Eurographics Symposium on Geometry Processing 2026
M. Skouras and M. Livesu
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 45 (2026), Number 5

Phong-Rodrigues Extrinsic Vector-Field Processing

Hongyi Liu[1] Oded Stein[2] Amir Vaxman[3] Mirela Ben-Chen[2] Misha Kazhdan[1]

[1]Johns Hopkins University, USA [2]Technion - Israel Institute of Technology, Israel [3]University of Edinburgh, UK

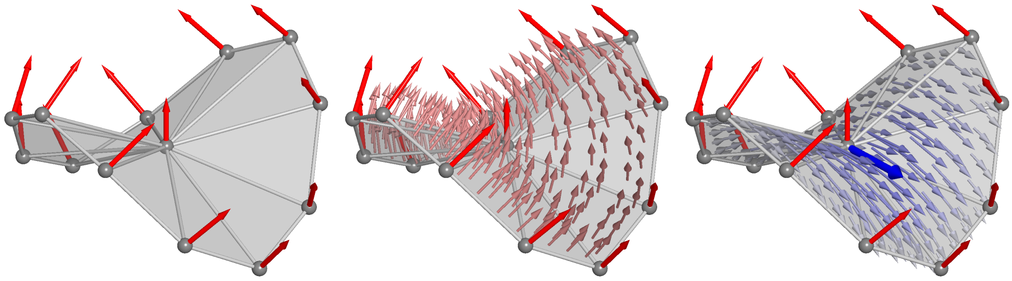

Figure 1: *To define a continuous basis for tangent vector-fields over a triangle mesh, we start with a triangle mesh with per-vertex normals (left). Using Phong interpolation, we obtain a continuous normal field over the mesh (center). Then, given a tangent vector at a vertex (dark blue), we extend it into the interior of the incident triangles by (1) transporting the vertex's tangent using the Rodrigues rotation taking the vertex's normal to the interior point's normal, and (2) scaling the transported tangent by the vertex's barycentric coordinate.*

Abstract

We introduce a new extrinsic discretization of tangent vector fields on triangle meshes that is continuous, with bounded derivatives that are continuous almost everywhere, supporting pointwise evaluation and integration of differential operators. We achieve this by building a continuous normal field over the mesh via Phong interpolation and using minimal Rodrigues rotations to transport vertex-based tangent vectors into triangle interiors. Unlike most existing discretizations, which typically sacrifice either continuity or the ability to evaluate derivatives pointwise, our approach supports both. Because it is pointwise evaluatable, and using the fact that the covariant derivative can be decomposed into its symmetric, antisymmetric, and scalar components, our discretization supports the construction of standard vector-field processing operators including the connection and Hodge Laplacians, Killing energy, divergence, curl, and the Lie bracket. This framework provides a simple and practical finite-element formulation for vector-field processing on meshes, supporting both integration-based operators and pointwise queries. To our knowledge, ours is the first discretization that jointly enables extrinsic continuous vector fields, bounded derivatives, and pointwise evaluation of this collection of operators.

CCS Concepts

• Computing methodologies → Mesh models; • Mathematics of computing → Quadrature; Differential calculus;

1. Introduction

Vector-field processing is an essential step in a broad class of geometry-processing applications, supporting the study of deformations of signals over a mesh and has been a corner-stone in a large body of work, including frame design, deformation, parallel transport, and optical flow. An essential first step is defining a discretization of vector-fields, enabling their representation, the study of their properties (e.g., smoothness), and the formulation of energies that have desired vector-fields as their minima.

A key challenge in discretizing vector-fields, and in particular in supporting the study of vector-field smoothness, is that the "natural" definition of vector-fields would represent them by values associated with faces, where a tangent space can be defined as the space perpendicular to the face normal. However, such representations are inherently discontinuous, due to the piecewise-constant triangle normals, resulting in unbounded derivatives across edges and making smoothness energies challenging to define.

These challenges have typically been addressed in one of two ways: Discrete Exterior Calculus (DEC) [Hir03] has been used to define discretizations of vector-fields and associated smoothness energies *without* explicitly constructing a basis. It provides a way to optimize vector-fields but only supports evaluation at specific locations and along specific directions. The second approach is to use a discontinuous vector-field basis (e.g. Nédélec/Whitney, Raviart-Thomas, or Crouzeix-Raviart [Hip99, Gat14, SJWG20]). These support point-wise evaluation and differentiation, but discontinuity along edges results in only some differential operators begin well-formulated in this representation.

An exception to this are the representations proposed by Knöppel *et al.* [KCPS13] and Liu *et al.* [LTGD16]. There, the authors propose extending an intrinsically represented tangent vector defined at a vertex into the incident triangles by a combination of parallel transport and linear scaling.

In our work we follow these approaches, replacing the intrinsic representation of tangent vectors with an explicit one, visualized in Fig. 1. To this end we assume we are given a triangle mesh with per-vertex normals as input (left), which we extend to a continuous normal field over the mesh (center). Then, given a tangent vector at a vertex, we extend that vector to a continuous vector-field over the mesh (right). Concretely, given a point in the interior of the incident triangle we define the tangent vector at that point by (1) using the Rodrigues rotation (taking the vertex's normal to the point's normal) to transport the tangent from the vertex to the point, and (2) scaling the transported tangent by the component of the point's barycentric coordinate associated with the vertex.

Our approach provides a simple expression for the vector-field basis that supports point-wise evaluation and differentiation. This, in turn, make it possible to define a covariant derivative that is point-wise evaluatable, enabling the construction of mass, stiffness (Hodge and Connection), and Killing energies. It also supports the point-wise evaluation of the Lie bracket of two vector-fields.

We demonstrate the efficacy of our approach in defining Connection, Hodge, and Killing energies, as well as in computing singularities, and estimating the Lie bracket. While subsets of these applications are supported by existing discretizations, the continuity of our vector-field basis makes it the first to support all of them.

2. Related Work

This work defines a continuous and pointwise-evaluatable extrinsic vector-field, together with its covariant derivative, for triangle meshes embedded in $\mathbb{R}^3$. This requires a choice of vector-field representation and a notion of transport between adjacent tangent spaces. Together these determine the covariant derivative, from which all other operators are derived. We review prior methods, focusing on continuity and pointwise-evaluability, and refer the reader to [VCD*16,dGDT16] for comprehensive reviews of vector-field representations and applications. A summary comparison is provided in Tab. 1.

	Rep.	C^0 a.e. and bounded	C^1 a.e. and bounded	DoFs
[KCPS13]	In.	✓	?	2 / vertex
[AOCBC15]	Ex.	✓	-	3 / vertex
[KCPS15]	In.	-	-	2 / vertex
[SSC19]	In.	-	-	2 / vertex
[LTGD16]	In.	✓	✓	2 / vertex
[Whi57]	In.	✓	-	1 / edge
[SWJG20]	In.	✓	-	2 / edge
[BV22]-PL	In.	✓	-	6 / face
Ours	Ex.	✓	✓	2 / vertex

Table 1: *Comparison of vector-field discretization methods by intrinsic/extrinsic representation, almost-everywhere continuity and boundedness of the vector field and its covariant derivative, and degrees of freedom. For [BV22] we describe the piecewise-linear (PL) version. (Though the discretization of Knöppel et al. [KCPS13] should have an evaluatable covariant derivative, the lack of closed-form expression makes this difficult in practice.)*

The most related work is the intrinsic discretization of Knöppel *et al.* [KCPS13] which defines a connection Laplacian on vertex-based tangent frames. These per-vertex frames are then extended into incident triangles via parallel transport along radii, yielding a connection Laplacian assembled in integrated form. Later works [KCPS15, SSC19] build on this representation, assembling the connection Laplacian purely from edge transport rotations without interpolating frames into triangle interiors, and are not pointwise-evaluatable.

As Liu *et al.* [LTGD16] noted, the connection Laplacian of Knöppel *et al.* [KCPS13] does not provide closed-form expressions for covariant derivatives, and first-order operators such as divergence and curl cannot be evaluated in their framework. Liu *et al.* [LTGD16] address this by additionally defining vertex-to-edge and vertex-to-triangle transition rotations, enabling explicit evaluation of the covariant derivatives. This requires introducing "impulse rotations" at chart crossings.

Degrees of freedom can also be placed on faces or edges. DEC provides a framework for discrete differential operators on triangle meshes, which can be realized in the finite-elements framework using the Whitney basis [WWT*06, FSDH07, DKT06]. Crouzeix-Raviart bases, with two degrees of freedom per edge encoding the components parallel and perpendicular to the edge, have also been proposed [Dje16, SWJG20]. Boksebeld and Vaxman [BV22] use a primal-dual decomposition, generalizing the connection Laplacian to higher order polynomials, with six degrees of freedom per face in the piecewise linear case. However, these three formulations all yield discontinuous vector fields.

While most methods are intrinsic, an extrinsic representation is given by Azencot *et al.* [AOCBC15] via differentiation of the ambient coordinates and projection onto the tangent plane. This builds on the operator view of vector-fields [ABCCO13] where the Lie bracket arises naturally from operator composition.

Our extrinsic approach follows Knöppel et al. [KCPS13] and Liu et al. [LTGD16] in placing degrees of freedom on vertices and extending vector fields into triangle interiors, but differs in two key ways. First, instead of using the intrinsic tangent frames derived from the triangle mesh, we define per-point tangent planes extrinsically via Phong interpolation, decoupling the tangent bundle from the mesh geometry. This choice makes ambient-space continuity possible, in contrast to prior methods that tend to focus on intrinsic continuity. Second, we define the covariant derivative by differentiating the extrinsic vector field in the ambient space and projecting onto the tangent plane, enabling the construction of a broad set of differential operators including the connection Laplacian, the Hodge Laplacian and the Lie bracket. This is similar to the approach of Azencot *et al.* [AOCBC15]. However, (1) the vector-fields we use are perpendicular to the normals by construction, and (2) our normals vary continuously over the surface.

3. Background

To define the covariant derivative and decompose it into its constituent parts, we leverage some basic linear algebra, which we briefly review here. We also recall the expression for the (minimal) Rodrigues rotation and introduce notation that we will use throughout the remainder of the discussion (Tab. 2).

3.1. Linear Algebra Review

We consider the covariant derivative of a vector-field v – a function assigning, to each point $p \in \mathcal{S}$ on the surface, an endomorphism on the tangent space at that point, $\nabla v|_p \in \text{End}(T_p\mathcal{S})$. Following earlier work, we compute an orthogonal decomposition of the covariant derivative into constituent components in order to extract the Killing, holomorphic, anti-holomorphic, etc. energies.

We review the linear algebra involved. Given a vector space V, we denote by V^* its *dual* – the space of linear functions on V. Given vector spaces V and W and a homomorphism $L : V \to W$, we denote by $L^* : W^* \to V^*$ the dual homomorphism defined in terms of the pull-back/composition, with $L^*(\phi) \equiv \phi \circ L$ for all $\phi \in W^*$. Finally, we denote an inner-product space as $\{V, B : V \to V^*\}$, with V the vector space, V^* its dual, and B the inner-product. Though the inner-product can be equivalently represented as a bilinear map, $B(v,w) \equiv [B(v)](w)$, its representation as a linear map (combined with the fact that inner-products are symmetric and positive-definite) makes the inverse $B^{-1} : V^* \to V$, well-defined.

Endomorphism Decomposition

As in previous work [dGLB*14, LTGD16], we decompose the space of endomorphisms into three orthogonal subspaces and consider the projection of individual endomorphisms onto the subspaces. We review the decomposition using a coordinate-free formulation. The subspaces are:

- $\mathcal{D}iv(V)$ – the space of scalar multiples: These describe infinitesimal isotropic scaling and, in the context of the covariant derivative, give the vector-field's divergence. (Represented with respect to any basis, these are scalar multiples of the identity.)
- $\overline{\mathcal{H}olom}(V)$ – the space of trace-free, self-adjoint operators: These describe infinitesimal, anisotropic, volume-preserving scaling and, in the context of the covariant derivative, give the vector-field's anti-holomorphic component. (Represented with respect to an orthonormal basis, these are symmetric matrices with vanishing trace.)
- $\mathcal{C}url(v)$ – the space of anti-self-adjoint operators: These describe infinitesimal rotations and, in the context of the covariant derivative, give the vector-field's curl. (Represented with respect to an orthonormal basis, these are skew-symmetric matrices.)

