

# Boundary-aware Neural Model Reduction for PDEs

LI LIAO\*, The University of Hong Kong, Hong Kong  
PENGFEI SHEN\*, The University of Hong Kong, Hong Kong  
YIFAN PENG†, The University of Hong Kong, Hong Kong




Fig. 1. **Unified Neural Model Reduction over Shape–Boundary Space.** **Left:** We demonstrate reduced-order elastodynamic simulation over a continuous joint parameter space. A single neural model captures variations in both boundary conditions and object geometry. From *left to right*, the Dirichlet (fixed) boundary region expands; from *near to far*, the object’s aspect ratio increases. All deformations are shown at the same specific simulation frame. **Right:** Reduced-order heat transfer analysis of a heatsink over a continuous joint parameter space. The base plane spans geometric shape parameters controlling fin layout, while the vertical axis corresponds to the Robin coefficient  $\sigma$ , modeling convective heat transfer with the environment. Eigenanalysis is performed on a periodic unit cell, with Dirichlet, Neumann, and Robin boundaries prescribed as indicated in the top-right inset.

Eigenanalysis of partial differential operators is essential in reduced-order modeling for physical simulation, providing eigenmode representations in elasticity, acoustics, and transient heat transport. While recent neural, mesh-free, and geometry-agnostic approaches enable differentiable eigenanalysis over continuously parameterized shape spaces, boundary conditions of eigenfunctions are limited to Neumann-type (natural) conditions, hindering their applicability in scenarios where boundary conditions must be precisely controlled or optimized. In this work, we focus on Laplace-type operators and extend shape space neural eigenanalysis to handle boundary conditions beyond natural Neumann settings. Building on the same variational, energy-based formulation, we show that Dirichlet, Robin, and mixed boundary conditions can be incorporated without altering the underlying eigenvalue optimization principle. In our formulation, boundary configurations—including boundary placement and Robin coefficients modeling boundary exchange processes—are treated as first-class parameters rather than fixed constraints. When combined with shape-parameterized domains, this leads to a *joint shape–boundary space* formulation, allowing eigenfunctions and spectra to be evaluated consistently across variations in both geometry and boundary configuration. We conduct experiments on

representative applications such as boundary-driven spectral optimization for rigid-walled cavity resonance tuning, reduced-order simulation with changing supports, and analysis of transient thermal behavior under varying boundary exchange conditions. By elevating boundary conditions from fixed constraints to operator-wise parameterization, our approach broadens the applicability of Laplace-type neural eigenanalysis to physical systems where boundary constraints serve as critical design and control variables.

CCS Concepts: • **Computing methodologies** → **Physical simulation.**

Additional Key Words and Phrases: Reduced-order modeling, Implicit neural representation, Boundary conditions, Eigenanalysis, Physical simulation

## ACM Reference Format:

Li Liao, Pengfei Shen, and Yifan Peng. 2026. Boundary-aware Neural Model Reduction for PDEs. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26)*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3799902.3811153>

## 1 Introduction

Reduced-order modeling (ROM) [Quarteroni and Rozza 2014] in physical simulation commonly relies on the spectral structure of partial differential operators, with eigenmodes serving as a basis in elasticity, acoustics, and heat transport [Hartwig et al. 2023; Ma et al. 2021; Molesky et al. 2018]. In this work, we focus on the boundary-aware neural eigenanalysis of the Laplace operator, while using these spectra as reduced bases or design signals in representative simulation settings. By projecting high-dimensional states onto a small number of eigenfunctions, ROM enables efficient simulation and design optimization at substantially reduced computational cost [Hesthaven et al. 2016; Levin et al. 2025].

\*Both authors contributed equally to this research.

† Corresponding author.

Authors’ Contact Information: Li Liao, The University of Hong Kong, Hong Kong, Hong Kong, [liliao24@connect.hku.hk](mailto:liliao24@connect.hku.hk); Pengfei Shen, The University of Hong Kong, Hong Kong, Hong Kong, [pengfei.shen@connect.hku.hk](mailto:pengfei.shen@connect.hku.hk); Yifan Peng, The University of Hong Kong, Hong Kong, Hong Kong, [evan.y.peng@gmail.com](mailto:evan.y.peng@gmail.com).



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGGRAPH Conference Papers '26, Los Angeles, CA, USA*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2554-8/26/07  
<https://doi.org/10.1145/3799902.3811153>

Traditional eigenanalysis pipelines are tightly coupled to mesh-based discretizations and fixed geometries. Any geometry change requires remeshing, operator reassembly, and eigenmode recomputation [Azizzadenesheli et al. 2024], hindering practical use of eigenfunction-based ROM in interactive and optimization-driven workflows. Physics-informed neural networks (PINNs) [Toscano et al. 2025] mitigate this via mesh-free, geometry-agnostic formulations that represent modes as continuous neural fields, optimized from variational principles without a fixed discretization [Benckroun et al. 2023; Modi et al. 2024]. Shape space eigenanalysis extends this to continuously parameterized shape families, enabling differentiable evaluation of eigenfunctions and spectra across geometry variations within a single model [Chang et al. 2025a; Liu et al. 2025].

Existing neural shape space eigenanalysis methods offer limited control over boundary conditions. In their Laplace eigenanalysis instantiation, current formulations derive eigenfunctions by minimizing Dirichlet energy via a Rayleigh–Ritz principle, enforcing only orthogonality and normalization. Consequently, the learned eigenfunctions satisfy natural (Neumann) conditions, with no mechanism to impose Dirichlet, Robin, or mixed conditions at the operator level. While adequate for certain applications, this restriction hinders the use of shape space eigenanalysis when boundary behavior must be precisely controlled.

In many physical systems, boundary conditions play a role that is qualitatively distinct from geometry variation. Supports, clamps, thermal exchange boundaries, and acoustic impedance interfaces modify the governing operator and can induce substantial spectral changes. Classical ROM pipelines handle such effects during simulation via constrained degrees of freedom or penalty-based formulations [Li et al. 2020; Modi et al. 2024]. Although Neumann eigenbases can still be reused, the resulting reduced models rely on the basis to approximate boundary conditions rather than encoding them directly. This separation becomes problematic in low-dimensional ROMs: with only a few modes, Neumann eigenfunctions lack sufficient expressiveness needed to capture strongly boundary-localized behavior, requiring higher-frequency modes that are expensive to compute due to accumulated errors caused by orthogonal projections and high-order mode mixing in shape space. Consequently, Neumann-only shape-space eigenbases are ill-suited for applications demanding accurate boundary control under tight mode budgets.

These observations motivate incorporating boundary conditions directly into the eigenanalysis. Variationally, Dirichlet and Robin conditions can be introduced by modifying the admissible function space or augmenting the energy with boundary terms, without altering the underlying minimization principle. Embedding such boundary-aware formulations into neural shape space eigenanalysis yields eigenfunctions that intrinsically satisfy boundary conditions, providing more efficient and expressive reduced bases. Moreover, boundary conditions are often key design variables: support placement (Dirichlet BC) and boundary exchange coefficients (Robin BC) are common optimization targets in acoustic, mechanical, and thermal design. A joint parameterization of geometry and boundary configuration enables boundary-driven analysis and optimization to be posed directly in operator space within a unified, differentiable framework.

In this work, we extend neural shape space eigenanalysis to support Dirichlet, Robin, and mixed boundary conditions within the same optimization formulation used for Neumann eigenanalysis. We focus on Laplace-type eigenproblems, where Dirichlet energy, admissible-space constraints, and Robin boundary energies form a consistent variational framework, which also extends to other operators by changing the target energy terms while keeping the same essential boundary treatment. By introducing a boundary-parameterized representation, we treat boundary configuration as a first-class input alongside shape parameters. This yields a joint shape–boundary space in which eigenfunctions and spectra vary smoothly across both geometry and boundary conditions, while preserving the mesh-free and geometry-agnostic advantages of prior methods.

