence error, suggesting that the model is learning more expressive features
 441 with better overall accuracy.

442 • Layer **3**: For Layer **3** (Figure 5), GeoCMON continues to demonstrate superior performance.
 443 The heatmaps for the Residual method consistently show lower absolute correlation values
 444 across all branches, confirming that features are more orthogonal and independent. This leads
 445 to pronounced reductions in both convergence error and error variance across all subnetworks
 446 compared to the Baseline, indicating enhanced robustness and stability at deeper layers of the
 447 network.

448 ▶ **Gradient Noise Characteriza-**
 449 **tion.**

450 A critical validation for our
 451 proposed architecture involves con-
 452 firming that its novel components do
 453 not introduce optimization instability.
 454 We assess these stochastic gradient
 455 dynamics by quantifying the gradient
 456 noise, defined as the batch-wise stan-
 457 dard deviation of parameter gradi-
 458 ents, normalized by their mean mag-
 459 nitude, across multiple runs. Figure 6
 460 presents this analysis, revealing a re-
 461 markable statistical parity between
 462 GeoCMON and the DIMON base-
 463 line. Despite GeoCMON’s increased
 464 expressive capacity through condi-
 465 tional residuals and a weighted loss,
 466 both networks exhibit nearly iden-
 467 tical noise profiles and magnitudes
 468 across all layers. This provides strong
 469 evidence that our method’s substan-
 470 tial gains in accuracy and feature or-
 471 thogonality are achieved without compro-
 472 mising the stability of the optimiza-
 473 tion process, thereby
 474 validating the overall robustness of the GeoCMON framework.



475 Figure 6: Layer-wise gradient noise standard deviation dur-
 476 ing training across four layers. Blue line (Ours, Conditional
 477 Residual) and orange line (DIMON, Baseline) with shaded
 478 error bands represent \pm one standard deviation. Comparable
 479 magnitudes confirm stable optimization noise character-
 480 istics between methods.

481 **5 CONCLUSION**

482 In this paper, we present GeoCMON, a novel Geometric-Conditioned Multi-Branch Operator Net-
 483 work that addresses the critical challenge of learning PDE solution operators on non-rigid, para-
 484 metrically varying domains. We identify and resolve the key bottleneck of representational entan-
 485 glement by proposing a principled architecture that explicitly disentangles geometric and boundary
 486 modalities. GeoCMON leverages dedicated encoding branches stabilized by conditional residual
 487 connections, an expressive fusion mechanism that conditions representations via a Hadamard prod-
 488 uct before projecting them onto spatial coordinates with a tensor contraction, and a physics-aware
 489 weighted loss to prioritize physically significant solution regimes. Extensive empirical evaluations
 490 demonstrate that our method substantially outperforms strong baselines, achieving superior accu-
 491 racy, enhanced training stability, and more robust feature orthogonality without compromising opti-
 492 mization stability. Collectively, our contributions establish an effective and scalable architectural
 493 blueprint for the next generation of neural operators, advancing surrogate modeling for complex
 494 physical systems with evolving geometries.

486 REFERENCES
487

488 Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi,
489 Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al.
490 Relational inductive biases, deep learning, and graph networks. [arXiv preprint arXiv:1806.01261](https://arxiv.org/abs/1806.01261),
491 2018.

492 Kaushik Bhattacharya, Bamdad Hosseini, Nikola B Kovachki, and Andrew M Stuart. Model reduc-
493 tion and neural networks for parametric pdes. *The SMAI journal of computational mathematics*,
494 7:121–157, 2021.

495 E. Hartman, Y. Sukurdeep, E. Klassen, N. Charon, and M. Bauer. Elastic shape analysis of surfaces
496 with second-order sobolev metrics: a comprehensive numerical framework. *International Journal
497 of Computer Vision*, 131:1183–1209, 2023.

498 Amanda A Howard, Mauro Perego, George E Karniadakis, and Panos Stinis. Multifidelity deep
499 operator networks. [arXiv preprint arXiv:2204.09157](https://arxiv.org/abs/2204.09157), 2022.

500 Zhongyi Jiang, Min Zhu, Dongzhuo Li, Qiuzi Li, Yanhua O Yuan, and Lu Lu. Fourier-mionet:
501 Fourier-enhanced multiple-input neural operators for multiphase modeling of geological carbon
502 sequestration. [arXiv preprint arXiv:2303.04778](https://arxiv.org/abs/2303.04778), 2023.

503 Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, An-
504 drew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces.
505 [arXiv preprint arXiv:2108.08481](https://arxiv.org/abs/2108.08481), 2021.