Formally, an endomorphism $L : V \to V$ is said to be *self-adjoint* if it commutes with the inner-product:

$$B \circ L = L^* \circ B.$$

Analogously, an endomorphism $L : V \to V$ is *anti-self-adjoint* if:

$$B \circ L = -L^* \circ B.$$

The subsets of self-adjoint and anti-self-adjoint endomorphisms form linear subspaces. Given an endomorphism $L : V \to V$, the orthogonal projections are, respectively:

$$L \mapsto \left(\frac{L + B^{-1} \circ L^* \circ B}{2}\right) \quad \text{and} \quad L \mapsto \left(\frac{L - B^{-1} \circ L^* \circ B}{2}\right).$$

Represented with respect to an orthonormal base, this amounts to averaging the matrix with its transpose (resp. negative transpose), resulting in symmetric (resp. skew symmetric) matrices.

The space of self-adjoint operators can be further decomposed into endomorphisms that are scalar multiples of the identity and endomorphisms with vanishing trace. Given a self-adjoint endomorphism $L \in \text{End}(V)$, the orthogonal projections onto the subspaces are, respectively:

$$L \mapsto \frac{\text{tr}(L)}{\dim(V)} \cdot \text{Id.} \quad \text{and} \quad L \mapsto L - \frac{\text{tr}(L)}{\dim(V)} \cdot \text{Id.}$$

In App. A we show that these maps are orthogonal projections (with respect to the canonical inner-product on the space of endmorphisms $\text{End}(V)$ induced by the inner-product on V) providing an orthogonal decomposition the space of endomorphisms as:

$$\text{End}(V) = \mathcal{D}iv(V) \oplus \overline{\mathcal{H}olom}(V) \oplus \mathcal{C}url(V).$$

3.2. Rodrigues Rotation

Given (non-antipodal) unit vectors $v, w \in S^{d-1}$, the Rodrigues rotation formula gives the minimal-angle rotation taking v to w:

$$R(v,w) \equiv \text{Id.} + (w \cdot v^\top - v \cdot w^\top) + \frac{(w \cdot v^\top - v \cdot w^\top)^2}{1 + \langle v, w \rangle} \in \text{SO}(d).$$

This expression is singular only when v and w are antipodal, in which case the rotation by 180° in *any* plane containing v will be minimal.) Otherwise, the coefficients of R are smooth functions of v and w, with derivatives that are readily computable.

3.3. Notation

Given a triangle with vertices $v_0, v_1, v_2 \in \mathbb{R}^3$, we denote the triangle's normal as:

$$n = \frac{(v_1 - v_0) \times (v_2 - v_0)}{|(v_1 - v_0) \times (v_2 - v_0)|}.$$

$R : S^2 \times S^2 \to \mathrm{SO}(3)$	Rodrigues rotation
$v_i \in \mathbb{R}^3$	Vertices of the triangle
$n_i \in S^2$	Normals of the triangle
$n \in S^2$	Normal defined by triangle $\{v_0, v_1, v_2\}$
$\mathbb{T} \subset \mathbb{R}^2$	Unit right triangle
$S^2 \subset \mathbb{R}^3$	Unit sphere
$T_p\mathbb{T} \simeq \mathbb{R}^2$	Tangent space at $p \in \mathbb{T}$
$T_nS^2 \subset \mathbb{R}^3$	Tangent space at $n \in S^2$
$\mathrm{End}(T_p\mathbb{T})$	The space of endomorphisms on $T_p\mathbb{T}$
$\Phi : \mathbb{T} \to \mathbb{R}^3$	Linear embedding of triangle $\{v_0, v_1, v_2\}$
$d\Phi\|_p : T_p\mathbb{T} \to T_nS^2$	Differential of Φ at $p \in \mathbb{T}$
$g\|_p : T_p\mathbb{T} \to T_p^*\mathbb{T}$	Metric on $T_p\mathbb{T}$ induced by $d\Phi\|_p$
$N : \mathbb{T} \to S^2$	Gauss map defined by vertex normals
$\widetilde{d\Phi}\|_p : T_p\mathbb{T} \to T_{N(p)}S^2$	Realization of $T_p\mathbb{T}$ in $T_{N(p)}S^2$
$\psi_i : \mathbb{T} \to \mathbb{R}$	Hat basis
$\mathbf{m}^\psi \in \mathbb{R}^{3\times 3}$	Element scalar mass matrix
$\mathbf{s}^\psi \in \mathbb{R}^{3\times 3}$	Element scalar stiffness matrix
$R_i : \mathbb{T} \to \mathrm{SO}(3)$	Transport from the i-th corner
$\overline{\omega}_i : \mathbb{T} \to \mathbb{R}^3$	Extrinsic tangent vector-field basis
$\omega_i \in \Gamma(T\mathbb{T})$	Pulled-back tangent vector-field basis
$\mathbf{m}^\omega \in \mathbb{R}^{3\times 3}$	Element vector mass matrix
$\mathbf{s}^\omega \in \mathbb{R}^{3\times 3}$	Element vector stiffness matrix
$\mathcal{V} \subset \mathbb{R}^3$	Mesh vertices
$\mathcal{T} \subset [0, \|\mathcal{V}\|)^3$	Mesh triangles

Table 2: *Notation for per-triangle system construction*

Computations are performed over the unit right triangle and unit sphere, $\mathbb{T} \subset \mathbb{R}^2$ and $S^2 \subset \mathbb{R}^3$:

$$\mathbb{T} = \left\{(s,t) \in [0,1]^2 \middle| s+t \le 1\right\} \quad \text{and} \quad S^2 = \left\{v \in \mathbb{R}^3 \middle| |v| = 1\right\}.$$

The tangent space at $p \in \mathbb{T}$ is equivalent to Euclidean 2-space and the tangent space at $v \in S^2$ is a subset of Euclidean 3-space:

$$T_p\mathbb{T} \simeq \mathbb{R}^2 \quad \text{and} \quad T_vS^2 \subset \mathbb{R}^3.$$

We denote by $\Phi : \mathbb{T} \to \mathbb{R}^3$ the linear embedding of the triangle:

$$\Phi(s,t) = (1-s-t)\cdot v_0 + s\cdot v_1 + t\cdot v_2$$

and we set $d\Phi|_p$ to be its differential at $p \in \mathbb{T}$:

$$d\Phi|_p : T_p\mathbb{T} \simeq \mathbb{R}^2 \to T_nS^2 \subset \mathbb{R}^3.$$

We set $g|_p$ to be the metric tensor on $T_p\mathbb{T}$ induced by the linear embedding of the unit right triangle:

$$g|_p = d\Phi|_p^\top \cdot d\Phi|_p.$$

We assume that we are given a Gauss Map:

$$N : \mathbb{T} \to S^2.$$

We denote by $\widetilde{d\Phi}|_p$ the map "realizing" the tangent space $T_p\mathbb{T}$ as a subspace of $\mathbb{R}^3$ perpendicular to $N(p)$ by first computing the differential of the embedding Φ and then applying the Rodrigues rotation taking the normal of the embedded triangle to the normal prescribed by the Gauss Map.

$$\widetilde{d\Phi}|_p \equiv R\big(n, N(p)\big) \circ d\Phi|_p : \mathbb{R}^2 \to \mathbb{R}^3.$$

We denote the "hat" functions as $\psi_i : \mathbb{T} \to \mathbb{R}$, with:

$$\psi_0(s,t) = 1-s-t, \quad \psi_1(s,t) = s, \quad \text{and} \quad \psi_2(s,t) = t,$$

We note that: (1) The differential $d\Phi|_p$, and hence the metric $g|_p$ is constant for all $p \in \mathbb{T}$; (2) The realization $\widetilde{d\Phi}|_p$ is an orthogonal transformation between $T_p\mathbb{T}$ with inner-product $g|_p$ and $T_{N(p)}S^2$ with inner-product obtained by restricting the Euclidean inner-product on R^3; (3) The inverse of $\widetilde{d\Phi}|_p$ is:

$$\widetilde{d\Phi}|_p^{-1} \equiv g^{-1} \cdot d\Phi|_p^\top \cdot R\big(N(p), n\big) \tag{1}$$

and is defined over all of $\mathbb{R}^3$ by projecting out the component parallel to $N(p)$.

4. Discretization

As our approach is based on defining a basis for function spaces on a mesh, composed of functions that are defined "piecewise" per triangle, it suffices to consider the case of a single triangle, as the global system can be constructed using finite-element assembly. To this end, we assume that we are given a triplet of oriented vertices $\{v_k, n_k\} \in \mathbb{R}^3 \times S^2$ (with $1 \le k \le 3$), and we extend the per-vertex normals to a Gauss map $N : \mathbb{T} \to S^2$ by Phong interpolation:

$$N(p) \equiv \frac{\sum_k \psi_k(p)\cdot n_k}{|\sum_k \psi_k(p)\cdot n_k|} \in S^2.$$

We assume that the Gauss map is nowhere equal to the negative of the triangle's normal, $N(p) \ne -n$, so that $\widetilde{d\Phi}|_p$ is well-defined.

4.1. Scalar Basis

Using the standard "hat" functions gives the mass, $\mathbf{m}^\psi \in \mathbb{R}^{3\times 3}$, and stiffness, $\mathbf{s}^\psi \in \mathbb{R}^{3\times 3}$, matrices:

$$\mathbf{m}^\psi_{ij} \equiv \int_{\mathbb{T}} \psi_i(p)\cdot\psi_j(p)\cdot\sqrt{\det(g|_p)}\,dp$$
$$\mathbf{s}^\psi_{ij} \equiv \int_{\mathbb{T}} \mathrm{tr}\left(g^{-1}\cdot d\psi_i|_p^\top\cdot d\psi_j|_p\right)\cdot\sqrt{\det(g|_p)}\,dp.$$

(The definition of the stiffness is consistent with Eq. 6, implicitly using the coordinate basis for $\mathbb{R}^3$ and leveraging the fact that, in this basis, the inner-product is represented by the identity matrix.)

4.2. Vector-Field Basis

To define a basis for vector-fields, we proceed as in the approach of Knöppel *et al.* [KCPS13], (1) defining an operator taking a tangent vector at a vertex and transporting it into the interior of incident triangles, (2) selecting a tangent frame at each vertex, and (3) blending the transported tangent vectors.

Fig. 2 visualizes our approach for a polygonal curve. Starting with per-vertex normals, Phong interpolation assigns normals to points in the interior of the edges. Then, given a tangent vector at the center vertex, the vector is transported using the rotation taking the center vertex's normal to the Phong normals – giving a vector-field. Finally, a continuous vector-field supported on the two edges is obtained by scaling with the "hat" function at the center vertex.

Transport Given a tangent vector $t_k \in \mathbb{R}^3$ at the k-th corner (perpendicular to n_k), we define its transport to $p \in \mathbb{T}$ by applying the Rodrigues rotation taking n_k to $N(p)$, to the tangent vector t_k. We

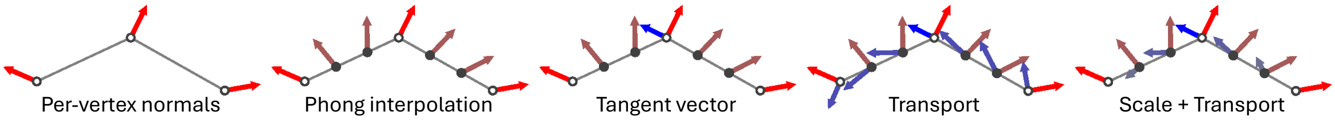

Figure 2: *Constructing a continuous tangent vector-field on a polyline by scaling and transporting a vector prescribed at the center vertex.*

denote by $R_k : \mathbb{T} \to \mathrm{SO}(3)$ the minimal-angle rotation taking the normal at the k-th corner to the interpolated normal at p:

$$R_k(p) \equiv R\big(n_k, N(p)\big).$$

Framing Vertices At the k-th corner we select framing vectors, $f_{2k}, f_{2k+1} \in \mathbb{R}^3$ perpendicular to the normal n_k. For simplicity, we choose the vectors to be orthonormal.