In summary, our contributions include:

- A unified neural eigenanalysis framework for the Laplace operator that incorporates Dirichlet, Robin, and mixed boundary conditions.
- A boundary-parameterized formulation that enables joint shape–boundary space eigenanalysis with differentiable dependence on boundary configurations.
- A boundary-aware reduced eigenbasis, used as scalar skinning weights, that improves accuracy in low-dimensional ROM settings under changing supports.
- Application demonstrations on boundary-driven air cavity resonance tuning, reduced-order simulation with varying shapes and supports, and transient thermal analysis with controllable boundary exchange coefficients and shape parameters.

## 2 Related Work

### 2.1 Reduced-Order Modeling for PDE-based Simulation

Reduced-Order Modeling accelerates PDE-based simulation by restricting high-dimensional fields to a low-dimensional subspace, enabling fast time stepping, inverse design, and interactive editing [Hesthaven et al. 2016; Li et al. 2023; Zong et al. 2023].

However, classical mesh-specific ROM pipelines rely on a fixed discretization, and the learned/constructed basis can be tightly coupled to given geometry, material layout, and boundary conditions [Kovachki et al. 2023], hindering their reuse when shapes or constraints change.

Relevant works on geometry- and discretization-agnostic ROM formulations with neural representations [Chang et al. 2023; Edavamadathil Sivaram et al. 2024; Yang et al. 2022] avoid explicit meshing by operating directly on continuous domains or implicit shape representations, enabling reduced simulation without constructing a fixed grid or mesh. While alleviating meshing difficulties and improving robustness for complex geometries, they remain restricted to single-shape settings.

Most recently, to extend eigenanalysis and reduced bases to shape families, Chang et al. [2025a] introduce a neural representation of eigenfunctions and eigenvalues over continuous shape space, enabling differentiable eigenanalysis across geometry variations. This line of work resolves challenges such as eigenvalue crossings and mode consistency by causal sorting, but is restricted to Neumann conditions and leaves boundary variation outside the scope of shape

parameterization. Meanwhile, neural basis representations have been studied in the context of heterogeneous materials and operators with sharp transitions. Recent lifting-based approaches [Chang et al. 2025b; Liu et al. 2025] explicitly encode gradient discontinuities in neural fields, allowing accurate representation of piecewise-smooth solutions and heterogeneous coefficients.

## 2.2 Boundary Conditions in Physics Simulation

Boundary conditions are essential in PDE-based physical models, with Dirichlet, Neumann, and Robin conditions arising naturally across various physical phenomena [Evans 2022]. Modifying these conditions alters the underlying operator, often inducing structural changes in spectrum and system behavior. Classical elastodynamic simulation enforces boundary conditions at the discretization level via constrained degrees of freedom, penalty forces, or projection operators applied to mesh-based systems [Bargteil et al. 2007; Bouaziz et al. 2014; Cutler et al. 2002]. While effective for fixed setups, these methods are tightly coupled to mesh topology and discretization and are difficult to extend to scenarios where boundary locations or types vary across instances [Hu et al. 2018; Schneider et al. 2018].

Within ROM frameworks, reduced bases are typically constructed without explicit awareness of boundary constraints [Chang et al. 2025a, 2023; Modi et al. 2024], and fixed boundary conditions are enforced in the reduced space via truncation, masking, or barrier-based constraint methods [Li et al. 2020]. Such strategies generally destroy basis optimality with respect to the constrained operator and can degrade accuracy when limited modes are retained.

## 2.3 Boundary Condition Enforcement in Neural Fields

*Dirichlet BC Enforcement.* Accurately enforcing Dirichlet (essential) boundary conditions in PINNs [Cuomo et al. 2022; Raissi et al. 2019] is challenging, as such constraints must hold point-wise while the solution is learned through optimization. A common *soft* strategy augments the loss with boundary penalties alongside interior PDE residuals [Aygün et al. 2023], but this trades boundary accuracy off interior convergence and requires careful weighting tuning to avoid slow convergence or boundary artifacts. Alternative *hard* strategies compose the network output with geometry-aware functions that vanish on the prescribed boundary, removing boundary residuals from the objective. Distance function-based formulations [Sukumar and Srivastava 2022] and related hard-constraint formulations [Wang et al. 2023] demonstrate that exactly enforcing essential boundary conditions could improve training stability. A systematic comparison across PINNs and variational PINNs [Berrone et al. 2023] further indicates that hard constraints generally outperform penalty methods in both robustness and accuracy.

This work is partially inspired by distance-function-based hard-constraint approaches. In our setting, however, Dirichlet boundaries evolve with geometry variation, making globally smooth distance-like functions costly to construct. We therefore enforce Dirichlet conditions using raw distance functions instead. Another line of work enforces Dirichlet conditions via architectural constraints embedded into network structures [Kumar and Yadav 2025; Zhou et al. 2021]. Instead of modifying outputs with distance functions, these approaches embed boundary information into basis functions

or activations, avoiding explicit distance computation and improving accuracy over standard MLPs [Kim et al. 2025]. However, such architectures typically rely on stable, readily expressed boundary parameterizations in the input space, limiting their applicability to problems with discretized or dynamically evolving mesh-based boundaries.

*Robin BC Enforcement.* Robin conditions in PINNs are typically imposed by adding boundary residuals to the loss alongside interior PDE terms [Raissi et al. 2019]. This soft enforcement couples multiple competing objectives in a single optimization problem, making performance sensitive to loss weighting. Before modern PINNs, neural-network-based solvers embedded boundary conditions into the trial function through analytical constructions or distance-based length factors, enabling exact satisfaction of Dirichlet and mixed conditions [Berg and Nyström 2018; Lagaris et al. 2000; McFall and Mahan 2009]. Although this removes boundary residual terms, it requires specialized trial function designs that become computationally expensive when shapes or boundary configurations vary.

A common alternative interprets Robin boundary conditions as natural boundary terms in a variational formulation. The Deep Ritz method minimizes the energy functional directly [E and Yu 2018], incorporating Neumann and Robin conditions without explicit boundary residual objectives. Related approaches employ penalty or Nitsche-type formulations to weakly impose essential boundary conditions in a variationally consistent manner [Liao and Ming 2021]. In this work, we adopt this variational viewpoint, where Robin boundary conditions arise naturally as part of the energy functional in eigenvalue optimization. This ties the training objective directly to the eigenvalue, enables efficient evaluation of boundary-dependent energy in shape-space settings, and avoids the loss-balancing issues inherent to strong-form PINNs.

## 3 Method

### 3.1 Variational Eigenanalysis with Boundary Parameters

Following prior work, spectral modes are derived via constrained variational minimization. For a scalar field  $u$  defined on  $\Omega(s)$ , we consider the Dirichlet energy

$$E_D[u] = \frac{1}{2} \int_{\Omega(s)} \|\nabla u(x)\|^2 dV, \quad (1)$$

subject to unit-norm and orthogonality constraints. In the absence of explicit boundary constraints, stationary points of  $E_D$  correspond to eigenfunctions under homogeneous Neumann boundary conditions.

The boundary parameter  $\beta$  induces a boundary decomposition as

$$\partial\Omega(s) = \partial\Omega_D(s, \beta) \dot{\cup} \partial\Omega_N(s, \beta) \dot{\cup} \partial\Omega_R(s, \beta), \quad (2)$$

where  $\partial\Omega_D$ ,  $\partial\Omega_N$ , and  $\partial\Omega_R$  denote Dirichlet, Neumann, and Robin regions, respectively. Each choice of  $(s, \beta)$  defines a distinct admissible function space and bilinear form, and thus a corresponding operator  $\mathcal{L}(s, \beta)$ .

As illustrated in Fig. 2, this variational formulation provides a unified mechanism for incorporating boundary conditions directly into the eigenanalysis pipeline.

