Learning Eigenstructures of Unstructured Data Manifolds

Roy Velich[1*] Arkadi Piven[1] David Bensaïd[1] Daniel Cremers[2,3] Thomas Dagès[1,2,3] Ron Kimmel[1]

In memory of Haïm Brezis (1944 – 2024).

Abstract

We introduce a novel framework that directly learns a spectral basis for shape and manifold analysis from unstructured data, eliminating the need for traditional operator selection, discretization, and eigensolvers. Grounded in optimal-approximation theory, we train a network to decompose an implicit approximation operator by minimizing the reconstruction error in the learned basis over a chosen distribution of probe functions. For suitable distributions, they can be seen as an approximation of the Laplacian operator and its eigendecomposition, which are fundamental in geometry processing. Furthermore, our method recovers in a unified manner not only the spectral basis, but also the implicit metric's sampling density and the eigenvalues of the underlying operator. Notably, our unsupervised method makes no assumption on the data manifold, such as meshing or manifold dimensionality, allowing it to scale to arbitrary datasets of any dimension. On point clouds lying on surfaces in 3D and high-dimensional image manifolds, our approach yields meaningful spectral bases, that can resemble those of the Laplacian, without explicit construction of an operator. By replacing the traditional operator selection, construction, and eigendecomposition with a learning-based approach, our framework offers a principled, data-driven alternative to conventional pipelines. This opens new possibilities in geometry processing for unstructured data, particularly in high-dimensional spaces.

1. Introduction

Differential geometry is at the core of shape and manifold analysis. By defining operators, one can encode the underlying geometry and compute intrinsic and extrinsic quantities, from geodesics and Gaussian curvature to mean curvature and normals. The specific choice of operator depends on the task at hand. For example, the Laplace-Beltrami operator (LBO) captures intrinsic geometric properties of the underlying manifold and governs heat diffusion. Hamiltonian operators, appearing in the famous Schrödinger equation, model wave motion and allow potential-modulated particle dynamics. The biharmonic operator encodes fourth-order behavior and is natural for smooth interpolation and shape-aware distances.

In practice, these operators are primarily used for their spectral decompositions, which serve as the computational foundation for geometry processing. For instance, LBO eigenvalues are intrinsic shape signatures [50, 51, 62, 102, 107], while LBO eigenfunctions power heat- and wave-kernel methods for smoothing and propagation [35, 55], define feature descriptors [7, 113, 123], and support matching via functional maps [93]. As such, these eigendecompositions are routinely required in geometry processing.

The standard pipeline for computing operator eigendecompositions on discrete data begins by choosing a discrete approximation that preserves key geometric properties such as positive semi-definiteness (PSD), symmetry, consistency, or weak formulations [100, 118, 135]. Then, explicitly construct a matrix acting as the discretized linear operator and solve a generalized eigenvalue problem using a classical numerical solver. Although effective on data sampled from surfaces in 3D, this pipeline designed for intrinsically two-dimensional manifolds does not scale gracefully to high-dimensional manifolds [118].

Here, we propose a different route. Rather than targeting a specific operator, explicitly discretizing and eigendecomposing it, we learn the spectral decomposition directly from the data. The learned spectral decomposition implicitly corresponds to an operator that is reconstructible a posteriori from the basis and associated eigenvalues. Building on optimal-approximation theory [3, 4, 24], we find the eigenbasis of an optimal-approximation operator, that minimizes the reconstruction error over a class of probe functions. Different probe classes induce different metrics and thus different operators. In particular cases, the learned operator approximates the LBO, but the proposed framework generalizes to more operators. While only the eigenbasis is learned, the eigenvalues are given directly as a by-product of the worst-case reconstruction error. This yields a learning methodology for spectral analysis that is data-driven and

*Corresponding author: royve@campus.technion.ac.il
[1]Technion - Israel Institute of Technology
[2]Technical University of Munich
[3]Munich Center for Machine Learning

avoids the need for explicit operator construction.

To summarize, our contributions are as follows:

- We learn spectral bases directly from point clouds, bypassing the explicit choice of an operator, its construction, and classical numerical eigensolvers. We ground the method in an optimal-basis representation theory and use it to propose a unique learning objective.
- We demonstrate on surfaces in 3D that the proposed approach provides eigenstructures and estimated metrics performing competitively with oracle discrete LBO baselines, while bypassing explicit operator construction.
- We show that the method scales from surfaces to higher-dimensional data manifolds, enabling scalable manifold learning where mesh-based pipelines are inapplicable and common graph-based ones are unreliable.
- We provide our code at https://github.com/royvelich/learning-eigenstructures.

2. Related Works

Our framework lies at the intersection of operator discretization and neural methods for eigenproblems.

Discrete operators. In geometry processing, discrete operators supply via spectral decomposition [28] the orthonormal basis that we actually use to process signals on meshes and graphs. For instance, filtering [15, 71], diffusion [35, 119], and convolution [27, 72] are usually implemented in this spectral basis. Notably, the community rarely applies the operator matrix directly. In practice, only the first eigenvectors are kept, which amounts to projecting a signal onto the lowest-energy subspace defined by the operator. This acts as a low-pass filter where only the (smooth) low-frequency components are retained. Because the operator defines what "energy", "smoothness", and "frequency" mean, different operators emphasize different properties of the manifold. This motivates the importance of choosing operators wisely in the field of geometry processing.

Perhaps the most important example, the LBO [12, 109] is ubiquitous in geometry processing [20–22, 26, 30, 123, 139, 145]. Discretizing the weak form gives the cotangent Laplacian [87, 100] – the reference in LBO discretization, but alternatives exist [32, 142]. Cotangent weights need well-triangulated meshes, i.e. Delaunay, to avoid negative edges violating the maximum principle [135]. They cannot be directly applied to unstructured data, like point clouds, requiring first wise meshings, e.g. with the tufted cover of the Robust Laplacian [118] enabling flips to Delaunay triangulations. While it extends to thin 3D volumes, this method does not scale beyond surfaces. For higher-dimensional manifolds, graph Laplacians [33, 65] are used, but they depend strongly on connectivity, e.g. local density, unlike the targeted smooth LBO [73]. Constructing a reliable high-dimensional generalization of the LBO is a challenge.

Learning Laplacians has become a popular line of research as the amount of data increases. Some works suggest learning the operator action [103] but lack eigendecompositions. Others learn eigenvalues but not eigenvectors [5]. Most learn a Laplacian matrix from data rather than heuristics [95, 141], but they still explicitly approximate the operator by assembling mass and stiffness matrices and then apply eigensolvers sensitive to these approximations. As they rely on triangle-based discretizations, they target surfaces and do not extend to higher-dimensionality. For implicit neural representations, the Laplacian can be built via Rayleigh quotients directly [140]. However, this assumes a two-dimensional surface and the Euclidean ambient metric, removing adaptivity to other metrics or dimensionalities. On the theoretical front, Laplacian estimation has progressed [29, 88, 98, 128], yet, translating these insights into practice remains largely unexplored.

Additionally, other discrete operators provide different insights. Changing the metric, via scaling (scale-invariant LBO [2, 20, 54, 117]), anisotropy [6, 17, 18, 106, 109], or asymmetry [9, 37–39, 91, 136], changes the LBO and its approximation. Beyond LBOs, both intrinsic [14, 31, 105] and extrinsic [134] alternatives are understudied.

Recent works have also focused on making spectral methods more robust to the arbitrary choice of eigenbasis for a given operator, either by defining spectral processing on eigenspaces [80] or by adaptive canonicalization [77]. However, these methods still begin from a predefined operator and its spectral decomposition.

Eigenproblems and neural networks. Recent efforts apply neural networks to operator eigenvalue problems, typically using variational or dynamical formulations. In [111], networks are trained to represent individual Laplacian eigenfunctions as continuous functions on parametrized domains, based on the Rayleigh quotient objective and finding eigenfunctions sequentially via Gram-Schmidt orthogonalization, an idea that can be extended to learning multiple eigenfunctions simultaneously [13]. However, this approach is limited to simplistic impractical domains, e.g. Euclidean. Another approach reformulates the eigenvalue problem as a fixed point of the operator's semi-group flow, training networks via forward-backward stochastic differential equations to handle high-dimensional problems, up to ten dimensions [56]. These methods fundamentally differ from the proposed approach: (1) they require explicit domain parameterization with global coordinates, (2) they assume flat geometry by taking Euclidean gradients via automatic differentiation, and (3) they must be trained from scratch for each specific domain, with no mechanism for generalization across different geometries. Thus, they cannot directly handle curved manifolds sampled as point clouds without manually specifying metric tensors

and managing coordinate singularities.