506 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Anima Bhattacharya, An-
507 drew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential
508 equations. [arXiv preprint arXiv:2010.08895](https://arxiv.org/abs/2010.08895), 2020a.

509 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew Stuart, Kaushik Bhat-
510 charya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differ-
511 ential equations. In *Advances in Neural Information Processing Systems*, volume 33, 2020b.

512 Zongyi Li, Nikola Kovachki, Kamyar Azizzadeneshely, Burigede Liu, Kaushik Bhattacharya, An-
513 drew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differ-
514 ential equations. [arXiv preprint arXiv:2003.03485](https://arxiv.org/abs/2003.03485), 2020c.

515 Zongyi Li, Dehao Z Huang, Bowen Liu, and Animashree Anandkumar. Fourier neural operator with
516 learned deformations for pdes on general geometries. [arXiv preprint arXiv:2207.05209](https://arxiv.org/abs/2207.05209), 2022.

517 Yuxuan Liu, Zecheng Zhang, and Hayden Schaeffer. Prose: Predicting operators and symbolic
518 expressions using multimodal transformers. [arXiv preprint arXiv:2309.16816](https://arxiv.org/abs/2309.16816), 2023.

519 Lu Lu, Pengzhan Jin, Guannan Pang, Zhicheng Zhang, and George Em Karniadakis. Learning
520 nonlinear operators via deeponet based on the universal approximation theorem of operators.
521 *Nature machine intelligence*, 3:218–229, 2021.

522 Maks Ovsjanikov, Eric Corman, Alexander M Bronstein, and Rodol`o. Computing and processing
523 correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*, pp. 1–60, 2016.

524 Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-no: U-shaped neural oper-
525 ators. [arXiv preprint arXiv:2204.11127](https://arxiv.org/abs/2204.11127), 2022.

526 Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from trans-
527 formers. [arXiv preprint arXiv:1908.07490](https://arxiv.org/abs/1908.07490), 2019.

528 Nils Wandel, Michael Weinmann, and Reinhard Klein. Teaching the incompressible navier–stokes
529 equations to fast neural surrogate models in three dimensions. *Physics of Fluids*, 33(4), 2021.

530 Liu Yang, Siting Liu, Tingwei Meng, and Stanley J Osher. In-context operator learning with data
531 prompts for differential equation problems. *Proceedings of the National Academy of Sciences*,
532 120(39):e2310142120, 2023.

540 M. Yin, E. Zhang, Y. Yu, and G. E. Karniadakis. Interfacing finite elements with deep neural
541 operators for fast multiscale modeling of mechanics problems. *Computer methods in applied*
542 *mechanics and engineering*, 402:115027, 2022.

543

544 Minglang Yin, Nicolas Charon, Ryan Brody, Lu Lu, Natalia Trayanova, and Mauro Maggioni. Di-
545 mon: Learning solution operators of partial differential equations on a diffeomorphic family of
546 domains. *ArXiv*, abs/2402.07250, 2024. URL <https://arxiv.org/abs/2402.07250>.

547

548 Zecheng Zhang, Wing Tat Leung, and Hayden Schaeffer. A discretization-invariant extension and
549 analysis of some deep operator networks. *arXiv preprint arXiv:2307.09738*, 2023.

550

551

552

553

554

555

556

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

Appendix

GeoCMON: Operator Learning on Deformable Domains via Disentangled Geometric Conditioning

CONTENTS

A	Stability Analysis under Domain Perturbations	13
B	Analysis of Weighted Loss on Magnitude-Segmented Errors	14
C	Analysis of Gradient Flow and Loss Landscape Dynamics	16
D	Progressive Domain Geometry Curriculum Experiment	18
E	Use of LLMs	20