Blending Transported Tangent Vectors We extend the framing vectors, $f_{2k}, f_{2k+1} \in \mathbb{R}^3$, to extrinsic vector-fields defined in the interior of the triangle, $\overline{\omega}_i : \mathbb{T} \to \mathbb{R}^3$, by transporting and blending:

$$\overline{\omega}_{2k+l}(p) \equiv \psi_k(p) \cdot R_k(p) \cdot f_{2k+l} \qquad \text{with } l \in \{0,1\} \tag{2}$$

We note that, by construction, these are perpendicular to the interpolated normals.

To obtain intrinsic vector-fields on the tangent space of the unit right triangle, we pull back using $\widetilde{d\Phi}|_p$:

$$\omega_i|_p \equiv \widetilde{d\Phi}|_p^{-1} \cdot \overline{\omega}_i|_p \in T_p\mathbb{T}.$$

To obtain the vector-fields' covariant derivatives, we recall that that the covariant derivative of an extrinsic vector-field can be computed by differentiating the field and projecting out the normal component. In particular, since $\widetilde{d\Phi}|_p^{-1}$ projects out the normal component, the intrinsic representation of the covariant derivative is obtained by differentiating Eq. 2 and pulling back:

$$\nabla\omega_i|_p \equiv \widetilde{d\Phi}|_p^{-1} \circ d\overline{\omega}_i|_p \in \mathrm{End}(T_p\mathbb{T}).$$

Using the six functions $\omega_i : \mathbb{T} \to \mathbb{R}^3$ we compute the mass and stiffness matrices, $\mathbf{m}^\omega, \mathbf{s}^\omega \in \mathbb{R}^{6\times 6}$ with:

$$\begin{aligned} \mathbf{m}^\omega_{ij} &\equiv \int_{\mathbb{T}} g|_p\left(\omega_i|_p, \omega_j|_p\right) \cdot \sqrt{\det(g|_p)}\, dp \\ \mathbf{s}^\omega_{ij} &\equiv \int_{\mathbb{T}} \mathrm{tr}\left(g|_p^{-1} \cdot (\nabla\omega_i|_p)^* \cdot g|_p \cdot (\nabla\omega_j|_p)\right) \cdot \sqrt{\det(g|_p)}\, dp. \end{aligned} \tag{3}$$

For the mass, the inner-product is computed with respect to the inner-product on the tangent space, $g|_p : T_p\mathbb{T} \to \mathbb{T}^*_p\mathbb{T}$. For the stiffness, we use the inner-product on the space of endomorphisms $\mathrm{End}(T_p\mathbb{T})$ induced by $g|_p$, reviewed in Eq. 6 of App. A.

4.3. Integration

For both the scalar and vector-field basis we compute the coefficients of the mass and stiffness matrices using numeric quadrature [Tay08]. For the mass matrices, this only requires that we can evaluate the basis functions at arbitrary locations in the triangle. For the stiffness matrix, this also requires that we are able to evaluate the differential of the basis functions. This is trivial for the scalar basis (since the differential is constant) and, given the simple expression of the Rodrigues rotation formula, is also straight-forward for the vector-field basis.

4.4. Finite Element Assembly

Given a triangle mesh, $\{\mathcal{V}, \mathcal{T}\}$, and given functionality for computing the system matrix associated to a triangle, $L : \mathcal{T} \to \mathbb{R}^{3K\times 3K}$, pseudocode for finite element assembly is summarized in Alg. 1: The matrix is initialized (step 1); Then, iterating over all triangles (step 2), the system matrix for the triangle is computed (step 3), and its coefficients are added to the corresponding entries of the system matrix defined over the mesh (steps 4-7).

Algorithm 1: FiniteElementAssembly

```
Require: $L : \mathcal{T} \to \mathbb{R}^{3K\times 3K}$
1: $\mathbf{L} \leftarrow \mathbf{0} \in \mathbb{R}^{|\mathcal{V}|K\times|\mathcal{V}|K}$
2: for $\tau = \{v_0, v_1, v_2\} \in \mathcal{T}$:
3:     $\mathbf{l} \leftarrow L(\tau)$
4:     for $m, n \in [0,3)$ and $k, l \in [0,K)$:
5:         $I \leftarrow v_m \cdot K + k, \quad J \leftarrow v_n \cdot K + l$
6:         $i \leftarrow m \cdot K + k, \quad j \leftarrow n \cdot K + l$
7:         $\mathbf{L}_{IJ} \leftarrow \mathbf{L}_{IJ} + \mathbf{l}_{ij}$
8: return L
```

For the scalar mass and stiffness matrices, which have one degree of freedom per vertex, we have $K = 1$. For the vector-field mass and stiffness matrices we have $K = 2$. In what follows, we abuse notation, using ψ_i (resp. ω_{2i} and ω_{2i+1}), with $i \in [0, |\mathcal{V}|)$, to denote the scalar (resp. vector-field) basis functions on the mesh, rather than at the corners of a single triangle.

5. Energies and Operators

While the previous section described the construction of the stiffness matrix for vector-fields, corresponding to the Dirichlet energy defined by the covariant derivative, a similar approach can be used to define other energies. We also show that our discretization can be used to define the Lie bracket operator.

The following discussion holds for all finite-elements discretization of vector-fields that support evaluation of their derivative. In particular, we discuss constructions of the Killing energy and definitions of the Lie bracket derived from the finite-elements proposed by Stein *et al.* [SWJG20] (ignoring discontinuities along edges). While in principle a similar approach can be used with the finite-elements of Knöppel *et al.* [KCPS13], the lack of a closed-form expression for the basis functions makes this challenging in practice. This issue is addressed in the later work of Liu *et al.* [LTGD16].

5.1. Energies

Using the decomposition from Sec. 3.1, we can factor the covariant derivative into orthogonal components consisting of scalar multiples of the identity, trace-free (self-adjoint) endomorphisms, and anti-self-adjoint endomorphisms. This, in turn, allows us to replace the covariant derivative in Eq. 3 with an individual component – allowing us to define classical energies used in vector-field processing such as:

$$\begin{aligned}
\text{Connection} &\longleftrightarrow \mathcal{D}\mathit{iv}(T_p\mathbb{T})\oplus\overline{\mathcal{H}\mathit{olom}}(T_p\mathbb{T})\oplus\mathcal{C}\mathit{url}(T_p\mathbb{T})\\
\text{Holom. / Hodge} &\longleftrightarrow \mathcal{D}\mathit{iv}(T_p\mathbb{T})\oplus\mathcal{C}\mathit{url}(T_p\mathbb{T})\\
\text{Anti-holom.} &\longleftrightarrow \overline{\mathcal{H}\mathit{olom}}(T_p\mathbb{T})\\
\text{Killing} &\longleftrightarrow \mathcal{D}\mathit{iv}(T_p\mathbb{T})\oplus\overline{\mathcal{H}\mathit{olom}}(T_p\mathbb{T})\\
\text{Divergence} &\longleftrightarrow \mathcal{D}\mathit{iv}(T_p\mathbb{T})\\
\text{Curl} &\longleftrightarrow \mathcal{C}\mathit{url}(T_p\mathbb{T})
\end{aligned}$$

We note that the spaces $\mathcal{D}\mathit{iv}(T_p\mathbb{T})$ and $\mathcal{C}\mathit{url}(T_p\mathbb{T})$ are one-dimensional, while the space $\overline{\mathcal{H}\mathit{olom}}(T_p\mathbb{T})$ is two-dimensional.

5.2. Lie Bracket

Our approach provides an expression for the covariant derivative that can be evaluated pointwise. In addition to being amenable to quadrature-based integration, it also enables the pointwise evaluation of the Lie bracket of two vector-fields. Concretely, given tangent vector-fields X and Y, we evaluate the Lie bracket at a point $p\in\mathbb{T}$ by taking the difference of the derivative of X along Y and the derivative of Y along X:

$$[X,Y](p)=\nabla X|_p\cdot Y(p)-\nabla Y|_p\cdot X(p).$$

We note that this is only an approximation of the Lie Bracket, since this relation is exact only when ∇ is the Levi-Civita Connection.

Global fitting

As our space of tangent vector-fields is not closed under the Lie bracket, the bracket of vector-fields X and Y cannot (in general) be expressed as a linear combination of our tangent vector-field basis. However, it is straightforward to compute the projection of the bracket $[X,Y]$ onto our space of vector-fields.

Given tangent vector-fields X and Y, we start by computing the "weak representation" of their Lie bracket, $\mathbf{b}\in\mathbb{R}^{2|\mathcal{V}|}$, by computing the integral of the dot-product of the bracket with the tangent vector-field basis functions (which only requires that the bracket be pointwise evaluatable):

$$\mathbf{b}_i\equiv\int g|_p\big(\omega_i(p),[X,Y](p)\big)\cdot\sqrt{\det(g|_p)}\,dp.$$

Then, we obtain the coefficients of the projection of the bracket onto the space of tangent vector-fields, $\mathbf{z}\in\mathbb{R}^{2|\mathcal{V}|}$, by solving:

$$\mathbf{M}^{\omega}\cdot\mathbf{z}=\mathbf{b}. \tag{4}$$

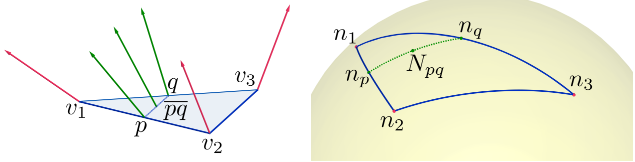

Figure 3: *Interpolated normals (green) along a segment on the triangle map onto a great circle on the Gauss map.*

6. Topology and N-Fields

6.1. Parallel transport and Gaussian curvature

As our vector-fields are continuous and we interpolate normals, their derivatives induce notions of per-triangle piecewise continuous parallel transport and Gaussian curvature. Consider a single triangle $\tau\in\mathcal{T}$ with normals $\{n_1,n_2,n_3\}$. Without loss of generality, consider the total rotation from and back to the first vertex:

$$\exp(\kappa_\tau\cdot n_1)=R(n_3,n_1)\cdot R(n_2,n_3)\cdot R(n_1,n_2).$$

Then, κ_τ is the *holonomy* of the triangle with the prescribed normals. Note that since our tangent bundle is decoupled from the geometry of the mesh, the total sum of curvature might have a different Euler characteristic than the mesh. For instance, it is 0 when the vertex normals are constant. By definition, it is however always a multiple of 2π. For our "natural" vertex normals, we expect it to match the Euler characteristic of the geometry when the mesh is decently-sampled, but we do not make this assumption.

The holonomy κ_τ can be computed directly as the angle defect of the geodesic triangle in the Gauss map [dC16]. We use the equivalent solid angle formula for the induced area:

$$\kappa_\tau=2\,\mathrm{atan2}\big(n_1\cdot(n_2\times n_3),\ 1+(n_1\cdot n_2)+(n_2\cdot n_3)+(n_3\cdot n_1)\big) \tag{5}$$

Remark 1 Given a triangle $\tau\in\mathbb{T}$, the normals along any straight segment within the triangle, $\overline{pq}\subset\tau$, all lie on a great circle (Fig. 3).

Proof Let $\alpha_k:\mathbb{T}\to\mathbb{R}$ be the barycentric coordinate functions. For any point $\overline{pq}(s)=(1-s)\cdot p+s\cdot q$ on the segment, the interpolated (un-normalized) normal is:

$$\tilde{N}\big(\overline{pq}(s)\big)=\sum_{k=1}^{3}\alpha_k\big(\overline{pq}(s)\big)\cdot n_k=(1-s)\cdot\tilde{n}_p+s\cdot\tilde{n}_q$$

where $\tilde{n}_p$ and $\tilde{n}_q$ are the (un-normalized) vectors obtained by linearly interpolating the normals from the corner of the triangle to p and q respectively. Thus, the normals along the segment $\overline{pq}$ all lie in the plane spanned by $\tilde{n}_p$ and $\tilde{n}_q$, and hence (since the span is independent of normalization) on a great circle. □

A direct implication is that the parallel transport between p and q is simply $R(n_p,n_q)$, where n_p and n_q are the unit normals at p and q obtained by interpolating from the corners and normalizing. More formally, the Gauss map of the interpolated normals is a part of a sphere, where segments map to geodesics (parts of great circles). As a consequence, the parallel transport along pq is the pullback of parallel transport along the sphere, as defined in Sec. 4.2.

