HodgeNet: Learning Spectral Geometry on Triangle Meshes

DMITRIY SMIRNOV and JUSTIN SOLOMON, Massachusetts Institute of Technology, USA

Fig. 1. Mesh segmentation results on the full-resolution MIT animation dataset. Each mesh in the dataset contains 20,000 faces (10,000 vertices). We show an example ground truth segmentation in the bottom-left. In contrast to previous works, which downsample each mesh by more than 10 $\times$, we efficiently process dense meshes both at train and test time.

Constrained by the limitations of learning toolkits engineered for other applications, such as those in image processing, many mesh-based learning algorithms employ data flows that would be atypical from the perspective of conventional geometry processing. As an alternative, we present a technique for learning from meshes built from standard geometry processing modules and operations. We show that low-order eigenvalue/eigenvector computation from operators parameterized using discrete exterior calculus is amenable to efficient approximate backpropagation, yielding spectral per-element or per-mesh features with similar formulas to classical descriptors like the heat/wave kernel signatures. Our model uses few parameters, generalizes to high-resolution meshes, and exhibits performance and time complexity on par with past work.

CCS Concepts: • **Computing methodologies** $\rightarrow$ **Shape analysis; Mesh geometry models; Neural networks.**

Additional Key Words and Phrases: Machine learning, meshes, operators

ACM Reference Format:

Dmitriy Smirnov and Justin Solomon. 2021. HodgeNet: Learning Spectral Geometry on Triangle Meshes. *ACM Trans. Graph.* 40, 4, Article 166 (August 2021), 11 pages. <https://doi.org/10.1145/3450626.3459797>

1 INTRODUCTION

Data-driven algorithms have altered the way we approach problem-solving in computer graphics. Machine learning tools garner top performance for tasks like image editing, user interaction, image synthesis, and layout, supported by large, well-curated datasets. Yet, while learning tools for areas like computational photography and rendering are widely adopted, another branch of graphics has been resistant to change: mesh-based shape analysis.

Authors' address: Dmitriy Smirnov, smirnov@mit.edu; Justin Solomon, jsolomon@mit.edu, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139, USA.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

0730-0301/2021/8-ART166

<https://doi.org/10.1145/3450626.3459797>

Numerous technical challenges preclude modern learning methods from being adopted for meshes. Deep learning—arguably the most popular recent learning methodology—relies on *regularity* of the data and *differentiability* of the objective function for efficiency. For example, convolutional neural network (CNN) training is built on high-throughput processing of images through convolution and per-pixel computations to obtain gradients with respect to network weights, required for stochastic gradient descent.

Meshes, a primary means of representing geometry in graphics, defy the considerations above. They come as sparse, irregular networks of vertices varying in number; the same piece of geometry easily can be represented by multiple meshes and at multiple resolutions/densities. Advances in graph neural networks (GNNs) have as a byproduct helped advance mesh processing, but typical graphs in geometry processing are fundamentally different from those in network science—vertices have low valence, are related through long chains of edges, can be connected in many roughly-equivalent ways, and can be deformed through rigid motions and isometries.

The end result is that mesh-based learning architectures often contort input data to make it compatible with existing learning toolkits. Restricting to GPU-parallel, regularly-structured computation is a vast limitation for mesh analysis. For example, while geometry processing algorithms frequently rely on inversion and eigenanalysis of sparse matrices, these operations are hardly compatible with deep learning. Instead, mesh-based learning algorithms differ from successful non-learning geometry processing algorithms, relying on easily differentiated/parallelized local operations.

In this paper, we ask whether we can invert this relationship: Rather than inventing new data streams for geometry processing to suit existing learning algorithms, can we develop learning methodologies from successful geometry processing techniques?

We target applications in shape analysis using a prevalent tool in that domain, spectral geometry. Myriad shape analysis algorithms follow a similar template, building a positive (semi)definite matrix whose sparsity pattern is inherited from the mesh and then using its spectral structure to infer information about meshed geometry.

Some examples include the Laplacian operator, the bilaplacian operator, the Dirac operator, and modal analysis. Our broad goal—also explored in some past work (see §2)—is to *learn* the entries of this operator as functions of local geometry.

Unlike past work, however, we observe that classical shape analysis relies on *near-zero eigenvalues* of these operators (and the corresponding eigenvectors); high-frequency information is discarded. This is a key reason why classical geometry processing algorithms involve sparse matrix inversion and partial computation of eigenvalues. Partial eigenvector computation from a sparse matrix, however, is incompatible with most existing learning pipelines, so learning algorithms that use low-order eigenanalysis typically precompute the relevant eigenvectors from a fixed operator. Approaches to operator learning work with operator-vector products (rather than inverting the operator), restrict to a pre-computed basis, or compute the full spectrum as a dense matrix, which is prohibitive for large meshes.

In this paper, we approximately differentiate through sparse operator construction for one class of operators motivated by discrete differential geometry. As a result, we can learn operators whose entries are functions of local geometry, which together modify the spectral structure of the operator—a global computation. Our method is competitive with existing mesh-based learning tools while being implemented from standard components of the geometry processing toolkit, and we show how to handle boundary conditions and vector-valued data.

We make some unconventional design decisions that resemble geometry processing rather than deep learning. For instance, our spectral computation and operator construction are implemented using sparse linear algebra on the CPU, and we implement geometry processing-specific strategies for data augmentation that promote resolution independence. These decisions do not hamper efficiency of our method relative to past work.

Contributions. We present a lightweight model for learning from triangle meshes, with or without boundary. Contributions include:

- a learnable class of sparse operators on meshes built from standard constructions in discrete exterior calculus;
- parallelizable algorithms for differentiating eigencomputation from these operators, including approximate backpropagation without sparse computation;
- end-to-end architectures for learning per-element or per-mesh features starting from mesh geometry without additional features;
- simple strategies for data augmentation and other practical techniques to improve performance of our method; and
- experiments demonstrating effectiveness in shape analysis tasks, including the generalization of our model to high-resolution meshes that are too dense to be compatible with related methods.

2 RELATED WORK

Machine learning from geometry is becoming a popular subfield of graphics and vision. Bronstein et al. [2017] provide a broad overview of challenges in this discipline; here, we focus on work directly related to our task of learning from meshed geometry.

2.1 Spectral Shape Analysis

Our method is built on ideas from spectral geometry, which captures shape properties through the lens of spectral (eigenvalue/eigenvector) problems. Wang and Solomon [2019] provide a comprehensive introduction to this approach to geometry processing.

The *Laplace–Beltrami* (or, *Laplacian*) operator is ubiquitous in spectral geometry processing. Most relevant to our work, numerous per-vertex and per-mesh features have been built from Laplacian eigenvalues and eigenvectors, including the global point signature [Rustamov 2007], the heat kernel signature [Sun et al. 2009], the wave kernel signature [Aubry et al. 2011], and the heat kernel map [Ovsjanikov et al. 2010]. These descriptors underlie algorithms for tasks as varied as symmetry detection [Ovsjanikov et al. 2008], correspondence [Ovsjanikov et al. 2012], shape recognition [Bronstein and Bronstein 2010; Reuter et al. 2006], and shape retrieval [Bronstein et al. 2011]—among countless others.

The Laplacian is popular given its multiscale sensitivity to intrinsic geometry, but recent work proposes replacements sensitive to other aspects of geometry like extrinsic deformation. Examples include the Dirac operator [Liu et al. 2017; Ye et al. 2018], modal analysis [Hildebrandt et al. 2012; Huang et al. 2009], the Hamiltonian [Choukroun et al. 2018], the curvature Laplacian [Liu and Zhang 2007], the concavity-aware Laplacian [Au et al. 2011; Wang et al. 2014], the volumetric Laplacian [Raviv et al. 2010], and the Dirichlet-to-Neumann operator [Wang et al. 2018a]. Other works add invariances to the Laplacian, e.g., to local scaling [Bronstein and Kokkinos 2010] or affine deformation [Raviv et al. 2011], while others incorporate local features like photometric information [Kovnatsky et al. 2011; Spagnuolo et al. 2012]. Nearly all these algorithms—with the notable exception of volumetric methods [Raviv et al. 2010; Wang et al. 2018a]—follow the same outline: Build an operator matrix whose sparsity pattern is inherited from the edges of a triangle mesh and construct features from its eigenvectors and eigenvalues; a widely-used strategy of *truncation* approximates spectral features using partial eigeninformation, usually the eigenvalues closest to 0.