3. Method

3.1. Foundations

Denote $\mathcal{M} \subset \mathbb{R}^d$ as a high-dimensional manifold, typically a curved low-dimensional surface embedded in $\mathbb{R}^d$, with functions $f \in \mathcal{F}(\mathcal{M}, \mathbb{R})$ and linear operators $L : \mathcal{F}(\mathcal{M}, \mathbb{R}) \to \mathcal{F}(\mathcal{M}, \mathbb{R})$ acting on them. When the manifold is discretely sampled by n points, scalar functions can be represented by vectors $f \in \mathbb{R}^n$ and operators as matrices $L \in \mathbb{R}^{n \times n}$. Here, we denote by $\langle \cdot, \cdot \rangle_L$ the L-weighted inner product $\langle f, g \rangle_L = f^\top L g$ for any *symmetric positive definite* (SPD) matrix L, with its induced norm $\|f\|_L = \sqrt{\langle f, f \rangle_L} = \sqrt{f^\top L f}$.

Optimal-approximation theory. Given a class $\mathcal{C}$ of signals on the discrete domain $\mathcal{M}$, a fundamental question arises: What is the optimal orthonormal basis for representing functions $f \in \mathcal{C}$? The answer depends on how we characterize the class of signals. A key insight [4, 24] is that many practical signal classes can be characterized by constraints of the form $\|f\|_L \leq 1$, where L is an SPD operator encoding prior knowledge about the signals. Remarkably, for any such constraint class, the optimal basis is given by the eigenvectors of L itself.

Theorem 3.1 (Min-Max Optimality [4, Theorem 2.1])**.** *Given a symmetric positive definite operator L with eigenvalues $0 < \lambda_1 \leq \cdots \leq \lambda_n$ and eigenvectors $e_1, \ldots, e_n$, the min-max approximation error*

$$\alpha_k \quad = \quad \min_{b=(b_1,\ldots,b_n)} \max_{\|f\|_L \leq 1} \left\| f - \sum_{i=1}^{k} \langle f, b_i \rangle b_i \right\|^2 ,$$

where b ranges over orthonormal bases, is minimized by the first k eigenvectors of L, i.e. $b_i = e_i \ \forall i \leq k$, with optimal value $\lambda_{k+1} = \frac{1}{\alpha_k}$. For simple spectrum, $\lambda_i < \lambda_{i+1} \ \forall i$, the basis b that is a solution for every k is unique up to signs.

The key insight is that for any SPD operator $L \in \mathbb{R}^{n \times n}$, the optimal orthonormal basis for the progressive k-term approximation, i.e. for every k, is uniquely given by the eigenvectors of L ordered by increasing eigenvalues. Crucially, the worst approximation error using the first k eigenvectors is $\frac{1}{\lambda_{k+1}}$, meaning that the $(k+1)$-th eigenvalue is inversely proportional to the worst maximum reconstruction error.

The min-max formulation (Theorem 3.1) optimizes over the set $\mathcal{C}_L = \{f \,;\, \|f\|_L \leq 1\}$. Instead of minimizing the worst-case reconstruction error, another natural approach to find optimal representations would be to maximize the captured variance. This alternative problem leads to Principal Component Analysis (PCA) on the same class of signals $\mathcal{C}_L$. Remarkably, we obtain the following result when performing PCA on uniformly distributed signals from $\mathcal{C}_L$.

Theorem 3.2 (Operator-Bounded PCA [4, Section 5])**.** *Given a symmetric positive definite operator L with eigenvalues $0 < \lambda_1 \leq \cdots \leq \lambda_n$, and eigenvectors $e_1, \ldots, e_n$, the PCA objective*

$$\min_{b=(b_1,\cdots,b_n)} \mathbb{E}_{f \sim \mathcal{U}\left(\|f\|_L \leq 1\right)} \left(\left\| f - \sum_{i=1}^{k} \langle f, b_i \rangle b_i \right\|^2 \right) ,$$

over orthonormal bases b is minimized by the eigenvectors of the covariance matrix $R_L = \mathbb{E}_{f \sim \mathcal{U}\left(\{\|f\|_L \leq 1\}\right)}[f f^\top]$, which are the first k eigenvectors of L, with eigenvalues (variances) $\lambda_1^{-1} \geq \cdots \geq \lambda_n^{-1}$. In other words: $R_L = L^{-1}$.

For a reminder why the expectation of the approximation error is PCA, i.e. iterative maximisation of the Rayleigh quotient of R_L, see Appendix A.1. Therefore, PCA on $\mathcal{C}_L$ yields the same eigenvectors and order as the min-max solution. The min-max optimization (Theorem 3.1) and its PCA counterpart (Theorem 3.2) are thus equivalent when applied to the same signal class. These dual formulations have useful practical implications that we exploit in our method.

A natural prior for signals f on a manifold is smoothness, implying that the Dirichlet energy $\|\nabla f\|_2$ is bounded. By Green's identity, $\|\nabla f\|_2 = \|f\|_\Delta$, where Δ is the (discrete) Laplace-Beltrami operator (LBO). Thus, bounding the Dirichlet energy leads to choosing Δ for L. To motivate our method, let us thus focus on the Laplacian operator and its construction. This will enable the derivation of our framework, generalizable both beyond Laplacian operators and to arbitrarily high-dimensional manifolds, from surfaces in $\mathbb{R}^3$ to image datasets.

Discrete Laplacians. The discrete Laplacian is traditionally constructed from two fundamental matrices, the mass matrix $M \in \mathbb{R}^{n \times n}$, which is a positive diagonal matrix encoding the local (metric-dependent) sampling density defining the manifold's Riemannian metric, and the stiffness matrix $S \in \mathbb{R}^{n \times n}$ encoding geometric relationships that discretize the continuous Laplace-Beltrami operator. For two-dimensional surfaces, M represents vertex areas and S contains cotangent weights [6, 19, 121]. For higher-dimensionality, the LBO is harder to approximate, thus M and S are constructed with schemes more sensitive to the connectivity of the relationship graph. In any case, the matrices M and S enable two common formulations for the discretized Laplacian: the unnormalized Laplacian $L = M^{-1}S$ and the symmetric normalized Laplacian $L_{\text{norm}} = M^{-\frac{1}{2}} S M^{-\frac{1}{2}}$. While both operators share identical eigenvalues, their eigenvectors differ. The unnormalized Laplacian's eigenvectors $\mathbf{v}_i \in \mathbb{R}^n$ are M-orthogonal, satisfying $\langle \mathbf{v}_i, \mathbf{v}_j \rangle_M = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise, whereas the normalized Laplacian's eigenvectors $\mathbf{q}_i \in \mathbb{R}^n$ are Euclidean-orthogonal, with $\langle \mathbf{q}_i, \mathbf{q}_j \rangle = \delta_{ij}$.

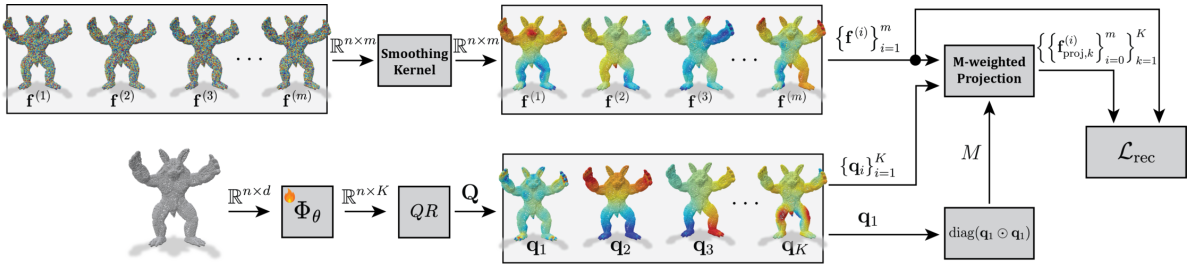

Figure 1. Overview of our neural framework to compute spectral bases directly from unstructured point clouds of any dimensionality, based on optimal-approximation theory, without first explicitly choosing, discretely approximating, and eigendecomposing an operator.

Both sets of eigenvectors are related by $\mathbf{v}_i = M^{-\frac{1}{2}}\mathbf{q}_i$, as $L(M^{-\frac{1}{2}}\mathbf{q}_i) = M^{-\frac{1}{2}}M^{-\frac{1}{2}}SM^{-\frac{1}{2}}\mathbf{q}_i = \lambda_i M^{-\frac{1}{2}}\mathbf{q}_i$.

