

# Shape Space Spectra

YUE CHANG, University of Toronto, Canada

OTMAN BENCHEKROUN, University of Toronto, Canada

MAURIZIO M. CHIARAMONTE, Meta Reality Labs Research, USA

PETER YICHEN CHEN, MIT CSAIL, USA

EITAN GRINSPUN, University of Toronto, Canada

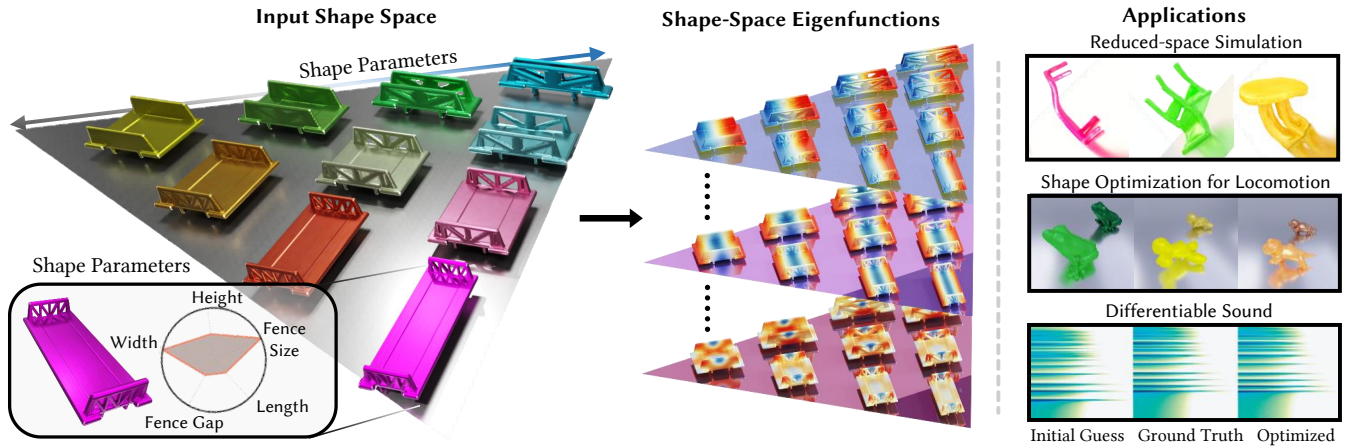


Fig. 1. Given a continuously-parameterized shape space (*left*), we introduce eigenanalysis parameterized over shape space (*center*), enabling applications such as subspace simulation and inverse design (*right*). Unlike eigenanalysis of a single shape, our shape-space eigenanalysis is readily differentiable with respect to shape parameters, enabling optimization objectives based on eigenmodes, such as locomotion or timbre. Our focus on shape space in turn requires new techniques for consistency of modes across shape space: compare the shading of shapes within one triangle (eigenfunctions varying over shape space) and across triangles (distinct modes revealed by eigenanalysis).

Eigenanalysis of differential operators, such as the Laplace operator or elastic energy Hessian, is typically restricted to a single shape and its discretization, limiting reduced order modeling (ROM). We introduce the first eigenanalysis method for continuously parameterized shape families. Given a parametric shape, our method constructs spatial neural fields that represent eigenfunctions across the entire shape space. It is agnostic to the specific shape representation, requiring only an inside/outside indicator function that depends on shape parameters. Eigenfunctions are computed by minimizing a variational principle over nested spaces with orthogonality constraints. Since eigenvalues may swap dominance at points of multiplicity, we jointly train multiple eigenfunctions while dynamically reordering them based on their eigenvalues at each step. Through *causal gradient filtering*, this reordering is reflected in backpropagation. Our method enables applications to operate over shape space, providing a single ROM that encapsulates vibration modes

Authors' addresses: Yue Chang, University of Toronto, Canada, changyue.chang@mail.utoronto.ca; Otman Benckekroun, University of Toronto, Canada, otman.benckekroun@mail.utoronto.ca; Maurizio M. Chiaramonte, Meta Reality Labs Research, USA, mchiaram@meta.com; Peter Yichen Chen, MIT CSAIL, USA, pyc@csail.mit.edu; Eitan Grinspun, University of Toronto, Canada, eitan@cs.toronto.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2025 Association for Computing Machinery.

0730-0301/2025/6-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

for all shapes, including previously unseen ones. Since our eigenanalysis is differentiable with respect to shape parameters, it facilitates eigenfunction-aware shape optimization. We evaluate our approach on shape optimization for sound synthesis and locomotion, as well as reduced-order modeling for elastodynamic simulation.

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

Additional Key Words and Phrases: Reduced-order modeling, Implicit neural representation, Computational design, Differentiable simulation

## ACM Reference Format:

Yue Chang, Otman Benckekroun, Maurizio M. Chiaramonte, Peter Yichen Chen, and Eitan Grinspun. 2025. Shape Space Spectra. *ACM Trans. Graph.* 1, 1 (June 2025), 16 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

From the aerodynamics of an airplane wing to the flexibility of a plastic fork, partial differential equations (PDEs) play a crucial role in physics-constrained product design. The eigenfunctions of PDE operators are essential for analyzing design solutions, as they help identify bending or fracture patterns, describe resonant frequencies, and encode geometric properties such as distance and curvature.

However, eigenfunctions are typically computed for a single geometry, and geometry modification requires recomputing the discrete operator and its eigendecomposition, which is nonlinear with respect to geometry. These compute-intensive steps hinder PDE-based shape optimization and interactive design tasks, which often

require numerous eigenfunction and derivative evaluations. A key challenge is that evaluating eigenfunctions during optimization is cost-prohibitive, and their derivatives with respect to shape parameters are not readily available.

We introduce a method and representation for eigenmodes over shape space, enabling efficient evaluation for any shape, including those unseen during training (see Fig. 1, Fig. 2, Fig. 4). Since our eigenfunctions and eigenvalues share continuous parameters with the corresponding geometry, differentiating eigenfunctions with respect to shape is straightforward. To the best of our knowledge, this is the first proposed method of its kind in the literature.

Our method is agnostic to the shape space representation (see Fig. 3), requiring only an inside/outside indicator function dependent on shape parameters. As a result, it applies to both manually defined shape interpolation and neural implicit representations trained over real-world datasets [Chen and Zhang 2019a].

Our first contribution is a variational method that generates eigenfunctions for single shape domains represented by neural fields. While the eigenfunction problem for PDE operators is well-established in the context of finite element discretizations, it remains largely unexplored for more generalized domains, such as those represented by neural implicit representations, even for single shapes. To address this, we introduce a projection module at the end of our neural field pipeline, ensuring that each eigenfunction's output lies within the classical eigenfunction constraint space. With these constraints *ensured by construction* in the output space of our network, we can compute the eigenfunctions by minimizing a suitable operator-aware loss functional.

Expanding the aforementioned method to a shape space presents its own unique challenges. For a single shape, these eigenfunctions are labeled, computed, and discussed in terms of the order of their dominance, corresponding to a monotonic sequence of eigenvalues. For a shape space, however, we can no longer think in terms of monotonicity because eigenvalues as functions over shape space often cross and exchange dominance (see Fig. 7-right and Fig. 8-right). Maintaining the right “topological structure” of eigenfunctions over shape space as they cross each other is critical to engineering design tasks that require differentiating eigenvalues or eigenfunctions with respect to shape parameters.

Therefore, our second contribution is a method that *jointly* optimizes eigenfunctions over shape space. This requires accounting




Fig. 3. *Eigenfunctions Across Different Representations* Our method is discretization-agnostic; we have demonstrated the calculation of eigenfunctions for point clouds, signed distance fields, tetrahedral meshes, and neural implicit representations.

for the exchange of dominance relationships over shape space and a special filtering step during back-propagation, which we refer to as *causal sorting* and *causal filtering*, respectively.

For the first time, we have a shape-dependent eigenfunction representation that correctly tracks crossovers of eigenvalues at points of multiplicity in shape space. This approach generalizes across families of shapes, allowing us to predict eigenfunctions for new shapes that were not seen during precomputation. With its novel dynamic reordering and gradient filtering, it encourages accurate reproduction of eigenfunctions, with consistent modes across different shapes, enabling applications such as single-model multiple-shape subspace physics, warm-starting PDE solutions across shapes, and inverse shape optimization for locomotion and sound profile.

## 2 RELATED WORK

### 2.1 Eigenfunctions of PDEs

Eigenfunctions of PDEs have a broad range of applications such as deformation [Benckroun et al. 2023; Hildebrandt et al. 2011; Jacobson et al. 2014; James and Pai 2002], fluid simulation [Cui et al. 2018; De Witt et al. 2012], locomotion [Kry et al. 2009; Nunes et al. 2012], sound analysis [Bharaj et al. 2015; Kac 1966; O’Brien et al. 2002], shape analysis [Mateus et al. 2008; Melzi et al. 2018; Ovsjanikov et al. 2012; Rampini et al. 2019; Sharma and Horaud 2010; Sun et al. 2023; Vallet and Lévy 2008], and geometric deep learning [Sharp et al. 2022; Smirnov and Solomon 2021].

The most straightforward approach to calculating the eigenfunctions for a given shape is to mesh the domain and perform eigendecomposition on the constructed discrete operator.

For the Laplace-Beltrami operator, despite the availability of various methods [Belkin et al. 2008, 2009; Bobenko and Springborn 2005; Fisher et al. 2006; Guezic et al. 1999; Liu et al. 2017; Pang et al. 2024; Sellán et al. 2019; Sharp and Crane 2020; Shimada and Gossard 1995], each with its own advantages, none perfect [Wardetzky et al. 2007], the most commonly adopted discrete operator is the cotangent matrix. For elastic energies, the commonly used discrete operator is the elastic Hessian, which varies depending on the specific elastic energy. Commonly used elastic energy models include linear elasticity [Sifakis and Barbic 2012], St. Venant-Kirchhoff [Barbič and James 2005], co-rotational (ARAP) [Rankin and Brogan 1986; Sorkine and Alexa 2007], and Neo-Hookean elasticity [Smith et al. 2018], to name a few. A generalized approach for calculating the Hessian matrix for a given elastic energy can be found in [Kim and Eberle 2022].



Fig. 2. *Elastic eigenfunctions over Shape Space*. With a single neural model, we can compute eigenfunctions for a bridge over the entire shape-space shown in Fig. 1.

After obtaining the discrete operator, the eigenvectors can trivially be calculated via eigensolvers [Arbenz et al. 2005; Duersch et al. 2018; Nasikun and Hildebrandt 2022]. Because all these approaches rely on a discrete operators and perform eigenanalysis on its representative matrix, they are tied to a *single discretization of one shape*. They cannot find eigenfunctions for general shape *families*.