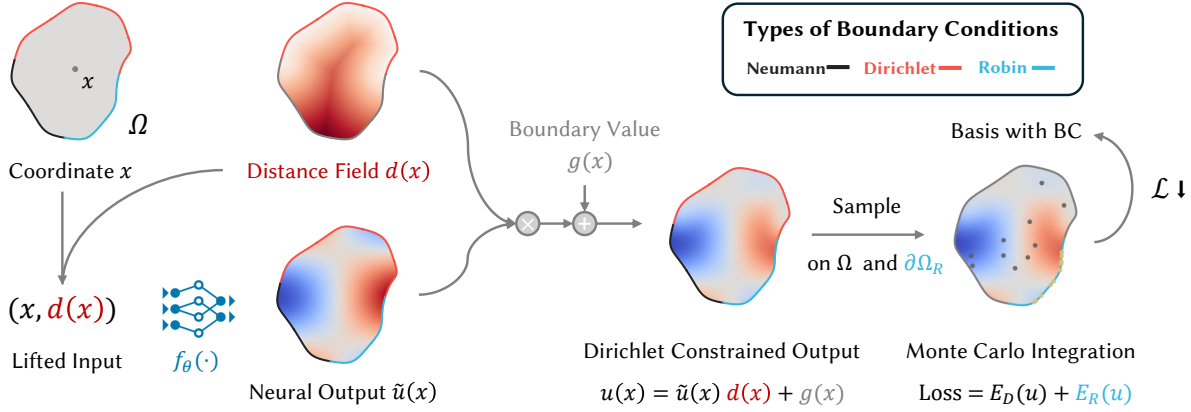


Fig. 2. **Pipeline of our extended boundary-aware neural modeling.** Given a domain  $\Omega$ , we augment spatial coordinates with a distance field  $d(x)$  to the Dirichlet boundary and evaluate a neural field  $\tilde{u}(x)$  on the lifted input  $(x, d(x))$ . The resulting constrained field  $u(x)$  satisfies Dirichlet conditions by construction, while Neumann and Robin conditions are handled through the variational loss. Energies are estimated by Monte Carlo integration over interior and boundary samples, producing boundary-aware basis functions.

### 3.2 Dirichlet Conditions as Essential Constraints

Dirichlet boundary conditions impose *essential constraints* that restrict the admissible function space. Rather than enforcing them via penalty terms, we embed Dirichlet conditions directly into the parameterization of the solution. Specifically, we represent the physical field  $u$  as

$$u(x) = \tilde{u}(x)d(x) + g(x), \quad (3)$$

where  $\tilde{u}$  denotes an unconstrained neural field,  $d(x)$  is a distance field function, satisfying  $d(x) = 0$  on  $\partial\Omega_D(s, \beta)$  and  $d(x) > 0$  in the interior, and  $g(x)$  is the prescribed Dirichlet boundary field. Both  $d$  and  $g$  are determined by  $(s, \beta)$ .

Since the variational energy depends on spatial gradients, the correct gradient of  $u$  is obtained via the chain rule

$$\nabla u = d \nabla \tilde{u} + \tilde{u} \nabla d + \nabla g. \quad (4)$$

This formulation precisely enforces Dirichlet constraints while preserving correct differential structure in the interior energy.

*Continuity Mismatch and Lifted Inputs.* For homogeneous elliptic operators with smooth coefficients, eigenfunctions are smooth in the interior of the domain. In the homogeneous setting, although the distance field  $d(x)$  is continuous everywhere, its gradient  $\nabla d(x)$  is discontinuous due to changes in closest-point projections along the Dirichlet boundary. By contrast, neural fields produce outputs with continuous spatial gradients. As a result, the product term  $\tilde{u}(x) \nabla d(x)$  in Eq. (4) inherits the non-smoothness of  $\nabla d(x)$ , leading to a mismatch with the expected interior gradient continuity of eigenfunctions in homogeneous domains.

To resolve this issue, we follow recent work on precise discontinuity modeling in neural fields [Chang et al. 2025b; Liu et al. 2025]. Instead of implicitly embedding discontinuities in network weights, we explicitly expose the source of non-smoothness to the network by augmenting its input with the distance field, expressed as

$$\tilde{u}(x) = f_\theta(x, d(x)). \quad (5)$$

This lifted representation enables the network to learn compensating behavior aligned with gradient discontinuities of  $d(x)$ , yielding eigenfunctions whose composite gradients remain consistent with the expected physical regularity.

### 3.3 Robin Conditions via Boundary Energy Operators

We adopt the standard variational formulation of Robin boundary conditions, in which the boundary operator is induced by an explicit surface energy term [E and Yu 2018].

*Mixed-boundary variational formulation.* Let the boundary be decomposed into disjoint subsets  $\partial\Omega = \partial\Omega_D \cup \partial\Omega_N \cup \partial\Omega_R$ . For a given shape parameter  $s$  and boundary parameter  $\beta$ , the Robin subset  $\partial\Omega_R$  is equipped with a boundary operator defined through the surface energy

$$E_R[u] = \int_{\partial\Omega_R} \left( \frac{1}{2} \sigma(x) u(x)^2 - \psi(x) u(x) \right) dS, \quad (6)$$

where  $\sigma(x)$  and  $\psi(x)$  are boundary fields parameterized by  $(s, \beta)$ .

The full variational objective is then defined as

$$E[u] = \frac{1}{2} \int_{\Omega} \|\nabla u\|^2 dV + E_R[u], \quad \text{subject to } \int_{\Omega} u^2 dV = 1, \quad (7)$$

together with orthogonality constraints for higher modes.

*Variational derivation of the Robin boundary operator.* Taking the first variation of the surface energy yields

$$\delta E_R = \int_{\partial\Omega_R} (\sigma(x) u(x) - \psi(x)) \delta u(x) dS. \quad (8)$$

Meanwhile, the interior Dirichlet energy produces the standard boundary flux term  $\int_{\partial\Omega} \partial_n u \delta u dS$  after integration by parts. For  $\partial\Omega_D$  and  $\partial\Omega_N$ , this term vanishes due to the constraints. While for  $\partial\Omega_R$ , the total boundary term in the first variation collects into

$$\int_{\partial\Omega_R} \left( \partial_n u(x) + \sigma(x) u(x) - \psi(x) \right) \delta u(x) dS. \quad (9)$$

Since  $\delta u$  is free on  $\partial\Omega_R$ , stationarity enforces the following:

$$\partial_n u(x) + \sigma(x) u(x) = \psi(x), \quad x \in \partial\Omega_R. \quad (10)$$

The Robin boundary condition arises as the Euler–Lagrange boundary operator induced by the surface energy in Eq. (6), rather than an externally imposed constraint. Including this Robin surface energy does not alter the fundamental Rayleigh–Ritz interpretation of the eigenproblem: eigenvalues remain interpretable as minimized energy levels, now augmented by the boundary energy term.

*Monte Carlo integration of volume and boundary energies.* We approximate the variational objective using Monte Carlo integration. When the loss consists of a single volume integral  $\int_{\Omega} f(x) dV$ , uniform sampling induces a constant density, and the estimator differs from the continuous energy only by a global factor  $|\Omega|$ , which does not affect the minimizer and therefore requires no explicit normalization. When both volume and boundary terms are present,  $\int_{\Omega} (\cdot) dV + \int_{\partial\Omega_R} (\cdot) dS$ , the two contributions are taken with respect to different measures. Their relative scaling is thus determined by the volume  $|\Omega|$  and surface area  $|\partial\Omega_R|$ , rather than by a tunable weighting coefficient. This preserves the variational structure of the continuous problem and avoids penalty-style hyperparameters.

*Why Robin does not degenerate to Dirichlet.* For any finite  $\sigma(x)$ , the Robin condition couples the boundary value  $u$  and its normal derivative  $\partial_n u$  without fixing  $u$  pointwise on  $\partial\Omega_R$ . As a result, the admissible function space and the associated operator differ fundamentally from those induced by Dirichlet conditions. Dirichlet behavior is recovered only in a singular limit (e.g.,  $\sigma \rightarrow \infty$ ), which alters both the conditioning and the operator family itself.

*Extension to other operators.* The same boundary-aware variational framework extends to other eigenproblems. Representative operators and energies are listed in supplementary Table S2. Dirichlet conditions continue to be imposed through the admissible space by the same distance-field parameterization and lifted-input construction since this mechanism depends on the essential boundary constraint. The operator-specific volume energy determines both the interior differential operator and its associated natural boundary flux, while Robin conditions are introduced by adding the corresponding boundary energy, as summarized in Table S2.