Other spectral methods use or produce *vectorial* data, working with operators that manipulate tangential fields. Vector diffusion operators move information along a manifold or surface while accounting for parallel transport [Sharp et al. 2019; Singer and Wu 2012]. The Killing operator also has been applied to intrinsic symmetry detection [Ben-Chen et al. 2010], segmentation [Solomon et al. 2011b], deformation [Claici et al. 2017; Solomon et al. 2011a], level set tracking [Tao et al. 2016], and registration/reconstruction [Chan and Wan 2013; Slavcheva et al. 2017]. These methods again analyze a sparse operator built from local features and mesh structure, although there is less agreement on the discretization of operators acting on vector-valued data [de Goes et al. 2016].

Spectral representations of geometry can be “complete” in the sense that a shape’s intrinsic structure or embedding can be reconstructed from the eigenvalues and eigenvectors of certain operators. For example, the discrete Laplacian determines mesh edge lengths [Zeng et al. 2012], and a modified operator adds the extrinsic information needed to obtain an embedding [Corman et al. 2017]. [Boscaini et al. 2015a; Corman et al. 2017; Cosmo et al. 2019] solve related inverse problems in practice.

Transitioning to the next section, an early machine learning method by Litman and Bronstein [2013] uses regression to learn spectral descriptors on meshes through learnable functions of Laplacian eigenvalues. This method does not learn the operator but rather the way per-vertex features are constructed from Laplacian eigenvalues. Henaff et al. [2015] propose a similar approach on graphs.

We attempt to generalize many of the methods above. Rather than defining a “bespoke” operator and, mapping from eigeninformation to features for each new task, however, we learn an operator from data.

2.2 Neural Networks on Meshes

Many papers propose algorithms for learning from meshes and other geometric representations. Here, we summarize past approaches for learning features from meshes, although specialized methods for mesh-based learning appear in tasks like generative modeling [Hertz et al. 2020; Liu et al. 2020], meshing [Sharp and Ovsjanikov 2020], and reconstruction [Gao et al. 2020; Hanocka et al. 2020].

Learning from graphs. Since triangle meshes are structured graphs, algorithms for learning from graphs inspired approaches to learning from meshes. Indeed, graph neural networks (GNNs) [Kipf and Welling 2017] are often used as baselines for geometric learning.

The graph analog of spectral geometry employs Laplacian matrices that act on per-vertex functions. Graph Laplacians provide a linear model for aggregating information between neighboring vertices. Spectral networks [Bruna et al. 2013] project per-vertex features onto a low-frequency Laplacian eigenbasis before applying a learned linear operator, followed by a per-vertex nonlinearity in the standard basis; convolution on images can be understood as a spectral filter, so these networks generalize image-based convolutional neural networks (CNNs). Subsequent work accelerated learning and inference from spectral networks, often using matrix functions in lieu of computing a Laplacian eigenbasis, e.g., via Chebyshev polynomials [Defferrard et al. 2016], random walks [Atwood and Towsley 2016], or rational functions [Levie et al. 2018].

Spatial domain. Many mesh-based learning methods operate in the “spatial domain,” relating vertices to their neighbors through constructions like local parameterization or tangent plane approximation. These methods often can be understood as GNNs with geometrically-motivated edge weights.

Starting with [Masci et al. 2015], many methods define convolution-like operations within local neighborhoods by parameterizing vertices and their k -ring neighborhoods. A challenge is how to orient the convolution kernel, since the tangent plane is different at every point; strategies include taking a maximum over all possible orientations [Masci et al. 2015; Sun et al. 2020], dynamically computing weights from neighboring features [Verma et al. 2018], aligning to principal curvatures [Boscaini et al. 2016], learning pseudo-coordinate functions represented as mixtures of Gaussians [Monti et al. 2017], projecting onto tangent planes [Tatarchenko et al. 2018], sorting nearby vertices based on feature similarity [Wang et al. 2018b], aligning to a 4-symmetry field [Huang et al. 2019], and weighting by normal vector similarity [Song et al. 2020] or directional curvature [He et al. 2020].

These and other methods must also define a means of representing localized convolution kernels. Many choices are available, including localized spectral filters [Boscaini et al. 2015b], B-splines [Fey et al. 2018], Zernike polynomials [Sun et al. 2020], wavelets [Schonsheck et al. 2018], and extrinsic Euclidean convolution [Schult et al. 2020].

Additional machinery is needed to compute vectorial features or relate tangent kernels at different vertices—a problem related to choosing a canonical orientation per vertex. Parallel transport is a choice motivated by differential geometry [Pan et al. 2018], which can be combined with circular harmonics [Wiersma et al. 2020] or pooling over multiple coordinates [Poulenard and Ovsjanikov 2018] to avoid dependence on a local coordinate system. Yang et al. [2020] employ locally flat connections for a similar purpose.

Simple GNN layers like [Kipf and Welling 2017] communicate information only among neighboring vertices. This small receptive field—inherited by several methods above—is a serious challenge for learning from meshes, which are sparse graphs for which a single such layer becomes more and more local as resolution increases. This issue creates dependency of performance on mesh resolution.

Mesh-based constructions. While it is valid to interpret meshes as graphs, this neglects the fact that meshes are highly-structured relative to graphs in other disciplines; a few learning algorithms leverage this additional structure to engineer mesh-specific convolutional-style layers. The popular MeshCNN architecture [Hanocka et al. 2019] learns edge features and performs pooling based on edge collapse operations. PD-MeshNet [Milano et al. 2020] augments the graph of mesh edges with the graph of dual edges capturing triangle adjacency, with pooling techniques inspired by mesh simplification and dynamic local aggregation using attention.

Global parameterization. Surface parameterization is a standard technique for texture mapping; some methods parameterize meshes into an image domain on which standard CNNs can be used for learning and inference from pushed-forward features. Sinha et al. [2016] pioneered this approach using geometry images [Gu et al. 2002] for parameterization. Maron et al. [2017] use seamless toric covers, conformally mapping four copies of a surface into a flat torus; this work was extended by Haim et al. [2019] to general covers to reduce distortion. Rendering-based techniques can also be understood as simple parameterizations onto the image plane, e.g., using panoramic [Sfikas et al. 2017; Shi et al. 2015] or multi-view [Kalogerakis et al. 2017; Su et al. 2015; Wei et al. 2016] projections.

Fixed operator methods. Some methods use operators on surfaces to construct convolution-like operations. Surface Networks [Kostrikov et al. 2018] use discrete Laplacian and Dirac operators as edge weights in GNNs. Yi et al. [2017] define kernels in Laplacian eigenbases, including spectral parameterizations of dilated convolutional kernels and transformer networks. Qiao et al. [2020] use Laplacian spectral clustering to define neighborhoods for pooling.

Learned operators. Some past methods learn relevant differential operators to a geometric learning task. Closest to ours, Wang et al. [2019] learn a parameterized sparse operator for geometry processing; see §4 for comparison of our operator to theirs. Their layers simulate iterations of algorithms like conjugate gradients by applying their operator, limiting its receptive field to the number of

Fig. 2. Data flow in HodgeNet; yellow boxes contain learnable parameters.

layers. In contrast, we explicitly perform eigendecomposition in our differentiable pipeline, allowing us to engineer the inductive bias inspired by the Hodge Laplacian. Similar discretizations are found in methods like [Eliasof and Treister 2020] for learning PDEs from data; this method uses algebraic multigrid to increase the receptive field.

Other. We mention a few other methods for learning from meshes that do not fall into the categories above. Xu et al. [2017] present a pipeline that combines purely local and mesh-wise global features; Feng et al. [2019] also propose extracting purely local features. Lim et al. [2018] apply recurrent neural networks (RNNs) to compute vertex features after unrolling local neighborhoods into prescribed spiral patterns. Deep functional maps [Litany et al. 2017] largely rely on precomputed features for geometric information, although some recent efforts bring this correspondence method closer to end-to-end [Donati et al. 2020; Sharma and Ovsjanikov 2020].

Concurrent and unreviewed work. Machine learning is a fast-paced discipline, with new papers released daily. Here, we acknowledge some “late-breaking” concurrent work.