On closed connected manifolds[1], these operators' nullspaces reveal important differences. While the vector of ones $\mathbf{1}$ lies in the nullspace of L, the vector $M^{\frac{1}{2}}\mathbf{1}$ lies in the nullspace of L_{norm}. This means the first eigenvector $\mathbf{q}_1$ of L_{norm} corresponding to eigenvalue zero is proportional to $M^{\frac{1}{2}}\mathbf{1}$, effectively encoding the square root of the sampling density weights (metric-sensitive). Should we want to learn directly the eigendecomposition of a Laplacian, the normalized formulation of L_{norm} offers a crucial advantage. Indeed, since the first eigenvector directly encodes the sampling weights of the discrete metric of the manifold, a trained neural network predicting the eigenbasis of L_{norm} simultaneously learns both the spectral decomposition and the manifold metric in a unified manner. This eliminates the need for separate modules to predict the mass matrix, simplifying the architecture while maintaining geometric consistency. Inspired by this remarkable property for Laplacians, which generalizes beyond them (see Appendix A.2), we designed our framework to similarly learn a basis, where we decided that the first eigenvector encodes the underlying metric and from which we can explicitly compute the metric-dependent area weights.

3.2. Neural Framework for Direct Spectral Bases

Exploiting the previous insights on optimal-approximation theory and discrete Laplacians, our goal is to learn directly a basis for geometry processing on discretely sampled data of arbitrary dimensions, bypassing the need to choose a specific operator, how to discretize it, and calls to sensitive eigensolvers on its discrete approximation. Should we choose an operator, like the Laplacian, a sub-goal would be to compute its eigendecomposition directly without first discretizing it (by combining mass and stiffness matrices) and post hoc eigendecomposition.

Given a point cloud $\mathcal{P}$ discretely sampling a manifold $\mathcal{M}$ with n points $\mathbf{p}_i \in \mathbb{R}^d$ for $1 \leq i \leq n$, we design a neural network $\Phi_\theta : \mathbb{R}^{n\times d} \to \mathbb{R}^{n\times K}$ to predict a K-dimensional feature vector for each point, where K is the number of basis vectors we wish to predict. These per-point predictions form a matrix $\Phi_\theta(\mathcal{P}) \in \mathbb{R}^{n\times K}$, on which we do QR decomposition, $\Phi_\theta(\mathcal{P}) = \mathbf{QR}$, where $\mathbf{Q} \in \mathbb{R}^{n\times K}$ has orthonormal columns and $\mathbf{R} \in \mathbb{R}^{K\times K}$ is upper triangular. By analogy with Laplacians, we can interpret $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_K]$, with i-th column $\mathbf{q}_i \in \mathbb{R}^n$, as the first K eigenvectors of a normalized operator, like L_{norm}, corresponding to eigenvalues $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \ldots, \lambda_K]^\top$, and ordered such that $0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_K$. For any $k \leq K$, we denote by $\mathbf{Q}_k = [\mathbf{q}_1, \ldots, \mathbf{q}_k] \in \mathbb{R}^{n\times k}$ the matrix containing the first $k \leq K$ columns of $\mathbf{Q}$.

Our pipeline computes a spectral basis directly, bypassing altogether first choosing, computing, and eigendecomposing an operator. Remarkably, we need not learn the metric M separately. Indeed, with our interpretation, the first eigenvector $\mathbf{q}_1$ directly encodes the mass matrix diagonal, by taking $M = \text{diag}(\mathbf{q}_1 \odot \mathbf{q}_1)$. We also need not learn separately the eigenvalues by exploiting the min-max theorem (Theorem 3.1) linking the eigenvalues to the maximum approximation error. During the forward pass, we generate random probe functions for the input point cloud $\mathcal{P}$ and project them onto our predicted truncated bases $\mathbf{Q}_k$ for all k. The maximum reconstruction error across these projections provides an estimate α_k, from which we compute $\lambda_{k+1} \approx \frac{1}{\alpha_k}$. This gives us eigenvalue estimates without requiring additional network parameters.

3.3. Learning by Optimal Approximations

Our framework (Fig. 1) operates on a batch of point clouds. Both at train or inference time, we progressively reconstruct probe functions by M-orthogonal projections onto the Euclidean-orthogonal estimated bases $\mathbf{Q}_k$ (see how in Appendix B.1) for increasing values of k and for each point cloud in the batch, as summarized in Algorithm 1.

Training. We train our model by learning a basis to optimally reconstruct probe functions. Depending on the choice of probe function distribution, we will be working implicitly with different operators L in the optimal reconstruc-

[1]This setting is standard. It is both common in practice for meshes, and is systematic for unstructured data like pointclouds, which is what we focus on, where potential manifold boundaries are both ill-defined and not provided. The optimal approximation theory generalizes to this case by focusing on the orthogonal of the nullspace.

Algorithm 1 Progressive Reconstruction Forward Pass

Input: Point cloud $\mathcal{P}$ with n points, number of probe functions m, number of eigenvectors K
Output: Reconstruction loss $\mathcal{L}_{\text{rec}}$ and maximum-error $e_{\max} \in \mathbb{R}^K$

Generate independently and uniformly m probe functions $\mathbf{f}^{(1)}, \mathbf{f}^{(2)}, \dots, \mathbf{f}^{(m)} \in \mathbb{R}^n$. By default, do this by iteratively applying Gaussian kernels on the k-nearset neighbor graph to independently and uniformly sampled signals of $\mathbb{R}^n$.
Compute a forward pass on the network $\Phi_\theta(\mathcal{P})$
Compute the QR-decomposition $\Phi_\theta(\mathcal{P}) = \mathbf{QR}$
for $k = 1$ to K **do**
    $\mathbf{Q}_k \leftarrow \mathbf{Q}$ truncated to the first k columns
    for $i = 1$ to m **do**
        $\mathbf{f}^{(i)}_{\text{proj},k} \leftarrow M$-projection of $\mathbf{f}^{(i)}$ onto $\mathbf{Q}_k$
        $e^{(i)}_k \leftarrow \|\mathbf{f}^{(i)} - \mathbf{f}^{(i)}_{\text{proj},k}\|_2^2$
    end for
    $i^{\max}_k \leftarrow \text{argmax}_i(e^{(i)}_k)$
    $e_{\max}.\text{append}(e^{(i^{\max}_k)}_k)$
end for
$\mathcal{L}_{\text{rec}} \leftarrow \frac{1}{mK}\sum_{i=1}^m \sum_{k=1}^K e^{(i)}_k$

tion theory (Theorems 3.1 and 3.2), leading to different estimated bases each having their own advantages. By default, we take inspiration from the Laplacian, which is the most commonly used operator, yet without computing it. Probe functions for the Laplacian should be smooth, with bounded Dirichlet energy, and uniformly distributed in this bounded set. However, computing the gradient requires the knowledge of the metric, which on unstructured data like high dimensional point clouds is not provided. We relax these constraints by generating probe functions from smoothing random functions with Gaussian kernels on the k-nearest neighbor graph of the data. The resulting distribution loosely resembles that of the constrained Laplacian, leading to similarly behaved bases, yet we neither aim for exact replica of the Laplacian nor or are we constrained to this operator or this choice of distribution.

Our training loss is based on optimal-reconstruction theory. Instead of the min-max formulation (Theorem 3.1), where for each k only one probe function in the batch (the worst approximated one) would be used to optimize the model, we switch to the equivalent PCA one (Theorem 3.2) averaging out the contribution of all probe functions. This leads to stabler optimization. Our proposed loss is thus the average reconstruction loss $\mathcal{L}_{\text{rec}}$ measuring the average quality of these progressive approximations

$$\mathcal{L}_{\text{rec}} = \frac{1}{mK}\sum_{i=1}^{m}\sum_{k=1}^{K} \|\mathbf{f}^{(i)} - \mathbf{f}^{(i)}_{\text{proj},k}\|_2^2, \tag{1}$$

where $\mathbf{f}^{(i)}, \mathbf{f}^{(i)}_{\text{proj},k} \in \mathbb{R}^n$ are the i-th probe function and its projection onto the first k estimated basis vectors. We train our model by backpropagating through the entire pipeline using $\mathcal{L}_{\text{rec}}$ as the sole training objective, enabling the network to learn the optimal basis through end-to-end gradient descent-based optimization. Note that we switched from the M-norm to the unweighted Euclidean norm. This hybrid approach, combining M-weighted projection with 2-norm error, creates an unsupervised mechanism for learning a density-related mass matrix and improves the training stability (see Appendix B.2).

Importantly, we never compute the implicit optimal reconstruction operator nor its eigenvalues during training. Nevertheless, they are easy to compute at inference time.