## 4 Experiments

*Implementation.* All experiments are trained and evaluated on an NVIDIA RTX 4090 GPU. We implement our method in PyTorch and optimize all network parameters using Adam. Our neural field backbone follows the SIREN implicit representation with sinusoidal activations and the standard initialization [Sitzmann et al. 2020]. Unless otherwise specified, we use an MLP with 12 hidden sine layers of width 64, followed by a final linear layer that produces the scalar output. For the Airplane cavity example (Fig. 5), we use a smaller MLP with 4 hidden sine layers of width 128.

During training, shape and boundary parameters are sampled uniformly at random from their prescribed ranges. For each sampled configuration, we estimate variational energies by Monte Carlo integration with uniformly sampled interior points and, when Robin conditions are present, separately sampled boundary points on the

Robin boundary. We use 65,536 samples for both interior and boundary integral across all experiments. We report training time statistics for representative examples in supplementary Table S1.

### 4.1 Analytic Agreement for Robin Conditions

We evaluate the agreement with analytic solutions using the Laplace eigenvalue problem on a rectangular domain  $\Omega = [0, a] \times [0, b]$  with Robin boundary conditions, defined as

$$\partial_n u + \sigma u = 0 \quad \text{on } \partial\Omega, \quad (11)$$

where  $\sigma$  is a boundary parameter controlling the boundary impedance.

We train a single boundary-parameterized model that represents a continuous family of operators indexed by  $\sigma$ . All experimental results are derived from *one reduced model*, with varying Robin coefficients  $\sigma$  supplied exclusively at inference time. We sweep  $\sigma$  over a fixed range while keeping the geometry and network parameters unchanged, and compare the predicted eigenfunctions against their analytic solution counterparts. Fig. 3 visualizes this comparison: columns correspond to increasing values of  $\sigma$ , with analytic and predicted eigenfunctions shown in the top and bottom, respectively.

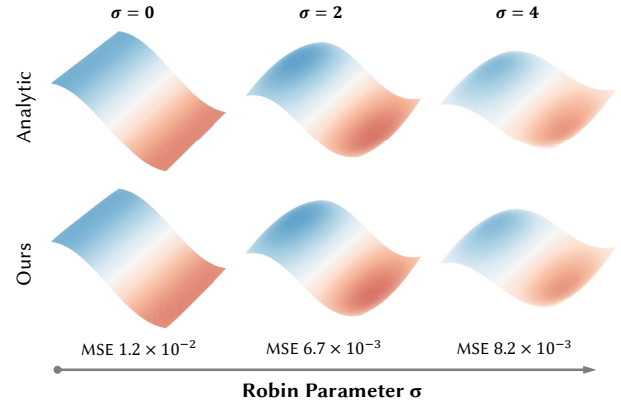


Fig. 3. **Agreement with analytic solutions on a rectangular domain under Robin boundary conditions.** The Robin coefficient  $\sigma$  varies from left to right. **Top:** analytic eigenfunctions. **Bottom:** predicted eigenfunctions, with mean squared error (MSE) relative to analytic solutions reported below.

### 4.2 Ablation: Lifted input

As discussed in Sec. 3.2, the observed continuity mismatch in our formulation is attributed not to the neural representation itself, but to the geometric term  $\nabla d(x)$ .

These discontinuities persist at the spatial gradient level, and therefore cannot be eliminated solely through any nonlinear mapping of the distance value.

To explicitly illustrate this behavior, Fig. 4 visualizes a representative 2D distance field together with its spatial gradient. Although  $d(x)$  appears smooth (continuous),  $\nabla d(x)$  exhibits distinct discontinuities that align with closest-point switching curves along the Dirichlet boundary. This observation motivates exposing the distance field explicitly as a network input, rather than forcing the neural field to implicitly absorb the induced non-smoothness.

We conduct a qualitative ablation study on a fixed planar domain with Dirichlet boundary. All models are trained using identical

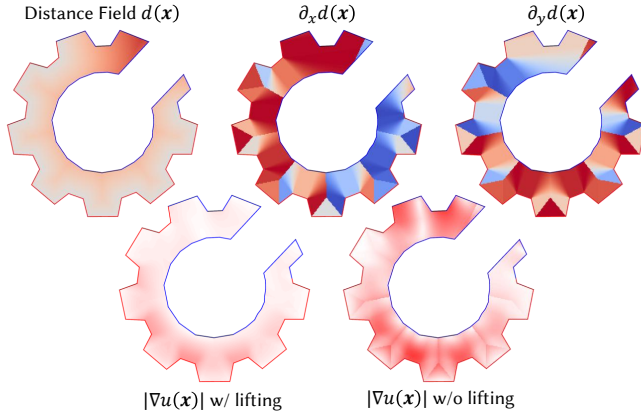


Fig. 4. **Distance-gradient discontinuity.** Top: a 2D distance field  $d(x)$  and its spatial gradient components  $\partial_x d(x)$ ,  $\partial_y d(x)$ . The red outline indicates the Dirichlet boundary. While  $d(x)$  is continuous, its gradient  $\nabla d(x)$  exhibits sharp discontinuities along closest-point switching curves. Bottom: gradient magnitude  $|\nabla u(x)|$  with and without distance-field lifting.

network architectures and optimization settings, differing only in the input parameterization. Specifically, we compare the resulting solution field  $u$  obtained from: (i) the reference solution; (ii) a formulation without lifted input; and (iii) our lifted formulation, where the distance to the Dirichlet boundary is provided as an additional network input.

This demonstrates that explicitly exposing boundary-induced geometric non-smoothness to the network effectively resolves a fundamental incompatibility. The network produces outputs with smooth spatial gradients, whereas the distance field  $d(x)$  exhibits discontinuities in  $\nabla d(x)$  at closest-point switching locations. When these two quantities are multiplied at the same spatial locations, the resulting field inherits the non-smoothness of  $d(x)$ , contaminating what should otherwise be a smooth interior representation. By lifting  $d(x)$  into the network input, the model can account for this structure directly rather than implicitly compensating for it through distorted interior predictions.

### 4.3 Boundary-Space Air Cavity Resonance Tuning

We evaluate our boundary-aware eigenanalysis on a fixed acoustic cavity of arbitrary geometry. Under the standard assumptions of small-amplitude acoustics in homogeneous air, the time-harmonic acoustic pressure satisfies a Helmholtz problem, so the corresponding eigenvalues determine cavity resonance frequencies. This setting isolates the influence of boundary operators from geometric variation: the cavity geometry remains fixed, while the resonance spectrum is controlled solely through boundary condition variation. Specifically, we introduce a localized pressure-release opening that slides along the boundary, inducing a family of operators parameterized by its position in boundary space.

Notably, rigid cavity walls impose a natural (Neumann) condition on the acoustic pressure field. In contrast, an opening connected to free space is approximated as a pressure-release boundary, and is

therefore modeled as a localized Dirichlet boundary patch. We introduce a single scalar boundary parameter  $\beta \in [0, 1]$  that controls the position of this opening along  $\partial\Omega$ , while the remainder of the boundary remains Neumann. This construction induces a one-parameter family of operators  $\mathcal{L}(\beta)$  acting on an identical geometric domain. The resonant medium here is the enclosed air, while the surrounding solid is assumed rigid and affects the system only through the cavity shape and opening configuration. Crucially, Dirichlet constraints are essential as they alter the admissible function space of the operator and can strongly affect the low-frequency cavity spectrum, particularly under strict low-rank spectral representations.