Sharp et al. [2020] propose a “learned diffusion layer” in which features are diffused along a geometric domain via the isotropic heat equation with learned amount of diffusion; they include diffusion time as a learnable parameter. Similarly to [Bruna et al. 2013], their diffusion is implemented in a fixed low-frequency Laplacian eigenbasis, computed during learning/inference. Additional features incorporate anisotropy via inner products of spatial gradients. Unlike our work, they use a prescribed Laplacian operator.

Other methods include [de Haan et al. 2020], which proposes anisotropic gauge-invariant kernels using a message passing scheme built from parallel transport; [Lahav and Tal 2020], an RNN-based approach employing random walks; [Schneider et al. 2020], which improves MeshCNN’s memory efficiency and resilience to class imbalance for medical applications; and [Budninskiy et al. 2020], which optimizes for a graph Laplacian parameterized by edge features.

3 OVERVIEW

Figure 2 gives an overview of our HodgeNet architecture for learning from triangle meshes. The boxes highlighted in yellow have

learnable parameters, while the remaining boxes are fixed computations.

Our goal is to learn an operator and associated spectral descriptor for a given learning-from-meshes task. As with most methods, the learning stage uses stochastic gradient descent to optimize for model parameters, which are fixed during inference.

Our model inputs a triangle mesh $M = (V, E, T)$ and constructs three objects:

- a combinatorial *differential* matrix $d \in \{-1, 0, 1\}^{|E| \times |V|}$,
- a diagonal *0-form Hodge star* matrix $\star_0 \in \mathbb{R}^{|V| \times |V|}$, and
- a diagonal *1-form Hodge star* matrix $\star_1 \in \mathbb{R}^{|E| \times |E|}$.

The matrix d is a fixed function of M , while $\star_0, \star_1$ are learnable functions of local neighborhoods around mesh vertices.

HodgeNet then computes the k eigenvectors x^i of the semidefinite Laplacian-type matrix $L = \star_0^{-1} d^T \star_1 d$ whose eigenvalues λ^i are closest to zero. Finally, per-vertex or per-mesh features are gathered from $\{(x^i, \lambda^i)\}$ using learnable formulas that generalize the form of popular spectral features like the heat kernel signature.

During training, we need to differentiate our loss function through the steps above. Most of the operations above are simple nonlinearities that can be differentiated using standard backpropagation methods. We show in §5 how to obtain approximate derivatives of the eigenproblem efficiently.

4 OPERATOR CONSTRUCTION

HodgeNet relies on a parameterized class of learnable operators whose entries are functions of local geometry. The basis of our construction, designed to encapsulate operator constructions in spectral geometry, resembles that proposed by Wang et al. [2019], with the key difference that we expose per-edge and per-vertex features using diagonal Hodge star operators; this difference greatly simplifies our backpropagation procedure in §5. In §4.2, we also show how to generalize this construction for vectorial operators.

4.1 Operator

Given an oriented manifold mesh $M = (V, E, T)$ (optionally with boundary) with vertices $V \subset \mathbb{R}^3$, edges $E \subset V \times V$, and oriented triangles $T \subset V \times V \times V$, HodgeNet constructs a positive (semi)definite operator matrix $L \in \mathbb{R}^{|V| \times |V|}$ whose spectral structure will be used for a mesh-based learning task.

Inspired by the factorization of the Laplacian in discrete exterior calculus [Desbrun et al. 2005], we parameterize L as a product:

$$L = \star_0^{-1} d^T \star_1 d. \quad (1)$$

Here, $d \in \{-1, 0, 1\}^{|E| \times |V|}$ is the differential operator given by

$$d_{ev} = \begin{cases} 1 & \text{if } v = v_2 \\ -1 & \text{if } v = v_1 \\ 0 & \text{otherwise,} \end{cases}$$

where $e = (v_1, v_2)$ is an oriented edge.

While d is determined by mesh topology, the diagonal Hodge star matrices $\star_0 \in \mathbb{R}^{|V| \times |V|}$ and $\star_1 \in \mathbb{R}^{|E| \times |E|}$ are learnable functions of local mesh geometry. To construct $\star_0, \star_1$, we input D per-vertex features $F \in \mathbb{R}^{|V| \times D}$. In our experiments, we use positions and normals as the per-vertex features, except when noted otherwise.

We detail the construction of these operators $\star_0(F)$, $\star_1(F)$ from F below.

Per-vertex features $\star_0$. Our construction of $\star_0(F)$ imitates area weight computation for discrete Laplacians. It takes place in two steps. First, we compute *per-triangle* features using a learnable function $f_\Phi : \mathbb{R}^{3D} \rightarrow \mathbb{R}$, where Φ contains the parameters of our model. To ensure positive (semi)definiteness for $\star_0(F)$ we square f_Φ . Finally, we gather features from triangles to vertices by summing and optionally adding a small constant ε (in practice, $\varepsilon = 10^{-4}$) to improve conditioning of the eigensystem. Overall, we can write our expression as follows:

$$(\star_0(F))_{vv} = \varepsilon + \sum_{t \sim v} f_\Phi(F_{v_1}, F_{v_2}, F_{v_3})^2 \quad (2)$$

where $t = (v_1, v_2, v_3) \in T$ is a triangle with vertices v_1, v_2, v_3 in counterclockwise order. This sum over t has a potentially different number of terms for each vertex, equal to the valence.

If $\varepsilon = 0$ and f_Φ^2 measures triangle area scaled by $1/3$, then $\star_0$ becomes the barycentric area weights matrix often used in finite elements and discrete exterior calculus. We give the details of our choice of functions f_Φ in §7. Squaring the inner part of (2) is one of many ways to make sure $(\star_0)_{vv} \geq 0$ and could be replaced, e.g., by ReLU activation, but we found empirically that this simple expression led to the best performance.

Per-edge features $\star_1$. The diagonal matrix $\star_1(F)$ contains per-edge features on its diagonal. Unlike (2), to compute $\star_1(F)$ we do not need to gather features from a variable-sized ring. Instead, we learn a function $g_\Phi : \mathbb{R}^{4D} \rightarrow \mathbb{R}$ and, for an interior edge $e = (v_1, v_2)$, compute

$$(\star_1(F))_{ee} = \varepsilon + g_\Phi(F_{v_1}, F_{v_2}, F_{v_3}, F_{v_4})^2, \quad (3)$$

where v_3 and v_4 are *opposite* the edge e as shown to the right. We order v_3 and v_4 so that (v_1, v_2, v_3) and (v_2, v_1, v_4) are all consistently oriented. We learn a separate function $\tilde{g}_\Phi(v_1, v_2, v_3)$ for boundary edges, since there is only one opposite vertex in this case.

If $\varepsilon = 0$ and g_Φ^2 gives the sum of interior angle cotangents at v_3 and v_4 , then $\star_1$ will be the famous cotangent Laplacian matrix common in geometry processing. While we have chosen to square the function g , thanks to conjugation by d in (1) this is sufficient but not necessary for positive (semi)definiteness of L , and indeed this design choice prevents us from exactly reproducing the cotangent Laplacian in the presence of obtuse triangles. Our architecture could easily be adjusted to allow for negative $(\star_1)_{ee}$ values and hence to reproduce the cotangent Laplacian operator, but the stability and ease of squaring g_Φ to ensure that L has no negative eigenvalues outweighed this largely theoretical consideration.

Discussion. Our parameterizations of L , $\star_0$, and $\star_1$ largely imitate the flow of information used to construct discrete Laplacian operators and related objects. They are readily incorporated into geometry processing pipelines and have familiar sparsity patterns encountered in this discipline.

It is worth acknowledging a few design decisions intended to simplify our framework at the cost of mathematical structure:

- Squaring g_Φ in (3) means we cannot reproduce the cotangent Laplacian operator for poorly-conditioned meshes with negative cotangent weights.
- We arbitrarily choose one of three possible cyclic orderings of the inputs to f_Φ in (2).
- Similarly, we arbitrarily choose among two orderings of the inputs to g_Φ in (3): (v_1, v_2, v_3, v_4) and (v_2, v_1, v_4, v_3) .

All three items above could be addressed at the cost of increasing the complexity of f_Φ, g_Φ , but building more general semidefiniteness conditions and/or order invariance did not bring practical benefit.