Inference. A feed-forward pass computes our optimal reconstruction basis for each point cloud in the inference batch. We can then easily compute the associated implicit optimal reconstruction operator or its eigenvalues for downstream geometric tasks. Thanks to the min-max formulation of optimal reconstruction theory (Theorem 3.1) the eigenvalues can be estimated from the worst-case reconstruction over all the probe functions at each spectral resolution $k \in \{1, \dots, K\}$ using $\lambda_{k+1} = \frac{1}{\max_i \|\mathbf{f}^{(i)} - \mathbf{f}^{(i)}_{\text{proj},k}\|_2^2}$. This provides theoretically-grounded eigenvalue estimates directly from the estimated basis rather than in parallel to it. Given the basis and its associated eigenvalues, we can then optionally explicitly reconstruct the implicit normalized symmetric operator by matrix multiplication $Q_K \Lambda_K Q_K^T$, or it unnormalized non-symmetric version $M^{-\frac{1}{2}} Q_K \Lambda_K Q_K^T M^{\frac{1}{2}}$. However, the operator matrix has little use in practice, as even its action is usually computed in its spectral basis. Thus, in our experiments, we never needed to recompute it explicitly. From the estimated normalized basis $\mathbf{q}_i$, we can also compute an unnormalized spectral basis $\mathbf{v}_i = M^{-\frac{1}{2}}\mathbf{q_i}$, which can be preferable downstream as it is more sampling invariant. For instance, the first vector $\mathbf{v}_1$ is constant regardless of the sampling of the the underlying smooth manifold.

4. Experiments

We demonstrate the versatility of our framework to compute spectral bases on arbitrary data. Starting from a sanity check in a toy 1D setting, with points sampled in $[0, 1] \subset \mathbb{R}$, we show that we can handle not only traditional 3D point clouds sampled on two-dimensional surfaces in $\mathbb{R}^3$, but also high-dimensional data in $\mathbb{R}^d$ such as image embeddings. Full details on the data, implementation, and further results are pushed to Appendix C.

4.1. Toy 1D Segment Manifold

As a sanity check, we illustrate our method on the simplest geometry: the unit interval $\Omega = [0, 1]$.

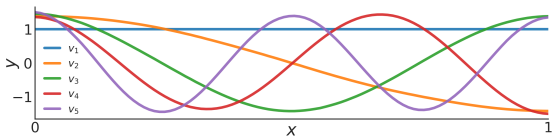

Figure 2. Learned eigenfunctions on [0,1] recover frequency-ordered harmonics resembling the Laplacian's spectrum.

Setting. We grid sample 100 points on $[0, 1]$, forming a point cloud. As in higher dimensions, random probe functions are first smoothened, without boundary constraints. Our learned extractor Φ_θ is a small MLP, which suffices in this single small point cloud setting.

Results. We plot in Fig. 2 the first five learned unnormalized spectral basis vectors $\mathbf{v}_i = M^{-\frac{1}{2}}\mathbf{q}_i$ after sign alignment. The network recovers a Fourier basis-like family with frequency-increasing harmonics: $\mathbf{v}_1$ is constant, while the other $\mathbf{v}_i$ exhibit i half-waves across the interval. This minimal example showcases how we can compute a frequency-ordered orthonormal basis resembling the Fourier basis, i.e. the Laplacian's eigenfunctions, just by optimising the approximation objective on basic probe functions. Importantly, this happens because the metric weights M, extracted from the first normalized basis vector $\mathbf{q}_1$, are sensible as they are correlated with intuition.

4.2. 2D Surface and 3D Volume Manifolds

We here show that our method can learn an approximation of the eigendecomposition of the most prevalent operator in shape analysis: the Laplace-Beltrami operator. By learning on single 3D point clouds (overfitting setting), we obtain highly accurate estimates. By learning on a wide collection of 3D point clouds from various datasets (generalization setting), we may obtain a foundation model that generalizes to unseen point clouds in $\mathbb{R}^3$ without retraining.

Datasets. The shapes in the overfitting setting are from [90]. In the generalization setting, we train our model on a wide collection of surface datasets to ensure broad generalization, ranging from protein structures to human scans [16, 66, 70, 75, 101]. We evaluate on reference shape analysis benchmarks [41, 48, 69, 76, 86, 90, 99, 122, 127, 146], spanning many challenges, e.g. deformations, topology variations, and different geometric characteristics. In each dataset, we dropped the mesh connectivity to work only with point clouds. Also, all shapes are scaled to fit within a unit sphere. We evaluate generalization to new manifold dimensionality by testing on 3D volumetric point clouds, computed by randomly sampling points inside the volume of shapes in [41], the model learned on surface point clouds.

Methods. We compare our neural framework (without connectivity) using a transformer [130] as learned extractor Φ_θ

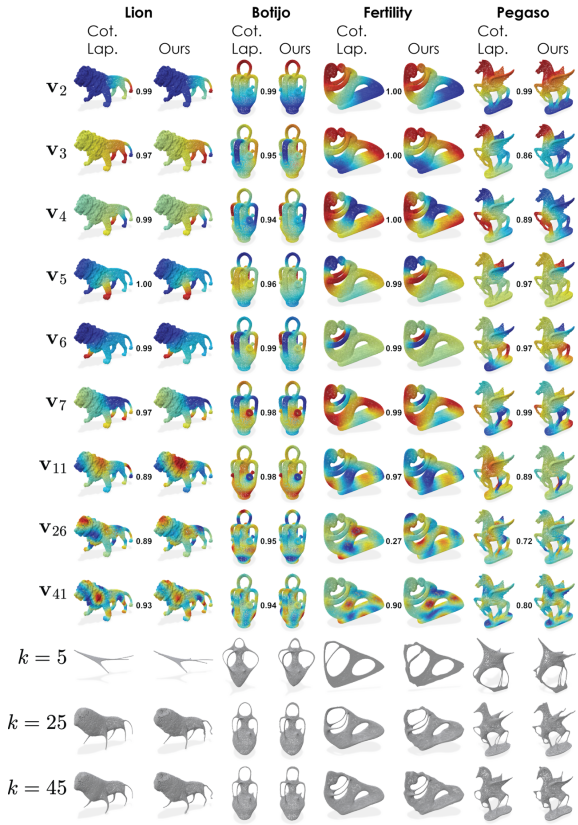

Figure 3. Unnormalized spectral basis (top) and xyz reconstruction from k basis vectors (bottom), using either the oracle cotangent Laplacian or our method (overfitting setting). Scalars are cosine similarities between basis vectors. We get similar if not more detailed reconstructions. More in Appendix C.

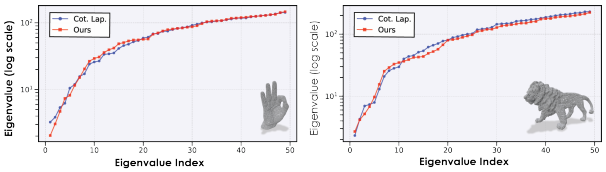

Figure 4. Eigenvalues of the oracle cotangent Laplacian and our estimated ones (overfitting setting). More in Appendix C.

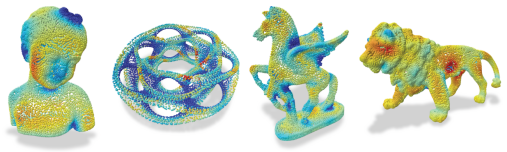

Figure 5. Estimated mass metric M from $\mathbf{q}_1$ (overfitting setting).

against the reference axiomatic oracle: the classical cotangent Laplacian (with oracle mesh connectivity).

Results – Overfitting setting. We plot the learned unnormalized eigenvectors $\mathbf{v}_i$ in Fig. 3. Remarkably, our eigenvectors, unsupervisedly learned solely using optimal-approximation theory and without knowledge of the underlying mesh, are almost identical to those computed for the

Table 1. Average cosine similarity between predicted and oracle eigenfunctions at different truncation levels k, and mean relative eigenvalue discrepancy (overfitting setting). More in Appendix C.

Shape	Image	$k \leq 10$	$k \leq 20$	$k \leq 50$	λ Discrepancy
Armadillo		0.968	0.967	0.773	0.200 ± 0.126
Bimba		0.964	0.945	0.822	0.093 ± 0.145
Botijo		0.972	0.955	0.813	0.153 ± 0.092
Elephant		0.979	0.866	0.687	0.105 ± 0.123
Fertility		0.874	0.866	0.720	0.083 ± 0.106
Kitten		0.993	0.988	0.981	0.088 ± 0.104
Laurent Hand		0.823	0.696	0.568	0.066 ± 0.078
Lion		0.951	0.908	0.822	0.067 ± 0.086
Pegaso		0.932	0.797	0.544	0.142 ± 0.140

cotangent Laplacian, which relies on the oracle mesh structure, with cosine similarity between them often close to 1 (see Tab. 1). Additionally, the eigenvalues extracted from the worst case errors provide a good approximation to those of the oracle (see Fig. 4). These results show that our neural unsupervised method can be used in an overfitting setting to get highly accurate spectral basis estimates that match those of the targeted Laplace-Beltrami oracle. On some shapes though, e.g. Pegaso, higher frequency basis vectors slightly diverge from the oracle. Yet by analyzing shape reconstruction (spectral filtering) results (see Fig. 3), projecting the xyz coordinates to the first k spectral vectors, we see that our model captures additional details lost in the oracle, providing better top k approximation. Thus not only can we imitate the reference oracle method, we can in some cases provide a superior version with better compressed information in the spectral basis. A cornerstone of our method is the unsupervised extraction of the metric weights M directly from the first estimated normalized vector $\mathbf{q}_1$, without knowledge of the mesh structure. We see in Fig. 5 that these estimated area weights are indeed well-behaved.