## 2.2 Shape Spaces and Neural Fields

The shape space in our paper refers to a family of shapes defined by a continuously parameterized shape code. Shape spaces can serve various purposes, including defining a solution space for shape optimization [Jin et al. 2024; Ma et al. 2021], visualizing physical properties [Schulz et al. 2017], and enabling shape representation and generalization [Chen and Zhang 2019a]. Common examples of such shape spaces in the graphics community include design spaces [Schulz et al. 2018], shape interpolation [Solomon et al. 2015], mesh Booleans [Liu et al. 2024; Yuan et al. 2024]. Unfortunately, different shapes usually entail vastly different discretizations, making it very difficult to smoothly navigate geometrically diverse shape spaces.

To make our method discretization-agnostic and able to integrate with large shape spaces, we make use of shape families defined by neural fields [Xie et al. 2021]. These neural fields parameterize a spatially dependent vector field through a neural network. Early seminal efforts by Chen and Zhang [2019b]; Mescheder et al. [2019]; Park et al. [2019] utilized this framework for encoding signed distance fields, wherein each distinct latent vector represents a unique geometry. Neural fields have since been extensively applied in various domains, including neural rendering [Mildenhall et al. 2020], 3D reconstruction [Wang et al. 2021; Yariv et al. 2020], geometry processing [Aigerman et al. 2022; Dodik et al. 2023; Mehta et al. 2022; Williamson and Mitra 2024; Yang et al. 2021], topology optimization [Zehnder et al. 2021], constitutive modeling [Li et al. 2023], and solving diverse PDE problems [Chang et al. 2023; Chen et al. 2022, 2023; Deng et al. 2023; Raissi et al. 2019].

Our work builds upon previous research and extends neural fields to model eigenfunctions for PDEs. Neural fields offer a significant advantage over traditional representations like meshes and grids when modeling eigenfunctions: the ability to model a large family of parameterized shapes. Specifically, neural fields enable the modeling of eigenfunctions across a family of shapes by using high-dimensional shape codes as inputs. Furthermore, unlike prior approaches that fit neural fields to pre-existing geometric data, our method trains eigenfunction neural fields in a geometry-informed manner *without* requiring any precomputed eigenfunctions as training data.

## 3 EIGENANALYSIS OF A SINGLE SHAPE

We begin with defining eigenfunctions for a single shape and then describe our method over shape space in Sec. 4. In both cases, we first discuss the variational perspective as a foundation, and then move on to implementation. We begin with a variational perspective on eigenanalysis of the Laplace operator for a specific, non-parametric volumetric shape in  $\mathbb{R}^n$ , before extending to elasticity and shape spaces.



Fig. 4. *Laplace Eigenfunctions for a Teapot Shape Space.* The eigenfunctions of the Laplace operator describe the low-frequency heat distributions. We demonstrate these eigenfunctions across different teapot shapes in the shape space.

### 3.1 Eigenanalysis: A Variational Perspective

Consider a compact subset  $\Omega \subset \mathbb{R}^n$  with a piecewise smooth boundary  $\partial\Omega$ . We are interested in studying the eigenfunctions of the Laplace operator  $\Delta$ , subject to appropriate boundary conditions on  $\partial\Omega$ . The Laplace operator is defined as

$$\Delta u = \nabla \cdot \nabla u = \sum_{i=1}^n \frac{\partial^2 u}{\partial x_i^2},$$

where  $u$  is a sufficiently smooth function defined on  $\Omega$ .

The dominant eigenfunction  $\phi_1(x)$  minimizes Dirichlet energy

$$E_D[\phi] = \frac{1}{2} \int_{\Omega} |\nabla \phi|^2 d\Omega, \quad (1)$$

among  $\mathcal{U} = \{f \in L^2(\Omega) \mid \|f\|_2 = 1\}$ , the unit-norm square-integrable functions in  $\Omega$ . Restricting the search to unit-norm functions (akin to taking the Rayleigh quotient) helps canonize the minimizer and equates eigenvalue to Dirichlet energy,  $\lambda_1 = E_D[\phi_1]$ .

*Peeling off eigenfunctions.* The subdominant eigenfunction *also* minimizes Dirichlet energy, but this time in the space  $\phi_1^\perp$ , the *orthogonal complement* to  $\phi_1$  in  $S$ . And so forth, in order of dominance: the  $i$ 'th eigenfunction  $\phi_i$  minimizes Dirichlet energy  $\lambda_i = E_D[\phi_i]$  in  $C_i = \text{span}\{\phi_1, \dots, \phi_{i-1}\}^\perp$ , the space orthogonal to earlier modes:

$$\phi_i = \arg \min_{\phi \in \mathcal{U} \cap C_i} E_D[\phi]. \quad (2)$$

In this iterative “peeling” procedure, echoing the Courant-Fischer-Weyl min-max principle, each dominant eigenfunction is found and “peeled away” revealing the complementary subspace containing the subdominant eigenspace, where the procedure is repeated. Because the spaces are nested, the minimizer of the “bigger” space cannot be greater than subsequent “smaller” spaces, therefore  $\lambda_1 \leq \lambda_2 \leq \dots$ .

Since we have not explicitly enforced any boundary conditions, the minimizers all satisfy the *natural* condition, which, for the Dirichlet energy, is the vanishing Neumann value  $\frac{\partial \phi_i}{\partial n} = 0$  on  $\partial\Omega$ , as shown in Fig. 5. We discuss avenues for essential boundary conditions in Sec. 7.






Fig. 5. *Boundary Condition*. Minimizing the Dirichlet energy on a shape naturally enforces a vanishing Neumann boundary condition. To illustrate this, we visualize the isolines of the eigenfunctions. Notably, the gradients along the boundary are nearly zero, highlighting the fulfillment of the Neumann condition.

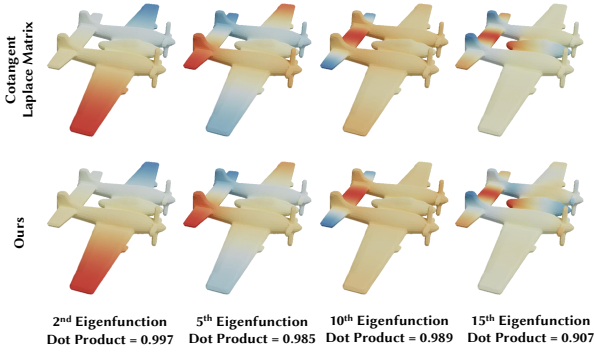


Fig. 6. *Comparison with eigenfunctions from cotangent Laplacians*. When training on a single shape, our method converges to results consistent with traditional eigenanalysis. We compare the eigenfunctions obtained using our approach with those derived from a cotan Laplace (linear finite element on triangle mesh) matrix. The accuracy of our method matches that of existing techniques that rely on matrix construction from point cloud sampling.

### 3.2 Implementation with Neural Fields

We implemented the optimization described in Equation 2 using neural fields. Our results show good agreement with prior discrete Laplace operators (see Fig. 6), boundary condition satisfaction (see Fig. 5), and SE(3) (or SE(2)) invariance (see Fig. 12).

Our implementation models each eigenfunction  $\phi_i$  as the composition of a corresponding neural field and a projection operator:

$$\phi_i = \mathcal{P}_i \circ \bar{\phi}_i. \quad (3)$$

The neural field  $\bar{\phi}_i$  is a multilayer perceptron (MLP) mapping domain position,  $\mathbf{x} \in \Omega$ , to field value,  $\bar{\phi}_i(\mathbf{x})$ . The projection operator  $\mathcal{P}_i : L_2 \rightarrow \mathcal{U} \cap C_i$  maps any field to a constraint-satisfying field.

Algorithm 1 performs eigenanalysis for the Laplace operator over a single shape. We describe the optimization, energy (loss) evaluation, and projection operator in turn.

Each neural field,  $\bar{\phi}_i$ , is trained by minimizing the loss  $\mathcal{L} = E_D[\phi_i]$ . After convergence, the next neural field is trained, in sequence. We estimate the domain integral with stochastic cubature:

$$\tilde{\mathcal{L}} = \tilde{E}_D[\phi_i] = \sum_{\mathbf{x}_j \in \mathcal{X}} |\nabla \phi_i(\mathbf{x}_j)|^2. \quad (4)$$

---

#### ALGORITHM 1: Optimization on a Single Shape for Laplace Eigenfunction

---

```

epoch = 0;
repeat
  X = {x1, ...};           ▶ Sample Domain Ω;
  φprev(X) = ones_like(X);
                        ▶ Hardcode the known eigenfunction;
  for i in range [0, k) do
    Evaluate Network  $\bar{\phi}_i(X)$ ;
    Calculate λ by doing the projection in Equation 5;
    Calculate φm+1P(X) by Equation 6;
    Calculate φ(X) by Equation 7;
    for x in X = {x1, ...} do
      Evaluate gradient  $\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}}$  by Equation 8;
      ▶ In practice, this is implemented using tensors;
    end
    Calculate loss  $\mathcal{L}$  by Equation 4;
    Backward loss;
    Concatenate φprev(X) with φi(X) and detach;
      ▶ Update for future orthogonal constraints;
  end
  epoch = epoch + 1;
until epoch = MaxEpoch;

```

---

Our implementation uniformly samples cubature points  $\mathcal{X} = \{\mathbf{x}_j \in \Omega\}$  via rejection sampling, drawing from a uniform distribution over an axis-aligned bounding volume, and rejecting samples based on the indicator (inside/outside query) function of  $\Omega$ . Since our implementation assumes only an indicator function, it is agnostic to the representation of  $\Omega$ .

The projection  $\mathcal{P}_i$  takes a general field  $\bar{\phi}_i$  and returns the closest field  $\phi_i \in \mathcal{U} \cap C_i$  satisfying the orthogonality and unit-norm conditions. The projection is composed of two steps: orthogonalization and normalization, which we represent via the dataflow schematic:

$$\mathcal{P}_i : \bar{\phi}_i \xrightarrow{\text{Gram-Schmidt}} \phi_i^P \xrightarrow{\text{normalize}} \phi_i.$$

*Orthogonalization (projection to  $C_i$ ).* Any two eigenfunctions  $\phi_i$  and  $\phi_j$ ,  $i \neq j$ , must be orthogonal,  $\int_{\Omega} \phi_i \phi_j d\Omega = 0$ , a constraint we achieve by Gram-Schmidt orthogonalization. Given a candidate (unconstrained) function  $\bar{\phi}_m$  not yet orthogonal to all previous eigenfunctions, we find its projection  $\phi_m^P$  onto the orthogonal subspace by subtracting from  $\bar{\phi}_m$  the component already spanned by the  $m-1$  dominating eigenfunctions  $\phi = (\phi_1, \dots, \phi_{m-1})$ . We seek the  $\lambda$ -weighted linear combination of dominating eigenfunctions  $\phi$  that best approximate  $\bar{\phi}_m$ ,

$$\lambda = \arg \min_{\lambda \in \mathbb{R}^{m-1}} \left\| \lambda^T \phi - \bar{\phi}_m \right\|_2, \quad (5)$$

where  $\|\cdot\|_2$  is the  $L_2$  norm, which, like all domain integrals, we estimate by uniform stochastic cubature. Next, we remove the component:

$$\phi_m^P = \bar{\phi}_m - \lambda^T \phi. \quad (6)$$

The gradient, required for training, follows by chain rule.