4.2 Vectorial operators

L discretizes operators that act on functions discretized using one value per vertex of a triangle mesh. We also can discretize operators acting on *vector-valued* functions with a value in $\mathbb{R}^k$ per vertex by adjusting our construction. For example, for planar triangle meshes and $k = 2$ we can reproduce the Killing operator described in [Claici et al. 2017; Solomon et al. 2011a]; for $k = 4$ we can mimic the Dirac operator used for shape analysis by Liu et al. [2017].

To extend to vectorial operators, we use a $k|E| \times k|V|$ block version of d whose blocks are given as follows:

$$d_{ev} = \begin{cases} I_{k \times k} & \text{if } v = v_2 \\ -I_{k \times k} & \text{if } v = v_1 \\ 0 & \text{otherwise,} \end{cases}$$

where $I_{k \times k}$ denotes the $k \times k$ identity matrix.

We generalize $f_\Phi : \mathbb{R}^{3D} \rightarrow \mathbb{R}^{k \times k}$ and $g_\Phi : \mathbb{R}^{4D} \rightarrow \mathbb{R}^{k \times k}$ to output $k \times k$ matrices. Then, we compute $\star_0$ and $\star_1$ as block diagonal matrices whose elements are as follows:

$$(\star_0)_{vv} = \varepsilon I_{k \times k} + \sum_{t \sim v} f_\Phi(F_{v_1}, F_{v_2}, F_{v_3})^\top f_\Phi(F_{v_1}, F_{v_2}, F_{v_3}) \quad (4)$$

$$(\star_1)_{ee} = \varepsilon I_{k \times k} + g_\Phi(F_{v_1}, F_{v_2}, F_{v_3}, F_{v_4})^\top g_\Phi(F_{v_1}, F_{v_2}, F_{v_3}, F_{v_4}). \quad (5)$$

These definitions generalize our scalar construction for the case of $k = 1$ and still lead to a semidefinite matrix $L = \star_0^{-1} d^\top \star_1 d \in \mathbb{R}^{k|V| \times k|V|}$.

5 DIFFERENTIABLE SPECTRAL ANALYSIS

Now that we have determined the form of our operator L , we turn the task of using its low-order eigenvalues and eigenvectors for learning tasks. The key challenge will be to differentiate through eigenvalue/eigenvector computation, a task we consider below. While general eigencomputation is extremely difficult for learning, we show how our particular form (1) for L facilitates backpropagation and reduces dependence on random-access computation.

Recall that a step of training requires evaluating the loss function $\mathcal{L}$ and its gradients with respect to the parameters Φ of the model. The loss function is evaluated in a *forward pass*, and the gradients are evaluated during *backpropagation*. We will perform both the forward and backward pass of our model on the CPU so as to take advantage of a sparse solver to compute a set of eigenvalues/eigenvectors for the operator L efficiently. While the computation of our features for our learning problem as well as the entirety of backpropagation

could be efficiently computed on the GPU, our model has sufficiently few parameters that we find it unnecessary to transfer data between GPU and CPU.

5.1 The HodgeNet Generalized Eigenproblem

Our architecture outputs features built from eigenvectors of L in (1). Recall that L —and, in particular, the Hodge stars $\star_0, \star_1$ —is built from a matrix $F \in \mathbb{R}^{|V| \times D}$ of per-vertex features:

$$L = \star_0(F)^{-1} d^\top \star_1(F) d.$$

Hence, our features are built from eigenvectors $x^i \in \mathbb{R}^{k|V|}$ satisfying

$$Lx^i = \lambda^i x^i \iff d^\top \star_1 d x^i = \lambda^i \star_0 x^i. \quad (6)$$

By construction, $d^\top \star_1 d \geq 0$ and $\star_0 \geq 0$, so $\lambda^i \geq 0$. By convention, we normalize eigenvectors to satisfy the condition $(x^i)^\top \star_0 x^j = \delta_{ij}$, possible thanks to symmetry of our operators.

To differentiate our per-vertex features, we need to differentiate the eigenvectors x^i and eigenvalues λ^i with respect to the parameters Φ of our learned functions f_Φ, g_Φ . The expressions in §4 for $(\star_0)_{vv}(F)$ and $(\star_1)_{ee}(F)$ are readily differentiated. Hence, for compatibility with the backpropagation algorithm for differentiation, we need to solve the following problem involving our loss function $\mathcal{L}$:

Given the partial derivatives $\partial \mathcal{L} / \partial \lambda^i$ and $\partial \mathcal{L} / \partial x_j^i$ for all i, j , compute the partial derivatives $\partial \mathcal{L} / \partial (\star_0)_{vv}$ and $\partial \mathcal{L} / \partial (\star_1)_{ee}$ for all $v \in V, e \in E$.

In words, given derivatives of the loss function with respect to the eigenvalues and eigenvectors of L , compute the derivatives of the loss function with respect to the Hodge stars.

In general, differentiating through eigenvalue problems is expensive. Libraries like TensorFlow and PyTorch allow for differentiation of computing the *full* spectrum of a matrix, but their implementations (1) cannot account for the sparsity structure of our mesh and (2) cannot target a few eigenvalues close to 0, which are typically the meaningful eigenvalues to compute in geometry processing applications. Solving the full eigenvalue problem is extremely expensive computationally, and storing a $k|V| \times k|V|$ matrix of eigenvectors is prohibitive.

Our pipeline addresses the issues above. We rely on CPU-based sparse eigensolvers during the forward pass of our network, solving (6) only for a subset of eigenvalues. This alleviates dependence on $k|V| \times k|V|$ dense matrices and instead only stores the $O(k|V|)$ nonzero entries.

5.2 Derivative Formulas

The vectorial operator L operates on vectors in $\mathbb{R}^k$ per vertex on a mesh. Following §4.2, we will use x_{ℓ}^i to refer to the ℓ -th element ($\ell \in \{1, \dots, k\}$) of entry v ($v \in V$) of the i -th eigenvector of L . We use $\star_{0\ell m}$ to refer to the element (ℓ, m) of the $k \times k$ block of $\star_0$ at vertex $v \in V$. More generally, we will use subscripts to refer to matrix elements and superscripts to index over eigenvalues.

Define the following tensors:

$$\begin{aligned} y_{\ell}^i &:= \sum_{v \in V} d_{ev} x_{\ell}^i \\ M_{ij} &:= \begin{cases} (\lambda^i - \lambda^j)^{-1} & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases} \\ N_{ij} &:= \begin{cases} \lambda^i / (\lambda^j - \lambda^i) & \text{if } i \neq j \\ -1/2 & \text{otherwise.} \end{cases} \end{aligned}$$

We compute y during the forward pass as dx and cache the result for use during backpropagation, since d is a sparse matrix.

Our algorithm relies on the following proposition:

PROPOSITION 5.1. *We can backpropagate derivatives of our loss function as follows:*

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \star_{0\ell m}} &= - \sum_i \frac{\partial \mathcal{L}}{\partial \lambda^i} \lambda^i x_{\ell}^i x_{wm}^i + \sum_{i \neq j} \frac{\partial \mathcal{L}}{\partial x_{vn}^i} N_{ij} x_{\ell}^j x_{wm}^i x_{vn}^j \\ \frac{\partial \mathcal{L}}{\partial \star_{1\ell m}} &= \sum_i \frac{\partial \mathcal{L}}{\partial \lambda^i} y_{\ell}^i y_{em}^i + \sum_{i \neq j} \frac{\partial \mathcal{L}}{\partial x_{vn}^i} M_{ij} x_{vn}^j y_{\ell}^i y_{em}^i \end{aligned} \quad (7)$$

Here, i, j index over the eigenvectors of L ; ℓ, n, m index over vector elements from 1 to k ; v, w are vertices of the mesh; and e is an edge.

We defer proof to the supplemental document, since it requires a fairly involved computation. That said, this proposition is roughly an application of standard derivative-of-eigenvalue formulae to our operator L in (1), which benefits from the fact that our differentiable parameters are in *diagonal* matrices $\star_0, \star_1$.

The expressions in (7) may appear complicated, but in reality they are efficiently computable. We have eliminated all sparse matrices and inverses from these formulas, which are readily implemented using a one-line call to Einstein summation functions in deep learning toolkits (e.g., `einsum` in PyTorch).