Results – Generalization setting. Here, our model is trained on a wide collection of surface point clouds. We plot learned unnormalized eigenvectors $\mathbf{v}_i$ and filter the xyz coordinates on unseen evaluation shapes, either surfaces or volumes (Fig. 6), which is a type of manifold never seen in training. Our model generalizes well to new eclectic types of shapes, yet with smaller precision than in the overfitting setting, demonstrating that our framework could provide unsupervised foundation models to compute spectral decompositions generalizing well beyond the training data.

Beyond Laplacians. We also implemented our method using other families of probes – piecewise constant, smooth polynomials, and Schrödinger-like functions – corresponding to non-LBO operators, see Fig. 7. Passing the burden from the choice of operator onto that of the probe family is a favorable trade-off, as generating simple families of probes is easier than operator discretization, especially in high dimensions. See Appendix C.4.2 for full details.

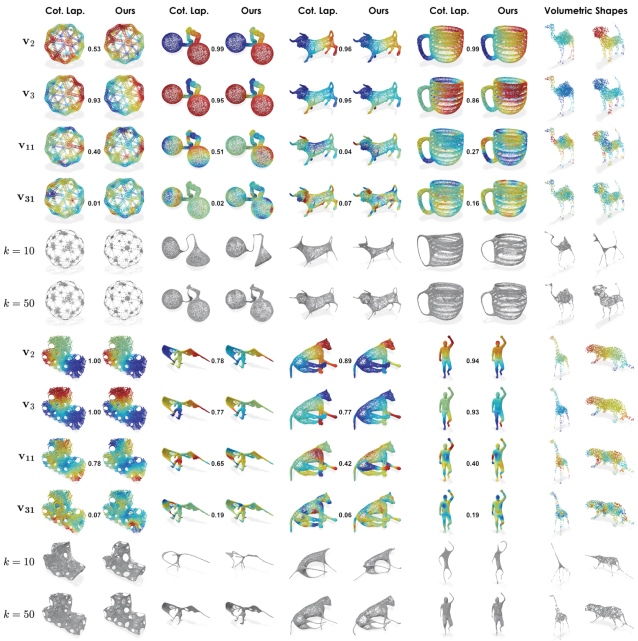

Figure 6. Unnormalized spectral basis $\mathbf{v}_1$ on unseen shapes, either surfaces (left) or volumes (right), when the model was trained on a wide collection of surface point clouds (generalization setting). Our model exhibits foundation-level generalization capabilities.

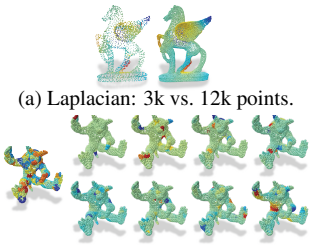

(a) Laplacian: 3k vs. 12k points.

(b) Piecewise: probe (left), basis (right).

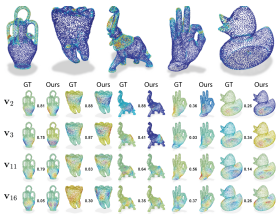

(c) Schrödinger: potential $V(x)$ (top), learned vs. GT (bot.).

Figure 7. Results for different probe families (overfitting setting).

4.3. High-Dimensional Manifolds

To demonstrate the generality of our approach, we experiment, beyond the classical 2D surface setting, on manifolds with high intrinsic dimensionality. Here, each image is a single point on the dataset manifold, with distances between images measured in a pretrained feature space.

Datasets. We use standard image classification datasets: STL10 [34], Imagenette [59], CIFAR100 [67], and Caltech256 [53] having 10, 10, 100, and 256 classes. Here, we do not mix datasets, but follow the train-test splits. As is standard, rather than working on raw pixel images,

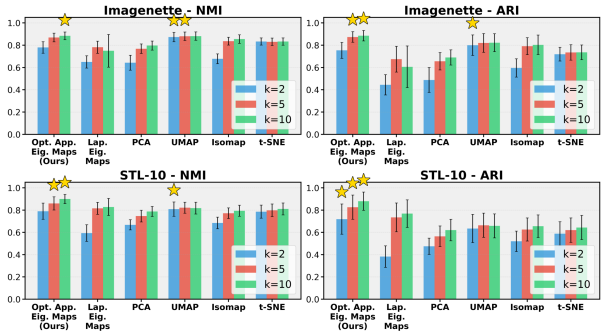

Figure 8. Average clustering performance over 50 runs of manifold learning methods on DINOv2 features of random data subsets (1500 images). Higher is better. More in Appendix C.

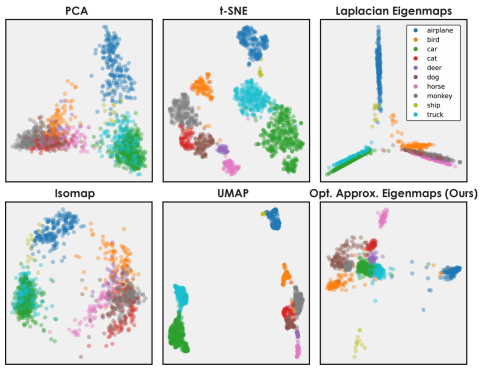

Figure 9. Manifold learning visualization of a random subset of STL10 by 2D embedding DINOv2 features. More in Appendix C.

we use reference feature embeddings, DINOv2 [92] and CLIP [104], as high-dimensional (yet lower dimensional than the raw pixel image) image coordinates in $\mathbb{R}^d$, with $d=768$ (resp. 512) for DINOv2 (resp. CLIP). Dataset manifolds are thus mapped to submanifolds of $\mathbb{R}^d$, and image distances are measured on embeddings.

Methods. Unlike 3D point clouds, high-dimensional data cannot be visualized and analyzed directly. Manifold learning addresses this by finding low-dimensional embeddings in $\mathbb{R}^k$, with $k \ll d$, that preserve pairwise dissimilarities to capture manifold structures. In particular, it enables 2D visualizations ($k = 2$) of high-dimensional data. For a short overview of manifold learning see Fig. 8. For classification data, it is widely accepted that better manifold learning methods provide better clustered embeddings correlating with the class labels. This can be both evaluated visually or with metrics such as the Normalized Mutual Information (NMI) [132] and Adjusted Rand Index (ARI) [60]. Our unnormalized spectral basis $\mathbf{v}_i$ naturally provides k-low-dimensional embeddings, which is our proposed manifold learning technique that we name Optimal-Approximation Eigenmaps. It is a direct generalization of Laplacian Eigenmaps [11], which uses the graph Laplacian spectral basis instead. We compare for various embedding dimensions $k \in \{2, 5, 10, 50\}$ our method against reference manifold learning techniques: PCA [58, 97], Isomap [126], Laplacian Eigenmaps [11], t-SNE [83, 129], and UMAP [85].

Results. Based on our results (see Figs. 8 and 9), our learned spectral basis consistently provides competitive or superior embeddings compared to baselines. These results showcase the strength of our neural spectral decomposition for capturing the intrinsic geometry of manifolds. In particular, our superiority compared to the similar Laplacian Eigenmaps suggests that our spectral basis and associated implicit operator are superior to those of the graph Laplacian, which is the standard for high-dimensional data.

5. Conclusion

We propose a neural framework to compute directly spectral bases on unstructured data of any dimensionality, bypassing the traditional need to first choose, discretize, and then eigendecompose an operator. Grounded in optimal-approximation theory, we find the orthonormal basis that optimally approximates constrained probe functions, with different probe distributions encoding different implicit operators. From the estimated basis, we can recover directly the manifold metric and compute explicitly the associated eigenvalues of the associated implicit operator. In a wide range of experiments covering one to three-dimensional manifolds, we showed that our method can recover the Laplace-Beltrami operator (LBO), which is the most useful operator in 3D geometry. However, our method scales in practice to high-dimensionality, unlike the LBO, and we show that it provides advantageous spectral representations via manifold learning experiments on image datasets. Our data-driven alternative to conventional pipelines paves the way for new avenues in geometry processing for unstructured data, especially in high-dimensions.