This allows us to measure holonomy on any sub-triangle simply by using Eq. 5 on the interpolated normals at the sub-triangle's vertices, making the algorithm of the next section simple to implement.

6.2. Singularities

Given a vector field ω, for every triangle $\tau \in \mathcal{T}$ [CDS10], the index of the field over the triangle is:

$$I_\tau^\omega = \frac{1}{2\pi}\left(\theta_\tau^\omega - \kappa_\tau\right),$$

where θ_τ^ω is the sum of the rotation angles of the field deviating from parallel transport, taken over the edges of τ. Since the field is continuous, singularities can appear anywhere (as in [BV22]). To find them, we rely on the fact that every segment within the triangle is a great circle on the Gauss map, so that the angle sum θ^ω and holonomy κ can be computed in closed form for any subtriangle. Thus, we can find the barycentric coordinates $\{\alpha_1, \alpha_2, \alpha_3\}$ of a singularity using bisection. We show examples in Fig. 4. The definition of the total index of a triangle is exact; we further assume that the pointwise singularities are as those of a standard PL field: either a single point per triangle, a full line, or the entire triangle (we omit the last two cases). We leave a concrete (dis)proof for future work.

To find where coordinate α_k crosses the singularity, Alg. 2 proceeds as follows: Assuming there is a singularity inside the triangle (step 1), the search interval is initialized (steps 2 and 3), and bisection is iterated until the interval length is within a prescribed tolerance (steps 4-9). Within each iteration, the mid-point is identified (step 5), the index of the triangle having k as one of its vertices and the positions α_k along the edges incident to vertex k as the other two vertices is computed (step 6), and the interval is refined to be the sub-interval with the singularity (steps 7 and 8). We use a tolerance of $\varepsilon = 10^{-7}$. We search using $k = 1$ and with $k = 2$, thereby locating the singularity within the triangle.

Algorithm 2: SingularityDetection

```
Require: $k \in \{1,2,3\}$,  $\varepsilon > 0$
1: if $I_\tau^\omega = 0$: return false
2: $\alpha_{\text{start}} \leftarrow 0$
3: $\alpha_{\text{end}} \leftarrow 1$
4: while $\alpha_{\text{end}} - \alpha_{\text{start}} \geq \varepsilon$:
5:     $\alpha_{\text{mid}} \leftarrow \frac{1}{2}(\alpha_{\text{start}} + \alpha_{\text{end}})$
6:     $I \leftarrow \text{SubTriangleIndex}(k, \alpha_{\text{mid}})$  (Eq.5)
7:     if $I \neq 0$:  $\alpha_{\text{end}} \leftarrow \alpha_{\text{mid}}$
8:     else:           $\alpha_{\text{start}} \leftarrow \alpha_{\text{mid}}$
9: return {true, $\alpha_{\text{mid}}$}
```

6.3. N-fields

Given a triangle $\tau \in \mathcal{T}$, we represent N-RoSy fields [VCD*16] using the orthonormal framing vectors f_{2k} and f_{2k+1} at the k-th vertex (Sec. 4.2), allowing us to define the power representation $X = \omega^N$ of an N-RoSy vector in each vertex's tangent plane. At each internal point p, we choose an arbitrary orthonormal frame, defining the connection $r_{k,p} \in \mathbb{C}$ as the representation of the frame $\{R(n_k, N(p)) \cdot f_{2k+l}\}_l$ with respect to the frame of p. Then, we get:

$$X_p = \sum_{k\in\tau} X_k \cdot \left(r_{k,p}\right)^N,$$

and the N-RoSy field at p as the roots $\sqrt[N]{X_p}$. Singularities are discovered as in Sec. 6.2, with two modifications: (1) we compute θ^X as the angle sum defined by X, rather than θ^ω, and (2) we use $N \cdot \kappa_\tau$ as the holonomy. The obtained integer indices are interpreted as integer multiples of $1/N$. This is standard practice in defining N-RoSy singularities [DVPSH14]. We show examples in Fig. 5. We note that PolyVectors can be defined similarly.

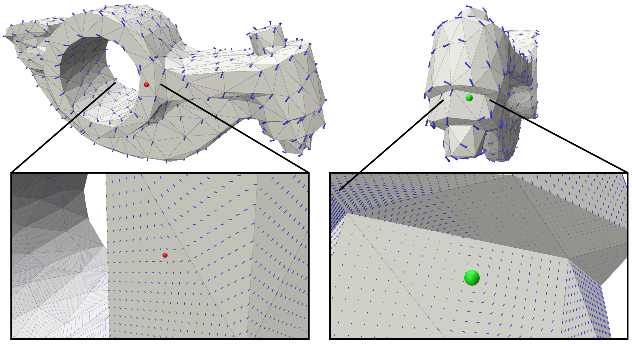

Figure 4: *We detect positive (green) and negative (red) singularities in the interior of faces.*

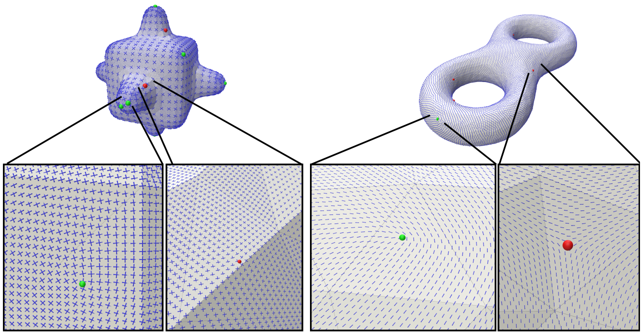

Figure 5: *Our framework allows the interpolation of N-fields, and locating their continuous singularities. Left: $N = 4$, right: $N = 2$.*

7. Evaluation

We evaluate our discretization of the space of tangent vector-fields in four ways: (1) We compare to existing approaches in standard interpolation applications; (2) We validate that our discretization preserves some basic invariants from the continuous theory; (3) We consider the spectral decomposition of the the connection Laplacian and associated operators; and (4) We demonstrate applications in computing the Lie bracket of vector-fields.

In our implementation, we use 3-point quadrature to integrate functions over a triangle. With the exception of the unit sphere we rescale all models to have unit area. We assign normals to vertices by computing the approximate limit surface under Loop subdivision [Loo87]. Concretely, for each vertex we compute the subdivision stencil, raised to the 10-th power, and apply it to the vertex and its one-ring neighbors. We then set the vertex's normal to the area-weighted average of the normals of the subdivided triangles.

With the exception of the spectral analysis on the sphere, we use the same geometry for all methods. In practice, this means that the method of Stein *et al.*, which assigns two degrees of freedom to each edge, has roughly three times as many degrees of freedom as the methods that assign two degrees of freedom to a vertex.

When ground-truth is available, we measure the error between the estimated solution, x and the ground-truth, x^* as the ratio:

$$E(x,x^*) \equiv \sqrt{\frac{\|x-x^*\|^2}{\|x\|^2+\|x^*\|^2}}$$

where, for a function/field f defined on the surface, "$\|f\|^2$" denotes the integrated square-norm over the mesh.

We visualize vector-fields using anisotropic diffusion [DPR00] in which initially randomly distributed noise at the vertices is smoothed in the direction of the flow-field. Similar to line integral convolution [CL93], the visualization captures the direction of the vector-field but not the sign. To address this, we overlay the visualization with stream-lines, shaded from black (start) to white (end).

To support follow-on research we provide an implementation in C++, using a functional-programming paradigm, that makes it straight-forward to represent fields over the mesh (https://github.com/mkazhdan/PhongRodriguesVF/). The implementation only requires that a user provide a functor whose input is a position on the mesh (defined by the triangle index and barycentric coordinates) and whose output is the field's value at that point. (These fields can themselves be constructed by evaluating vector-fields and their covariant derivatives to construct the field's values, as in the case of computing the Lie bracket.) Using quadrature, the implementation integrates the fields, performing the finite-element assembly required to construct system matrices and vectors for targeted vector-field-processing applications.

We compare our approach to that of Knöppel *et al.* [KCPS13], Stein *et al.* [SWJG20], and Sharp *et al.* [SSC19], with implementations provided by the authors, [SC*19, CS13] We also compare to the discretization obtained using the Whitney basis [Whi57]. We could not compare to Liu *et al.* [LTGD16] as a public implementation is not available.

7.1. Example Applications

Sparse Vector-Field Interpolation

Given a sparse subset of vertices $\mathcal{W} \subset \mathcal{V}$ and target tangent vectors at those vertices $\{t_w\}_{w\in\mathcal{W}}$, the stiffness matrix can be used to solve for the as-smooth-as-possible interpolating vector-field:

$$\mathbf{x} = \underset{\mathbf{y}\in\mathbb{R}^{3|\mathcal{V}|}}{\operatorname{argmin}} \left(\mathbf{y}^\top \cdot \mathbf{S} \cdot \mathbf{y}\right), \quad \text{s.t.} \quad \mathbf{y}_w = t_w, \forall w \in \mathcal{W}$$

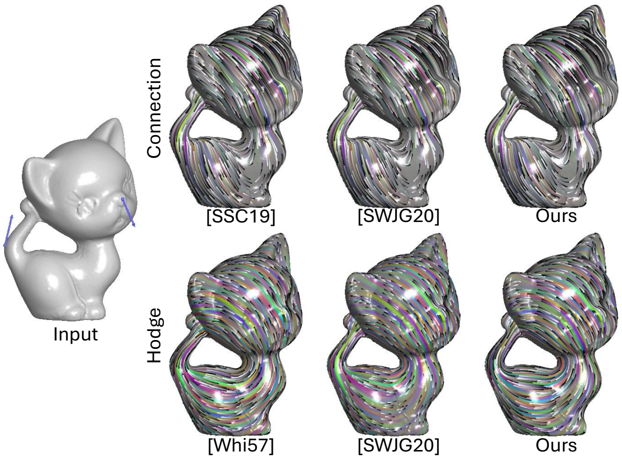

Figure 6: *Visualization of smooth interpolation of sparse vectors (draw in blue) using smoothness energies defined by the Connection (top) and Hodge (bottom) Laplacians.*

with $\mathbf{S}$ a stiffness matrix.

As an example, Fig. 6 shows a model with two tangent vector constraints (blue) and the computed vector-fields obtained with $\mathbf{S}$ defined by the connection (top) and Hodge (bottom) energies. As the model has genus one, there are two harmonic vector-fields and Hodge interpolation gives the vector-field circulating around the hole. In contrast, the connection energy, which is non-singular, penalizes the vector-field wherever it is non-zero, resulting in an interpolant that gets smaller away from the constraints.

The top row compares our approach to that of Sharp *et al.* (second column) and Stein *et al.* (third column). As the figure shows, all three approaches produce a similar sparse interpolant.

Since the approach of Stein *et al.* provides a pointwise evaluation of the covariant derivative, we can similarly define a Hodge energy by considering the holomorphic component of the covariant derivative (ignoring the unbounded discontinuity at edges). Interpolation results obtained using that approach are shown in the bottom row. We also show results obtained using the DEC formulation of the Hodge Laplacian associated with the Whitney 1-form basis [Whi57, Bos98, DKT06]. We find that our results are more consistent with those obtained using the Whitney basis. For the approach of Stein *et al.* it was necessary to add a connection energy regularizer, as otherwise the system was not well-conditioned. This is consistent with the authors' incorporation of a connection regularizer to remove spurious minima in the Killing energy.

For a more detailed empirical discussion, please see App. B.