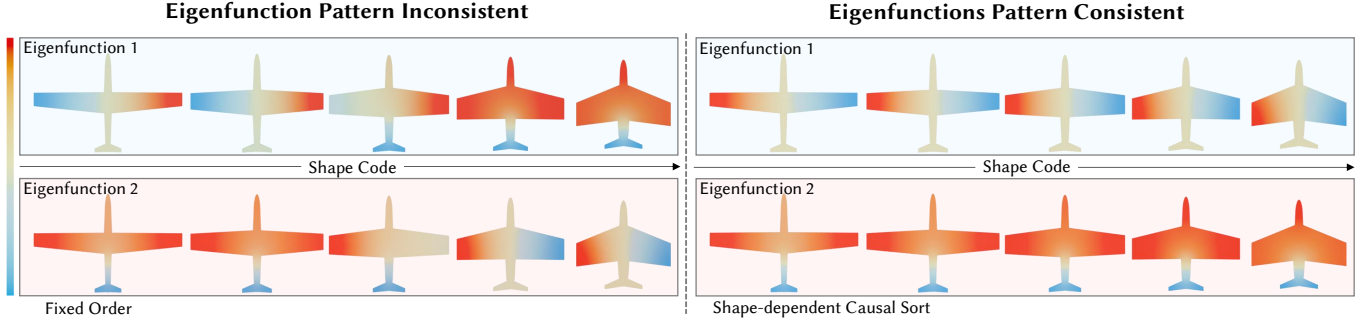


Fig. 7. The wing shape varies over shape space. (left) **Ablation**: eigenfunctions 1 & 2 swap modal patterns at the point of eigenvalue multiplicity (see Fig. 8). (right) **Our approach**: shape-dependent causal sorting improves mode consistency across shape space.

*Normalization (projection to  $\mathcal{U}$ ).* We enforce the unit-norm constraint by normalizing  $\phi^P(\mathbf{x})$ :

$$\phi(\mathbf{x}) = \frac{\phi^P(\mathbf{x})}{\|\phi^P\|_2}, \quad (7)$$

again estimating the  $L_2$  norm by uniform stochastic cubature.

To accelerate gradient computation, we approximate the gradient of the norm by treating it as independent of its argument, a similar approach to [Lévy and Zhang 2010; Williamson and Mitra 2024].

$$(\partial/\partial\mathbf{x})\phi \approx \|\phi^P\|_2^{-1}(\partial/\partial\mathbf{x})\phi^P. \quad (8)$$

Now, we have successfully constructed a set of continuous functions that are orthogonal to each other and possess a unit norm.

*Implementation details for a single Shape.* We use ADAM [Kingma and Ba 2017] to optimize the neural network weights, and implement our method in PyTorch, evaluating spatial gradients such as  $\nabla\bar{\phi}_i(\mathbf{x})$  using PyTorch’s autodiff. Theory suggests that to minimize bias in the optimization, each domain integral should be estimated using an independently sampled cubature set, however, we found that using one cubature set per epoch, for all integrals, provides good results.

*Extension to Elasticity.* Our method can also be applied to elasticity. In this case, the resulting eigenfunctions (elastic modes) are vector-valued fields,  $\phi_i(\mathbf{x}) \mapsto \mathbb{R}^3$ . We implemented linear elastic energy [Sifakis and Barbic 2012]. The eigenfunctions in this context are minimizers of the energy functional:

$$E_e[\phi] = \frac{1}{2} \int_{\Omega} \mu |\nabla\phi + \nabla\phi^T|_F^2 + \frac{\lambda}{2} \text{Tr}^2(\nabla\phi + \nabla\phi^T) d\Omega, \quad (9)$$

Here,  $\phi$  represents the vector-valued eigenfunction, and  $\nabla\phi$  is the deformation gradient (or Jacobian), expressed as a  $3 \times 3$  matrix. The term  $\nabla\phi + \nabla\phi^T$  is known as the small strain tensor. The Frobenius norm is denoted by  $|\cdot|_F$ , and  $\mu$  and  $\lambda$  are the Lamé coefficients. We estimate  $\tilde{E}_e \approx E_e$  using uniform stochastic cubature.

Crucially, other than substituting  $\tilde{E}_D$  with  $\tilde{E}_e$ , and reconsidering the known modes (discussed below), our theory and implementation are unchanged, highlighting the broader applicability of the method.

*Handling Known Modes.* When subject to Neumann boundary conditions, the dominant mode ( $\lambda_1 = 0$ ) of the Laplace operator is always the constant function, irrespective of domain geometry. While

the peeling implementation described above is able to find this first trivial mode, this is wasted computation, since the result is already known. Therefore, we hard-code the known geometry-independent eigenfunction analytically, and represent the remaining, geometry-dependent modes using neural fields. For the elasticity operator, we hard-code the zero eigenvalue modes corresponding to rigid translations and rotations.

#### 4 EIGENANALYSIS OVER SHAPE SPACE

We are ready to dive into training over shape spaces. Let the domain  $\{\Omega^g | g \in \mathcal{D}\}$  be parameterized by a geometry code  $g \in \mathcal{D}$  drawn from a shape space  $\mathcal{D}$ . Since choosing a point  $g \in \mathcal{D}$  fixes the shape of the domain  $\Omega^g$ , we could use the eigenanalysis of a single shape from §3 to determine the eigenfunctions  $\phi_i^g(\mathbf{x})$ .

On the other hand, by explicitly leaving  $g \in \mathcal{D}$  as a free parameter, we can think of eigenfunctions  $\phi_i^g(\mathbf{x}) \equiv \phi_i(g, \mathbf{x})$  as spatial fields parameterized over shape space, whose cross-sections at some  $g = g_j$  correspond to the eigenfunctions over  $\Omega^{g_j}$ . Classical results from perturbation theory [Kato 1980] state that for small, smooth perturbations of the domain, eigenvalues and eigenfunctions of elliptic operators (e.g., the Laplacian) vary smoothly, or even analytically. This opens the door to a reduced-order or parametrized representation of the eigenfunctions across the entire shape space  $\mathcal{D}$ , that is, to efficient learning and representation of eigenfunctions  $\phi_i^g(\mathbf{x})$  across a continuous family of domain geometries.

##### 4.1 Shape-Space Eigenanalysis: A Variational Perspective

A first, albeit misguided, attempt to generalize the method discussed in §3 might be to modify the domain of integration to be the product of shape space and the spatial domain. The first eigenfunction  $\phi_1$  minimizes the integral of Dirichlet energy over shape space  $\int_{\mathcal{D}} E_D[\phi_1^g] d\mathcal{D}$ , the second eigenfunction  $\phi_2$  does so restricted to  $\phi_1^\perp$ , and so forth. But such peeling would inherently yield eigenfunctions that maintain a fixed dominance relationship across shape space, which, as we are about to see, is damaging and unnecessary.

We present a didactic example in Fig. 7-left, where the shape space represents a family of airplanes of changing wing thicknesses  $g$ . As the shape code changes, the vertically oriented eigenfunctions (varying from the left wingtip to the right wingtip) and the horizontally oriented eigenfunctions (varying from the tail to the

aircraft’s body) swap dominance. If we also plot the eigenvalues of the two eigenfunctions, as shown in Fig. 8-left, we observe that the crossover occurs at a point of geometric symmetry and eigenvalue multiplicity.

In this example, the classical view of eigenfunctions ordered by dominance leads to defining the blue curve that (by construction) always dominates the red curve throughout shape space Fig. 8-left (“Fixed Order”). Observe how the blue curve represents an eigenfunction that is only *piecewise* smooth over design space, with a kink at the crossover; the same for the red curve. Such a kink is undesirable for neural field training and its applications, consuming more network capacity, slowing convergence, and leading to numerical challenges or failures in applications that harness the smoothness and differentiability of eigenfunctions with respect to  $g$ . Worse, such discontinuities increase in number and topological complexity with the dimension of design space and the number of eigenfunctions.

Fortunately, these discontinuities are an unnecessary, fictitious fabrication arising only from clinging to the single shape mindset. Indeed, Kato [1980] argues instead for viewing eigenvalues as an *unordered set*, whose subscript indexing merely provides a unique identifier, not an ordering. Now eigenvalues and eigenfunctions are smooth even over crossings at points of multiplicity. This is depicted in Fig. 7-right and Fig. 8-right (“Shape-Dependent Causal Sort”), wherein the eigenvalue curve colored blue corresponds to the eigenfunction that is smooth and horizontally-oriented *throughout* shape space; likewise, the red curve corresponds to the vertically-oriented mode smoothly-varying over shape space. Fig. 13 depicts the same kind of comparison for a complex shape space.

We therefore turn to a variational principle that *jointly* considers a *set* of eigenfunctions minimizing the sum of their Dirichlet energies integrated over shape space,

$$\arg \min_{\phi_0 \dots \phi_k} \sum_{i=0}^k \int_{\mathcal{D}} E_D[\phi_i^g] dg, \quad \text{subject to “orthogonality.”} \quad (10)$$

This energy is a straightforward extension of the single shape case, and the story would end here if eigenvalues never crossed. The key remaining ingredient is the enforcement of orthogonality relationships, which becomes nontrivial because dominance relationships vary over shape space.

## 4.2 Implementation Using Neural Fields

To achieve the desirable construction of Fig. 8-right, we must forgo sequential peeling of an ordered sequence and instead jointly optimize an unordered set of eigenfunctions. As depicted in the same figure, our optimization must allow different functions to dominate in different regions of shape space.

This novel direction presents three interwoven technical challenges: (1) We seek to find multiple eigenfunctions governed by coupled energy minimization principles across a descending chain of spaces, (2) Each minimization problem is constrained to a subspace determined by the solution of all preceding minimizations, establishing a causal relation, and (3) The ordering of these causal relations varies across the shape space.

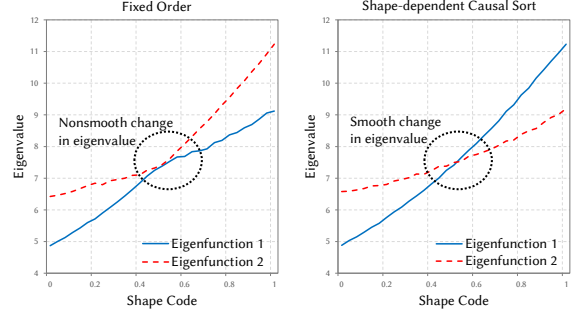


Fig. 8. We visualize the eigenvalues corresponding to Fig. 7. The x-axis shape code matches the middle panel of Fig. 7, following the same shape variation. The (in)consistency of modal patterns in Fig. 7 can be understood by examining the eigenvalues as functions over shape space. (left) **Ablation**: Comparing the eigenvalue plot to the eigenfunctions depicted in Fig. 7-left, the mode swap occurs at the point of multiplicity, where the eigenvalue curve kinks. (right) **Ours**: Our eigenmodes are as consistent as possible across shape space, and our eigenvalue curves are smoother, crossing as appropriate at points of multiplicity.