Limitations and future work. Training our method is computationally expensive, requiring time and GPUs, although this can be improved by optimizing our code with efficient compilation and CUDA implementations. Due to constraints, we focused on few tasks, yet we intend to explore further downstream applications, especially in high-dimensions, such as using our spectral basis in graph neural networks, or learning a more general foundation model that would apply to data of any dimensionality. At the heart of our method, the choice of probe functions imposes an underlying metric and implicit operator. By default we smoothen random functions on the kNN graph, giving similar results to the Laplacian. Yet, to best approximate it we need to manually tune distribution hyperparameters per shape, as we do in the overfitting setting. In future work, we will learn probe distributions and explore the approximation of more operators beyond the traditional Laplacians.

References

[1] Yonathan Aflalo and Ron Kimmel. Spectral multidimensional scaling. *Proceedings of the National Academy of Sciences*, 110(45):18052–18057, 2013. 5

[2] Yonathan Aflalo, Ron Kimmel, and Dan Raviv. Scale invariant geometry for nonrigid shapes. *SIAM Journal on Imaging Sciences*, 6(3):1579–1597, 2013. 2

[3] Yonathan Aflalo, Haim Brezis, and Ron Kimmel. On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences*, 8(2):1141–1160, 2015. 1

[4] Yonathan Aflalo, Haïm Brezis, Alfred Bruckstein, Ron Kimmel, and Nir Sochen. Best bases for signal spaces. *Comptes Rendus. Mathématique*, 354(12):1155–1167, 2016. 1, 3

[5] Yulin An and Enrique del Castillo. An ai approach for learning the spectrum of the laplace-beltrami operator. *arXiv preprint arXiv:2507.07073*, 2025. 2

[6] Mathieu Andreux, Emanuele Rodola, Mathieu Aubry, and Daniel Cremers. Anisotropic laplace-beltrami operators for shape analysis. In *European conference on computer vision*, pages 299–312. Springer, 2014. 2, 3

[7] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 1:1626–1633, 2011. 1

[8] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 3

[9] Thomas Barthelmé. A natural finsler-laplace operator. *Israel Journal of Mathematics*, 196(1):375–412, 2013. 2

[10] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. *Advances in neural information processing systems*, 14, 2001. 5

[11] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003. 8, 5

[12] Eugenio Beltrami. *Saggio di interpretazione della geometria non-euclidea*. Stab. Tip. De Angelis, 1868. 2

[13] Ido Ben-Shaul, Leah Bar, Dalia Fishelov, and Nir Sochen. Deep learning solution of the eigenvalue problem for differential operators. *Neural Computation*, 35(6):1100–1134, 2023. 2

[14] David Bensaïd, Amit Bracha, and Ron Kimmel. Partial shape similarity by multi-metric hamiltonian spectra matching. In *Scale Space and Variational Methods in Computer Vision: 9th International Conference, SSVM 2023, Santa Margherita Di Pula, Italy, May 21–25, 2023, Proceedings*, page 717–729, Berlin, Heidelberg, 2023. Springer-Verlag. 2

[15] David Bensaïd, Noam Rotstein, Nelson Goldenstein, and Ron Kimmel. Partial matching of nonrigid shapes by learning piecewise smooth functions. In *Computer Graphics Forum*, page e14913. Wiley Online Library, 2023. 2

[16] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3794–3801, 2014. 6, 2

[17] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. *Advances in neural information processing systems*, 29, 2016. 2

[18] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Michael M Bronstein, and Daniel Cremers. Anisotropic diffusion descriptors. In *Computer Graphics Forum*, pages 431–441. Wiley Online Library, 2016. 2

[19] Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Levy. Polygon mesh processing. pages 97–122, 2010. 3

[20] Amit Bracha, Oshri Halimi, and Ron Kimmel. Shape Correspondence by Aligning Scale-invariant LBO Eigenfunctions. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2020. 2

[21] Amit Bracha, Thomas Dagès, and Ron Kimmel. On unsupervised partial shape correspondence. In *Proceedings of the Asian Conference on Computer Vision*, pages 4488–4504, 2024.

[22] Amit Bracha, Thomas Dagès, and Ron Kimmel. Wormhole loss for partial shape matching. In *Advances in Neural Information Processing Systems*, pages 131247–131277. Curran Associates, Inc., 2024. 2, 5

[23] Matthew Brand. Nonrigid embeddings for dimensionality reduction. In *European Conference on Machine Learning*, pages 47–59. Springer, 2005. 5

[24] Haïm Brezis and David Gómez-Castro. Rigidity of optimal bases for signal spaces. *Comptes Rendus. Mathématique*, 355(7):780–785, 2017. 1, 3

[25] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5): 1168–1172, 2006. 5

[26] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 2

[27] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *International Conference on Learning Representations (ICLR)*, 2013. 2

[28] Augustin-Louis Cauchy. Sur l'équation à l'aide de laquelle on détermine les inégalités séculaires des mouvements des planètes. *Oeuvres Complètes (IIeme Série)*, 9:174–195, 1829. 2

[29] Frédéric Chazal, Ilaria Giulini, and Bertrand Michel. Data driven estimation of laplace-beltrami operator. *Advances in Neural Information Processing Systems*, 29, 2016. 2

[30] Gengxiang Chen, Xu Liu, Qinglu Meng, Lu Chen, Changqing Liu, and Yingguang Li. Learning neural operators on riemannian manifolds. *National Science Open*, 3 (6):20240001, 2024. 2

[31] Yoni Choukroun, Gautam Pai, and Ron Kimmel. Sparse approximation of 3d meshes using the spectral geometry of the hamiltonian operator. *J. Math. Imaging Vis.*, 60(6): 941–952, 2018. 2

[32] Ming Chuang, Linjie Luo, Benedict J Brown, Szymon Rusinkiewicz, and Michael Kazhdan. Estimating the laplace-beltrami operator by restricting 3d functions. In *Computer graphics forum*, pages 1475–1484. Wiley Online Library, 2009. 2

[33] Fan RK Chung. *Spectral graph theory*. American Mathematical Soc., 1997. 2

[34] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. 7

[35] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Special Issue: Diffusion Maps and Wavelets. 1, 2, 5

[36] Ronald R Coifman, Stephane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the national academy of sciences*, 102(21): 7426–7431, 2005. 5

[37] Thomas Dagès, Michael Lindenbaum, and Alfred M Bruckstein. Metric convolutions: A unifying theory to adaptive image convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13974–13984, 2025. 2

[38] Thomas Dagès, Simon Weber, Ya-Wei Eileen Lin, Ronen Talmon, Daniel Cremers, Michael Lindenbaum, Alfred M Bruckstein, and Ron Kimmel. Finsler multi-dimensional scaling: Manifold learning for asymmetric dimensionality reduction and embedding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 25842–25853, 2025. 5

[39] Thomas Dagès, Simon Weber, Daniel Cremers, and Ron Kimmel. Harnessing data asymmetry: Manifold learning in the Finsler world. *arXiv preprint arXiv:2603.11396*, 2026. 2, 5

[40] David L Donoho and Carrie Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100 (10):5591–5596, 2003. 5

[41] Roberto M Dyke, Yu-Kun Lai, Paul L Rosin, Stefano Zappalà, Seana Dykes, Daoliang Guo, Kun Li, Riccardo Marin, Simone Melzi, and Jingyu Yang. SHREC'20: Shape correspondence with non-isometric deformations. *Computers & Graphics*, 92:28–43, 2020. 6, 2, 3

[42] Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y Zeevi. The farthest point strategy for progressive image sampling. *IEEE transactions on image processing*, 6(9):1305–1315, 1997. 3

[43] William A Falcon. Pytorch lightning. *GitHub*, 2019. 3

[44] Bruce Fischl, Martin I Sereno, and Anders M Dale. Cortical surface-based analysis: Ii: inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9(2):195–207, 1999. 5

[45] Edward B Fowlkes and Colin L Mallows. A method for comparing two hierarchical clusterings. *Journal of the American statistical association*, 78(383):553–569, 1983. 6

[46] Clement Fuji Tsang, Maria Shugrina, Jean Francois Lafleche, Or Perel, Charles Loop, Towaki Takikawa, Vismay Modi, Alexander Zook, Jiehan Wang, Wenzheng Chen, Tianchang Shen, Jun Gao, Krishna Murthy Jatavallabhula, Edward Smith, Artem Rozantsev, Sanja Fidler, Gavriel State, Jason Gorski, Tommy Xiang, Jianing Li, Michael Li, and Rev Lebaredian. Kaolin: A pytorch library for accelerating 3d deep learning research. 3