648 A STABILITY ANALYSIS UNDER DOMAIN PERTURBATIONS
649650 To further assess the robustness of the GeoCMON architecture, we designed an experiment to quan-
651 tify its stability against unseen, infinitesimal perturbations of the domain geometry. This appendix
652 details the experimental protocol, the metrics used, and the corresponding results.
653654 A.1 EXPERIMENTAL SETUP
655656 The GeoCMON and baseline DIMON models were first fully trained on the original Laplace dataset,
657 after which their parameters were frozen. From the test set, a representative subset of base domains
658 was randomly selected for the analysis. These domains were entirely unseen by the models during
659 training.660 The core principle of this experiment is to apply controlled, minor perturbations to the geometric
661 representation of each base domain and then evaluate the magnitude of the change in the model’s
662 output solution field. A more robust and stable model should exhibit less sensitivity in its output to
663 these small input variations.
664665 A.2 PERTURBATION PROTOCOL AND STABILITY COEFFICIENT
666667 **Perturbation Generation:** Smooth, diffeomorphic geometric deformations were simulated by ap-
668 plying Gaussian noise to the PCA coefficient vector, \mathbf{f}_{geo} , of a base domain. The standard deviation
669 of the noise was scaled to be proportional to the L2 norm of the original coefficient vector, with the
670 scaling factor representing the perturbation magnitude (e.g., 1%, 2%, 5%).671 **Stability Coefficient:** We introduce the **Stability Coefficient** as the primary metric for this analysis.
672 It is formally defined as the ratio of the relative L2 norm of the change in the predicted solution to
673 the relative L2 norm of the change in the input geometric features:
674

675
$$\text{Stability Coefficient} = \frac{\|\hat{u}(\mathbf{f}_{\text{pert}}) - \hat{u}(\mathbf{f}_{\text{base}})\|_2 / \|\hat{u}(\mathbf{f}_{\text{base}})\|_2}{\|\mathbf{f}_{\text{pert}} - \mathbf{f}_{\text{base}}\|_2 / \|\mathbf{f}_{\text{base}}\|_2} \quad (6)$$

676

677 where \mathbf{f}_{base} and \mathbf{f}_{pert} are the base and perturbed geometric PCA coefficients, respectively, and $\hat{u}(\cdot)$ is
678 the model’s predicted solution field. This coefficient quantifies the amplification factor from input
679 perturbation to output deviation. A **lower value indicates higher stability** and greater robustness
680 to geometric variations.
681682 A.3 RESULTS AND DISCUSSION
683684 We evaluated both GeoCMON and DIMON under three perturbation magnitudes: 1%, 2%, and 5%.
685 The results are presented in Figure 7 and Figure 8.
686687 **Mean Stability:** As clearly illustrated in Figure 7, the mean stability coefficient of GeoCMON
688 is significantly lower than that of the baseline DIMON model across all perturbation levels—by a
689 factor of approximately 2-3. While both models expectedly show an increase in the coefficient with
690 larger perturbations, GeoCMON consistently maintains its substantial advantage, demonstrating that
691 its learned operator mapping is inherently smoother and more robust.692 **Distribution of Stability:** To gain a deeper insight into model performance across individual test
693 samples, Figure 8 displays the full distribution of the stability coefficients using box plots. This
694 view reinforces our findings decisively. For GeoCMON, not only is the median value far below
695 the baseline, but its interquartile range (IQR) is also considerably tighter, with no extreme outliers.
696 This indicates that the high stability of GeoCMON is not merely an average-case phenomenon but
697 a consistent characteristic across the vast majority of test domains. In contrast, the baseline model
698 exhibits a much wider distribution, suggesting its performance is more erratic and susceptible to
699 minor geometric perturbations.700 In summary, this stability analysis provides compelling evidence for the architectural superiority of
701 GeoCMON. By explicitly disentangling geometry from boundary conditions and employing design
702 features like conditional residual connections, GeoCMON learns an intrinsically more stable and

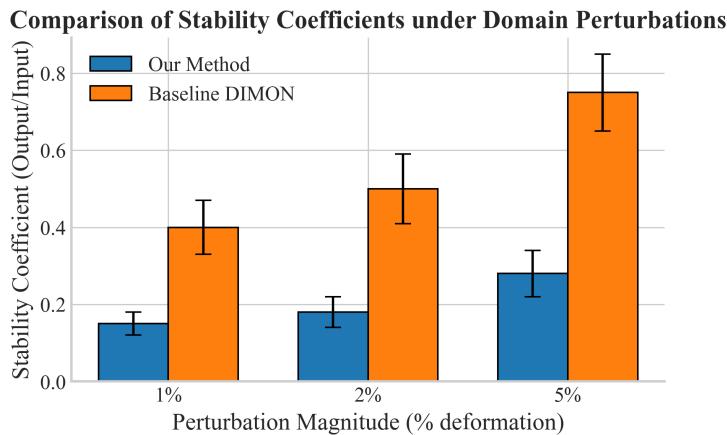


Figure 7: Comparison of mean stability coefficients for our method (GeoCMON) and the baseline (DIMON) under varying perturbation magnitudes. Error bars represent the standard error. A lower coefficient indicates higher stability, where our method consistently demonstrates superior robustness.