5.3 Derivative Approximation

Here we briefly address one challenge using Proposition 5.1 to differentiate HodgeNet. Recall from §5.1 that we compute an incomplete set of eigenvectors of L , far fewer than the largest possible number. This choice is reasonable for constructing a loss function, which will only depend on this low-order eigenstructure. However, (7) requires *all* eigenvectors of L to evaluate the sums over j .

We use a simple strategy to address this issue. During the forward pass we compute and cache more eigenvalues/eigenvectors than are needed to evaluate $\mathcal{L}$; in practice, we use $2\times$ (see §8.5 for validation). Then, in backpropagation we truncate the sums over j in (7) to include only these terms. A straightforward argument reveals that the resulting gradient approximation still yields a descent direction for $\mathcal{L}$.

The first term in each sum is computable from exclusively the partial set of eigenvalues, implying we can exactly differentiate $\mathcal{L}$ with respect to the eigenvalues λ^i ; our approximation is only relevant to the eigenvectors.

6 FROM EIGENVECTORS TO FEATURES

Recall that our broad task is to design a learnable mapping from meshes to task-specific features. So far, we have designed a learnable

operator from mesh geometry and provided a means of differentiating through its eigenvectors/eigenvalues. It is tempting to use the eigenvectors as per-vertex features, but this is not a suitable choice: The choice of sign $\pm x_i$ for each eigenvector is arbitrary.

We return to geometry processing for inspiration. Classical shape descriptors built from operator eigenfunctions circumvent the sign issue by *squaring* the Laplacian eigenfunctions pointwise. For instance, the heat kernel signature [Sun et al. 2009], wave kernel signature [Aubry et al. 2011], and general learned kernels [Litman and Bronstein 2013] take the form

$$\sum_i f(\bar{\lambda}^i) \psi^i(p)^2,$$

where $\bar{\lambda}^i$ is the i -th eigenvalue and ψ^i is the i -th eigenvector of the Laplacian. The fact that ψ^i is squared alleviates sign dependence. Similarly, for eigenfunctions of the vectorial *vector Laplacian*, sign-agnostic features can be computed from the outer product of the pointwise vector and itself $\psi^i(p)\psi^i(p)^\top \in \mathbb{R}^{k \times k}$ (see, e.g., [Singer and Wu 2012, eq. (3.13)]).

Generalizing the hand-designed features above, we construct a sign-agnostic learnable per-vertex feature as follows. Take m to be the number of eigenvectors of L we will use to compute features, and take n to be the number of output features. We learn a function $h_\Phi : \mathbb{R} \rightarrow \mathbb{R}^n$ and construct a matrix $H \in \mathbb{R}^{m \times n}$ whose columns are $h(\lambda^i)$ for $i \in \{1, \dots, m\}$. Then, for $j \in \{1, \dots, n\}$, the j -th output feature at vertex $v \in V$, denoted G_v^j , is given by:

$$G_v^j := \sum_i H_{ij} \cdot (x_v^i)(x_v^i)^\top,$$

where x_v^i denotes the i -th eigenvector of L evaluated at vertex v as a $k \times 1$ column vector. We omit the 0 eigenvalue corresponding to the constant eigenfunction. We give our form for h_Φ in §7.

Having computed per-vertex features G_v , we optionally follow a standard max pooling approach to obtain per-face features

$$G_f = \max_{v \sim f} G_v,$$

or per-mesh features

$$G_M = \max_{v \in V} G_v$$

depending on the learning task at hand. We map these features to the desired output dimensions d using a learned function $o_\Phi : \mathbb{R}^n \rightarrow \mathbb{R}^d$.

7 ADDITIONAL DETAILS AND PARAMETERS

We model each of f_Φ , g_Φ , h_Φ , o_Φ as an MLP with batch normalization and Leaky ReLU nonlinearity [Maas et al. 2013] before each hidden layer. f_Φ and g_Φ each consist of four hidden layers, each of size 32; h_Φ consists of four hidden layers, each of size n ; and o_Φ consists of two hidden layers, each of size 32, except for the classification task, where the layers have 64 units. In all our experiments, we set vector dimensionality $k = 4$, output feature size $n = 32$, and number of eigenpairs used $m = 32$. We use an additional 32 eigenpairs for improved derivative approximation, as described in §5.3.

We train our network using the optimizer AdamW [Loshchilov and Hutter 2019] with a batch size of 16 and learning rate of 0.0001. We use gradient clipping with maximum norm of 1.0 to stabilize training. We implement our pipeline in PyTorch, using SciPy eigsh

with ARPACK for solving our sparse eigenproblem and libigl for mesh processing. We train our models on 128 2.5 GHz CPUs.

8 EXPERIMENTS

We demonstrate the efficacy of our method on several shape analysis tasks and provide experiments justifying some of our parameter and design choices. We also compare to state-of-the-art methods developed for learning on meshes. Other geometric deep learning approaches tend to use GNNs, ignoring the structure and relying on multiple layers to aggregate global data, whereas our method uses spectral geometry to infer global information from local features.

8.1 Mesh Segmentation

We train our network for the task of mesh segmentation on four datasets—the Human Body dataset [Maron et al. 2017] and the vase, chair, and alien categories of the Shape COSEG dataset [Wang et al. 2012]—optimizing the standard cross entropy loss. We use the same version of the Human Body dataset as in [Hanocka et al. 2019; Milano et al. 2020], which is downsampled to 2000 faces per mesh. We evaluate on the test set of the Human Body dataset, and generate a random 85%-15% train-test split for each Shape COSEG category, as in [Hanocka et al. 2019; Milano et al. 2020]. We train for 100 epochs (about 3 hours), randomly decimating each input mesh to a resolution of 1000-2000 faces and randomly applying anisotropic scaling of up to 5% in each dimension. We then fine-tune by training for 100 more epochs without decimation or scaling. In the case of the Human Body dataset, where meshes are not canonically rotated, we also apply random rotations as data augmentation to the training set. We center each mesh about the vertex center of mass and rescale to fit inside the unit sphere.

Table 1. Segmentation accuracy on the Human Body test set.

Method	# Parameters	Accuracy
Ours	31,720	85.03%
PD-MeshNet [Milano et al. 2020]	173,728	85.61%
MeshCNN [Hanocka et al. 2019]	2,279,720	85.39%

Table 2. Area-weighted segmentation accuracy on Human Body test set.

Method	Accuracy
Ours	86.48%
PD-MeshNet [Milano et al. 2020]	86.45%

Table 3. Test segmentation accuracy on Shape COSEG.

Method	Vases	Chairs	Aliens
Ours	90.30%	95.68%	96.03%
PD-MeshNet [Milano et al. 2020]	95.36%	97.23%	98.18%
MeshCNN [Hanocka et al. 2019]	92.36%	92.99%	96.26%

Fig. 3. Segmentation results on the Shape COSEG dataset. Meshes shown above are randomly selected from the test set for each category.

Table 4. Area-weighted test segmentation accuracy on Shape COSEG.

Method	Vases	Chairs	Aliens
Ours	94.38%	99.22%	97.97%
PD-MeshNet [Milano et al. 2020]	97.49%	97.86%	98.66%

We report segmentation accuracies in Tables 1 and 3 and area-weighted segmentation accuracies in Tables 2 and 4. For fair comparison, as in [Milano et al. 2020], we report accuracies based on “hard” ground-truth segmentation face labels for MeshCNN [Hanocka et al. 2019] rather than “soft” edge labels; see [Milano et al. 2020, Supplementary Material, Section H] for details regarding the segmentation metrics. Our method obtains results comparable to state-of-the-art for each dataset while requiring significantly fewer learnable parameters. We also show our learned segmentations on the entire Human Body test set in Figure 4 and on a sampling of the Shape COSEG test sets in Figure 3.

8.2 High-Resolution Mesh Segmentation

In contrast to earlier works, our method is capable of training on dense, non-decimated mesh data. We demonstrate this by training a segmentation model on the MIT animation dataset [Vlasic et al. 2008], where each mesh consists of 20,000 faces. We pre-initialize our model with the segmentation model trained on the Human Body dataset above and train for an additional 30 epochs, approximately 4 hours. The pre-initialization allows us

Table 5. Mesh segmentation accuracies for five random test splits of the full-resolution MIT animation dataset.