Vector Heat

We also consider the Vector Heat method for parallel transport of vectors [SSC19]. The method is implemented using both vector heat diffusion (using the connection Laplacian) and scalar heat diffusion (using the co-tangent Laplacian). The result is a vector-field on the surface whose value within a Voronoi region of a con-

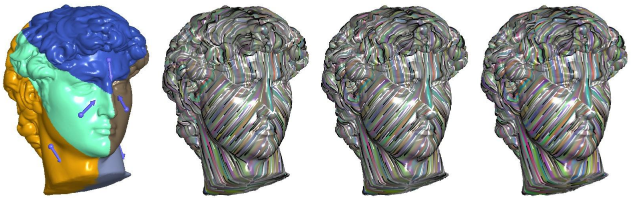

Input [SSC19] [SWJG20] Ours

Figure 7: *Input vector-field constraints, with the associated Voronoi partition of the geometry (left), and the results of applying the Vector Heat method using different discretizations of the connection Laplacian.*

straint point is the parallel transport of the constraint vector along the shortest path geodesic.

Fig. 7 shows the results for five constraint vectors placed on the surface of the David model. The visualization shows the initial constraint vectors and associated Voronoi regions (left) as well as the output vector-field. As with sparse vector-field interpolation, we compare our results (right) to those of Sharp *et al.* (center left) and Stein *et al.* (center right). Empirically, we find almost no difference between the approaches. This is consistent with the earlier observation that, when using the connection stiffness for sparse vector-field interpolation, the approaches give similar results.

7.2. Rotation Invariance

As with the discretizations of Knöppel *et al.* and Stein *et al.*, our function space is closed under pointwise rotation by a fixed angle in the tangent plane. In particular, letting $\mathbf{M},\mathbf{T},\overline{\mathbf{T}},\overline{\mathbf{A}},\mathbf{J}\in\mathbb{R}^{2|\mathcal{V}|\times 2|\mathcal{V}|}$ be the matrices corresponding to the:

$\mathbf{M}$ mass,
$\mathbf{T}$ stiffness of component $\mathscr{D}iv(T_p\mathbb{T})$,
$\overline{\mathbf{T}}$ stiffness of component $\overline{\mathscr{H}olom}(T_p\mathbb{T})$,
$\overline{\mathbf{A}}$ stiffness of component $\mathscr{C}url(T_p\mathbb{T})$, and
$\mathbf{J}$ 90° counter-clockwise rotation around each vertex's normal,

we validate that the mass matrix and trace-less component are both invariant under counter-clockwise rotation by 90° and that the scalar-multiple of the identity and anti-self-adjoint components are related to each other by 90° counter-clockwise rotation.

For the models in Fig. 10, we show the relative errors in Tab. 3. The table validates the expected properties of the system matrices under counter-clockwise rotation by 90°. The tiny relative errors for the mass matrix are expected, as the vector-field representation is extrinsic and the Euclidean inner-product is invariant under rotation. For the other energies, the relative error is also small, with the largest errors for the "bunny" and "icosa" models. We believe this is due to the quality of the triangulations (both models were obtained by applying Marching Cubes [LC87] to an implicit representation). In particular, considering the distribution of triangle aspect ratios in Fig. 11, we find that the error is strongly correlated with triangulation quality. (As $\mathbf{J}\cdot\mathbf{J}=\mathbf{J}^\top\cdot\mathbf{J}^\top=-\mathbf{Id}$, it follows that $\mathbf{T}=\mathbf{J}^\top\cdot\overline{\mathbf{A}}\cdot\mathbf{J}$ if and only if $\overline{\mathbf{A}}=\mathbf{J}^\top\cdot\mathbf{T}\cdot\mathbf{J}$. This was empirically confirmed by noting that the relative errors are nearly identical.)

Model	$\frac{\|\mathbf{M}-\mathbf{J}^\top\cdot\mathbf{M}\cdot\mathbf{J}\|_F}{\|\mathbf{M}+\mathbf{J}^\top\cdot\mathbf{M}\cdot\mathbf{J}\|_F}$	$\frac{\|\overline{\mathbf{T}}-\mathbf{J}^\top\cdot\overline{\mathbf{T}}\cdot\mathbf{J}\|_F}{\|\overline{\mathbf{T}}+\mathbf{J}^\top\cdot\overline{\mathbf{T}}\cdot\mathbf{J}\|_F}$	$\frac{\|\mathbf{T}-\mathbf{J}^\top\cdot\overline{\mathbf{A}}\cdot\mathbf{J}\|_F}{\|\mathbf{T}+\mathbf{J}^\top\cdot\overline{\mathbf{A}}\cdot\mathbf{J}\|_F}$
Bunny	1.5×10^{-16}	1.3×10^{-10}	1.0×10^{-10}
Eight	1.5×10^{-16}	2.2×10^{-16}	2.3×10^{-16}
Elephant	1.4×10^{-16}	1.6×10^{-15}	1.3×10^{-15}
Genus-6	1.4×10^{-16}	4.6×10^{-16}	3.7×10^{-16}
Icosa	1.4×10^{-16}	4.2×10^{-10}	3.1×10^{-10}

Table 3: *Invariance under 90°-rotation, computed using the ratio of the Frobenius norm of the difference to the Frobenius norm of the sum.*

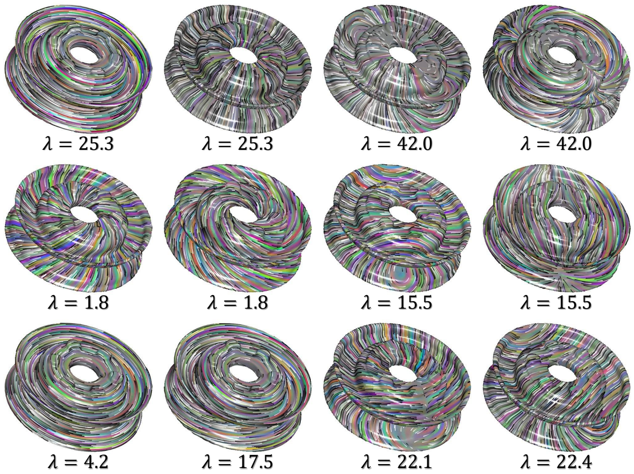

Figure 8: *Visualization of The four smallest eigenvectors of the Connection (top), Hodge (middle), and Killing (bottom) energies.*

7.3. Spectral Analysis

Spectral Vector-Fields

Using the energies from Sec. 5.1, we solve the generalized eigenproblem $\mathbf{S}\cdot\mathbf{x}=\lambda\cdot\mathbf{M}\cdot\mathbf{x}$ where $\mathbf{M}$ is the mass matrix and $\mathbf{S}$ is one of the connection, Hodge, and Killing stiffness matrices.

Fig. 8 shows the Pulley model and the smallest four eigenvectors of each energy, along with the associated eigenvalues. We note that:

- Since the connection and Hodge energies use equal parts of $\mathscr{D}iv(T_p\mathbb{T})$ and $\mathscr{C}url(T_p\mathbb{T})$, their eigenvectors come in pairs, related by 90° counter-clockwise rotation (see above).
- Since the model has genus one, the two smallest eigenvectors of the Hodge Laplacian closely match the harmonic vector-fields and their associated eigenvalues are close to zero.
- Since the model is nearly a surface of revolution, the smallest eigenvector of the Killing energy corresponds to the flow around the axis of revolution, and its associated eigenvalue is noticeably smaller than that of the next eigenvector. (Though the second eigenvector also corresponds to a vector-field circulating about the axis of rotational symmetry, the directions of flow are reversed on the top and bottom halves of the pulley.)

For visualization, we further decompose the eigenspaces of the connection and Hodge energies. To this end, we set $\mathbf{M}\in\mathbb{R}^{2|\mathcal{V}|\times 2|\mathcal{V}|}$

to be the mass matrix and $\mathbf{T} \in \mathbb{R}^{2|\mathcal{V}|\times 2|\mathcal{V}|}$ to be the stiffness matrix associated with component $\mathcal{D}iv(T_p\mathbb{T})$ – defining the quadratic energy giving the integrated squared divergence of a vector-field. Then, for a particular eigenspace, X, of dimension $2k$, letting $\mathbf{X}^{2|\mathcal{V}|\times 2k}$ be the matrix whose columns are the eigenvectors, we define matrices $\mathbf{m},\mathbf{t} \in \mathbb{R}^{2k\times 2k}$ with:

$$\mathbf{m} = \mathbf{X}^\top \cdot \mathbf{M} \cdot \mathbf{X} \qquad \text{and} \qquad \mathbf{t} = \mathbf{X}^\top \cdot \mathbf{T} \cdot \mathbf{X}.$$

Solving the generalized eigenvalue problem:

$$\mathbf{t} \cdot \mathbf{x} = \lambda \cdot \mathbf{m} \cdot \mathbf{x}$$

we obtain an orthonormal basis for the eigenspace X graded by divergence energy. In particular, letting $\mathbf{x}_1, \dots \mathbf{x}_{2k}$ be the computed eigenvectors, we order the eigenvectors as k pairs, with the first vector minimizing squared divergence and the second its 90° counter-clockwise rotation:

$$\{\mathbf{X}\cdot\mathbf{x}_1, \mathbf{J}\cdot\mathbf{X}\cdot\mathbf{x}_1, \dots, \mathbf{X}\cdot\mathbf{x}_k, \mathbf{J}\cdot\mathbf{X}\cdot\mathbf{x}_k\}.$$

Comparison to Stein et al.

In principle, the formulation of Stein *et al.* [SWJG20] can also be used to define a Killing energy. However, as discussed in that work, the direct definition of such an energy exhibits spurious high-frequency minima (likely due to aliasing arising from the discontinuous nature of the basis). This limitation can be mitigated by adding a connection energy regularizer to the Killing energy. However, that introduces the additional challenge of tuning the regularization weight. Please see App. C for additional discussion.

Spectral Analysis on the Sphere

To study the spectral decomposition of the connection Laplacian defined over the sphere, we generate ten random triangulations of the unit sphere in two ways: For the first, we randomly sample points on the sphere and triangulate the points by computing the convex hull. For the second, we randomly sample points on the ellipsoid with semi-axis lengths $(1,4,1)$, triangulate the points by computing the convex hull, and rescale the points to have unit-norm. We then compute the mass and connection Laplacian matrices, $\mathbf{M}$ and $\mathbf{S}$ and solve the generalized eigenvalue problem:

$$\mathbf{S} \cdot \mathbf{x} = \lambda \cdot \mathbf{M} \cdot \mathbf{x}$$

to obtain the first 240 eigenvalues.

As the n-th eigen-space of the connection Laplacian is $(4n+2)$-dimensional, with associated eigenvalue $n(n+1)-1$, we evaluate the connection Laplacian by measuring the difference of the estimated eigenvalues from the ground-truth.

Fig. 9 compares our results to those of Knöppel *et al.* [KCPS13], Stein *et al.* [SWJG20], and Sharp *et al.* [SSC19]. For the visualization, we plot the ratio $|ev_c - ev_{gt}|/|ev_c + ev_{gt}|$, with ev_c the computed eigenvalue and ev_{gt} the ground-truth. Since the mass matrix of Sharp *et al.* is not obtained using a finite-elements approach, it may not be positive definite when the mesh is not Delaunay, and we apply intrinsic Delaunay triangulation [FSBS06] before computing their system matrices. For all methods, we discretize the connection Laplacian over a sphere sampled with $60K$ points. For the method of Stein *et al.*, we also show results for a sphere sampled with $20K$ points, so that the number of degrees of freedom match (two degrees of freedom per edge for Stein *et al.*, versus two degrees of freedom per vertex for the other methods).

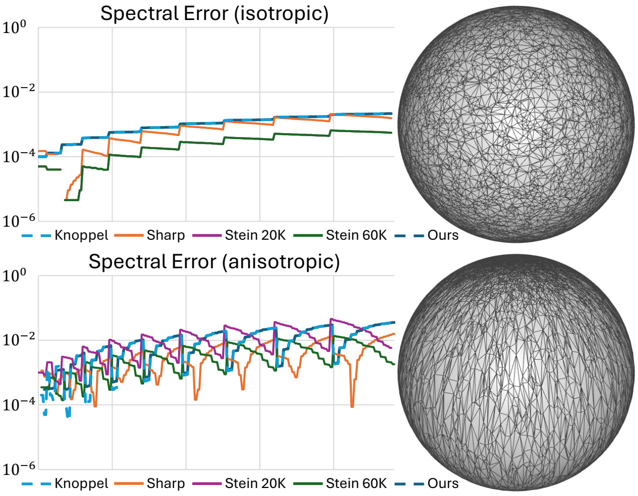

Figure 9: *The differences between the analytic eigenvalues and the estimated ones, averaged over ten random tessellations of the unit sphere, as a function of eigenvalue index. The top chart plots results when the sphere is sampled isotropically. The bottom plots results for the more challenging case of anisotropic sampling. (The associated eigenvalues are $\{1,5,11,19,29,41,55,71,89,109\}$.) Representative tessellations of the sphere are shown on the right.*

For both the isotropically (top) and anisotropically (bottom) sampled sphere, our results are indistinguishable from those of Knöppel *et al.* and are comparable to those of the method of Stein *et al.* when using $20K$ samples. (Using a higher resolution of $60K$ samples, the discretization of Stein *et al.* gives better results, as expected.) While also comparable to the method of Sharp *et al.* for the anisotropically sampled sphere, we find that the spectrum of Sharp *et al.*'s Laplacian is closer to that of the ground-truth in the case of anisotropic sampling. We believe that this is due to the use of the intrinsic Delaunay triangulation, which has the effect of undoing the anisotropic triangulation.