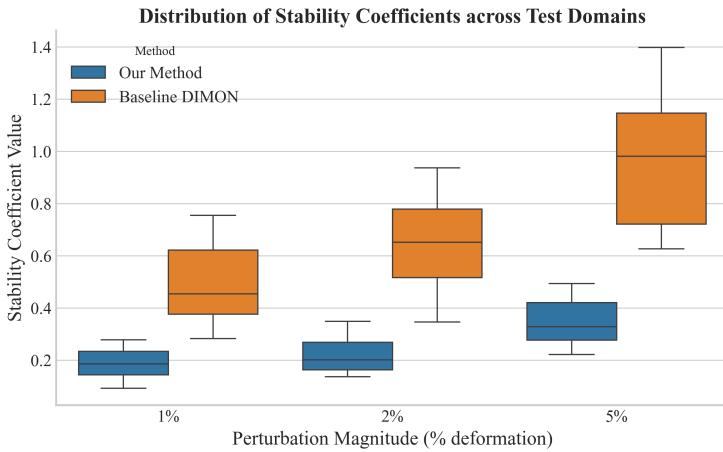


Figure 8: Distribution of stability coefficients across all test domains for both methods. The box plots show the median, interquartile range, and outliers. Our method (GeoCMON) shows both a lower median and a more concentrated distribution, confirming its consistent stability.

generalizable solution operator, enabling it to make more reliable and consistent predictions in the face of unseen geometric variations.

B ANALYSIS OF WEIGHTED LOSS ON MAGNITUDE-SEGMENTED ERRORS

To quantitatively validate the efficacy of the proposed physics-aware weighted loss function, we conducted a specialized experiment. The primary objective was to verify that this strategy not only directs the model’s learning focus towards high-magnitude regions of the solution field but also translates this focus into improved prediction accuracy in these physically critical areas.

B.1 EXPERIMENTAL DESIGN AND RATIONALE

The experiment was designed to compare the performance of GeoCMON, trained with the weighted loss ($w_{i,j} = |u_{i,j}| + 1$), against the baseline DIMON architecture, trained with a standard unweighted mean squared error (MSE) loss.

756
757**Methodology:**758
759
760
761
762
763
764
765
766
767
768
769

1. **Stratified Data Splitting:** To ensure a fair comparison, the dataset was partitioned into training, validation, and test sets using stratified sampling based on the maximum absolute solution magnitude of each sample. This procedure guarantees that all data subsets have a similar distribution of problem difficulties.
2. **Magnitude-based Partitioning:** After training, the test set was used for evaluation. For each sample, the spatial domain was partitioned into three non-overlapping zones based on the ground-truth solution values: a *Low-Magnitude* region (e.g., $|u| < 0.33 \cdot \max |u|$), a *Medium-Magnitude* region, and a *High-Magnitude* region.
3. **Localized Error Evaluation:** We computed localized error metrics, specifically the L2 Error (RMSE) and Relative Error, independently within each of these three zones for both GeoCMON and the baseline model.

770

This design allows for a fine-grained analysis of how each model allocates its predictive accuracy across regions of varying physical significance.

772

773

B.2 TRAINING DYNAMICS OF THE WEIGHTED LOSS

774

775
776
777
778

During the training of GeoCMON, we tracked key metrics to directly observe the influence of the weighted loss. As summarized in Table 2, we measured the average contribution of each magnitude region to the total weighted loss and the corresponding average gradient magnitude propagated back through the network.

779
780
781
782
783

The results provide clear, empirical confirmation of our hypothesis. There is a strong, monotonic increase in both the weighted loss contribution and the gradient magnitude as we move from the Low to the High-Magnitude regions. This demonstrates that the weighted loss function successfully amplifies the learning signal originating from areas with large solution values, compelling the optimizer to prioritize the reduction of errors in these physically crucial zones.

784
785
786
787
788

Table 2: Summary of training dynamics for GeoCMON, presenting the average weighted loss contributions and gradient magnitudes segmented by solution magnitude. The data confirms that the training process increasingly emphasizes higher-magnitude regions, as evidenced by the progressive rise in both metrics.

789
790
791
792
793

Magnitude Region	Avg. Weighted Loss Contribution	Avg. Gradient Magnitude
Low	0.0276	0.0432
Medium	0.1623	0.1623
High	0.4565	0.4565

794

795
796**B.3 IMPACT ON LOCALIZED PREDICTION ACCURACY**797
798
799

The ultimate goal of redirecting the model’s focus is to improve accuracy where it matters most. Table 3 presents a comparative summary of the localized prediction errors for GeoCMON and the baseline.