[47] Thorben Funke, Tian Guo, Alen Lancic, and Nino Antulov-Fantulin. Low-dimensional statistical manifold embedding of directed graphs. In *8th International Conference on Learning Representations (ICLR 2020)*, pages 2018–2035. Curran, 2020. 5

[48] Daniéla Giorgi, Silvia Biasotti, and Laura Paraboschi. Watertight models track. *Shape Retrieval Contest*, 2007. 6, 2

[49] Teofilo F Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical computer science*, 38: 293–306, 1985. 3

[50] Carolyn Gordon and David Webb. You can't hear the shape of a drum. *American Scientist*, 84(1):46–55, 1996. 1

[51] Carolyn Gordon, David L Webb, and Scott Wolpert. One cannot hear the shape of a drum. *Bulletin of the American Mathematical Society*, 27(1):134–138, 1992. 1

[52] Jianping Gou, Xia Yuan, Ya Xue, Lan Du, Jiali Yu, Shuyin Xia, and Yi Zhang. Discriminative and geometry-preserving adaptive graph embedding for dimensionality reduction. *Neural Networks*, 157:364–376, 2023. 5

[53] Gregory Griffin, Alex Holub, Pietro Perona, et al. Caltech-256 object category dataset. Technical report, Technical Report 7694, California Institute of Technology Pasadena, 2007. 7

[54] Oshri Halimi and Ron Kimmel. Self functional maps. In *2018 International Conference on 3D Vision (3DV)*, pages 710–718. IEEE, 2018. 2

[55] David K. Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011. 1

[56] Jiequn Han, Jianfeng Lu, and Mo Zhou. Solving high-dimensional eigenvalue problems using deep neural networks: A diffusion monte carlo like approach. *Journal of Computational Physics*, 423:109792, 2020. 2

[57] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016. 3

[58] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933. 8, 5

[59] Jeremy Howard. Imagenette: A smaller subset of imagenet, 2019. 7

[60] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985. 8, 6

[61] Parvaneh Joharinad, Hannaneh Fahimi, Lukas Silvester Barth, Janis Keck, and Jürgen Jost. Isumap: Manifold learning and data visualization leveraging vietoris-rips filtrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 17699–17706, 2025. 5

[62] Mark Kac. Can one hear the shape of a drum? *The american mathematical monthly*, 73(4P2):1–23, 1966. 1

[63] Jungeum Kim and Xiao Wang. Inductive global and local manifold approximation and projection. *Transactions on Machine Learning Research*, 2024. 5

[64] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 2

[65] Gustav Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148 (12):497–508, 1847. 2

[66] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 6, 2

[67] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 7

[68] Joseph B Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29(1):1–27, 1964. 5

[69] Zorah Lähner, Emanuele Rodolà, Michael M Bronstein, Daniel Cremers, Oliver Burghard, Luca Cosmo, Alexander Dieckmann, Reinhard Klein, Y Sahillioğlu, et al. SHREC'16: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, pages 55–60. Eurographics Association, 2016. 6, 2

[70] Florent Langenfeld, Tunde Aderinwale, Feryal Windal, Mahmoud Melkemi, Ekpo Otu, Reyer Zwiggelaar, David Hunter, Yonghuai Liu, Léa Sirugue, Huu-Nghia H. Nguyen, Tuan-Duy H. Nguyen, Vinh-Thuyen Nguyen–Truong, Charles Christoffer, Danh Le, Hai-Dang Nguyen, Minh-Triet Tran, Matthieu Montès, Woong-Hee Shin, Genki Terashi, Xiao Wang, Daisuke Kihara, Halim Benhabiles, Karim Hammoudi, and Adnane Cabani. SHREC 2021: Surface-based Protein Domains Retrieval. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2021. 6, 2

[71] Thibault Lescoat, Hsueh-Ti Derek Liu, Jean-Marc Thiery, Alec Jacobson, Tamy Boubekeur, and Maks Ovsjanikov. Spectral mesh simplification. In *Computer Graphics Forum*, pages 315–324. Wiley Online Library, 2020. 2

[72] Ron Levie, Wei Huang, Lorenzo Bucci, Michael Bronstein, and Gitta Kutyniok. Transferability of spectral graph convolutional neural networks. *Journal of Machine Learning Research*, 22(272):1–59, 2021. 2

[73] Bruno Lévy. Laplace-beltrami eigenfunctions towards an algorithm that" understands" geometry. In *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*, pages 13–13. IEEE, 2006. 2

[74] Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, and Soumith Chintala. Pytorch distributed: experiences on accelerating data parallel training. *Proc. VLDB Endow.*, 13(12):3005–3018, 2020. 5

[75] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12706–12716, 2021. 6, 2

[76] Z. Lian, J. Zhang, S. Choi, H. ElNaghy, J. El-Sana, T. Furuya, A. Giachetti, R. A. Guler, L. Lai, C. Li, H. Li, F. A. Limberger, R. Martin, R. U. Nakanishi, A. P. Neto, L. G. Nonato, R. Ohbuchi, K. Pevzner, D. Pickup, P. Rosin, A. Sharf, L. Sun, X. Sun, S. Tari, G. Unal, and R. C. Wilson. Non-rigid 3D Shape Retrieval. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2015. 6, 2

[77] Ya-Wei Eileen Lin and Ron Levie. Adaptive canonicalization with application to invariant anisotropic geometric networks. In *The Fourteenth International Conference on Learning Representations*, 2026. 2

[78] Ya-Wei Eileen Lin, Ronald R Coifman, Gal Mishne, and Ronen Talmon. Hyperbolic diffusion embedding and distance for hierarchical representation learning. In *International Conference on Machine Learning*, pages 21003–21025. PMLR, 2023. 5

[79] Ya-Wei Eileen Lin, Ronald R Coifman, Gal Mishne, and Ronen Talmon. Tree-wasserstein distance for high dimensional data with a latent feature hierarchy. *arXiv preprint arXiv:2410.21107*, 2024. 5

[80] Ya-Wei Eileen Lin, Ronen Talmon, and Ron Levie. Equivariant machine learning on graphs with nonlinear spectral filters. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 2

[81] Yongming Liu. Curvature augmented manifold embedding and learning. *arXiv preprint arXiv:2403.14813*, 2024. 5

[82] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 3

[83] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (Nov):2579–2605, 2008. 8, 5

[84] Manuel Martín-Merino and Alberto Muñoz. Visualizing asymmetric proximities with SOM and MDS models. *Neurocomputing*, 63:171–192, 2005. 5

[85] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv*, 2018. 8, 5

[86] Simone Melzi, Riccardo Marin, Emanuele Rodolà, Umberto Castellani, Jing Ren, Adrien Poulenard, P Ovsjanikov, et al. SHREC'19: matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval*, pages 1–8. The Eurographics Association, 2019. 6, 2

[87] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H Barr. Visualization and mathematics III. *Mathematics and Visualization*, pages 35–57, 2003. 2

[88] Hrushikesh N Mhaskar and Ryan O'Dowd. Learning on manifolds without manifold learning. *Neural Networks*, 181:106759, 2025. 2

[89] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed precision training. In *International Conference on Learning Representations*, 2018. 3

[90] Ashish Myles, Nico Pietroni, and Denis Zorin. Robust field-aligned global parametrization. *ACM Transactions on Graphics*, 33(4):1–14, 2014. 6, 2

[91] Shin-ichi Ohta and Karl-Theodor Sturm. Heat flow on finsler manifolds. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 62(10):1386–1433, 2009. 2

[92] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 8

[93] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics*, 31(4):1–11, 2012. 1

[94] Gautam Pai, Alex Bronstein, Ronen Talmon, and Ron Kimmel. Deep isometric maps. *Image and Vision Computing*, 123:104461, 2022. 5

[95] Bo Pang, Zhongtian Zheng, Yilong Li, Guoping Wang, and Peng-Shuai Wang. Neural laplacian operator for 3d point clouds. *ACM Transactions on Graphics*, 43(6):1–14, 2024. 2

[96] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 3

[97] Karl Pearson. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11):559–572, 1901. 8, 5

[98] John Wilson Peoples and John Harlim. A higher order local mesh method for approximating laplacians on unknown manifolds. *arXiv preprint arXiv:2405.15735*, 2024. 2