Hodge Laplacian

Considering just the holomorphic component of the covariant derivative (i.e. the sum of the projections onto $\mathcal{D}iv(T_\mathbf{p}\mathbb{T})$ and $\mathcal{C}url(T_\mathbf{p}\mathbb{T})$ we obtain a quadratic Hodge energy.

We measure the similarity between the spectra of our holomorphic energy and the spectra of the co-tangent Laplacian. We compare with the spectra obtained using the Hodge Laplacian computed using the Whitney basis.

Fig. 10 shows results for a number of shapes, with different genuses. We plot the ratio $|ev_h - ev_c|/|ev_h + ev_c|$ where ev_h is the eigenvalue of the Hodge Laplacian and ev_c is the associated eigenvalue of the co-tangent Laplacian, for the first 200 eigenvalues. Concretely, for a genus g surface, we compare eigenvalues $2(g+i)$ and $2(g+i)+1$ of the Hodge Laplacian with eigenvalue $i+1$ of

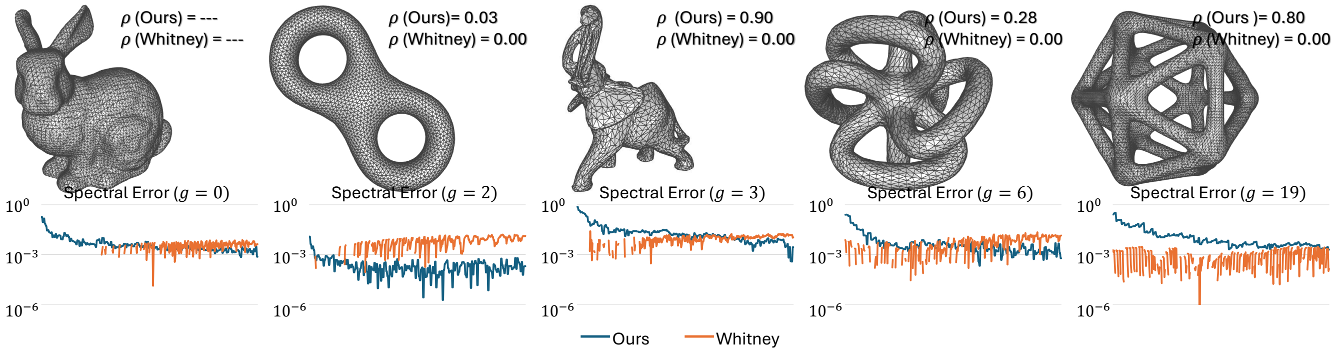

Figure 10: *The relative difference between the eigenvalue of the Hodge Laplacian and the co-tangent Laplacian, computed using the holomorphic part of the connection Laplacian and the discretization using the Whitney basis. The plot also gives the genus, g, of the model and the ratio, ρ, of eigenvalues* $2g$ *and* $2g+1$.

the co-tangent Laplacian, accounting for the $2 \cdot g$ harmonic vector-fields and the constant scalar field in the kernels of the Laplacians.

In addition to comparing the spectra, the plots also give the ratio, ρ, of the $2g$-th and $(2g+1)$-st eigenvalues of the Hodge Laplacian. Since we expect to have $2g$ harmonic vector-fields, this ratio measures the extent to which the Hodge Laplacian fails to identify the harmonic vector-fields.

As the Whitney 1-forms are consistent with the structure-preserving DEC discretization, they exactly identify the harmonics ($\rho = 0$). Furthermore, since the differentials of the eigenvectors of the co-tangent Laplacian are themselves eigenvectors of the Hodge Laplacian, every eigenvalue of the co-tangent Laplacian will necessarily be an eigenvalue of the Hodge Laplacian obtained using the DEC formulation. However, since the Whitney 1-forms are not closed under rotation by 90° degrees, the eigenvalues obtained using the DEC discretization do not come in pairs and it is *not* the case that the $(i+1)$-st eigenvalue of the co-tangent Laplacian will appear as either the $(2g+2i)$-th or $(2g+2i+1)$-st eigenvalue of the Hodge Laplacian discretized using DEC. Thus, while the DEC discretizations tend to give errors where every other one is initially zero, this ceases to hold at higher frequencies.

More generally, we found that at lower frequencies, the spectrum of the Whitney basis discretizations does a better job of matching the spectrum of the co-tangent Laplacian, while at higher frequencies our discretization does better.

To understand the behavior of our discretizations under refinement, we repeated the experiment comparing the spectrum of the Hodge Laplacian to the spectrum of the co-tangent Laplacian, performing multiple passes of Loop subdivision [Loo87]. We found that both our and the Whitney energies improved with refinement, and the extent of improvement appeared to be comparable. We also found that the trend of the Whitney basis outperforming our discretization at lower frequencies persists. Unfortunately, the ability of our discretization to distinguish the harmonic vector-fields did not improve consistently under refinement.

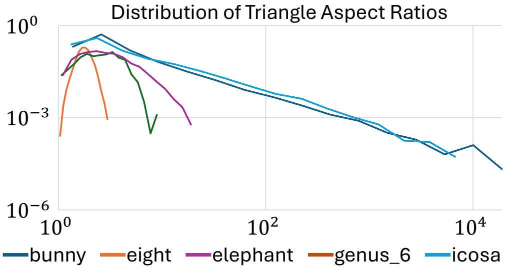

Figure 11: *Distribution of triangle aspect ratios for the models visualized in Fig. 10.*

7.4. Lie Bracket

To evaluate our discretization of the Lie bracket, we compare to three approaches. The first two define the bracket indirectly, representing derivations in terms of functional maps [ABCCO13, AOCBC15]. The third is obtained analogously to ours, differentiating the finite-elements representation of Stein *et al.* [SWJG20] to obtain the covariant derivative and evaluating the expression $[X,Y](p) = \nabla X|_p \cdot Y(p) - \nabla Y|_p \cdot X(p)$. For both finite-elements discretizations the obtained bracket $[X,Y]$ is not in the span of the basis and we obtain the least-squares fit by solving the linear system in Eq. 4 using the respective mass matrices. For our approach this requires approximating integrals using quadrature. For Stein *et al.*'s discretization the bracket is piecewise linear, and integration over the interior of triangles is computed in closed-form.

We evaluate the discretizations of the Lie bracket on two tessellations of the sphere. For the first we recursively subdivide and refine the vertices of an icosahedron. For the second we randomly sample points on the sphere and compute the convex hull. Both tessellations contain roughly 10K vertices.

As input we use vector-fields defined in the ambient 3D space. We start by generating random vector-fields $\tilde{X}, \tilde{Y} : \mathbb{R}^3 \to \mathbb{R}^3$, whose Fourier coefficients are band-limited complex exponentials with uniformly randomly assigned Fourier coefficients. (Antipodal frequencies are assigned conjugate values so that the coordinate func-

	b	[ABCCO13]	[AOCBC15]	[SWJG20]	Ours
icosa	2	3.5×10^{-3}	9.5×10^{-4}	$\mathbf{1.0\times10^{-4}}$	7.7×10^{-4}
	5	8.0×10^{-3}	1.5×10^{-3}	$\mathbf{3.7\times10^{-4}}$	2.4×10^{-3}
	10	2.6×10^{-2}	3.7×10^{-3}	$\mathbf{7.3\times10^{-4}}$	8.9×10^{-3}
	20	9.6×10^{-2}	1.1×10^{-1}	$\mathbf{1.4\times10^{-3}}$	3.5×10^{-2}
CH	2	1.7×10^{-1}	3.0×10^{-1}	$\mathbf{9.6\times10^{-3}}$	1.3×10^{-2}
	5	1.2×10^{-1}	3.1×10^{-1}	$\mathbf{1.5\times10^{-2}}$	1.6×10^{-2}
	10	1.3×10^{-1}	4.0×10^{-1}	$\mathbf{2.2\times10^{-2}}$	3.4×10^{-2}
	20	2.0×10^{-1}	4.6×10^{-1}	$\mathbf{2.9\times10^{-2}}$	9.6×10^{-2}

Table 4: *Relative error in estimating the Lie bracket for of two vector-fields for different band-widths, b, and different discretizations of the sphere, “icosa” vs. “CH”.*

tions are real-valued.) Setting π to be the map projecting vector-fields onto the tangent space of the unit sphere:

$$[\pi(Z)](p) \equiv Z(p) - p\cdot\langle Z(p), p\rangle,$$

we compute the projected vector-fields $X=\pi(\tilde{X})$ and $Y=\pi(\tilde{Y})$. and obtain their Lie bracket as:

$$[X,Y] \equiv \pi(\nabla X\cdot Y - \nabla Y\cdot X)$$

with ∇ the Euclidean gradient applied to the coordinates of the vector field. Because $\tilde{X}$ and $\tilde{Y}$ are linear combinations of trigonometric functions, and because the projection π has a simple expression, the input vector-fields X and Y, and the ground-truth bracket $[X,Y]$ can be evaluated in closed-form.

Tab. 4 gives the relative errors for the estimated bracket, computed for band-widths $b\in\{2,5,10,20\}$, using the two tessellations of the sphere. (For [AOCBC15] we computed the brackets using spectral dimensions varying from 10 to 1000 and give the smallest error.)

As expected, the table shows that discretization accuracy deteriorates as vector-fields become higher frequency and as the tessellation becomes less uniform. In addition, we see that while all methods perform comparably on the subdivided icosahedron, the operator-based approaches are less robust in the case of non-uniform sampling. Interestingly, despite the discontinuity of their representation, we find that the finite-elements discretization of Stein *et al.* provides the best results. (It is possible that some of the performance is due to the fact that Stein *et al.*'s basis has approximately 3× the numbers of degrees of freedom.)

Fig. 12 visualizes the results, showing the bracket estimated using the four different approaches (first four columns) as well as the analytic solution (right), for vector-fields $\tilde{X}$ and $\tilde{Y}$ with band-width $b=10$. On the uniform tessellation of the sphere obtained by subdividing the icosahedron (top) all methods match the ground-truth. However, on the non-uniform tessellation the brackets computed using the operator-based approaches produce tangent vectors whose magnitudes do not match those of the ground-truth. (Note that both the seeding of the initial noise for anisotropic diffusion, and the seeding of flow-lines at triangle centers are tessellation dependent. As such, while the visualizations in the rightmost column represent the same ground-truth vector-field, they appear different.)

Performance Under Refinement We also evaluate the performance of our method under refinement. To this end we construct tessellations of the torus by randomly sampling points in the periodic plane, computing the Delaunay triangulation, and mapping to the torus in 3D with inner (resp. outer) radius 2.0 (resp. 1.0):

$$\Phi(s,t) = \big(\cos(s), 0, \sin(s)\big)\cdot\big(2+\cos(t)\big) + \big(0, \sin(t), 0\big).$$

As above, we synthesize band-limited vector-fields with uniformly random Fourier coefficients, $X,Y:\mathbb{R}^2\to\mathbb{R}^2$, this time in the parameterization domain. Leveraging the metric-independence of the Lie bracket, we obtain the ground-truth by analytically computing $[X,Y]:\mathbb{R}^2\to\mathbb{R}^2$ and using the differential of the parameterization Φ to map it to 3D. We compare this to the bracket obtained by applying our method to the 3D vector-fields obtained by applying the differential of the embedding, $d\Phi$, to X and Y.