To address these challenges, we extend the single shape algorithm with three novel, interconnected concepts, all three *strictly required* to obtain eigenfunctions analytic over shape space:

- (1) **Joint training**: As we have seen, sequential peeling produces an incorrect structure because no single ordering of eigenfunctions is valid across all of the shape space. Therefore, we *must* learn  $n$  eigenfunctions *jointly*.
- (2) **Gradient causal filtering**: As we shall see, a naïve approach to joint learning suffers from an action-reaction artifact, whereby functions earlier in the causal chain are affected by an orthogonality constraint that should only affect functions subsequent in the chain. To address this, our backpropagation filters the gradient to enforce the causality of the orthogonality constraint.
- (3) **Shape-dependent causal sorting**: Since the ordering of causal relations cannot be predetermined and indeed varies over shape space, we determine the order dynamically. At each evaluation of the loss function, we re-establish the causal ordering of orthogonality constraints based on the relative dominance of eigenvalues.

**Necessity of these three advances**: These three interwoven concepts are all *required* for our method to achieve, for the first time, a shape-dependent eigenfunction representation that correctly tracks crossovers of eigenvalues at points of multiplicity.

We show that our dynamic reordering encourages smoother eigenfunctions across different shape (see Fig. 7-right) and facilitating eigenfunction-dependent shape design.

As in §3.2, we represent each eigenfunction  $\phi_i$  as the composition of a corresponding neural field and a projection operator,  $\phi_i = \mathcal{P}_i \circ \bar{\phi}_i$ . However, this time the neural field maps *both* shape code,  $g \in \mathcal{D}$ , and domain position,  $x \in \Omega$ , to field value,  $\bar{\phi}_i(g, x)$ ; and the

projection operator,  $\mathcal{P}_i$ , which restrict the field to the constraint-satisfying function space  $\mathcal{U} \cap C_i$ , must now account for the variation of dominance relations over shape space.

Algorithm 2 performs eigenanalysis for the Laplace operator over shape space. We describe the optimization, energy (loss) evaluation, and projection operator in turn.

**4.2.1 Energy Evaluation.** A set of neural fields  $\{\bar{\phi}_0, \dots, \bar{\phi}_k\}$  is jointly trained by minimizing the loss (recalling Eq. (10))

$$\mathcal{L} = \sum_{i=0}^k \int_{\mathcal{D}} E_D[\phi_i^g] d\mathbf{g}. \quad (11)$$

We estimate integrals over shape space and spatial domains via uniform stochastic cubature (see §3.2), uniformly drawing a finite set of shapes  $\mathcal{G} = \{g_j \in \mathcal{D}\}$ , yielding the discretized loss

$$\tilde{\mathcal{L}} = \sum_{i=0}^k \sum_{g_j \in \mathcal{G}} \tilde{E}_D[\phi_i^{g_j}]. \quad (12)$$

The computation of the loss includes evaluation of  $\tilde{E}_D$ , wherein the gradient  $\nabla \equiv \nabla_{\mathbf{x}}$  is a spatial gradient calculated solely with respect to spatial coordinates (recall Eq. (4)). The implementation for elasticity simply replaces  $\tilde{E}_D$  with  $\tilde{E}_e$  (see §3.2).

**4.2.2 Joint Training and Gradient Causal Filtering.** Joint training of eigenfunctions complicates the enforcement of orthogonality constraints. Since the weights of both the dominating and dominated neural fields are now optimization variables, backpropagation produces an undesirable artifact. The gradient of the orthogonalization step has components along both the dominating and dominated fields. This is the “action-reaction” principle of a constraint force. A naïve backpropagation of this constraint gradient disrespects the direction of causality, with the dominated field “pushing back” the dominating field away from its optimum.

The inset figure illustrates the action-reaction artifact. Given a dominating function  $\phi_1$  and a dominated function  $\phi_2$  subject to the constraint  $\phi_2 \in \phi_1^\perp$ , there are two ways to decrease the Dirichlet energy of the dominated function, either (1) by making adjustments within the orthogonal subspace  $\phi_1^\perp$ , or (2) by making adjustments to  $\phi_1$  so as to modify the admissible space  $\phi_1^\perp$ . The first option (“action”) respects causality, whereas the latter (“reaction”) does not; yet both arise from differentiating the constraint  $\phi_2 \in \phi_1^\perp$  with respect to  $\phi_1$  and  $\phi_2$ .

To eliminate the causality-violating reaction, we must *not* differentiate the constraint with respect to the dominating function. We call this causality-enforcing ignoring of a gradient term **gradient causal filtering**. Such filtering could be used to enforce any causal constraint relationship in joint training of neural networks, and we use it to protect the causality of the descending chains of orthogonal spaces. We implement the filtering using `detach()` in PyTorch. After detachment, the gradient of the orthogonality




Fig. 9. To illustrate the importance of *gradient causal filtering*, we compare Laplacian eigenfunctions over the unit interval as produced by: (top) **Ground truth**: analytical evaluation of sinusoids satisfying the vanishing Neumann condition; (middle) **Ours**: variational eigenanalysis including gradient causal filtering; (bottom) **Ablation**: our approach excluding gradient causal filtering.

constraint is considered *only* with respect to the dominated eigenfunction. Figure 9 presents an ablation study for gradient causal filtering.

**Speedup of single shape training.** We train jointly because it is the only viable path to producing correct results over shape space, but as a side bonus, the joint approach trains 2× to 4× faster, too, compared to a sequential approach.

**4.2.3 Shape-Dependent Causal Sorting.** Since fixing the dominance order of the eigenfunctions is undesirable (recall §4.1), our optimization allows ordering to vary over shape space. To achieve the desirable results of Fig. 7-right, we determine the dominance order *dynamically* at each optimization step, by comparing the eigenvalues (equivalently, Dirichlet energy) of the two eigenmodes for some point in shape space. Wherever the horizontally-oriented eigenfunction has the smaller eigenvalue, it dominates the vertically-oriented eigenfunction, and vice-versa.

Therefore, the order of projection in  $\mathcal{P}$  is determined by the eigenvalue and varies for each shape. Since we integrate over shape space using stochastic cubature, this amounts to determining the dominance chain at each sample point in shape space by sorting the eigenvalues (i.e., Dirichlet energy) of the eigenfunctions. Algorithm 2 first evaluates the energy for each unit-norm eigenfunction, then sorts to determine the causal ordering, and then applies this order to construct the projection operator.

With dynamic sorting, we produce the desirable results of Fig. 7-right and 8-right, where each eigenfunction and eigenvalue evolves smoothly over shape space.

## 5 EXPERIMENTS

### 5.1 Training Time & Statistics

We summarize the training time and network parameters for all examples in Table. 1, all data reported was obtained on an AMD Ryzen 9 7950X CPU and an NVIDIA GeForce RTX 4090 GPU.

### 5.2 Agreement with Discrete Operator Eigenfunctions

In Fig. 10, we evaluated the relative differences between our method and the cotangent matrix as well as the analytical solution. For



**ALGORITHM 2:** Training Over a Shape Space

---

```

epoch = 0;
repeat
  g = Sample Geometry Code;
  ▶ The geometry code is an explicit shape parameter (e.g., width)
  or a latent variable (e.g., from an auto-decoder);
  x = Sample Domain  $\Omega$ ;
   $E_d = []$ ;
  for i in range [0, k) do
    Evaluate Network  $\bar{\phi}_i(x, g)$  and gradient  $\frac{\partial \bar{\phi}_i(x, g)}{\partial x}$ ;
     $E_d.append(\frac{1}{\sum \bar{\phi}_i(x, g)^2} \sum (\frac{\partial \bar{\phi}_i(x, g)}{\partial x})^2)$ ;
    ▶ Evaluate Dirichlet energy on normalized gradients;
  end
  index, sortedEigenvalue = sort( $E_d$ )
  for idx in range [0, k) do
    i = index[idx] ▶ Use the order from the sorted index;
    Do projection and optimization same as Algorithm 1;
  end
  epoch = epoch + 1;
until epoch = MaxEpoch;

```

---




Fig. 10. Our method’s error is comparable with the results of Sharp and Crane [2020], who present a discretization-agnostic Laplacian eigenanalysis on a single shape. The different plotlines correspond to the execution of our method with different neural network configurations, with ‘PE’ indicating the maximum frequency of positional encoding and ‘MLP’ indicating the width of the MLP.




Fig. 11. For the 1D domain where the analytic solution is known, our method produces results that closely match the analytical solution.

the first 20 eigenfunctions, our method’s eigenfunctions exhibit an average difference of approximately 6% compared to those from the cotangent matrix, increasing to 20% as the number of eigenfunctions grows to 50. This error is comparable with Sharp and Crane [2020], where both our method and theirs calculate Laplace eigenfunctions without the need for mesh connectivity. For the 1D domain, where




Fig. 12. **Rotation invariance:** eigenmodes produced by our method are consistent with respect to rotation of the domain.

the analytical solution is known, our method consistently produces results that closely match the analytical solution, as shown in Fig. 11.

### 5.3 Rotation Invariance of Neural Eigenfunctions

We retrained the network on rotations of the same shape with different angles. Fig. 12 shows our method provides eigenfunctions under rotational perturbations that are highly consistent with the base unrotated configuration. Specifically, for rotations of 30 and 60 degrees, the average error over the first 30 modes between the rotated and original domains is 0.22% and 0.25%, respectively. For context, variations within the 0.2% – 0.3% range frequently occur due to re-initializing network weights before training, indicating no significant bias for rotations. The 30th eigenfunction from the three training results is visualized in this figure.

### 5.4 Ablation Study on Causal Sorting

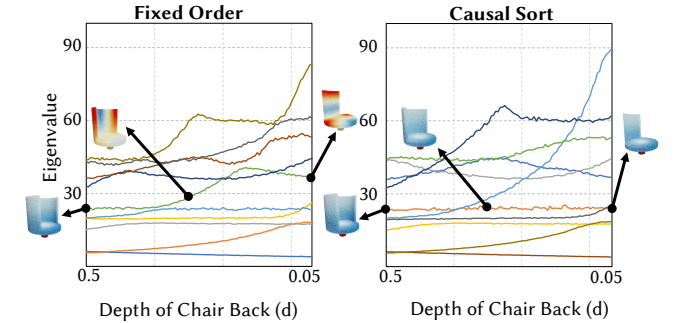


Fig. 13. **Eigenvalue Crossover.** Plot of eigenvalues versus geometry code for the chair shape space parameterized by radius ( $r$ ), height ( $h$ ), and depth ( $d$ ). comparing the naïve fixed-order method against shape-dependent causal sorting. Observe that both plots span a similar set of eigenvalues, but with different topology, with our proposed naïve baseline (*left*) surpassed by our proposed method (*right*), the latter more accurately resolving crossing eigenfunctions. Inherent to resolving an infinite function space with a finite number of eigenfunctions over shape space is that the highest-eigenvalue resolved eigenfunctions may include components of two or more unresolved eigenfunctions, in this case, as evident by the kink in the dark blue eigenfunction 7. For this plot,  $r = 0.1$ ,  $h = 0.8$ , and  $d$  varies.