For each boundary configuration, we compute the lowest  $K = 16$  eigenpairs of the resulting mixed-boundary Laplace operator, forming a boundary-parameterized spectrum  $\lambda(\beta) = (\lambda_1(\beta), \dots, \lambda_K(\beta))$ . Since the corresponding frequencies satisfy  $f_i(\beta) = \frac{c}{2\pi} \sqrt{\lambda_i(\beta)}$ , matching the eigenvalue spectrum is equivalent to matching the low-frequency cavity resonance spectrum. Thus, given a target spectrum  $\lambda^{\text{gt}} \in \mathbb{R}^K$ , we optimize the boundary parameter by minimizing the following:

$$\beta^* = \arg \min_{\beta} \|\lambda(\beta) - \lambda^{\text{gt}}\|_2^2. \quad (12)$$

*Results and Insights.* As the Dirichlet opening traverses the boundary, the low-frequency spectrum undergoes smooth yet pronounced deformation, despite the fixed geometry. Across all experiments, the first 16 eigenvalues vary continuously with respect to the boundary parameter  $\beta$ , thereby forming a low-dimensional spectral manifold. Optimizing within boundary space yields spectra that closely match the target through boundary manipulation alone. These results mirror behavior observed in shape-parameterized spectral spaces, demonstrating that boundary conditions can serve as an effective low-dimensional control axis for the resonance spectrum in this air cavity setting.

### 4.4 Reduced-Order Elastic Simulation

We evaluate the practical impact of our boundary-aware eigenbases in reduced-order elastic simulation. We adopt the subspace deformation model of Modi et al. [2024], where the displacement field  $\phi$  is constructed from a basis  $\{u_j\}_{j=1}^K$  and affine handles  $\mathbf{q}$ :

$$\phi(\mathbf{X}, \mathbf{q}) = \sum_{j=1}^K u_j(\mathbf{X}) \mathbf{q}_j \begin{bmatrix} \mathbf{X} \\ 1 \end{bmatrix}. \quad (13)$$

We use the Neo-Hookean energy model, implicit time integration, and standard Projected Newton method for the elastic simulation. We project all terms onto the reduced subspace and solve the densified system in the reduced space of handle weights  $\mathbf{q}$ .

Here, the simulation quality depends critically on the basis’s ability to satisfy boundary conditions. Most existing neural spectral methods [Chang et al. 2025a] produce Neumann eigenbases that do not inherently satisfy Dirichlet constraints. To enforce fixed boundaries with such bases, prior works often employ barrier functions during optimization [Li et al. 2020; Modi et al. 2024]. Following this, we adopt the log-barrier energy

$$E_{\text{barrier}} = -\mu \sum_{i \in \mathcal{B}} \log \left( 1 - \frac{\|\mathbf{u}_i - \mathbf{u}_{\text{bc},i}\|}{d_{\text{max}}} \right). \quad (14)$$

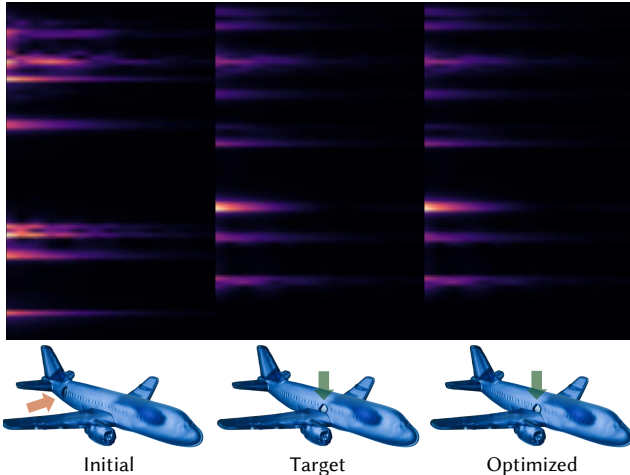


Fig. 5. **Boundary-space cavity resonance matching** on a fixed rigid-walled air cavity. We visualize the low-frequency spectrum (first  $K = 16$  eigenvalues) and corresponding boundary configurations for three cases: an initial boundary configuration (init.), the target spectrum (target), and the optimized result obtained by adjusting the boundary parameter  $\beta$ . The optimized configuration yields a spectrum that more closely matches the target than the initial guess (see arrows), even though the cavity geometry remains unchanged.




Fig. 6. **Reduced-order elastic simulation initialized via modal deformations.** We visualize the dynamic evolution of a soft object simulated using our boundary-aware eigenbasis. The simulations are initialized by applying a transformation projected onto specific eigenmodes: Mode 1, Mode 8, and all modes. The resulting motion, shown as time-lapsed overlays, demonstrates that the simulation evolves naturally from diverse initial conditions while strictly preserving the fixed boundary constraints (green).

while suffering from a fundamental conflict: when the reduced subspace is small, it is often kinematically impossible for a linear combination of unconstrained Neumann modes to satisfy the fixed boundary. Thus, the solver must “fight” the basis, leading to boundary drift, locking, or artificial stiffness. In contrast, our method generates bases that satisfy Dirichlet conditions *by construction*. This ensures the reduced simulation evolves strictly within the admissible kinematic subspace, eliminating the need for boundary barrier terms. As shown in Fig. 12, using the same reduced dimension ( $K = 12$ ), our method maintains exact boundary constraints and produces deformations that closely match the full-space ground truth, whereas the Neumann baseline exhibits visible violations.

*Simulation in Shape–Boundary Space.* A key advantage of our approach is the ability to parameterize dynamics over a continuous joint space of geometry and boundary conditions. We train a single reduced model over varying shape parameters and boundary codes. Fig. 13 shows that this model produces smooth, physically plausible deformation responses as both the object geometry and the boundary constraints (e.g., location of fixed regions) are varied interactively, without re-computation or mesh-dependent setups.

#### 4.5 Thermal Decay in Shape–Boundary Space

We study transient heat conduction in finned solids as a testbed for joint modeling of geometric variation and boundary conditions. The solid domain  $\Omega(\theta)$  represents a heat sink geometry parameterized by a two-dimensional shape code  $\theta$ , which controls the height and thickness of the fins. Heat transport within the solid follows the diffusion equation  $\partial_t T = \alpha \Delta T$ .

The base of the heat sink is assumed to be in perfect contact with a temperature-controlled plate and is therefore held at a fixed temperature, modeled as a Dirichlet boundary condition. Due to symmetry of the periodic fin structure, we also impose Neumann conditions on the cross-sectional planes of the base. All remaining external surfaces exchange heat with ambient air via convection, modeled by a Robin boundary condition  $\partial_n T + \sigma T = 0$ . Here,  $\sigma$  is proportional to the convective heat transfer coefficient, which corresponds to the strength of heat exchange with the environment. This setup defines a coupled shape–boundary parameter space, illustrated in the right part of Fig. 1.

To analyze transient behavior, we first compute the steady-state temperature field  $T_{ss}$  satisfying the full mixed boundary-value problem. We then introduce a homogenized variable  $v = T - T_{ss}$ , which obeys homogeneous Dirichlet conditions on the base and homogeneous Robin conditions on the remaining boundary. Importantly, the eigenbasis considered in our method is defined for this homogenized problem and governs only the *transient* component of the temperature field, rather than the steady state itself. Under this formulation,  $v$  admits a modal expansion in terms of eigenfunctions of the Laplacian with mixed boundary conditions, with eigenvalues  $\lambda_i$  determining the exponential decay rates of the corresponding thermal modes. Fig. 9 shows the temporal evolution of the temperature field for two representative parameter settings.

Treating both geometry  $\theta$  and boundary parameter  $\sigma$  as explicit inputs allows us to examine how transient thermal behavior varies across configurations. In particular, the leading eigenvalues provide a compact summary of global cooling characteristics, while higher modes capture more localized heat dissipation patterns.

Fig. 7 visualizes how the diffusion spectrum changes jointly with geometry and boundary conditions. For each fixed Robin coefficient  $\sigma$ , the leading eigenvalue varies smoothly as the shape parameter modifies fin proportions, while different  $\sigma$  curves exhibit consistent spectral shifts under identical geometry. Given that, we can optimize the shape and boundary parameters to achieve a desired spectral target, as illustrated in the right part of Fig. 7.