Fig. 13 (bottom left) shows the relative errors between the estimated and ground-truth brackets, computed for varying band-widths $b\in[1,50]$ and averaged over 10 tessellations. As expected, error increases with frequency and decreases consistently as we increase the resolution of the tessellation from $10K$ vertices to $160K$. The figure also shows an example of two vector-fields, X and Y, computed at a bandwidth of $b=10$ using $80K$ vertices (top) and the estimated Lie bracket $[X,Y]$ (bottom right).

8. Conclusion and Future Directions

8.1. Future work

Our work is distinct from other directional-field methods in that it decouples the tangent bundle from the geometry of the mesh. We demonstrated that this unlocks a potential for smooth fields with PL vertex-based normals, and we believe it has the potential for more flexible choices of tangent bundle representation. In the immediate future, we would like to consider extensions of our work using second-order Lagrange elements for discretization, and using a hierarchy akin to that in Subdivision Exterior Calculus [dGDMD16] to support computation over a smoother surface. More broadly, we would like to explore incorporating structure preservation into our discretization by building cochain sequences and optimizing the choice of per-vertex normals to improve the quality of the discretization. Additionally, it may be possible to extend our work to vector-field processing with general normal maps [TWBO03], allowing for detailed field editing on coarse meshes. Finally, we would like to develop polar representations [CDS10, WV25] for our fields, to be able to explicitly control singularities.

8.2. Discussion

Rotation Like the discretizations of Knöppel *et al.* [KCPS13] and Stein *et al.* [SWJG20], assuming that the triangle mesh is orientable, our discretization is closed under the action of counter-clockwise rotation by 90° in the tangent plane. In particular, this suggests that the discretization does not exhibit a preference for divergence-free vs. curl-free vector-fields. This is in contrast to the Whitney 1-form basis that only exhibit curl in the interior of triangles, and whose 90° counter-clockwise rotations are not in the space of Whitney 1-forms.

Figure 12: *Visualization of the brackets computed for a frequency* $b = 10$ *function using a regular tessellation obtained by subdividing an icosahedron (top) and an irregular tessellation obtained by randomly sampling points on the unit sphere and computing the convex hull (bottom).*

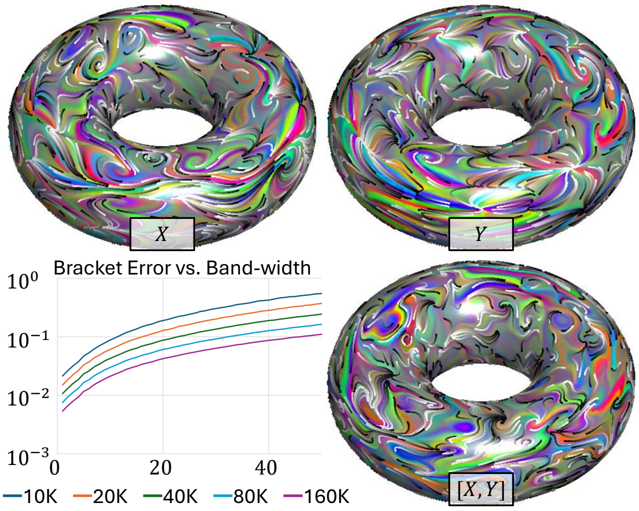

Figure 13: *The relative error between the estimated Lie bracket of two vector-fields on a torus and the analytic solution, given as a function of the band-width of the vector-fields. The plots show that error increases with the frequency of the vector-fields, and decreases as the resolution is increased.*

Relation to Scalar Gradients As with the discretization of Knöppel *et al.* [KCPS13], our vector-fields are continuous across the mesh. An implication of this is that they cannot represent the gradients of functions expressed in terms of the hat-basis, as those gradients are piecewise constant.

Implicit Geometry In our discretization, we implicitly treat the defined $\widetilde{d\Phi}|_p$ operator as the differential of an embedding, using it to define the mapping from the tangent space of the unit right triangle to the tangent space of the sphere. However, there may not be any Φ whose differential is $\widetilde{d\Phi}$. For example, the rows of $\widetilde{d\Phi}|_p$ may not be curl-free. One immediate implication is that the associated second fundamental form, $\widetilde{d\Phi}|_p^\top \cdot dN|_p$, need not be symmetric. It would be interesting to consider other choices for $\widetilde{d\Phi}|_p$, including ones that do not define an orthogonal transformation between $T_p\mathbb{T}$ and $T_{N(p)}S^2$, thereby inducing a non-constant Riemmanian metric within each triangle.

Acknowledgements

Mirela Ben-Chen acknowledges the support of the Israel Science Foundation (grant No. 1073/21).

References

[ABCCO13] Azencot O., Ben-Chen M., Chazal F., Ovsjanikov M.: An operator approach to tangent vector field processing. *Computer Graphics Forum 32*, 5 (2013), 73–82. 2, 11, 12

[AOCBC15] Azencot O., Ovsjanikov M., Chazal F., Ben-Chen M.: Discrete derivatives of vector fields on surfaces – an operator approach. *ACM Trans. Graph. 34*, 3 (May 2015). 2, 3, 11, 12

[Bos98] Bossavit A.: *Computational electromagnetism : variational formulations, complementarity, edge elements / Alain Bossavit.* Electromagnetism. Academic Press, San Diego, 1998. 8

[BV22] Boksebeld I., Vaxman A.: High-order directional fields. *ACM Trans. Graph. 41*, 6 (2022). 2, 7

[CDS10] Crane K., Desbrun M., Schröder P.: Trivial connections on discrete surfaces. *Computer Graphics Forum (SGP) 29*, 5 (2010), 1525–1533. 7, 12

[CL93] Cabral B., Leedom L. C.: Imaging vector fields using line integral convolution. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1993), SIGGRAPH '93, Association for Computing Machinery, p. 263–270. 8

[CS13] Crane K., Schröder P.: fieldgen. URL: `https://github.com/GeometryCollective/fieldgen/`. 8

[dC16] DO CARMO M. P.: *Differential Geometry of Curves and Surfaces: Revised and Updated Second Edition*. Dover Publications, New York, United States, 2016. 6

[dGDMD16] DE GOES F., DESBRUN M., MEYER M., DEROSE T.: Subdivision exterior calculus for geometry processing. *ACM Trans. Graph. 35*, 4 (July 2016). 12

[dGDT16] DE GOES F., DESBRUN M., TONG Y.: Vector field processing on triangle meshes. In *ACM SIGGRAPH 2016 Courses* (New York, NY, USA, 2016), SIGGRAPH '16, Association for Computing Machinery. 2

[dGLB*14] DE GOES F., LIU B., BUDNINSKIY M., TONG Y., DESBRUN M.: Discrete 2-tensor fields on triangulations. *Computer Graphics Forum 33*, 5 (2014), 13–24. 3

[Dje16] DJERBETIAN A.: *Tangent Vector Fields on Triangulated Surfaces – An Edge-Based Approach*. Master's thesis, Technion, 2016. 2

[DKT06] DESBRUN M., KANSO E., TONG Y.: Discrete differential forms for computational modeling. In *ACM SIGGRAPH 2006 Courses* (New York, NY, USA, 2006), SIGGRAPH '06, Association for Computing Machinery, p. 39–54. 2, 8

[DPR00] DIEWALD U., PREUßER T., RUMPF M.: Anisotropic diffusion in vector field visualization on euclidean domains and surfaces. *IEEE Transactions on Visualization and Computer Graphics 6* (Apr. 2000), 139–149. 8

[DVPSH14] DIAMANTI O., VAXMAN A., PANOZZO D., SORKINE-HORNUNG O.: Designing n-polyvector fields with complex polynomials. *Comput. Graph. Forum 33*, 5 (Aug. 2014), 1–11. 7

[FSBS06] FISHER M., SPRINGBORN B., BOBENKO A. I., SCHRODER P.: An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. In *ACM SIGGRAPH 2006 Courses* (New York, NY, USA, 2006), SIGGRAPH '06, Association for Computing Machinery, p. 69–74. 10

[FSDH07] FISHER M., SCHRÖDER P., DESBRUN M., HOPPE H.: Design of tangent vector fields. *ACM Trans. Graph. 26*, 3 (July 2007), 56–es. 2

[Gat14] GATICA G.: *A Simple Introduction to the Mixed Finite Element Method*. SpringerBriefs in Mathematics, 2014. 2

[Hip99] HIPTMAIR R.: Canonical construction of finite elements. *Mathematics of Computation 68*, 228 (1999), 1325–1346. 2

[Hir03] HIRANI A.: *Discrete Exterior Calculus*. PhD thesis, California Institute of Technology, 2003. 2

[KCPS13] KNÖPPEL F., CRANE K., PINKALL U., SCHRÖDER P.: Globally optimal direction fields. *ACM Trans. Graph. 32*, 4 (July 2013). 2, 4, 5, 8, 10, 12, 13

[KCPS15] KNÖPPEL F., CRANE K., PINKALL U., SCHRÖDER P.: Stripe patterns on surfaces. *ACM Trans. Graph. 34* (2015). 2

[LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1987), SIGGRAPH '87, Association for Computing Machinery, p. 163–169. 9

[Loo87] LOOP C.: *Smooth Subdivision Surfaces Based on Triangles*. Master's thesis, January 1987. 8, 11

[LTGD16] LIU B., TONG Y., GOES F. D., DESBRUN M.: Discrete connection and covariant derivative for vector field analysis and design. *ACM Trans. Graph. 35* (Mar. 2016). 2, 3, 5, 8

[SC*19] SHARP N., CRANE K., ET AL.: Geometrycentral: A modern c++ library of data structures and algorithms for geometry processing. 8

[SJWG20] STEIN O., JACOBSON A., WARDETZKY M., GRINSPUN E.: A smoothness energy without boundary distortion for curved surfaces. *ACM Trans. Graph. 39* (Mar. 2020). 2

[SSC19] SHARP N., SOLIMAN Y., CRANE K.: The vector heat method. *ACM Trans. Graph. 38*, 3 (2019). 2, 8, 10, 16

[SWJG20] STEIN O., WARDETZKY M., JACOBSON A., GRINSPUN E.: A simple discretization of the vector dirichlet energy. *Computer Graphics Forum 39*, 5 (2020), 81–92. 2, 5, 8, 10, 11, 12, 16

[Tay08] TAYLOR M. A.: Asymmetric cubature formulas for polynomial integration in the triangle and square. *J. Comput. Appl. Math. 218*, 1 (Aug. 2008), 184–191. 5

[TWBO03] TASDIZEN T., WHITAKER R., BURCHARD P., OSHER S.: Geometric surface processing via normal maps. *ACM Transactions on Graphics (TOG) 22*, 4 (2003), 1012–1033. 12

[VCD*16] VAXMAN A., CAMPEN M., DIAMANTI O., PANOZZO D., BOMMES D., HILDEBRANDT K., BEN-CHEN M.: Directional field synthesis, design, and processing. *Computer Graphics Forum 35*, 2 (2016), 545–572. 2, 7

[Whi57] WHITNEY H.: *Geometric Integration Theory*. Princeton University Press, 1957. 2, 8

[WV25] WANG J. B., VAXMAN A.: Power-linear polar directional fields. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers* (2025), SIGGRAPH Conference Papers '25. 12

[WWT*06] WANG K., WEIWEI, TONG Y., DESBRUN M., SCHRÖDER P.: Edge subdivision schemes and the construction of smooth vector fields. *ACM Trans. Graph. 25*, 3 (July 2006), 1041–1048. 2

Appendix A: Coordinate-Free Decomposition of Endomorphisms

Given an inner-product space $\{V, B : V \to V^*\}$, with B symmetric and positive definite, we would like to show that the maps taking an endomorphism $L \in \mathrm{End}(V)$ into $\mathcal{Div}(V)$, $\overline{\mathcal{Holom}}(V)$, and $\mathcal{Curl}(V)$ are orthogonal projections.

To show that a mapping is a *projection* onto a subspace requires showing (1) that the image of the map is the subspace and (2) that the map acts as the identity on the subspace. To show that the projection is *orthogonal* it suffices to show that it is self-adjoint.