Split #	Accuracy
1	90.57%
2	86.90%
3	90.02%
4	89.07%
5	90.15%

to save training time by avoiding training a model from scratch: our model trained on low-resolution mesh data is able to capture some triangulation-invariant features, making this transfer learning possible. We train on five randomly sampled 95%-5% train-test splits and achieve a mean accuracy of 89.34%. We report segmentation accuracies for each split in Table 5 and render Split 1 in Figure 1 and Split 2 in Figure 5.

8.3 Mesh Classification

We evaluate our method on mesh classification on the downsampled version of the SHREC dataset [Lian et al. 2011], as in [Hanocka et al. 2019; Milano et al. 2020], optimizing the standard cross entropy loss. We report our results on two different splits of the dataset—*Split 10*, where each of the 30 shape categories is randomly split into 10 test and 10 training examples, and *Split 16*, where each category is split into 4 test and 16 training examples—in Table 6. We train for 100 epochs using decimation to 400–500 faces, anisotropic scaling, and random rotations as data augmentation and then fine-tune for another 100 epochs for *Split 16* and 200 epochs for *Split 10*.

Table 6. Classification accuracy on the SHREC test set.

Method	Split 16	Split 10
Ours	99.17%	94.67%
PD-MeshNet [Milano et al. 2020]	99.7%	99.1%
MeshCNN [Hanocka et al. 2019]	98.6%	91.0%

8.4 Dihedral Angle Prediction

As a stress test, we demonstrate that our method is capable of learning an operator that is sensitive to extrinsic geometry. To this end, we propose a synthetic dataset for dihedral angle regression. Previous methods that rely on computing a Laplacian would necessarily fail at this task, as they are only aware of intrinsic structure.

We take a regular mesh of a flat square consisting of 100 faces and crease it down the center at a random angle $\theta \in [0, 2\pi]$, as shown. Our network learns a two-dimensional vector per mesh, and we optimize cosine distance to the ground truth θ . We use the same hyperparameters as for the other experiments with a batch size of 32. For this experiment, we only use vertex positions as the input features—we do not provide normals. After training for just 15 minutes, we are able to predict the angle with an average error of 0.17°.

8.5 Ablation

We perform an ablation study to justify some the design and parameter choices in our architecture. In Table 7, we report test accuracy on the Shape COSEG vases dataset after 100 epochs of training (without fine-tuning). The accuracy degrades when we do not provide normals as part of the input mesh features, when we do not cache any additional eigenpairs for improved derivative approximation, when we reduce the vector dimensionality k , when we reduce the learned feature size n , or when we use fewer eigenpairs m for feature computation.

Fig. 4. Mesh segmentation results on the Human Body test set.

Fig. 5. Test set mesh segmentations of Split 2 of the full-resolution MIT animation dataset.

Table 7. Ablation study of our parameter choices on segmentation of the Shape COSEG vases dataset.

Model	Accuracy
full	87.78%
no normals	86.26%
no additional eig.	87.44%
$k = 2$	79.02%
$m = 16$	86.34%
$n = 8$	87.08%

9 DISCUSSION AND CONCLUSION

HodgeNet has many features that make it an attractive alternative for learning from meshes. During inference, its structure resembles that of most spectral geometry processing algorithms: construct a useful operator, and compute features from its spectrum. Our model is lightweight in the sense that the learnable functions act only on local neighborhoods, yet our model has a global receptive field thanks to the eigenvector computation. It has relatively few parameters and can be evaluated efficiently on the CPU.

This exploratory work suggests many avenues for future research. The most obvious next step is to extend our model to tetrahedral meshes for volumetric problems; we do not anticipate any major issues with this extension. We also can use our method’s connection to DEC to make learnable versions of other discrete differential operators, e.g. ones acting on k -forms for $k \geq 1$, and we can consider other applications of learning on meshes like generative modeling.

Our work also reveals some insight into other learning problems. Our architecture could easily be applied to graphs rather than triangle meshes, essentially by mildly changing the parameterization of $\star_1$ and taking $\star_0$ to be the identity matrix; we hence anticipate that there may be some applications to network analysis and other graph learning problems. Our lightweight differentiation strategy for eigenvectors may also prove useful in other contexts demanding eigenstructure of large matrices.

From the broadest perspective, our work demonstrates one of many potential applications of differentiable sparse linear algebra. While our derivative approximations and specially-formulated operator provide one way to circumvent development of a general framework for combining deep learning and linear algebra, a framework coupling sparse linear algebra to deep learning toolkits would enable a vast set of modeling choice and applications currently hamstrung by available architectures.

ACKNOWLEDGMENTS

The MIT Geometric Data Processing group acknowledges the generous support of Army Research Office grant W911NF2010168, of Air Force Office of Scientific Research award FA9550-19-1-031, of National Science Foundation grant IIS-1838071, from the CSAIL Systems that Learn program, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, from a gift from Adobe Systems, from an MIT.nano Immersion Lab/NCSoft Gaming Program seed grant, and from the Skoltech-MIT Next Generation Program. This work was also supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374.

REFERENCES

- James Atwood and Don Towsley. 2016. Diffusion-convolutional neural networks. In *Proc. NeurIPS*. 1993–2001.
- Oscar Kin-Chung Au, Youyi Zheng, Menglin Chen, Pengfei Xu, and Chiew-Lan Tai. 2011. Mesh segmentation with concavity-aware fields. *IEEE Transactions on Visualization and Computer Graphics* 18, 7 (2011), 1125–1134.
- Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV workshops*. 1626–1633.
- Mirela Ben-Chen, Adrian Butscher, Justin Solomon, and Leonidas Guibas. 2010. On discrete Killing vector fields and patterns on surfaces. In *Computer Graphics Forum*.