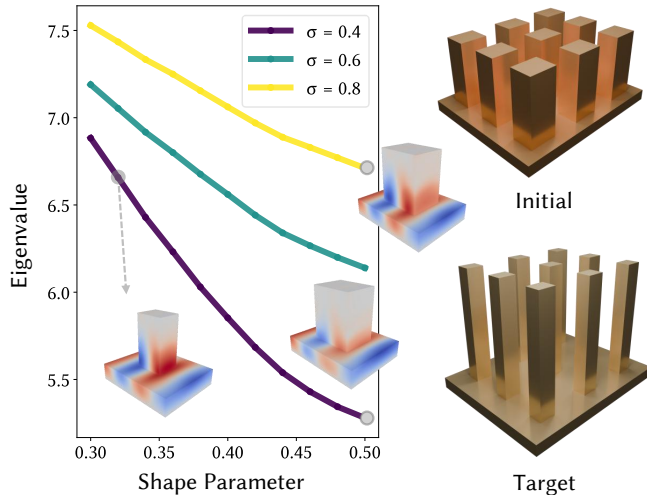


Fig. 7. **Spectral variation in shape-boundary space for a finned heat sink.** Left: leading Robin-Laplacian eigenvalue versus a geometric shape parameter, shown for multiple Robin coefficients  $\sigma$ . Insets visualize the corresponding dominant eigenmodes at selected configurations. Right: representative initial and target geometries.

## 5 Discussion

We have presented a boundary-aware neural model reduction framework that extends shape space eigenanalysis to mixed boundary conditions. Our complete instantiation focuses on mixed-boundary Laplace-type eigenproblems. By parameterizing boundary configurations together with geometry, the model learns eigenfunctions over a joint shape-boundary operator space, while the lifted input improves distance-based Dirichlet enforcement. Experiments show that the boundary-aware basis provides a compact, constraint-consistent ROM space under varying boundary configurations.

*Beyond the mixed-boundary Laplacian.* The boundary treatment itself is not tied to the scalar Laplacian. The main requirement is a variational eigenproblem with a well-defined volume energy and boundary terms. Under this view, extending the method mainly requires deriving the operator-specific volume energy and boundary flux. Dirichlet enforcement remains unchanged through the distance-field admissible-space parameterization, while Robin conditions are incorporated by adding the corresponding surface energy whose variation balances the boundary flux terms. Examples are listed in Table S2.

*Limitations and Future Work.* Our current implementation relies on accurate distance fields to enforce Dirichlet conditions. While efficient for parameterized shape families, computing high-quality differentiable distance fields for topologically complex or dynamically changing boundaries remains challenging and is orthogonal to this work. Furthermore, covering high-dimensional joint shape-boundary spaces requires more training data and network capacity than single-instance learning. Future work includes hybrid neural-explicit boundary representations for complex topologies and extensions to non-self-adjoint operators arising in fluid dynamics and dissipative systems.

## Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (62322217), the Innovation and Technology Fund of Hong Kong (MHP/313/24), and the Research Grants Council of Hong Kong (GRF 17208023).

## References

- Atakan Aygun, Romit Maulik, and Ali Karakus. 2023. Physics-informed neural networks for mesh deformation with exact boundary enforcement. *Engineering Applications of Artificial Intelligence* 125 (2023), 106660. doi:10.1016/j.engappai.2023.106660
- Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. 2024. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics* 6, 5 (May 2024), 320–328. doi:10.1038/s42254-024-00712-5
- Adam W. Bargteil, Chris Wojtan, Jessica K. Hodgins, and Greg Turk. 2007. A finite element method for animating large viscoplastic flow. In *ACM SIGGRAPH 2007 Papers* (San Diego, California) (SIGGRAPH '07). Association for Computing Machinery, New York, NY, USA, 16–es. doi:10.1145/1275808.1276397
- Otman Benckroun, Jiayi Eris Zhang, Siddhartha Chaudhuri, Eitan Grinspun, Yi Zhou, and Alec Jacobson. 2023. Fast Complementary Dynamics via Skinning Eigenmodes. *ACM Trans. Graph.* 42, 4, Article 106 (July 2023), 21 pages. doi:10.1145/3592404
- Jens Berg and Kaj Nyström. 2018. A unified deep artificial neural network approach to partial differential equations in complex geometries. *Neurocomputing* 317 (2018), 28–41. doi:10.1016/j.neucom.2018.06.056
- S. Berrone, C. Canuto, M. Pintore, and N. Sukumar. 2023. Enforcing Dirichlet boundary conditions in physics-informed neural networks and variational physics-informed neural networks. *Heliyon* 9, 8 (Aug. 2023), e18820. doi:10.1016/j.heliyon.2023.e18820
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective dynamics: fusing constraint projections for fast simulation. *ACM Trans. Graph.* 33, 4, Article 154 (July 2014), 11 pages. doi:10.1145/2601097.2601116
- Yue Chang, Otman Benckroun, Maurizio M. Chiaramonte, Peter Yichen Chen, and Eitan Grinspun. 2025a. Shape Space Spectra. *ACM Trans. Graph.* 44, 4, Article 121 (July 2025), 16 pages. doi:10.1145/3731148
- Yue Chang, Peter Yichen Chen, Zhecheng Wang, Maurizio M. Chiaramonte, Kevin Carlberg, and Eitan Grinspun. 2023. LICROM: Linear-Subspace Continuous Reduced Order Modeling with Neural Fields. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 111, 12 pages. doi:10.1145/3610548.3618158
- Yue Chang, Mengfei Liu, Zhecheng Wang, Peter Yichen Chen, and Eitan Grinspun. 2025b. Lifting the Winding Number: Precise Discontinuities in Neural Fields for Physics Simulation. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers* (SIGGRAPH Conference Papers '25). Association for Computing Machinery, New York, NY, USA, Article 25, 11 pages. doi:10.1145/3721238.3730597
- Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. 2022. Scientific Machine Learning Through Physics-Informed Neural Networks: Where we are and What's Next. *J. Sci. Comput.* 92, 3 (Sept. 2022), 62 pages. doi:10.1007/s10915-022-01939-z
- Barbara Cutler, Julie Dorsey, Leonard McMillan, Matthias Müller, and Robert Jagrow. 2002. A procedural approach to authoring solid models. *ACM Trans. Graph.* 21, 3 (July 2002), 302–311. doi:10.1145/566654.566581
- Weinan E and Bing Yu. 2018. The Deep Ritz Method: A Deep Learning-Based Numerical Algorithm for Solving Variational Problems. *Communications in Mathematics and Statistics* 6, 1 (March 2018), 1–12. doi:10.1007/s40304-018-0127-z
- Venkataram Edavamadathil Sivaram, Tzu-Mao Li, and Ravi Ramamoorthi. 2024. Neural Geometry Fields For Meshes. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 29, 11 pages. doi:10.1145/3641519.3657399
- Lawrence C. Evans. 2022. *Partial Differential Equations*. American Mathematical Society. Google-Books-ID: Ott1EAAAQBAJ.
- Florine Hartwig, Josua Sassen, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. 2023. An Elastic Basis for Spectral Shape Correspondence. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 58, 11 pages. doi:10.1145/3588432.3591518
- Jan S Hesthaven, Gianluigi Rozza, and Benjamin Stamm. 2016. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. Springer International Publishing. doi:10.1007/978-3-319-22470-1
- Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral meshing in the wild. *ACM Trans. Graph.* 37, 4, Article 60 (July 2018), 14 pages. doi:10.1145/3197517.3201353
- Bongseok Kim, Jiahao Zhang, and Guang Lin. 2025. BEKAN: Boundary condition-guaranteed evolutionary Kolmogorov-Arnold networks with radial basis functions