Figure 13 plots eigenvalues with respect to geometry code for the shape space of Figure 14. Our method resolves eigenvalue crossings, whereas the proposed naïve baseline does not. The exception is

Table 1. Timing statistics. We report the training and inference times for all examples presented in this paper, along with the number of eigenfunctions and the parameters per MLP. The inference time represents the duration required for a single evaluation of all eigenfunctions at 30k cubature points. We assume a uniform, fixed density for all examples.

Example	Training Time	Inference Time(ms)	Eigenfunction Count	Parameters per MLP
Bridge (Fig. 1)	5.0h	6	25	5.2k
Teapot (Fig. 4)	4.2h	8	30	4.7k
Airplane (Fig. 6)	69s	4	15	15k
Round Chair (Fig. 14, Fig. 15)	4.6h	3	10	15k
Sofa (Fig. 14)	4.0h	3	10	15k
Simulation on Wide Range of Shapes (Fig. 19)	35h	5	15	13k
Differentiable Sound (Fig. 23)	6.9h	11	32	4.6k
Sound - Little Star (Fig. 20, Fig. 21)	7.0h	11	25	5.1k
Walking robot (Fig. 24, Fig. 25)	1.5h	8	30	4.7k
Walking Animals (Fig. 22)	19h	7	30	5.0k

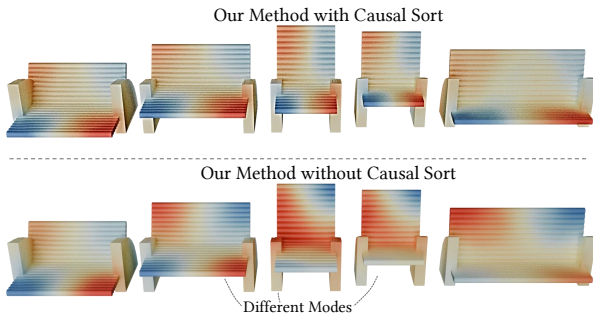


Fig. 14. We visualized an eigenfunction from a simple chair shape space with and without causal sorting as we navigate the space. Our causal sorting approach preserves the mode pattern throughout shape space.

the kink in the “highest” eigenvalue (eigenmode 7), which exposes a fundamental challenge in eigenanalysis over shape space (see §Sec. 7).

### 5.5 Ablation Study on Consistent Modes

The importance of resolving eigenvalue crossings is illuminated by studying the mode of an eigenfunction as it evolves over shape space. We have already shown that resolving crossings leads to consistent mode patterns across shape space. In addition to that, as depicted in Fig. 14 (top row), the absence of shape-dependent causal sorting leads to significant, rapid changes in the eigenfunction modes as the shape evolves over the kinks corresponding to unresolved eigenvalue crossings. By resolving crossings (bottom row), our method obtains a smooth evolution of eigenfunctions.

## 6 RESULTS & APPLICATIONS

### 6.1 Reduced-space Simulation for Shape Families

*Interactive Visualization of Physical Properties.* Our shape-space eigenfunctions enable us to track how physical quantities evolve across the shape space, facilitating fast previews of physical properties. This significantly simplifies the often tedious and iterative shape design process. To demonstrate this, we used the shape space shown in Fig. 1, which is a 5-dimensional shape space that controls

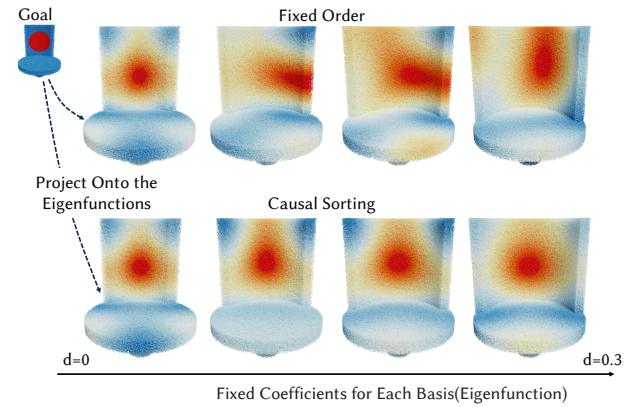


Fig. 15. *Physics Property Transfer.* A target heat distribution is projected onto the basis functions and transferred across the shape space. Causal sorting helps preserve initial distribution of the physical property.




Fig. 16. *Interactive Visualization of Deformation.* The elastic modes of the bridge shape space enable interactive visualization of deformations across various shapes.

the width, length, height, fence size, and fence gap of a bridge. We assume a uniform and fixed density for the bridge and all subsequent examples. As shown in Fig. 16, Fig. 17 and Fig. 18, the deformation and heat distribution initially defined on a simulated bridge can be seamlessly extrapolated to another bridge with a much larger width at runtime.

For instance, the solution  $u(\mathbf{x})$  of a PDE computed over a shape  $\mathbf{g}$  can be transferred to another shape  $\mathbf{g}'$ , providing an excellent warm start  $u'(\mathbf{x})$  for solving the PDE on the new shape: simply




Fig. 17. *Error and Runtime Comparison Between Our Method and Full-Space Simulation* We further evaluated the error in our reduced-space simulation by comparing it with a full-space simulation. Specifically, we simulated two bridges with fixed boundaries under gravitational loading, following the setup in Fig. 16. The full-space simulation runs at 311.0 ms per frame. Our method achieves a runtime of 9.6 ms per frame, yielding a 32.4× speedup, with a mean squared error of 2.4% in the final displacement. The norm of the displacement is shown using color coding.




Fig. 18. *Interactive Visualization of Heat*. Once we solve the PDE for one shape, the solution can be transferred across the shape space, warm-starting the PDE solver over another shape. Without causal sorting, we observe an order of magnitude slower convergence, due to missed crossings. By ensuring consistent mode-shapes via causal sorting, our method enables faster convergence.

project onto the eigenbasis,  $\alpha_i = \int_{\Omega_g} u(\mathbf{x}) \phi_i^g(\mathbf{x}) d\mathbf{x}$ , to obtain a low-dimensional embedding of the solution, then reconstruct on the new shape,  $u'(\mathbf{x}) = \sum_i \alpha_i \phi_i^g(\mathbf{x})$ . As illustrated in Figure 18, causal sort proves an order of magnitude lower error (and correspondingly faster convergence) using this warm start.

The reason for this is that the mode with causal sorting maintains greater consistency across the shape space, resulting in better preservation of the physical property distribution. As demonstrated in Fig. 15, with causal sorting, the symmetric distribution remains intact as the shape changes. In contrast, without causal sorting, the distribution shifts unevenly across the chair’s back.

**Direct Simulation on Generated Shapes.** Recent methods represent shapes through mappings from a latent space to implicit representations, such as occupancy fields [Chen and Zhang 2019a] and signed distance fields [Park et al. 2019]. By interpolating within this latent space, new shapes can be generated, and our method facilitates direct and efficient simulation on these generated models without requiring spatial domain meshing. We leveraged a pretrained model from Chen and Zhang [2019a], which provides a 256-dimensional

shape space, training our approach on specific shape-space coordinates and testing it on additional generated shapes outside the training set.

We implemented the method of Modi et al. [2024], which simulates corotational elasticity dynamics in a discretization-agnostic manner. They leverage the reduced basis proposed by Benckekroun et al. [2023], a linear blend skinning basis with skinning weights set to Laplace eigenfunctions. In their formulation, the displacement in the reduced space simulation is expressed as a weighted sum of affine transformations applied to the shape’s rest positions, with the rigid mode captured by a single scalar weight for the entire shape. These two works build a reduced model for a single shape. With our eigenanalysis, it is now possible to build a reduced model for a continuous family of shapes.

As shown in Figure 19, we built one reduced model to simulate over 250 shapes from 13+ ShapeNet categories, including shapes outside the training set. With our single model, switching between these diverse shapes is as simple as selecting a new shape code, with no need for domain (re)meshing or retraining of neural networks. The resulting simulation effectively captures local deformations, as demonstrated by the deformed chair/desk legs due to collisions with the plane, shown in the right.

To train this example, we first built the shape space, and then performed eigenanalysis. The shape space was trained as an implicit decoder mapping a 256-dimensional coordinate to an occupancy field [Chen and Zhang 2019a]. Eigenanalyzing this shape space yielded the Laplace eigenfunctions that serve as the scalar weight functions for the skinning eigenmodes [Benckekroun et al. 2023]. While this example uses a learned shape space, as we discussed earlier, our approach applies also to hand-designed shape spaces (e.g., parametric CAD models).

Traditionally, simulating a diverse set of objects requires meshing each individual generated shape and performing eigenmode calculations for discrete operators (matrices) on the resulting meshes. Depending on the meshing algorithm, creating a mesh can take anywhere from around 0.02 seconds [Geuzaine and Remacle 2020] to over 20 seconds [Hu et al. 2018], and the subsequent eigenanalysis further adds to the computational burden. For instance, at a resolution of 30k vertices, standard eigendecomposition using SciPy [Virtanen et al. 2020] takes over 6 seconds to compute the first 15 modes. In contrast, our method eliminates the need for meshing and directly computes eigenfunctions, enabling fast simulations across diverse shapes. It calculates eigenfunction for the same number of cubature points in just 0.005 seconds, significantly reducing both computational time and memory requirements.

Compared to prior discretization-agnostic methods such as Modi et al. [2024], our pretraining time for this example is longer at 35 hours. However, our method generalizes to a broader range of shapes. While Modi et al. [2024] require retraining for each individual shape, our approach supports evaluation on a continuous family of shapes, making it more efficient when handling over 200 query shapes. Additionally, our method has a significant advantage in memory consumption, as its memory cost remains constant regardless of the number of shapes, whereas their method incurs a linear increase in memory cost as the number of shapes grows.






Fig. 19. *Reduced-space Simulation*. With just one training session, the model can represent the basis for hundreds of shapes, enabling direct reduced-space simulations. Using the pretrained occupancy field from [Chen and Zhang 2019a], we constructed a shape space mapping a 256-dimensional coordinate to an occupancy field. After training on this shape space, our eigenfunctions can represent the skinning eigenmodes for all shapes within it. This allows for direct simulation of any shape from the model without requiring meshing over the spatial domain or retraining neural networks.

Table 2. Quantitative comparisons on shapes from DiffSound [Jin et al. 2024]. We evaluated the mean absolute error (MAE) by averaging the differences between predictions and ground truth across varying thicknesses. Our method enables pretraining on shape families, which is difficult for DiffSound due to its requirement for remeshing in each optimization step. This allows us to achieve a comparable MAE to DiffSound while significantly reducing query time.