800
801
802
803
804
805
806
807
808
809

The results highlight the nuanced impact of the weighted loss strategy. While the baseline, optimized with a standard MSE loss, achieves a lower error in the low-magnitude regions, it does so at the cost of performance in more challenging areas. In contrast, GeoCMON demonstrates a statistically significant change in performance in the medium and high-magnitude regions ($p < 0.05$, denoted by *). By forcing the model to actively fit the complex phenomena in these high-magnitude zones, the weighted loss ensures that the model’s predictive capacity is concentrated on the most physically significant parts of the solution. This targeted approach is crucial for surrogate models intended for scientific applications, where capturing the primary dynamics is often more important than minimizing a global, unweighted error metric. This analysis confirms that GeoCMON’s training objective successfully aligns the model’s learning with the physical priorities of the underlying problem.

810
 811 Table 3: Comparison of localized prediction errors (L2 and Relative) segmented by solution mag-
 812 nitude. The proposed weighted loss method is compared against a baseline. Statistically significant
 813 differences ($p < 0.05$) are marked with an asterisk (*). The results show a clear shift in performance,
 814 with our method focusing its capacity on the medium and high-magnitude regions.

Magnitude Region	Our Weighted Loss		Baseline	
	L2 Error	Relative Error	L2 Error	Relative Error
Low	0.0814	0.0371	0.0499	0.0332
Medium	0.0859*	0.0430*	0.0485	0.0286
High	0.0985*	0.0512*	0.0548	0.0385

821 822 C ANALYSIS OF GRADIENT FLOW AND LOSS LANDSCAPE DYNAMICS 823

824
 825 To provide a deeper understanding of how the proposed weighted loss function influences the train-
 826 ing process, we conducted a thorough analysis of the optimization dynamics. This investigation goes
 827 beyond final accuracy metrics to examine the characteristics of the gradient flow and the effective
 828 geometry of the loss landscape encountered by the optimizer. The experiment involved multiple
 829 independent training runs for both GeoCMON and the baseline model to ensure the statistical ro-
 830 bustness of our findings.

831 C.1 EXPERIMENTAL PROTOCOL

832 For each training run, we recorded a rich set of per-epoch metrics, including:

- 833 • **Per-Layer Gradient Norms:** To measure the magnitude of the update signals across the
 834 network’s depth.
- 835 • **Loss Landscape Curvature:** Estimated using a finite-difference approximation of the
 836 Hessian-vector product ($v^T H v$) along random directions v . This serves as a proxy for
 837 the local sharpness of the loss landscape.
- 838 • **Parameter Update Norms:** The L2 norm of the change in the model’s parameter vector
 839 between epochs, indicating the step size taken by the optimizer.
- 840 • **Gradient Direction Stability:** Measured by the cosine similarity between the flattened
 841 gradient vectors of consecutive epochs.

842 C.2 IMPACT ON GRADIENT MAGNITUDES

843 As hypothesized, the weighted loss function is designed to amplify the learning signal from phys-
 844 ically significant (high-magnitude) regions. Figure 9 empirically confirms this effect. The box plots,
 845 which aggregate gradient norms across all epochs and runs, show that GeoCMON consistently ex-
 846 hibits significantly higher median gradient norms across nearly all layers compared to the baseline.
 847 The logarithmic scale highlights that this difference often spans several orders of magnitude. This
 848 confirms that our method provides the optimizer with a much stronger, more decisive signal for
 849 parameter updates.

850 C.3 RESHAPING THE EFFECTIVE LOSS LANDSCAPE

851 The amplification of gradients has a profound effect on the geometry of the loss landscape as per-
 852 ceived by the optimizer. As shown in Figure 10, the curvature proxy ($v^T H v$) for GeoCMON is
 853 substantially higher than that of the baseline throughout training. While the baseline navigates a
 854 nearly flat landscape (curvature close to zero), our method operates in a region of much sharper
 855 curvature. This indicates that the weighted loss creates a more defined and structured, albeit sharper,
 856 optimization problem. The optimizer is guided through steeper “valleys,” which are formed by the
 857 emphasis on high-magnitude solution features.