[99] D. Pickup, X. Sun, P. L. Rosin, R. R. Martin, Z. Cheng, Z. Lian, M. Aono, A. Ben Hamza, A. Bronstein, M. Bronstein, S. Bu, U. Castellani, S. Cheng, V. Garro, A. Giachetti, A. Godil, J. Han, H. Johan, L. Lai, B. Li, C. Li, H. Li, R. Litman, X. Liu, Z. Liu, Y. Lu, A. Tatsuma, and J. Ye. SHREC'14 track: Shape retrieval of non-rigid 3d human models. In *Proceedings of the 7th Eurographics workshop on 3D Object Retrieval*. Eurographics Association, 2014. 6, 2

[100] Ulrich Pinkall and Konrad Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993. 1, 2

[101] Adrien Poulenard, Marie-Julie Rakotosaona, Yann Ponty, and Maks Ovsjanikov. Effective rotation-invariant point cnn with spherical harmonics kernels. In *2019 International Conference on 3D Vision (3DV)*, pages 47–56. IEEE, 2019. 6, 2

[102] Murray H Protter. Can one hear the shape of a drum? revisited. *Siam Review*, 29(2):185–197, 1987. 1

[103] B Quackenbush and PJ Atzberger. Geometric neural operators (gnps) for data-driven deep learning in non-euclidean settings. *Machine Learning: Science and Technology*, 5(4): 045033, 2024. 2

[104] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 8

[105] Arianna Rampini, Irene Tallini, Maks Ovsjanikov, Alex M Bronstein, and Emanuele Rodolà. Correspondence-free region localization for partial shape similarity via hamiltonian spectrum alignment. In *2019 International Conference on 3D Vision (3DV)*, pages 37–46. IEEE, 2019. 2

[106] Dan Raviv and Ron Kimmel. Affine invariant non-rigid shape analysis. *Int. J. Comput. Vision, submitted*, 2015. 2

[107] Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. Laplace–beltrami spectra as 'shape-dna' of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006. Symposium on Solid and Physical Modeling 2005. 1

[108] Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CoNLL)*, pages 410–420, 2007. 6

[109] Steven Rosenberg. *The Laplacian on a Riemannian manifold: an introduction to analysis on manifolds*. Number 31. Cambridge University Press, 1997. 2

[110] Guy Rosman, Michael M Bronstein, Alexander M Bronstein, and Ron Kimmel. Nonlinear dimensionality reduction by topologically constrained isometric embedding. *International Journal of Computer Vision*, 89(1):56–68, 2010. 5

[111] Conor Rowan, John Evans, Kurt Maute, and Alireza Doostan. Solving engineering eigenvalue problems with neural networks using the rayleigh quotient. *arXiv*, 2025. 2

[112] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000. 5

[113] Raif M. Rustamov. Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation. In *Geometry Processing*. The Eurographics Association, 2007. 1

[114] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998. 5

[115] Ariel Schwartz and Ronen Talmon. Intrinsic isometric manifold learning with application to localization. *SIAM Journal on Imaging Sciences*, 12(3):1347–1391, 2019.

[116] E.L. Schwartz, A. Shaw, and E. Wolfson. A numerical solution to the generalized mapmaker's problem: Flattening nonconvex polyhedral surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(9):1005–1008, 1989. 5

[117] Matan Sela, Yonathan Aflalo, and Ron Kimmel. Computational caricaturization of surfaces. *Computer Vision and Image Understanding*, 141:1–17, 2015. 2

[118] Nicholas Sharp and Keenan Crane. A laplacian for nonmanifold triangle meshes. *Computer Graphics Forum*, 39 (5):69–80, 2020. 1, 2, 8

[119] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)*, 41(3): 1–16, 2022. 2

[120] Roger N Shepard. The analysis of proximities: multidimensional scaling with an unknown distance function. i. *Psychometrika*, 27(2):125–140, 1962. 5

[121] O Sorkine, D Cohen-Or, Y Lipman, M Alexa, C Rössl, and H P Seidel. Laplacian surface editing. *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 175–184, 2004. 3

[122] Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Transactions on Graphics*, 23(3): 399–405, 2004. 6, 2

[123] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum*, pages 1383–1392. Wiley Online Library, 2009. 1, 2

[124] Ryota Suzuki, Ryusuke Takahama, and Shun Onoda. Hyperbolic disk embeddings for directed acyclic graphs. In *International Conference on Machine Learning*, pages 6066–6075. PMLR, 2019. 5

[125] Jian Tang, Jingzhou Liu, Ming Zhang, and Qiaozhu Mei. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th international conference on world wide web*, pages 287–297, 2016. 5

[126] Joshua B Tenenbaum, Vin de Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. 8, 5

[127] Elia Moscoso Thompson, Silvia Biasotti, Andrea Giachetti, Claudio Tortorici, Naoufel Werghi, Ahmad Shaker Obeid, Stefano Berretti, Hoang-Phuc Nguyen-Dinh, Minh-Quan Le, Hai-Dang Nguyen, Minh-Triet Tran, Leonardo Gigli, Santiago Velasco-Forero, Beatriz Marcotegui, Ivan Sipiran, Benjamin Bustos, Ioannis Romanelis, Vlassis Fotis, Gerasimos Arvanitis, Konstantinos Moustakas, Ekpo Otu, Reyer Zwiggelaar, David Hunter, Yonghuai Liu, Yoko Arteaga, and Ramamoorthy Luxman. SHREC 2020: Retrieval of digital surfaces with similar geometric reliefs. *Computers & Graphics*, 91:199–218, 2020. 6, 2

[128] Nicolás García Trillos, Chenghui Li, and Raghavendra Venkatraman. Minimax rates for the estimation of eigenpairs of weighted laplace-beltrami operators on manifolds. *arXiv preprint arXiv:2506.00171*, 2025. 2

[129] Laurens Van Der Maaten. Accelerating t-sne using tree-based algorithms. *The journal of machine learning research*, 15(1):3221–3245, 2014. 8

[130] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv*, 2017. 6

[131] Jarkko Venna, Jaakko Peltonen, Kristian Nybo, Helena Aidos, and Samuel Kaski. Information retrieval perspective to nonlinear dimensionality reduction for data visualization. *Journal of Machine Learning Research*, 11(2), 2010. 5

[132] Nguyen Xuan Vinh, Julien Epps, and James Bailey. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proceedings of the 26th annual international conference on machine learning*, pages 1073–1080, 2009. 8, 6

[133] Brian A Wandell, Suelika Chial, and Benjamin T Backus. Visualization and measurement of the cortical surface. *Journal of cognitive neuroscience*, 12(5):739–752, 2000. 5

[134] Yu Wang, Mirela Ben-Chen, Iosif Polterovich, and Justin Solomon. Steklov spectral geometry for extrinsic shape analysis. *ACM Transactions on Graphics (TOG)*, 38(1):1–21, 2018. 2

[135] Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. Discrete Laplace operators: no free lunch. In *Symposium on Geometry processing*, page 37. Aire-la-Ville, Switzerland, 2007. 1, 2

[136] Simon Weber, Thomas Dagès, Maolin Gao, and Daniel Cremers. Finsler-Laplace-Beltrami operators with application to shape analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3131–3140, 2024. 2

[137] Kilian Weinberger, Benjamin Packer, and Lawrence Saul. Nonlinear dimensionality reduction by semidefinite programming and kernel matrix factorization. In *International Workshop on Artificial Intelligence and Statistics*, pages 381–388. PMLR, 2005. 5

[138] Kilian Q Weinberger and Lawrence K Saul. Unsupervised learning of image manifolds by semidefinite programming. *International journal of computer vision*, 70:77–90, 2006. 5

[139] Aaron Wetzler, Yonathan Aflalo, Anastasia Dubrovina, and Ron Kimmel. The Laplace-Beltrami operator: a ubiquitous tool for image and shape processing. In *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, pages 302–316. Springer, 2013. 2

[140] Romy Williamson and Niloy J. Mitra. Neural geometry processing via spherical neural surfaces. *Computer Graphics Forum*, 44(2):e70021, 2025. 2

[141] Oguzhan Yigit and Richard C Wilson. Lbonet: Supervised spectral descriptors for shape analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2

[142] Ao Zhang, Qing Fang, Peng Zhou, and Xiao-Ming Fu. Topology-controlled laplace–beltrami operator on point clouds based on persistent homology. *Graphical Models*, 139:101261, 2025. 2

[143] Zhenyue Zhang and Jing Wang. MLLE: Modified locally linear embedding using multiple weights. *Advances in neural information processing systems*, 19, 2006. 5

[144] Zhenyue Zhang and Hongyuan Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM journal on scientific computing*, 26(1): 313–338, 2004. 5

[145] Hongyu Zhou and Zorah Lähner. Laplace-Beltrami operator for gaussian splatting. *arXiv preprint arXiv:2502.17531*, 2025. 2

[146] Qingnan Zhou and Alec Jacobson. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797*, 2016. 6, 2