Object	Target Thickness, DiffSound[Jin et al. 2024] / Ours										MAE↓	
	0.3		0.4		0.5		0.6		0.7			
Bunny	0.304	0.299	0.407	0.400	0.508	0.497	0.608	0.597	0.709	0.699	0.0073	0.0016
Armadillo	0.338	0.284	0.456	0.403	0.590	0.501	0.696	0.592	0.730	0.691	0.0623	0.0075
Bulbasaur	0.308	0.298	0.411	0.398	0.512	0.500	0.614	0.597	0.718	0.695	0.0125	0.0023
Squirtle	0.312	0.296	0.416	0.394	0.520	0.495	0.624	0.595	0.718	0.707	0.0177	0.0054




Fig. 20. *Elastic Mode Visualization*. We visualized the final shapes along with the elastic modes, using color coding to represent the norm of the  $\mathbb{R}^3$ -valued eigenfunction. Note that the vibration is concentrated in the bowl of the glass, with almost no vibration occurring in the foot or stem.

## 6.2 Differentiable Modal Sound Synthesis

Our method enables fast shape optimization with eigenvalue-based objectives, such as finding a shape whose vibration modes and frequencies when struck produce a desired sound. We implemented the damped modal sound synthesizer and optimization approach proposed by Jin et al. [2024].

They study the Volumetric Thickness Inference problem, which seeks to estimate the wall thickness  $t$  of a hollow shape based on its




Fig. 21. *Little Star*. Our method enables the optimization of shapes to achieve target sounds. For instance, we optimize a family of wine glasses so that tapping their edges produces the melody of the song 'Little Star.'

sound when struck. We briefly recall their shape space definition: Given one solid shape, define a family of hollow shapes parameterized by wall thickness  $g = t \in [0.15, 0.8]$ , by thresholding the signed distance field (SDF) of the solid shape, such that  $x$  is in the




Fig. 22. *Shape Optimization for Locomotion Across a Broader Range of Characters.* We extended our approach to a more complex shape space derived from interpolations between 12 animals. This enhanced shape space enables the optimization of more complex and diverse shapes.

shape if  $-t < \text{SDF}(\mathbf{x}) < 0$ . As a precomputation, we eigenanalyze the elastic energy Hessian for the shape space.

We implement their query: Given a target eigenvalue (vibration frequency)  $\lambda_i^{\text{gt}}$ , we seek the shape  $\mathbf{g}^*$  that minimizes the difference in eigenvalue:

$$\mathbf{g}^* = \arg \min_{\mathbf{g}} \sum_{i=1}^N \left\| \lambda_i(\mathbf{g}) - \lambda_i^{\text{gt}} \right\|^2, \quad (13)$$

where  $\lambda(\mathbf{g})$  is the elastic energy, because the eigenfunctions are unit norm.

We follow their testing protocol: Select a target shape from four different shapes, select a thickness from  $\{0.3, 0.4, 0.5, 0.6, 0.7\}$ , synthesize the desired (ground truth) sound based on the first  $N = 32$  eigenvalues, optimize Eq. (13) and compare the thickness and spectrogram of the result  $\mathbf{g}^*$ . As depicted in Fig. 23 the optimized shape and its spectrogram align more closely with the target shape and sound after optimization. As shown in Table. 2, our method achieves comparable error to Jin et al. [2024], however, while they report about two hours per query, our approach completes in under three minutes, a  $40\times$  speedup. Our discretization-agnostic approach avoids the costly remeshing required at each optimization step of their approach. Our method requires pretraining on the shape space, unlike Jin et al. [2024]. The precomputation time for our approach is approximately 7 hours per shape space, making it more efficient when handling 4 or more queries.

*Shape Optimization for Pitch.* We perform shape optimization with the goal of designing shapes that produce specific pitches. By optimizing the geometry of wine glasses and tapping them, we generate sounds corresponding to a predefined sequence of notes.

We constructed a 4-dimensional shape space for glasses, where the shape parameters control the radius of the bowl, the radius of the foot, the length of the stem, and a global scaling factor. We optimized the eigenvalue for a chosen low-frequency mode, which is visualized in Fig. 20.

To ensure the chosen mode dominates the resulting sound, we set its damping factor,  $d$ , to 12.5, which is half the damping factor of the other modes; here  $d$  is the damping factor in the model of Jin et al’s Eqs. (9) and (10) [Jin et al. 2024]. We minimize the loss (see Eq. (13)) over  $N = 1$  mode to obtain  $\mathbf{g}^*$ .

To generate an engaging melody, we set the target eigenvalues  $\lambda^{\text{gt}}$  to correspond to the square of frequencies of the notes: ‘do’ (261.63 Hz), ‘re’ (293.66 Hz), ‘mi’ (329.63 Hz), ‘fa’ (349.23 Hz), ‘so’ (392.00 Hz), and ‘la’ (440.00 Hz). Each optimization requires 500 gradient descent steps and takes less than 50 seconds to complete.

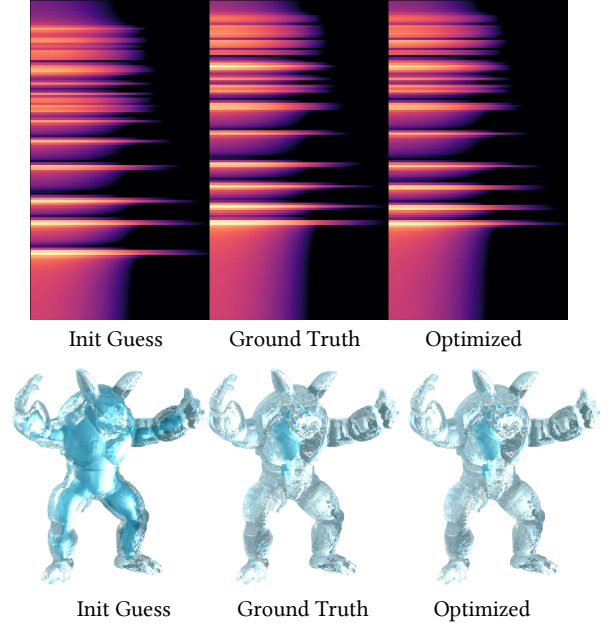


Fig. 23. *Volumetric Thickness Inference.* We optimize the thickness of the shape to produce sounds that better match the target sound. The spectrogram of the resulting sound and the corresponding shapes are visualized, showing that the optimized results are closer to the goal compared to the initial guess.

As shown in Figure 21, we successfully obtain shapes that achieve the target sound through optimization.

### 6.3 Shape Optimization for Locomotion

Inspired by work showing the utility of eigenfunctions as an actuation signal, our method allows us to take locomotion optimization one step further, and ask questions regarding shape optimization for locomotion. Inspired by [Benckroun et al. 2024], we define a controlled actuation force on a character using the eigenfunctions:

$$\mathbf{d} = \mathbf{D}(\mathbf{g})\mathbf{a}(t),$$

where the matrix  $\mathbf{D}(\mathbf{g}) \in \mathbb{R}^{3n \times 3r}$  contains the eigenfunctions of our elastic operator, specifically done by setting the columns  $3i + 1 \dots 3i + 3$  to  $[\phi_i^g, \mathbf{0}_n, \mathbf{0}_n]^T$ ,  $[\mathbf{0}_n, \phi_i^g, \mathbf{0}_n]^T$ ,  $[\mathbf{0}_n, \mathbf{0}_n, \phi_i^g]^T$  respectively, for  $0 \leq i < r$  [von Tycowicz et al. 2013].

To move the character forward, we also added an explicit damping friction force  $f_c \in \mathbb{R}^{3r}$  described by Benchekroun et al. [2024], penalizing the relative velocity between the character and the ground,

$$f_c = \mu J^c |_T \dot{z},$$

where  $\mu$  is the damping coefficient, and the contact Jacobian  $J^c|_T$  maps the reduced velocity  $\dot{z}$  to the full-space contact tangential velocities before projecting them back into the reduced space.

We couple this actuation signal with a differentiable simulation and run it for 200 timesteps and obtain the final full-space displacement field  $u^*(g) = D(g)z^*$ . Then, we optimize for the shape code parameters  $g^*$  by maximizing the norm of  $u^*$ :

$$g^* = \arg \max_g (\|u^*(g)\|^2)$$

We first designed a 3-dimensional shape space for the walking robot. Starting from an initial guess within the shape space, we perform gradient-based optimization. The final shape obtained from our optimization moves faster than the initial guess, as demonstrated in Figure 24. Note that this process is particularly challenging with traditional methods, as it requires computing the gradient of the eigenfunctions (which are used to design the controllers for the actuation forces) with respect to the shape parameters. This is difficult using traditional eigenanalysis based on matrices, and even more so for shapes with algebraic multiplicity.

The change in modes across the shape space can impact the optimization process. Specifically, when modes shift, the gradient with respect to the shape space increases in magnitude, which can complicate the later stages of optimization. In Figure 25(a), we visualize the shape-space norm gradient. Without causal sorting, these abrupt mode changes lead to a larger gradient, making the optimization process more challenging. In contrast, our proposed causal sort ensures that the modes transition smoothly across the shape space, as shown on the right, reducing this issue and enabling a more stable and efficient optimization process.

To demonstrate the impact of omitting causal sorting on optimization convergence, we performed shape optimization with and without causal sorting, using the same initial guess: width = 0.7, length = 0.7, and leg thickness = 0.3, optimized for 30 steps. As shown on the right of Fig. 25(a), the final solution is better with causal sorting. To further demonstrate this, we optimized the walking robot's shape with different initial guesses in the shape space and performed shape optimization, as shown in Figure 25(b). Without causal sorting, two distinct final solutions are found, each with different movement patterns (in one, the robot walks sideways, and in the other, it walks vertically), with one solution corresponding to a local minimum. However, with causal sorting, only one consistent solution is obtained, with a stable walking pattern.

Our optimization approach can be extended to operate within a more intricate shape space derived from signed distance field (SDF)-based interpolation of 12 distinct animal shapes. This interpolation constructs a continuous 12-dimensional shape space, enabling smooth transitions between different animal characters and offering a rich variety of geometries for optimization. As illustrated in Figure 22, our method effectively optimizes these animal characters within the shape space.






Fig. 24. *Shape Optimization for Walking Robot.* We optimize the design of a walking robot for a fixed controller, resulting in a robot that walks 18× faster than the initial design.



(a) Visualization of gradient norm scaling in shape space. As shown on the left, fixed-order eigenfunction calculations result in mode changes at crossovers, leading to large gradient norms and making optimization more challenging. By contrast, the middle shows how our method, with causal sorting, achieves smoother eigenfunction transitions across the shape space, resulting in smaller gradient norms. The consistent eigenfunction can lead to a better final solution of the optimization, as shown on the right.



(b) Starting from different initial guesses within the shape space of a walking robot, we observed two distinct final solutions with differing movement patterns when causal sorting was not applied. This occurs because the actuation force relies on eigenfunctions, and mode changes in eigenfunctions lead to variations in movement. By using causal sorting, we achieve consistent final solutions.

Fig. 25. The mode changes at crossovers significantly increase the gradient norm in shape space, as shown in 25(a). This creates an unsmooth gradient landscape, making optimization more challenging. Without causal sorting, as illustrated in 25(b), the process resulted in two different final solutions. In contrast, causal sorting encourages consistent final solutions.

## 7 DISCUSSION AND FUTURE WORK

With these promising results, there remain areas for improvement and expansion.