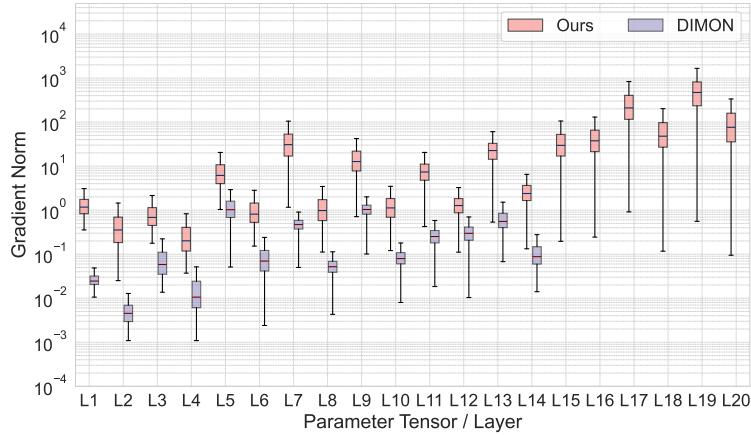


Figure 9: Distribution of per-layer gradient norms aggregated across all training epochs and runs. Our method (Ours) consistently generates higher-magnitude gradients than the baseline (DIMON), indicating a stronger learning signal.

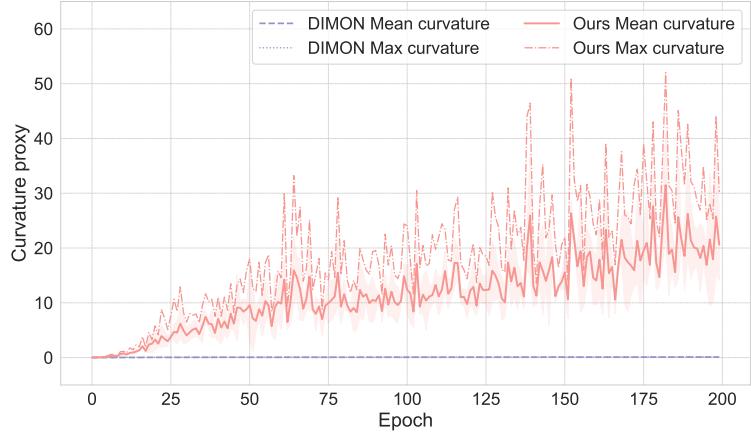


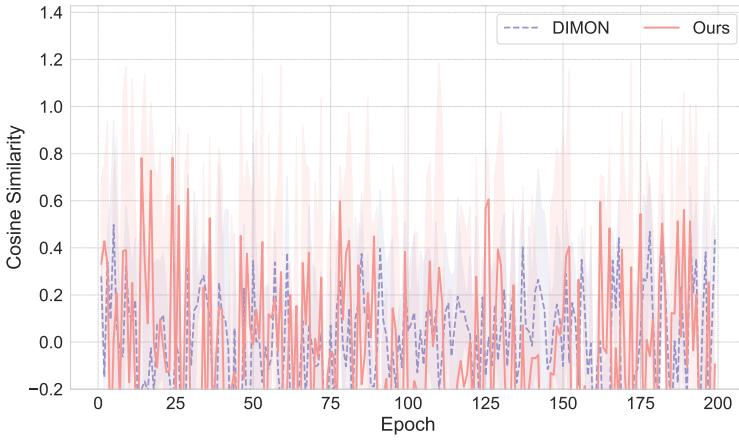
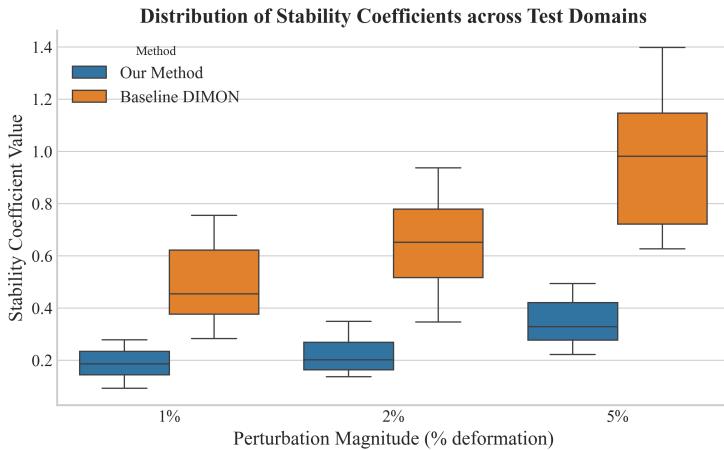
Figure 10: Evolution of the loss landscape curvature proxy ($v^T Hv$) during training. Our method operates in a significantly sharper region of the loss landscape compared to the nearly flat landscape of the baseline.