- for solving PDE problems. doi:10.48550/arXiv.2510.03576 arXiv:2510.03576 [cs].
- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Aizzadenesheli, Kaushik Bhatnagary, Andrew Stuart, and Anima Anandkumar. 2023. Neural operator: learning maps between function spaces with applications to PDEs. *J. Mach. Learn. Res.* 24, 1, Article 89 (Jan. 2023), 97 pages.
- Harender Kumar and Neha Yadav. 2025. Deep learning algorithms for solving differential equations: a survey. *Journal of Experimental & Theoretical Artificial Intelligence* 37, 4 (2025), 609–648. doi:10.1080/0952813X.2023.2242356
- I.E. Lagaris, A.C. Likas, and D.G. Papageorgiou. 2000. Neural-network methods for boundary value problems with irregular boundaries. *IEEE Transactions on Neural Networks* 11, 5 (Sep. 2000), 1041–1049. doi:10.1109/72.870037
- David IW Levin, Peter Yichen Chen, and Eitan Grinspun. 2025. Reduced-Order Modeling for Physical Simulation: From the Classical to the Neural. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Frontiers (SIGGRAPH Frontiers '25)*. Association for Computing Machinery, New York, NY, USA, Article 15, 2 pages. doi:10.1145/3736539.3737842
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.* 39, 4, Article 49 (Aug. 2020), 20 pages. doi:10.1145/3386569.3392425
- Xuan Li, Yu Fang, Lei Lan, Huamin Wang, Yin Yang, Minchen Li, and Chenfanfu Jiang. 2023. Subspace-Preconditioned GPU Projective Dynamics with Contact for Cloth Simulation. In *SIGGRAPH Asia 2023 Conference Papers (Sydney, NSW, Australia) (SA '23)*. Association for Computing Machinery, New York, NY, USA, Article 1, 12 pages. doi:10.1145/3610548.3618157
- Yulei Liao and Pingbing Ming. 2021. Deep Nitsche Method: Deep Ritz Method with Essential Boundary Conditions. *Communications in Computational Physics* 29, 5 (June 2021), 1365–1384. doi:10.4208/cicp.OA-2020-0219 arXiv:1912.01309 [math].
- Mengfei Liu, Yue Chang, Zhecheng Wang, Peter Yichen Chen, and Eitan Grinspun. 2025. Precise Gradient Discontinuities in Neural Fields for Subspace Physics. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 26, 11 pages. doi:10.1145/3757377.3763810
- Pingchuan Ma, Tao Du, John Z. Zhang, Kui Wu, Andrew Spielberg, Robert K. Katzschmann, and Wojciech Matusik. 2021. DiffAqua: a differentiable computational design pipeline for soft underwater swimmers with shape interpolation. *ACM Trans. Graph.* 40, 4, Article 132 (July 2021), 14 pages. doi:10.1145/3450626.3459832
- Kevin Stanley McFall and James Robert Mahan. 2009. Artificial Neural Network Method for Solution of Boundary Value Problems With Exact Satisfaction of Arbitrary Boundary Conditions. *IEEE Transactions on Neural Networks* 20, 8 (Aug 2009), 1221–1233. doi:10.1109/TNN.2009.2020735
- Vismay Modi, Nicholas Sharp, Or Perel, Shinjiro Sueda, and David I. W. Levin. 2024. Simplicit: Mesh-Free, Geometry-Agnostic Elastic Simulation. *ACM Trans. Graph.* 43, 4, Article 117 (July 2024), 11 pages. doi:10.1145/3658184
- Sean Molesky, Zin Lin, Alexander Y. Piggott, Weiliang Jin, Jelena Vucković, and Alejandro W. Rodriguez. 2018. Inverse design in nanophotonics. *Nature Photonics* 12, 11 (Nov. 2018), 659–670. doi:10.1038/s41566-018-0246-9
- Alfio Quarteroni and Gianluigi Rozza. 2014. *Reduced Order Methods for Modeling and Computational Reduction*. Springer International Publishing. doi:10.1007/978-3-319-02090-7
- M. Raissi, P. Perdikaris, and G.E. Karniadakis. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378 (2019), 686–707. doi:10.1016/j.jcp.2018.10.045
- Teseo Schneider, Yixin Hu, Jérémie Dumas, Xifeng Gao, Daniele Panozzo, and Denis Zorin. 2018. Decoupling simulation accuracy from mesh quality. *ACM Trans. Graph.* 37, 6, Article 280 (Dec. 2018), 14 pages. doi:10.1145/3272127.3275067
- Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. 2020. Implicit neural representations with periodic activation functions. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 626, 12 pages.
- N. Sukumar and Ankit Srivastava. 2022. Exact imposition of boundary conditions with distance functions in physics-informed deep neural networks. *Computer Methods in Applied Mechanics and Engineering* 389 (2022), 114333. doi:10.1016/j.cma.2021.114333
- Juan Diego Toscano, Vivek Oommen, Alan John Varghese, Zongren Zou, Nazanin Ahmadi Daryakenari, Chenxi Wu, and George Em Karniadakis. 2025. From PINNs to PIKANs: recent advances in physics-informed machine learning. *Machine Learning for Computational Science and Engineering* 1, 1 (March 2025), 15. doi:10.1007/s44379-025-00015-1
- Jiaji Wang, Y.L. Mo, Bassam Izzuddin, and Chul-Woo Kim. 2023. Exact Dirichlet boundary Physics-informed Neural Network EPINN for solid mechanics. *Computer Methods in Applied Mechanics and Engineering* 414 (2023), 116184. doi:10.1016/j.cma.2023.116184
- Bangbang Yang, Chong Bao, Junyi Zeng, Hujun Bao, Yinda Zhang, Zhaopeng Cui, and Guofeng Zhang. 2022. NeuMesh: Learning Disentangled Neural Mesh-Based Implicit Field for Geometry and Texture Editing. In *Computer Vision – ECCV 2022*, Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer Nature Switzerland, Cham, 597–614. doi:10.1007/978-3-031-19787-1\_34
- Kaixiong Zhou, Xiao Huang, Daochen Zha, Rui Chen, Li Li, Soo-Hyun Choi, and Xia Hu. 2021. Dirichlet Energy Constrained Learning for Deep Graph Neural Networks. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 21834–21846. [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/b6417f112bd27848533e54885b66c288-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/b6417f112bd27848533e54885b66c288-Paper.pdf)
- Zeshun Zong, Xuan Li, Minchen Li, Maurizio M. Chiaramonte, Wojciech Matusik, Eitan Grinspun, Kevin Carlberg, Chenfanfu Jiang, and Peter Yichen Chen. 2023. Neural Stress Fields for Reduced-order Elastoplasticity and Fracture. In *SIGGRAPH Asia 2023 Conference Papers (Sydney, NSW, Australia) (SA '23)*. Association for Computing Machinery, New York, NY, USA, Article 78, 11 pages. doi:10.1145/3610548.3618207



Fig. 8. **Dirichlet boundary-aware basis functions with interior slicing.** A sequence of neural basis functions computed under Dirichlet boundary conditions, with the constrained boundary highlighted in yellow. Each domain is sliced to expose a volumetric cross-section. The neutral white color denotes zero. The smooth transition toward white near the Dirichlet boundary visually confirms the constraint enforced by the basis construction.




Fig. 9. **Transient heat evolution in shape-boundary space.** Setting 1 (top row) combines a tall, narrow fin with strong convection, while Setting 2 (bottom row) uses a short, wide fin with weak convection. The temperature field is initialized close to the bottom temperature and evolves toward the steady state. Snapshots from left to right show different shape-boundary configurations lead to distinct decay rates and spatial heat distributions.




Fig. 10. **Qualitative ablation of lifted input using a 2D scalar field,** visualized as a 3D height surface. **Left:** reference smooth eigenfunction. **Middle:** neural field without lifted input, exhibiting visible artifacts aligned with distance-gradient discontinuity. **Right:** lifted input formulation, which restores surface smoothness and closely matches the reference.