For a single shape represented by a triangle mesh, eigenanalysis of the Cotan matrix is faster than stochastic gradient descent optimization. Our method's comparative advantage therefore lies in its generalization to shape spaces and general shape representations.

As the number of eigenfunctions increases, so does their spatial frequency, and accurate training of higher frequencies is a known challenge for neural fields. Since our mathematical framework and optimization approach are not specific to neural fields, and it would






Fig. 26. *Gradient Indicators of Higher-Frequency Crossovers.* The norm of shape space gradients (the gradient with respect to the geometry code  $|\frac{\partial \phi(\mathbf{x}, \mathbf{g})}{\partial \mathbf{g}}|^2$ ) shows a significant increase at the crossover, therefore can be used as an indicator of the crossover from higher-frequency eigenfunctions.

be interesting to apply the variational formulation, causal gradient filtering, and causal sorting to the optimization of eigenfunctions in other representations, including representations that easily extend to higher frequencies.

We have currently limited our scope to Neumann boundary conditions. Eigenfunctions subject to Dirichlet or Robin boundary conditions could be handled in the future by using PINN techniques for enforcing essential boundary conditions, such as blending with a weighting function that ensures constraint satisfaction by construction; alternatively, by introducing additional penalization terms that contribute to the loss function [Berrone et al. 2023]. In the latter case, causal gradient filtering could be leveraged to ensure that the Dirichlet penalty dominates the variational energy corresponding to the PDE operator.

Our implementation of stochastic cubature presently gives uniform weight to all samples. This is appropriate if the cubature is drawn from a uniform distribution. However, certain domain shapes may benefit from nonuniform cubature such as importance sampling. In that case, the number of samples per shape would potentially differ, potentially biasing the training to emphasize some regions of shape space over others, further necessitating a normalization of cubature weights per shape.

Truncating an infinite-dimensional space with a finite number of subspaces inherently leads to missing information. In our case, this limitation can result in crossovers from higher-frequency eigenfunctions that cannot be fully captured due to the finite number of eigenfunctions used (recall eigenfunction 7 of Figure 13). While with a finite budget we do not offer a fundamental way around this, we can identify missed crossings by evaluating shape space gradients  $\frac{\partial \phi(\mathbf{g}, \mathbf{x})}{\partial \mathbf{g}}$ . As shown in Figure 26, there is a significant increase in the gradient norm at points where there is a crossover from a higher-frequency eigenfunction, which serves as a reliable indicator of these crossovers. In applications, then, the few eigenfunctions with missed crossings may be omitted, or, less aggressively, they may be included with knowledge of their piecewise nature. For example, in our application to transfer of PDE solutions across shapes, such an eigenfunction would be included if the two shapes lie on the same side of the crossing.

In the future, it would be interesting to explore continuation methods [Allgower and Georg 2003] that seek to “build out” the remainder of the incomplete eigenfunction curves without arbitrarily increasing the total eigenfunction budget.

The gradients with respect to the geometry code may have further uses. For instance, by incorporating a shape space gradient norm into the training objective, it may be possible to achieve even smoother evolving functions among shape families.

Shape-dependent representations of eigenfunctions appear to be a promising and versatile tool for working with eigenfunctions in continuously-parameterized shape spaces, broadening the scope of reduced physical models and accelerating inverse design problems. We believe that they also open the door to other applications in geometry processing (e.g., spectral processing methods), computer vision (e.g., shape correspondence), computational physics (e.g., fluid and deformable solid coupling), and indeed in all fields that leverage eigenanalysis of PDE operators.

## ACKNOWLEDGMENTS

We thank Max Wardetzky for valuable discussions about the Min-Max Theorem and the spectral properties of the Laplace operator. We are also grateful to Mengfei Liu for narrating the video and to Vishnu Nittoor for proofreading the paper. We thank Abhishek Madan, David Levin, Jonathan Panuelos, Silvia Sellán, and Zhecheng Wang for their helpful feedback on the title. We would also like to thank our lab system administrator, John Hancock, and our financial officer, Xuan Dam, for their invaluable administrative support in making this research possible. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), grant RGPIN-2021-03733.

## REFERENCES

- Noam Aigerman, Kunal Gupta, Vladimir G Kim, Siddhartha Chaudhuri, Jun Saito, and Thibault Groueix. 2022. Neural Jacobian Fields: Learning Intrinsic Mappings of Arbitrary Meshes. *arXiv preprint arXiv:2205.02904* (2022).
- Eugene L. Allgower and Kurt Georg. 2003. *Introduction to Numerical Continuation Methods*. Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9780898719154> arXiv:https://epubs.siam.org/doi/pdf/10.1137/1.9780898719154
- Peter Arbenz, Ulrich L. Hetmaniuk, Richard B. Lehoucq, and Raymond S. Tuminaro. 2005. A comparison of eigensolvers for large-scale 3D modal analysis using AMG-preconditioned iterative methods. *Internat. J. Numer. Methods Engrg.* 64, 2 (2005), 204–236. <https://doi.org/10.1002/nme.1365> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.1365
- Jernej Barbič and Doug L. James. 2005. Real-Time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph.* 24, 3 (July 2005), 982–990. <https://doi.org/10.1145/1073204.1073300>
- Mikhail Belkin, Jian Sun, and Yusu Wang. 2008. Discrete laplace operator on meshed surfaces. In *Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry* (College Park, MD, USA) (SCG '08). Association for Computing Machinery, New York, NY, USA, 278–287. <https://doi.org/10.1145/1377676.1377725>
- Mikhail Belkin, Jian Sun, and Yusu Wang. 2009. Constructing Laplace operator from point clouds in Rd. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms* (New York, New York) (SODA '09). Society for Industrial and Applied Mathematics, USA, 1031–1040.
- Otman Benckekroun, Kaixiang Xie, Hsueh-Ti Derek Liu, Eitan Grinspun, Sheldon Andrews, and Victor Zordan. 2024. Actuators A La Mode: Modal Actuators for Soft Body Locomotion. In *SIGGRAPH Asia 2024 Conference Papers* (SA '24). Association for Computing Machinery, New York, NY, USA, Article 53, 10 pages. <https://doi.org/10.1145/3680528.3687638>
- Otman Benckekroun, 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 (jul 2023), 21 pages. <https://doi.org/10.1145/3592404>

- 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 (2023), e18820. <https://doi.org/10.1016/j.heliyon.2023.e18820>
- Gaurav Bharaj, David I. W. Levin, James Tompkin, Yun Fei, Hanspeter Pfister, Wojciech Matusik, and Changxi Zheng. 2015. Computational design of metallophone contact sounds. *ACM Trans. Graph.* 34, 6, Article 223 (Nov. 2015), 13 pages. <https://doi.org/10.1145/2816795.2818108>
- Alexander I. Bobenko and Boris A. Springborn. 2005. A Discrete Laplace–Beltrami Operator for Simplicial Surfaces. *Discrete & Computational Geometry* 38 (2005), 740–756. <https://api.semanticscholar.org/CorpusID:708910>
- Yue Chang, Peter Yichen Chen, Zhecheng Wang, Maurizio M. Chiamonte, 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. <https://doi.org/10.1145/3610548.3618158>
- Honglin Chen, Rundui Wu, Eitan Grinspun, Changxi Zheng, and Peter Yichen Chen. 2022. Implicit Neural Spatial Representations for Time-dependent PDEs. *arXiv preprint arXiv:2210.00124* (2022).
- Peter Yichen Chen, Jinxu Xiang, Dong Heon Cho, Yue Chang, G A Pershing, Henrique Teles Maia, Maurizio M Chiamonte, Kevin Thomas Carlberg, and Eitan Grinspun. 2023. CROM: Continuous Reduced-Order Modeling of PDEs Using Implicit Neural Representations. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=FUORztG8Og>
- Zhiqin Chen and Hao Zhang. 2019a. Learning Implicit Fields for Generative Shape Modeling. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019).
- Zhiqin Chen and Hao Zhang. 2019b. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5939–5948.
- Qiaodong Cui, Pradeep Sen, and Theodore Kim. 2018. Scalable laplacian eigenfluids. *ACM Trans. Graph.* 37, 4, Article 87 (July 2018), 12 pages. <https://doi.org/10.1145/3197517.3201352>
- Tyler De Witt, Christian Lessig, and Eugene Fiume. 2012. Fluid simulation using Laplacian eigenfunctions. *ACM Trans. Graph.* 31, 1, Article 10 (Feb. 2012), 11 pages. <https://doi.org/10.1145/2077341.2077351>
- Yitong Deng, Hong-Xing Yu, Diyang Zhang, Jiajun Wu, and Bo Zhu. 2023. Fluid Simulation on Neural Flow Maps. *ACM Trans. Graph.* 42, 6, Article 248 (2023).
- Ana Dodik, Oded Stein, Vincent Sitzmann, and Justin Solomon. 2023. Variational Barycentric Coordinates. *ACM Transactions on Graphics* (2023). <https://doi.org/10.1145/3618403>
- Jed A. Duersch, Meiyue Shao, Chao Yang, and Ming Gu. 2018. A Robust and Efficient Implementation of LOBPCG. *SIAM Journal on Scientific Computing* 40, 5 (2018), C655–C676. <https://doi.org/10.1137/17M1129830>
- Matthew Fisher, Boris Springborn, Alexander I. Bobenko, and Peter Schroder. 2006. An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) (SIGGRAPH '06). Association for Computing Machinery, New York, NY, USA, 69–74. <https://doi.org/10.1145/1185657.1185668>
- Christophe Geuzaine and Jean-Francois Remacle. 2020. *Gmsh*. <http://gmsh.info/>
- A. Guezic, F. Bossen, G. Taubin, and C. Silva. 1999. Efficient compression of non-manifold polygonal meshes. In *Proceedings Visualization '99 (Cat. No.99CB37067)*. 73–512. <https://doi.org/10.1109/VISUAL.1999.809870>
- Klaus Hildebrandt, Christian Schulz, Christoph Von Tycowicz, and Konrad Polthier. 2011. Interactive surface modeling using modal analysis. *ACM Trans. Graph.* 30, 5, Article 119 (Oct. 2011), 11 pages. <https://doi.org/10.1145/2019627.2019638>
- 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. <https://doi.org/10.1145/3197517.3201353>
- Alec Jacobson, Zhigang Deng, Ladislav Kavan, and JP Lewis. 2014. Skinning: Real-time Shape Deformation. In *ACM SIGGRAPH 2014 Courses*.
- Doug L. James and Dinesh K. Pai. 2002. DyRT: dynamic response textures for real time deformation simulation with graphics hardware. *ACM Trans. Graph.* 21, 3 (July 2002), 582–585. <https://doi.org/10.1145/566654.566621>
- Xutong Jin, Chenxi Xu, Ruohan Gao, Jiajun Wu, Guoping Wang, and Sheng Li. 2024. DiffSound: Differentiable Modal Sound Rendering and Inverse Rendering for Diverse Inference Tasks. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 128, 12 pages. <https://doi.org/10.1145/3641519.3657493>
- Mark Kac. 1966. Can One Hear the Shape of a Drum? *The American Mathematical Monthly* 73, 4 (1966), 1–23. <http://www.jstor.org/stable/2313748>
- Tosio Kato. 1980. *Perturbation theory for linear operators*. Springer Berlin, Heidelberg. see page 124 of 623 pages. <https://doi.org/10.1007/978-3-642-66282-9>
- Theodore Kim and David Eberle. 2022. Dynamic deformables: implementation and production practicalities (now with code!). In *ACM SIGGRAPH 2022 Courses* (Vancouver, British Columbia, Canada) (SIGGRAPH '22). Association for Computing Machinery, New York, NY, USA, Article 7, 259 pages. <https://doi.org/10.1145/3532720.3535628>
- Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs.LG]* <https://arxiv.org/abs/1412.6980>
- P.G. Kry, L. Reveret, F. Faure, and M.-P. Cani. 2009. Modal Locomotion: Animating Virtual Characters with Natural Vibrations. *Computer Graphics Forum* 28, 2 (2009), 289–298. <https://doi.org/10.1111/j.1467-8659.2009.01368.x>
- Bruno Lévy and Hao (Richard) Zhang. 2010. Spectral mesh processing. In *ACM SIGGRAPH 2010 Courses* (Los Angeles, California) (SIGGRAPH '10). Association for Computing Machinery, New York, NY, USA, Article 8, 312 pages. <https://doi.org/10.1145/1837101.1837109>
- Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2023. Neural Metamaterial Networks for Nonlinear Material Design. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–13.
- Hsueh-Ti Derek Liu, Maneesh Agrawala, Cem Yuksel, Tim Omernick, Vinith Misra, Stefano Corazza, Morgan McGuire, and Victor Zordan. 2024. A Unified Differentiable Boolean Operator with Fuzzy Logic. In *ACM SIGGRAPH 2024 Conference Papers* (Denver, CO, USA) (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, Article 109, 9 pages. <https://doi.org/10.1145/3641519.3657484>
- Yong-Jin Liu, Dian Fan, Chun-Xu Xu, and Ying He. 2017. Constructing Intrinsic Delaunay Triangulations from the Dual of Geodesic Voronoi Diagrams. *ACM Trans. Graph.* 36, 2, Article 15 (apr 2017), 15 pages. <https://doi.org/10.1145/2999532>
- 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 Transactions on Graphics (TOG)* 40, 4 (2021), 132.
- Diana Mateus, Radu Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. 2008. Articulated shape matching using Laplacian eigenfunctions and unsupervised point registration. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. 1–8. <https://doi.org/10.1109/CVPR.2008.4587538>
- Ishit Mehta, Mamohan Chandraker, and Ravi Ramamoorthi. 2022. A level set theory for neural implicit evolution under explicit flows. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*. Springer, 711–729.
- S. Melzi, E. Rodolà, U. Castellani, and M. M. Bronstein. 2018. Localized Manifold Harmonics for Spectral Shape Analysis. *Computer Graphics Forum* 37, 6 (2018), 20–34. <https://doi.org/10.1111/cgf.13309>
- Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4460–4470.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2020. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*. Springer, 405–421.
- Vismay Modi, Nicholar Sharp, Or Perel, David I. W. Levin, and Shinjiro Sueda. 2024. Simplicits: Mesh-Free, Geometry-Agnostic, Elastic Simulation. *arXiv preprint* (2024).
- Ahmad Nasikun and Klaus Hildebrandt. 2022. The Hierarchical Subspace Iteration Method for Laplace–Beltrami Eigenproblems. *ACM Trans. Graph.* 41, 2, Article 17 (jan 2022), 14 pages. <https://doi.org/10.1145/3495208>
- Rubens F. Nunes, Joaquim B. Cavalcante-Neto, Creto A. Vidal, Paul G. Kry, and Victor B. Zordan. 2012. Using natural vibrations to guide control for locomotion. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Costa Mesa, California) (I3D '12). Association for Computing Machinery, New York, NY, USA, 87–94. <https://doi.org/10.1145/2159616.2159631>
- James F. O'Brien, Chen Shen, and Christine M. Gatchalian. 2002. Synthesizing sounds from rigid-body simulations. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (San Antonio, Texas) (SCA '02). Association for Computing Machinery, New York, NY, USA, 175–181. <https://doi.org/10.1145/545261.545290>
- Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.* 31, 4, Article 30 (jul 2012), 11 pages. <https://doi.org/10.1145/2185520.2185526>
- Bo Pang, Zhongtian Zheng, Yilong Li, Guoping Wang, and Peng-Shuai Wang. 2024. Neural Laplacian Operator for 3D Point Clouds. *ACM Trans. Graph.* 43, 6, Article 239 (Nov. 2024), 14 pages. <https://doi.org/10.1145/3687901>
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Maziar Raissi, Paris Perdikaris, and George 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.