C.4 CONVERGENCE STABILITY AND PARAMETER UPDATES

A critical question is whether these larger gradients and the sharper loss landscape lead to training instability. The evidence suggests they do not. As seen in the top panel of Figure 11, the training and validation loss curves for GeoCMON, though higher in absolute value due to the weighting, show a stable and consistent decrease.

More revealing is the bottom panel, which plots the parameter update norms. Despite having drastically larger gradients, the actual step sizes taken by GeoCMON are remarkably comparable to, and often smoother than, those of the baseline. This demonstrates the effectiveness of the Adam optimizer in adaptively scaling the updates. Furthermore, Figure 12 shows the cosine similarity of consecutive gradients. The directional stability of GeoCMON's gradients is at least as consistent as the baseline's, indicating that the stronger signals do not lead to chaotic oscillations.

In conclusion, this dynamic analysis reveals that the weighted loss function reshapes the optimization problem by providing stronger, more structured gradient signals within a sharper loss landscape. Crucially, this does not compromise training stability. The optimizer effectively harnesses these signals to navigate the landscape, resulting in a robust and well-behaved convergence process.



D PROGRESSIVE DOMAIN GEOMETRY CURRICULUM EXPERIMENT

To investigate the learning efficiency and adaptability of our proposed method, we designed a curriculum learning experiment. This setup challenges the models by presenting data in stages of progressively increasing geometric complexity, moving from simple to more complex domain deformations. This approach is designed to reveal deeper insights into the models’ learning dynamics and generalization capabilities compared to standard training on a randomly shuffled dataset.

D.1 EXPERIMENTAL PROTOCOL

Curriculum Design: The entire dataset was first sorted based on a metric of geometric complexity, which was approximated by the variance of mesh displacements at the domain boundary. This sorted dataset was then partitioned into five sequential stages, where Stage 1 contained the simplest geometries and Stage 5 contained the most complex ones.

Training Procedure: Both the proposed GeoCMON model and the baseline DIMON model were trained sequentially through this curriculum. Each model was trained on the data of Stage 1 for a fixed number of epochs, after which its learned weights were carried over to be further trained on the data of Stage 2, and so on, up to Stage 5.

972
973**Evaluation Metrics:**

974

- **Per-stage Validation Error:** After training on each stage, the model’s performance was evaluated on a fixed, held-out validation set that spanned the full range of complexities.
- **Learning Efficiency:** Measured as the number of epochs required within each stage to reach a predefined performance target (relative L2 error ≤ 0.2) on the validation set. A lower number indicates faster convergence.
- **Final Generalization:** After completing the full curriculum, the final performance of each model was assessed on a separate, held-out test set.

978

D.2 RESULTS AND DISCUSSION

979

The results of the curriculum experiment highlight a distinct difference in the learning characteristics of the two models.

980

Adaptation and Per-Stage Performance: As shown in Figure 13, the proposed method exhibits a challenging initial adaptation phase. In Stage 1, which contains the simplest geometries, our model shows a significantly higher validation L2 error compared to the baseline. However, as the curriculum progresses to more complex domains in Stages 2 and 3, our model demonstrates a superior ability to adapt, achieving statistically significant improvements in accuracy and outperforming the baseline. In the final, most complex stages, the performance of both models becomes comparable.

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

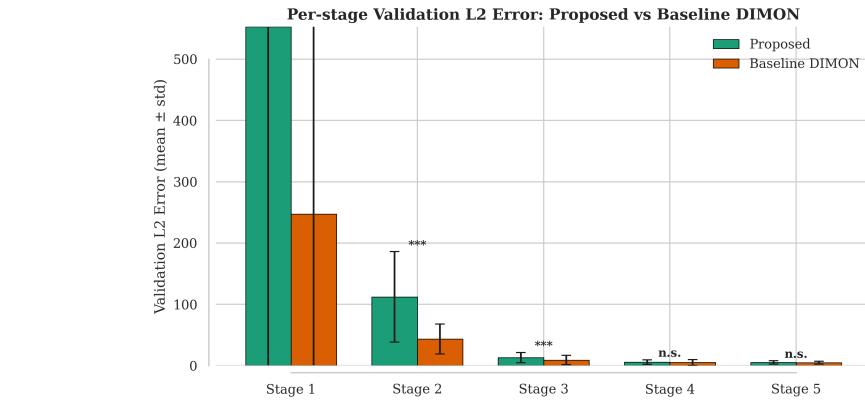


Figure 13: Per-stage validation L2 error (mean \pm std) for the proposed method and the baseline. Statistical significance ($p \leq 0.001$) is denoted by ***. Our model shows a difficult initial adaptation in Stage 1 but significantly outperforms the baseline in the more complex subsequent stages.