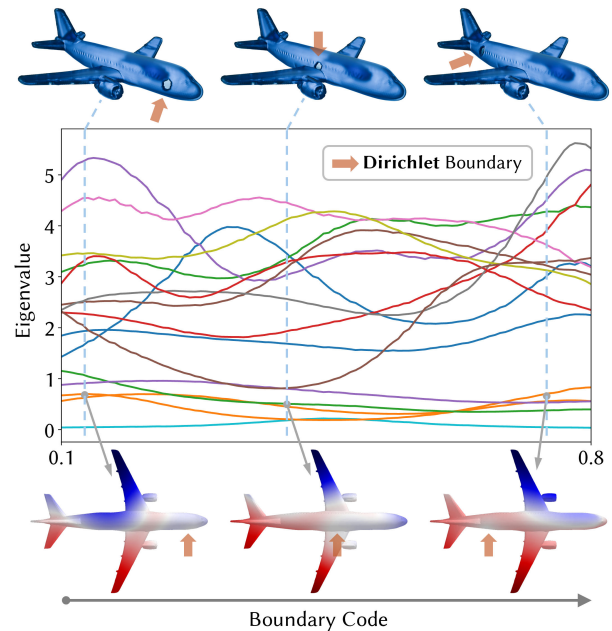


Fig. 11. **Boundary-space spectral tuning on a fixed air cavity.** The geometry is fixed, while a localized pressure-release opening is modeled as a Dirichlet boundary patch that slides along  $\partial\Omega$ , parameterized by a scalar boundary code  $\beta \in [0.1, 0.8]$ . The plot shows the lowest  $K = 16$  eigenvalues as functions of  $\beta$ , exhibiting smooth yet significant spectral variation induced purely by boundary condition changes. Insets visualize representative boundary configurations and corresponding pressure eigenmodes, with arrows indicating the Dirichlet region.

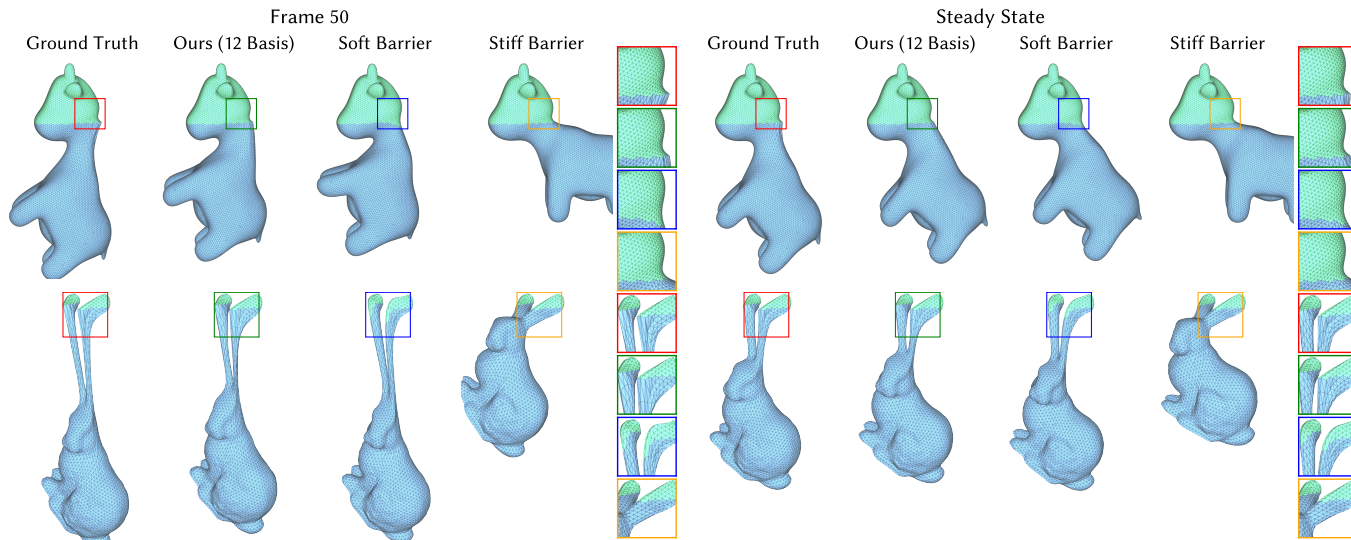


Fig. 12. **Comparison of reduced-order elastic simulations.** We evaluate simulation fidelity on two objects using different bases ( $K = 12$ ). From left to right: Full-space FEM (Ground Truth); our neural Dirichlet-aware basis; Neumann basis with soft barrier ( $\mu = 1$ ,  $d_{max} = 100$ ); and Neumann basis with stiff barrier ( $\mu = 10$ ,  $d_{max} = 10$ ). Close-up views highlight the conflict in prior methods: soft barriers fail to pin the boundary (blue boxes), while stiff barriers induce artificial locking that restricts deformation. Our method satisfies Dirichlet constraints exactly by construction, matching the ground truth without requiring penalty tuning.

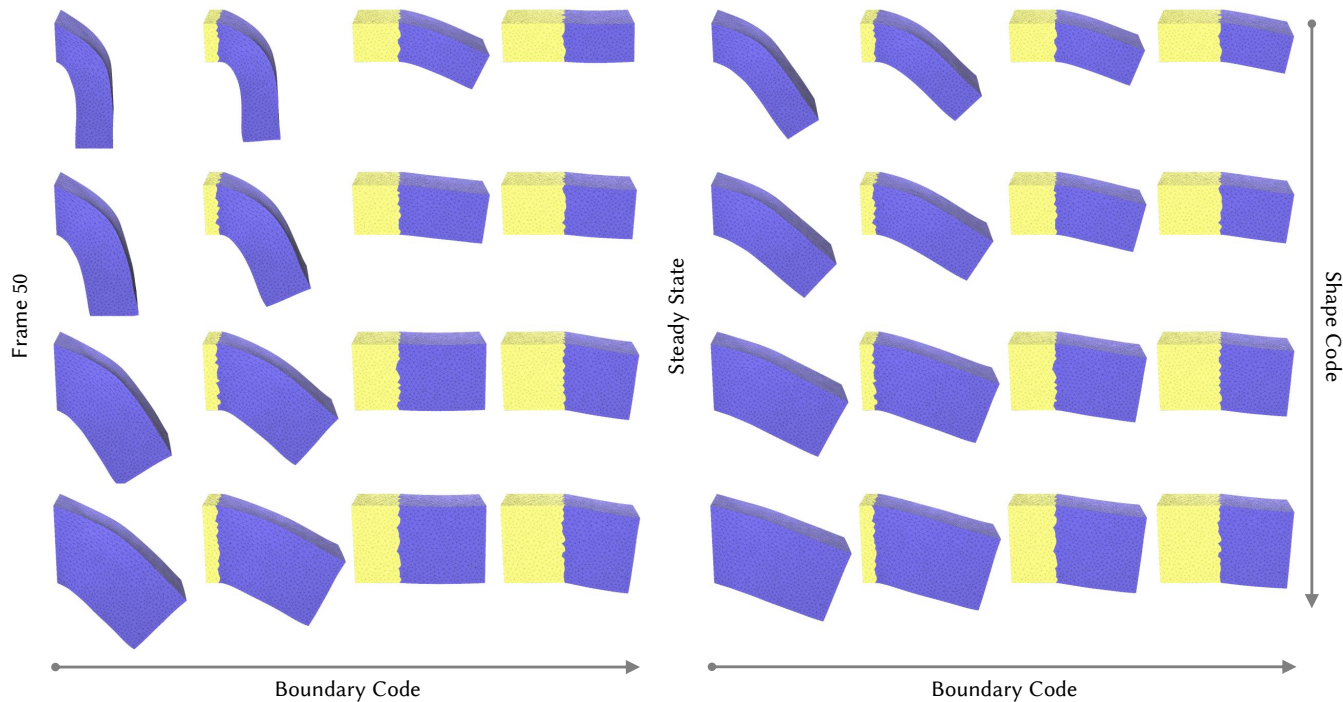


Fig. 13. **Unified Reduced-Order Simulation in Shape–Boundary Space.** We visualize deformation results at Frame 50 (left block) and Steady State (right block). In each block, the boundary parameter varies from left to right, extending the fixed region (yellow) to cover a larger portion of the surface. From top to bottom, the shape code varies, altering the cuboid’s aspect ratio. All simulations are performed using a *single* reduced-order model trained over the joint space, which adapts smoothly to simultaneous changes in geometry and boundary conditions.