- Arianna Rampini, Irene Tallini, Maks Ovsjanikov, Alexander M. Bronstein, and Emanuele Rodolà. 2019. Correspondence-Free Region Localization for Partial Shape Similarity via Hamiltonian Spectrum Alignment. *CoRR* abs/1906.06226 (2019). arXiv:1906.06226 <http://arxiv.org/abs/1906.06226>
- C. C. Rankin and F. A. Brogan. 1986. An Element Independent Corotational Procedure for the Treatment of Large Rotations. *Journal of Pressure Vessel Technology* 108, 2 (05 1986), 165–174. <https://doi.org/10.1115/1.3264765> arXiv:[https://asmedigitalcollection.asme.org/pressurevesseltech/article-pdf/108/2/165/5534963/165\\_1.pdf](https://asmedigitalcollection.asme.org/pressurevesseltech/article-pdf/108/2/165/5534963/165_1.pdf)
- Adriana Schulz, Harrison Wang, Eitan Grinspun, Justin Solomon, and Wojciech Matusik. 2018. Interactive exploration of design trade-offs. *ACM Trans. Graph.* 37, 4, Article 131 (July 2018), 14 pages. <https://doi.org/10.1145/3197517.3201385>
- Adriana Schulz, Jie Xu, Bo Zhu, Changxi Zheng, Eitan Grinspun, and Wojciech Matusik. 2017. Interactive design space exploration and optimization for CAD models. *ACM Trans. Graph.* 36, 4, Article 157 (July 2017), 14 pages. <https://doi.org/10.1145/3072959.3073688>
- Silvia Sellán, Herng Yi Cheng, Yuming Ma, Mitchell Dembowski, and Alec Jacobson. 2019. Solid Geometry Processing on Deconstructed Domains. *Computer Graphics Forum* (2019). <https://doi.org/10.1111/cgf.13592>
- Avinash Sharma and Radu Horaud. 2010. Shape matching based on diffusion embedding and on mutual isometric consistency. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*. 29–36. <https://doi.org/10.1109/CVPRW.2010.5543278>
- Nicholas Sharp, Souhaib Attai, Keenan Crane, and Maks Ovsjanikov. 2022. Diffusion-Net: Discretization Agnostic Learning on Surfaces. *ACM Trans. Graph.* 41, 3, Article 27 (mar 2022), 16 pages. <https://doi.org/10.1145/3507905>
- Nicholas Sharp and Keenan Crane. 2020. A Laplacian for Nonmanifold Triangle Meshes. *Computer Graphics Forum (SGP)* 39, 5 (2020).
- Kenji Shimada and David C. Gossard. 1995. Bubble mesh: automated triangular meshing of non-manifold geometry by sphere packing. In *Proceedings of the Third ACM Symposium on Solid Modeling and Applications* (Salt Lake City, Utah, USA) (SMA '95). Association for Computing Machinery, New York, NY, USA, 409–419. <https://doi.org/10.1145/218013.218095>
- Eftychios Sifakis and Jernej Barbic. 2012. FEM simulation of 3D deformable solids: a practitioner's guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses* (Los Angeles, California) (SIGGRAPH '12). Association for Computing Machinery, New York, NY, USA, Article 20, 50 pages. <https://doi.org/10.1145/2343483.2343501>
- Dmitriy Smirnov and Justin Solomon. 2021. HodgeNet: Learning Spectral Geometry on Triangle Meshes. *SIGGRAPH* (2021).
- Breannan Smith, Fernando De Goes, and Theodore Kim. 2018. Stable Neo-Hookean Flesh Simulation. *ACM Trans. Graph.* 37, 2, Article 12 (March 2018), 15 pages. <https://doi.org/10.1145/3180491>
- Justin Solomon, Fernando de Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. 2015. Convolutional wasserstein distances: efficient optimal transportation on geometric domains. *ACM Trans. Graph.* 34, 4, Article 66 (July 2015), 11 pages. <https://doi.org/10.1145/2766963>
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (Barcelona, Spain) (SGP '07). Eurographics Association, Goslar, DEU, 109–116.
- Mingze Sun, Shiwei Mao, Puhua Jiang, Maks Ovsjanikov, and Ruqi Huang. 2023. Spatially and Spectrally Consistent Deep Functional Maps. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. 14451–14461. <https://doi.org/10.1109/ICCV51070.2023.01333>
- B. Vallet and B. Lévy. 2008. Spectral Geometry Processing with Manifold Harmonics. *Computer Graphics Forum* 27, 2 (2008), 251–260. <https://doi.org/10.1111/j.1467-8659.2008.01122.x> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8659.2008.01122.x>
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Christoph von Tycowicz, Christian Schulz, Hans-Peter Seidel, and Klaus Hildebrandt. 2013. An efficient construction of reduced deformable objects. *ACM Trans. Graph.* 32, 6, Article 213 (Nov. 2013), 10 pages. <https://doi.org/10.1145/2508363.2508392>
- Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *Advances in Neural Information Processing Systems* 34 (2021), 27171–27183.
- Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. 2007. Discrete laplace operators: no free lunch. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (<conf-loc>, <city>Barcelona</city>, <country>Spain</country>, </conf-loc>) (SGP '07). Eurographics Association, Goslar, DEU, 33–37.
- Romy Williamson and Niloy J. Mitra. 2024. Neural Geometry Processing via Spherical Neural Surfaces. arXiv:2407.07755 [cs.GR] <https://arxiv.org/abs/2407.07755>
- Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. 2021. Neural Fields in Visual Computing and Beyond. *arXiv preprint arXiv:2111.11426* (2021).
- Guandao Yang, Serge Belongie, Bharath Hariharan, and Vladlen Koltun. 2021. Geometry Processing with Neural Fields. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. 2020. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems* 33 (2020), 2492–2502.
- Haocheng Yuan, Adrien Bousseau, Hao Pan, Quancheng Zhang, Niloy J. Mitra, and Changjian Li. 2024. DiffCSG: Differentiable CSG via Rasterization. In *SIGGRAPH Asia 2024 Conference Papers (SA '24)*. Association for Computing Machinery, New York, NY, USA, Article 9, 10 pages. <https://doi.org/10.1145/3680528.3687608>
- Jonas Zehnder, Yue Li, Stelian Coros, and Bernhard Thomaszewski. 2021. Ntopo: Mesh-free topology optimization using implicit neural representations. *Advances in Neural Information Processing Systems* 34 (2021), 10368–10381.