- Vol. 29. 1701–1711.
- Davide Boscaini, Davide Eynard, Drosos Kourounis, and Michael Bronstein. 2015a. Shape-from-operator: Recovering shapes from intrinsic operators. In *Computer Graphics Forum*, Vol. 34. 265–274.
- Davide Boscaini, Jonathan Masci, Simone Melzi, Michael Bronstein, Umberto Castellani, and Pierre Vanderghenst. 2015b. Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. In *Computer Graphics Forum*, Vol. 34. 13–23.
- Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. 2016. Learning shape correspondence with anisotropic convolutional neural networks. In *Proc. NeurIPS*.
- Alexander Bronstein, Michael Bronstein, Leonidas Guibas, and Maks Ovsjanikov. 2011. Shape Google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics* 30, 1 (2011), 1–20.
- Michael Bronstein and Alexander Bronstein. 2010. Shape recognition with spectral distances. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 5 (2010), 1065–1071.
- Michael Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vanderghenst. 2017. Geometric deep learning: going beyond Euclidean data. *IEEE Signal Processing Magazine* 34, 4 (2017), 18–42.
- Michael Bronstein and Iasonas Kokkinos. 2010. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR*. 1704–1711.
- Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2013. Spectral networks and locally connected networks on graphs. *arXiv:1312.6203* (2013).
- Max Budninskiy, Ameer Abdelaziz, Yiyang Tong, and Mathieu Desbrun. 2020. Laplacian-optimized diffusion for semi-supervised learning. *Computer Aided Geometric Design* 79 (2020), 101864.
- Ken Chan and Justin Wan. 2013. Reconstruction of missing cells by a Killing energy minimizing nonrigid image registration. In *International Conference of the IEEE Engineering in Medicine and Biology Society*. 3000–3003.
- Yoni Choukroun, Gautam Pai, and Ron Kimmel. 2018. Sparse approximation of 3D meshes using the spectral geometry of the Hamiltonian operator. *Journal of Mathematical Imaging and Vision* 60, 6 (2018), 941–952.
- Sebastian Claić, Mikhail Bessmeltsev, Scott Schaefer, and Justin Solomon. 2017. Isometry-aware preconditioning for mesh parameterization. In *Computer Graphics Forum*, Vol. 36. 37–47.
- Etienne Corman, Justin Solomon, Mirela Ben-Chen, Leonidas Guibas, and Maks Ovsjanikov. 2017. Functional characterization of intrinsic and extrinsic geometry. *ACM Transactions on Graphics* 36, 2 (2017), 1–17.
- Luca Cosmo, Mikhail Panine, Arianna Rampini, Maks Ovsjanikov, Michael Bronstein, and Emanuele Rodolà. 2019. Isospectralization, or how to hear shape, style, and correspondence. In *Proc. CVPR*. 7529–7538.
- Fernando de Goes, Mathieu Desbrun, and Yiyang Tong. 2016. Vector field processing on triangle meshes. In *ACM SIGGRAPH 2016 Courses*. 1–49.
- Pim de Haan, Maurice Weiler, Taco Cohen, and Max Welling. 2020. Gauge Equivariant Mesh CNNs: Anisotropic convolutions on geometric graphs. *arXiv:2003.05425* (2020).
- Michaël Defferrard, Xavier Bresson, and Pierre Vanderghenst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. *Proc. NeurIPS* 29 (2016), 3844–3852.
- Mathieu Desbrun, Anil Hirani, Melvin Leok, and Jerrold E Marsden. 2005. Discrete exterior calculus. *arXiv preprint math/0508341* (2005).
- Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. 2020. Deep geometric functional maps: Robust feature learning for shape correspondence. In *Proc. CVPR*. 8592–8601.
- Moshe Eliasof and Eran Treister. 2020. DiffGCN: Graph convolutional networks via differential operators and algebraic multigrid pooling. *Proc. NeurIPS* (2020).
- Yutong Feng, Yifan Feng, Haoxuan You, Xibin Zhao, and Yue Gao. 2019. MeshNet: Mesh neural network for 3D shape representation. In *Proc. AAAI*, Vol. 33. 8279–8286.
- Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. 2018. SplineCNN: Fast geometric deep learning with continuous B-spline kernels. In *Proc. CVPR*. 869–877.
- Jun Gao, Wenzheng Chen, Tommy Xiang, Alec Jacobson, Morgan McGuire, and Sanja Fidler. 2020. Learning deformable tetrahedral meshes for 3D reconstruction. *Proc. NeurIPS* 33 (2020).
- Xianfeng Gu, Steven Gortler, and Hugues Hoppe. 2002. Geometry images. In *Proc. SIGGRAPH*. 355–361.
- Niv Haim, Nimrod Segol, Heli Ben-Hamu, Haggai Maron, and Yaron Lipman. 2019. Surface networks via general covers. In *Proc. ICCV*. 632–641.
- Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. 2019. MeshCNN: A network with an edge. *ACM Transactions on Graphics* 38, 4 (2019), 1–12.
- Rana Hanocka, Gal Metzer, Raja Giryes, and Daniel Cohen-Or. 2020. Point2Mesh: A self-prior for deformable meshes. *ACM Transactions on Graphics* 39, 4 (July 2020).
- Wenchong He, Zhe Jiang, Chengming Zhang, and Arpan Man Sainju. 2020. CurvaNet: Geometric deep learning based on directional curvature for 3D shape analysis. In *Proc. KDD*. 2214–2224.
- Mikaël Henaff, Joan Bruna, and Yann LeCun. 2015. Deep convolutional networks on graph-structured data. *arXiv:1506.05163* (2015).
- Amir Hertz, Rana Hanocka, Raja Giryes, and Daniel Cohen-Or. 2020. Deep geometric texture synthesis. *ACM Transactions on Graphics* 39, 4 (July 2020).
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2012. Modal shape analysis beyond Laplacian. (2012).
- Jingwei Huang, Haotian Zhang, Li Yi, Thomas Funkhouser, Matthias Nießner, and Leonidas Guibas. 2019. TextureNet: Consistent local parametrizations for learning from high-resolution signals on meshes. In *Proc. CVPR*. 4440–4449.
- Qi-Xing Huang, Martin Wicke, Bart Adams, and Leonidas Guibas. 2009. Shape decomposition using modal analysis. In *Computer Graphics Forum*, Vol. 28. 407–416.
- Evangelos Kalogerakis, Melinos Averkiou, Subhansu Maji, and Siddhartha Chaudhuri. 2017. 3D shape segmentation with projective convolutional networks. In *Proc. CVPR*. 3779–3788.
- Thomas Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proc. ICLR*.
- Ilya Kostrikov, Zhongshi Jiang, Daniele Panozzo, Denis Zorin, and Joan Bruna. 2018. Surface networks. In *Proc. CVPR*. 2540–2548.
- Artiom Kovnatsky, Michael Bronstein, Alexander Bronstein, and Ron Kimmel. 2011. Photometric heat kernel signatures. In *International Conference on Scale Space and Variational Methods in Computer Vision*. 616–627.
- Alon Lahav and Ayellet Tal. 2020. MeshWalker: Deep mesh understanding by random walks. *arXiv:2006.05353* (2020).
- Ron Levie, Federico Monti, Xavier Bresson, and Michael Bronstein. 2018. CayleyNets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Transactions on Signal Processing* 67, 1 (2018), 97–109.
- Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoué, H. V. Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, and D. Vandermeulen. 2011. SHREC '11 Track: Shape Retrieval on Non-rigid 3D Watertight Meshes. In *Eurographics Workshop on 3D Object Retrieval*, H. Laga, T. Schreck, A. Ferreira, A. Godil, I. Pratikakis, and R. Veltkamp (Eds.). The Eurographics Association. <https://doi.org/10.2312/3DOR/3DOR11/079-088>.
- Isaak Lim, Alexander Dielen, Marcel Campen, and Leif Kobbelt. 2018. A simple approach to intrinsic correspondence learning on unstructured 3D meshes. In *Proc. ECCV*.
- Or Litany, Tal Remez, Emanuele Rodolà, Alex Bronstein, and Michael Bronstein. 2017. Deep functional maps: Structured prediction for dense shape correspondence. In *Proc. ICCV*. 5659–5667.
- Roe Litman and Alexander Bronstein. 2013. Learning spectral descriptors for deformable shape correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 1 (2013), 171–180.
- Hsueh-Ti Derek Liu, Alec Jacobson, and Keenan Crane. 2017. A Dirac operator for extrinsic shape analysis. In *Computer Graphics Forum*, Vol. 36. 139–149.
- Hsueh-Ti Derek Liu, Vladimir Kim, Siddhartha Chaudhuri, Noam Aigerman, and Alec Jacobson. 2020. Neural subdivision. *ACM Transactions on Graphics* 39, 4 (July 2020).
- Rong Liu and Hao Zhang. 2007. Mesh segmentation via spectral embedding and contour analysis. In *Computer Graphics Forum*, Vol. 26. 385–394.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg6RiCqY7>.
- A.L. Maas, A.Y. Hannun, and A.Y. Ng. 2013. Rectifier Nonlinearities Improve Neural Network Acoustic Models. In *Proceedings of the International Conference on Machine Learning*. Atlanta, Georgia.
- Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir Kim, and Yaron Lipman. 2017. Convolutional neural networks on surfaces via seamless toric covers. *ACM Transactions on Graphics* 36, 4 (2017).
- Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vanderghenst. 2015. Geodesic convolutional neural networks on Riemannian manifolds. In *Proc. ICCV*. 37–45.
- Francesco Milano, Antonio Loquercio, Antoni Rosinol, Davide Scaramuzza, and Luca Carlone. 2020. Primal-Dual Mesh Convolutional Neural Networks. In *Proc. NeurIPS*.
- Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, and Michael Bronstein. 2017. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *Proc. CVPR*.
- Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics* 31, 4 (2012), 1–11.
- Maks Ovsjanikov, Quentin Mérigot, Facundo Mémoli, and Leonidas Guibas. 2010. One point isometric matching with the heat kernel. In *Computer Graphics Forum*, Vol. 29. 1555–1564.
- Maks Ovsjanikov, Jian Sun, and Leonidas Guibas. 2008. Global intrinsic symmetries of shapes. In *Computer Graphics Forum*, Vol. 27. 1341–1348.
- Hao Pan, Shilin Liu, Yang Liu, and Xin Tong. 2018. Convolutional neural networks on 3D surfaces using parallel frames. *arXiv:1808.04952* (2018).
- Adrien Poulenard and Maks Ovsjanikov. 2018. Multi-directional geodesic neural networks via equivariant convolution. *ACM Transactions on Graphics* 37, 6 (2018),