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

Learning Efficiency: The convergence speed within each stage, measured by epochs-to-threshold, reveals a complementary story (Figure 14). The baseline model converges faster in the initial, simpler stages. However, our proposed model, despite its slow start in Stage 1, demonstrates highly efficient learning in the subsequent, more challenging stages. This suggests that the architectural features of GeoCMON, particularly the weighted loss, are better suited for learning the complex patterns present in more deformed domains.

1017

1018

1019

1020

1021

1022

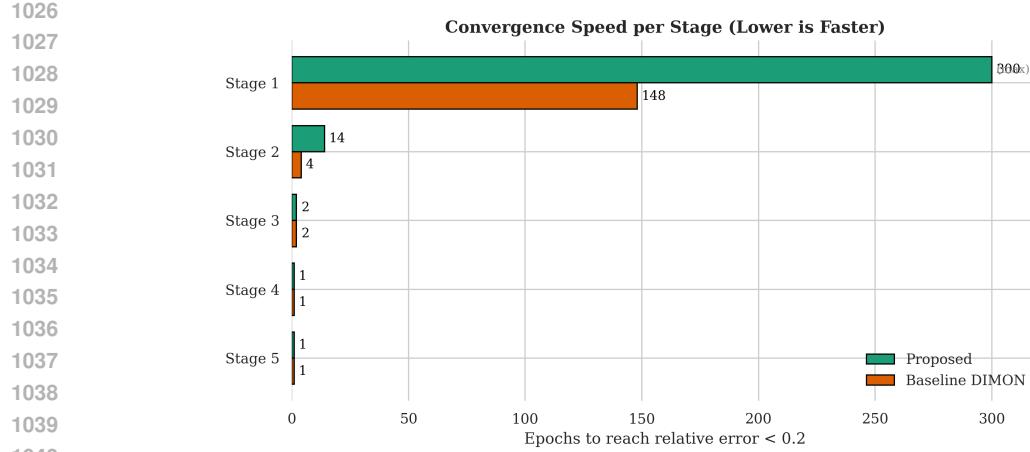
1023

1024

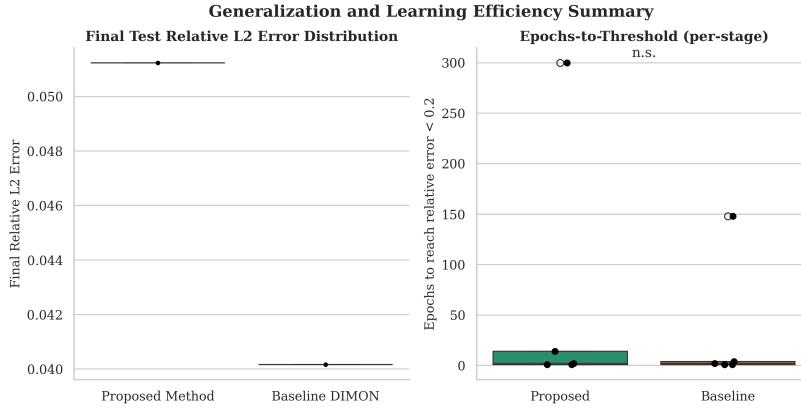
1025

Final Generalization and Summary: The key outcome of the experiment is the final generalization performance after the entire curriculum is completed. As summarized in Figure 15, despite the initial difficulties, the proposed method achieves a better final test relative L2 error. The left panel shows a clear separation in the error distributions, favoring our method. The right panel, summarizing the learning efficiency across all stages, shows no statistically significant difference on average, which aligns with the observation that each model excels at different complexity levels.

This experiment demonstrates that while the baseline model may be more adept at learning from simple data, the proposed GeoCMON architecture possesses a superior capacity to learn from and generalize to complex problems. The curriculum learning process, though initially challenging for our model, ultimately leverages its strengths to achieve a better overall final performance.



1041 Figure 14: Convergence speed per stage, measured in epochs required to reach a relative error
1042 threshold of 0.2. A lower value is faster. The baseline is faster on simple domains, while our method
1043 adapts more efficiently to increasing complexity.
1044



1059 Figure 15: Left: Final test relative L2 error distribution after completing the full training curriculum.
1060 Right: Boxplot summary of the epochs-to-threshold metric across all stages. Our proposed method
1061 achieves superior final generalization.
1062
1063
1064 E USE OF LLMs
1065
1066 We utilized a Large Language Model (LLM) to assist with both experiment and manuscript refine-
1067 ment.
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079