- 1–14.
- Yi-Ling Qiao, Lin Gao, Paul Rosin, Yu-Kun Lai, Xilin Chen, et al. 2020. Learning on 3D meshes with Laplacian encoding and pooling. *IEEE Transactions on Visualization and Computer Graphics* (2020).
- Dan Raviv, Michael Bronstein, Alexander Bronstein, and Ron Kimmel. 2010. Volumetric heat kernel signatures. In *Proceedings of the ACM Workshop on 3D Object Retrieval*. 39–44.
- Dan Raviv, Michael Bronstein, Alexander Bronstein, Ron Kimmel, and Nir Sochen. 2011. Affine-invariant diffusion geometry for the analysis of deformable 3D shapes. In *CVPR 2011*. 2361–2367.
- Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. 2006. Laplace–Beltrami spectra as ‘shape-DNA’ of surfaces and solids. *Computer-Aided Design* 38, 4 (2006), 342–366.
- Raif Rustamov. 2007. Laplace–Beltrami eigenfunctions for deformation invariant shape representation. In *Symposium on Geometry Processing*. 225–233.
- Lisa Schneider, Annika Niemann, Oliver Beuing, Bernhard Preim, and Sylvia Saalfeld. 2020. MedMeshCNN–Enabling MeshCNN for medical surface models. *arXiv:2009.04893* (2020).
- Stefan Schonsheck, Bin Dong, and Rongjie Lai. 2018. Parallel transport convolution: A new tool for convolutional neural networks on manifolds. *arXiv:1805.07857* (2018).
- Jonas Schult, Francis Engelmann, Theodora Kontogianni, and Bastian Leibe. 2020. DualConvMesh-Net: Joint geodesic and Euclidean convolutions on 3D meshes. In *Proc. CVPR*. 8612–8622.
- Konstantinos Sfikas, Theoharis Theoharis, and Ioannis Pratikakis. 2017. Exploiting the PANORAMA representation for convolutional neural network classification and retrieval. *Proc. 3DOR* 6 (2017), 7.
- Abhishek Sharma and Maks Ovsjanikov. 2020. Weakly supervised deep functional map for shape matching. In *Proc. NeurIPS*.
- Nicholas Sharp, Souhaib Attaki, Keenan Crane, and Maks Ovsjanikov. 2020. Diffusion is all you need for learning on surfaces. *arXiv:2012.00888* (2020).
- Nicholas Sharp and Maks Ovsjanikov. 2020. PointTriNet: Learned triangulation of 3D point sets. In *Proc. ECCV*.
- Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. The vector heat method. *ACM Transactions on Graphics* 38, 3 (2019), 1–19.
- Baoguang Shi, Song Bai, Zhichao Zhou, and Xiang Bai. 2015. DeepPano: Deep panoramic representation for 3-d shape recognition. *IEEE Signal Processing Letters* 22, 12 (2015), 2339–2343.
- Amit Singer and H.-T. Wu. 2012. Vector diffusion maps and the connection Laplacian. *Communications on Pure and Applied Mathematics* 65, 8 (2012), 1067–1144.
- Ayan Sinha, Jing Bai, and Karthik Ramani. 2016. Deep learning 3D shape surfaces using geometry images. In *Proc. ECCV*. 223–240.
- Miroslava Slavcheva, Maximilian Baust, Daniel Cremers, and Slobodan Ilic. 2017. Killing-Fusion: Non-rigid 3D reconstruction without correspondences. In *Proc. CVPR*. 1386–1395.
- Justin Solomon, Mirela Ben-Chen, Adrian Butscher, and Leonidas Guibas. 2011a. As-Killing-as-possible vector fields for planar deformation. In *Computer Graphics Forum*, Vol. 30. 1543–1552.
- Justin Solomon, Mirela Ben-Chen, Adrian Butscher, and Leonidas Guibas. 2011b. Discovery of intrinsic primitives on triangle meshes. In *Computer Graphics Forum*, Vol. 30. 365–374.
- An Ping Song, Xin Yi Di, Xiao Kang Xu, and Zi Heng Song. 2020. MeshGraphNet: An effective 3D polygon mesh recognition With topology reconstruction. *IEEE Access* 8 (2020), 205181–205189.
- M. Spagnuolo, M. Bronstein, A. Bronstein, and A. Ferreira. 2012. Affine-invariant photometric heat kernel signatures. *Eurographics* (2012), 39–46.
- Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. 2015. Multi-view convolutional neural networks for 3D shape recognition. In *Proc. ICCV*. 945–953.
- Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. 2009. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer Graphics Forum*, Vol. 28. 1383–1392.
- Zhiyu Sun, Ethan Rooke, Jerome Charton, Yusen He, Jia Lu, and Stephen Baek. 2020. ZerNet: Convolutional neural networks on arbitrary surfaces via Zernike local tangent space estimation. In *Computer Graphics Forum*.
- Michael Tao, Justin Solomon, and Adrian Butscher. 2016. Near-isometric level set tracking. In *Computer Graphics Forum*, Vol. 35. 65–77.
- Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. 2018. Tangent convolutions for dense prediction in 3D. In *Proc. CVPR*. 3887–3896.
- Nitika Verma, Edmond Boyer, and Jakob Verbeek. 2018. FeastNet: Feature-steered graph convolutions for 3D shape analysis. In *Proc. CVPR*. 2598–2606.
- Daniel Vlasic, Ilya Baran, Wojciech Matusik, and Jovan Popović. 2008. Articulated mesh animation from multi-view silhouettes. In *ACM SIGGRAPH 2008 papers*. 1–9.
- Hao Wang, Tong Lu, Oscar Kin-Chung Au, and Chiew-Lan Tai. 2014. Spectral 3D mesh segmentation with a novel single segmentation field. *Graphical Models* 76, 5 (2014), 440–456.
- Pengyu Wang, Yuan Gan, Panpan Shui, Fenggen Yu, Yan Zhang, Songle Chen, and Zhengxing Sun. 2018b. 3D shape segmentation via shape fully convolutional networks. *Computers & Graphics* 70 (2018), 128–139.
- Yunhai Wang, Shmulik Asafi, Oliver Van Kaick, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. 2012. Active co-analysis of a set of shapes. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 1–10.
- Yu Wang, Mirela Ben-Chen, Iosif Polterovich, and Justin Solomon. 2018a. Steklov spectral geometry for extrinsic shape analysis. *ACM Transactions on Graphics* 38, 1 (2018), 1–21.
- Yu Wang, Vladimir Kim, Michael Bronstein, and Justin Solomon. 2019. Learning geometric operators on meshes. In *Representation Learning on Graphs and Manifolds (ICLR workshop)*.
- Yu Wang and Justin Solomon. 2019. Intrinsic and extrinsic operators for shape analysis. In *Handbook of Numerical Analysis*. Vol. 20. 41–115.
- Lingyu Wei, Qixing Huang, Duygu Ceylan, Etienne Vouga, and Hao Li. 2016. Dense human body correspondences using convolutional networks. In *Proc. CVPR*. 1544–1553.
- Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt. 2020. CNNs on surfaces using rotation-equivariant features. *ACM Transactions on Graphics* 39, 4 (2020).
- Haotian Xu, Ming Dong, and Zichun Zhong. 2017. Directionally convolutional networks for 3D shape segmentation. In *Proc. ICCV*. 2698–2707.
- Yuqi Yang, Shilin Liu, Hao Pan, Yang Liu, and Xin Tong. 2020. PFCNN: Convolutional neural networks on 3D surfaces using parallel frames. In *Proc. CVPR*. 13578–13587.
- Zi Ye, Olga Diamanti, Chengcheng Tang, Leonidas Guibas, and Tim Hoffmann. 2018. A unified discrete framework for intrinsic and extrinsic Dirac operators for geometry processing. In *Computer Graphics Forum*, Vol. 37. 93–106.
- Li Yi, Hao Su, Xingwen Guo, and Leonidas Guibas. 2017. SpecSyncCNN: Synchronized spectral CNN for 3D shape segmentation. In *Proc. CVPR*. 2282–2290.
- Wei Zeng, Ren Guo, Feng Luo, and Xianfeng Gu. 2012. Discrete heat kernel determines discrete Riemannian metric. *Graphical Models* 74, 4 (2012), 121–129.