We review how inner-products $B_V : V \to V^*$ and $B_W : W \to W^*$ on vector spaces V and W define an inner-product on the space of homomorphisms between V and W. In particular, this allows us to define an inner-product on the space of endomorphisms $\mathrm{End}(V)$, which is needed to show that a projection is orthogonal.

Then, we step through the derivations showing that the map onto the subspace of self-adjoint operators:

$$L \mapsto \left(\frac{L + B^{-1} \circ L^* \circ B}{2}\right)$$

is an orthogonal projection. (A similar argument shows that the other maps are orthogonal projections as well.)

Inner-Products on Linear Maps

Given inner-product spaces $\{V, B_V\}$ and $\{W, B_W\}$, there is a canonical inner-product on the space of linear maps between V and W, defined in terms of the trace of an endomorphism on V. Concretely, for linear maps $L, M \in \mathrm{Hom}(V, W)$:

$$\langle L, M \rangle \equiv \mathrm{tr}\left(B_V^{-1} \circ L^* \circ B_W \circ M\right). \tag{6}$$

This definition is independent of the choice of bases for V and W. We note that the argument to the trace is itself an endomorphism on V. Concretely: M maps V to W; B_W maps W to W^*; L^* maps W^* to V^*; and B_V^{-1} maps V^* back to V. Thus, the trace is well-defined.

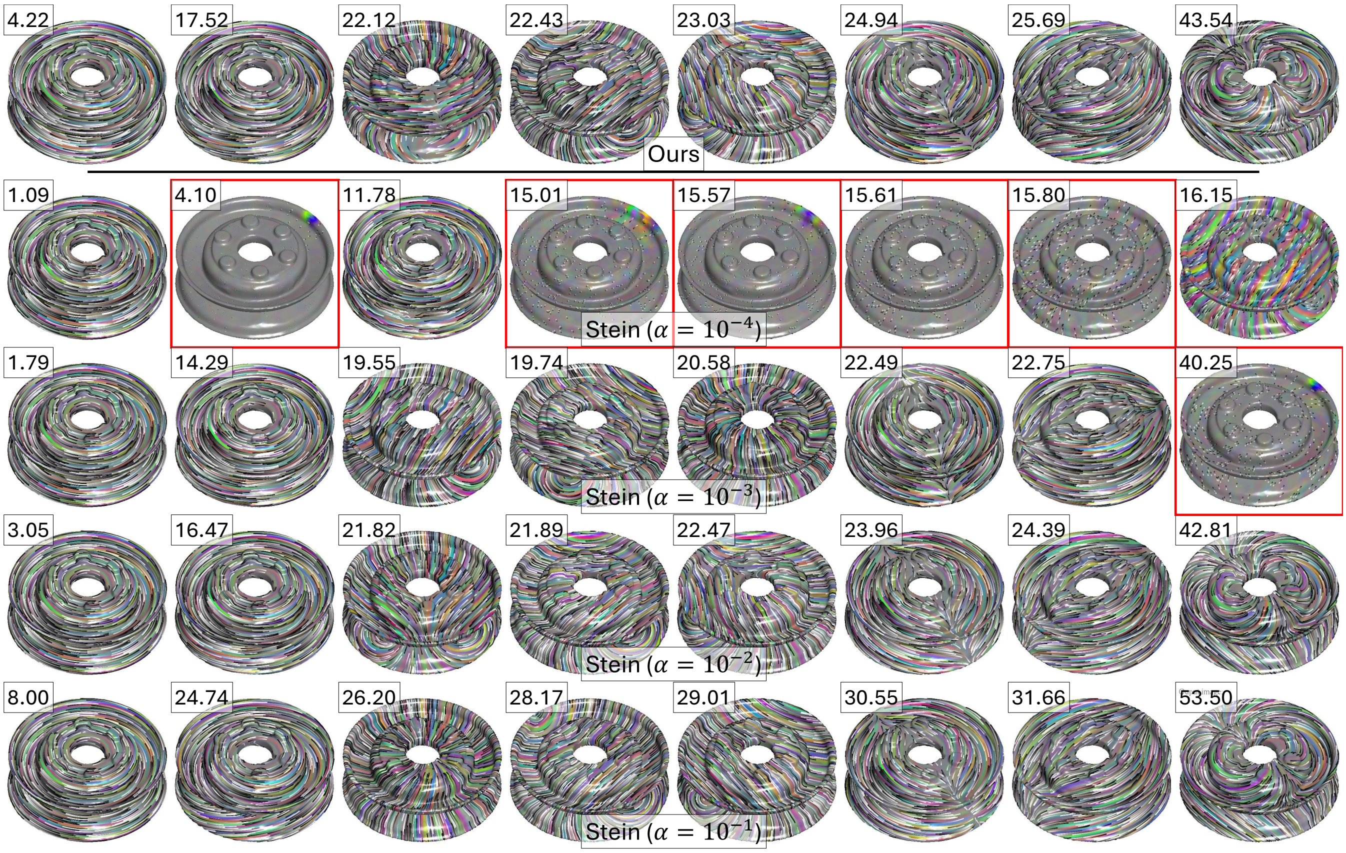

Figure 14: *Spectra of the Killing energy: Comparing the smallest eigenvectors obtained using our approach (top row) with those obtained Stein et al.'s method (second through fifth rows), with different regularization weights* α. *Spurious minima are highlighted in red.*

In particular, when operators are represented w.r.t. orthonormal bases on V and W, so that the inner-products B_V and B_W are represented by the identity matrices, the expression reduces to

$$\mathrm{tr}(\mathbf{L}^\top \cdot \mathbf{M}) \equiv \langle \mathbf{L}, \mathbf{M} \rangle_F,$$

the Frobenius inner-product of the matrices $\mathbf{M}$ and $\mathbf{L}$ associated to M and L respectively.

Orthogonal Projection

We begin by showing that the map:

$$\begin{aligned} \pi : \mathrm{End}(V) &\to \mathrm{End}(V) \\ L &\mapsto \frac{L + B^{-1} \circ L^* \circ B}{2} \end{aligned}$$

is a projection onto the subspace of self-adjoint endomorphism and then show that it is orthogonal.

Projection

To show that π is a projection we show that its image lies in the subspace of self-adjoint operators and that it acts as the identity on that subspace.

Image of π Noting that that the dual of the composition is the composition of the dual in reversed order (i.e. $(L \circ M)^* = M^* \circ L^*$) and that B is symmetric (i.e. $B = B^*$ and $(B^{-1})^* = B^{-1}$):

$$\begin{aligned} B \circ \pi(L) = B \circ \left(\frac{L + B^{-1} \circ L^* \circ B}{2} \right) &= \left(\frac{B \circ L + L^* \circ B}{2} \right) \\ &= \left(\frac{L^* \circ B + B \circ L}{2} \right)^* = \big(B \circ \pi(L)\big)^* = \pi(L)^* \circ B. \end{aligned}$$

Thus, the image of π is self-adjoint.

Identity on the subspace In the case that the endomorphism L is self-adjoint – i.e. that $B \circ L = L^* \circ B$ – we have:

$$\pi(L) = \left(\frac{L + B^{-1} \circ L^* \circ B}{2} \right) = \left(\frac{L + B^{-1} \circ B \circ L}{2} \right) = L.$$

Orthogonality

To show that π is an orthogonal projection, we show that it is self-adjoint. That is, for all $L, M \in \mathrm{End}(V)$ we show that:

$$\langle \pi(L), M \rangle = \langle \pi(M), L \rangle,$$

	[SSC19]	[SWJG20]	Ours
Without unfolding	0.335	0.468	**0.332**
With unfolding	0.214	0.222	**0.208**

Table 5: *Relative smoothness of the interpolating vector-fields obtained using a smoothness energy defined by the Connection Laplacian, and visualized in Fig. 6 (top).*

with the inner-product on endomorphisms induced from the inner-product on V (Eq. 6). Expanding twice the left-hand-side, we get:

$$\begin{aligned} 2\cdot\langle\pi(L),M\rangle &= \langle L,M\rangle + \langle B^{-1}\circ L^{*}\circ B, M\rangle \\ &= \langle L,M\rangle + \operatorname{tr}\left(B^{-1}\circ\left(B^{-1}\circ L^{*}\circ B\right)^{*}\circ B\circ M\right) \\ &= \langle L,M\rangle + \operatorname{tr}\left(B^{-1}\circ\left(B\circ L\circ B^{-1}\right)\circ B\circ M\right) \\ &= \langle L,M\rangle + \operatorname{tr}\left(L\circ M\right). \end{aligned}$$

As this is symmetric in L and M, the projection π is self-adjoint.

Appendix B: Comparing Sparse Interpolation

Although a ground-truth solution for the sparse interpolation problem described in Sec. 7.1 is not available, one can consistently measure the smoothness of the different results. To this end, we sample the different solutions at the faces to obtain a per-triangle representation of the vector-field, and measure the difference between vectors assigned to triangles on opposite sides of an edge.

For our approach, the per-triangle vectors are obtained by directly evaluating the basis functions to obtain extrinsic vectors. For the methods of Stein *et al.* and for the Whitney basis, this is done by first evaluating the basis functions at the centers of the triangles to obtain an intrinsic representation of the tangent vectors, and then using the triangles' embeddings to obtain extrinsic vectors tangent to the triangles. For the method of Sharp *et al.*, we interpret the per-vertex coefficients as extrinsic vectors perpendicular to the vertices' normals and compute their average.

We measure the relative smoothness by taking the sum of squared differences between the vectors associated to triangles incident on an edge, weighted by the ratio of the (absolute) primal-to-dual edge lengths. The error is normalized by the area-weighted magnitude of the vector-field.

We compare the vector-fields in two ways: (1) We directly compute the weighted sum of squared differences; and (2) Before computing the contribution of an edge, we first perform a "hinge-unfolding" to bring the two incident triangles into a common plane and compare the aligned vectors. Because our method does not generate tangent vectors perpendicular to the triangle's normal, we first apply the Rodrigues rotation taking the interpolated normal to the triangle's normal before performing the hinge-unfolding. Similarly, for the method of Sharp *et al.*, we obtain vectors perpendicular to the triangle's normal by applying the Rodrigues rotation mapping the vertex normals to the triangles' normals, before averaging to get the per-triangle vector.

Tab. 5 gives the relative smoothness for the different methods visualized in Fig. 6 (top), indicating that our method consistently generates a smooth interpolant. In addition, we find that the hinge-unfolding provides a lower smoothness energy for all methods, which is expected as vectors are not forced to point in different directions when the associated triangles meet at a sharp edge.

As the smoothness mimics the Connection Laplacian – e.g. it associates a non-zero energy to harmonic vector-fields – we do not believe it provides a meaningful measure of the quality of the interpolated field. (However, in this case as well, the relative smoothness for the vector-field computed using our is method is lower than that computed using either the Whitney basis or the method of Stein *et al.*)

Appendix C: Bracket Comparison

Fig. 14 compares the smallest eigenvectors obtained using our discretization of the Killing energy (first row), with the eigenvectors obtained using Stein *et al.*'s discretization. As noted in that work, directly using the Killing energy defined by their discretization is not robust due to spurious minimizers. They address this by adding an additional connection energy regularizer to the Killing energy. The different rows show the eigenvectors obtained using Stein *et al.*'s discretization for progressively larger values of regularization weight α. The visualization also gives the associated eigenvalue in the top left corner of each figure.

While a small regularization weight $\alpha = 10^{-4}$ effectively generates the as Killing-as-possible vector-field associated with the near rotational symmetry of the pulley, the problem with spurious minima is evident in subsequent eigenvectors. This is mitigated by increasing the regularization weight to $\alpha = 10^{-3}$, though even in this case one sees artifacts in the right-most column. While further increasing the regularization weight to $\alpha = 10^{-2}$ produces results similar to ours, this can have the detrimental effect of biasing the spectral decomposition away from Killing vector-fields in preferences of vector-fields that are smoother.

This highlights a limitation observed by Stein *et al.* – while their discretization provides properties desirable for the connection Laplacian (e.g. linear reproduction), the underlying Crouzeix-Raviart basis is discontinuous and can result in artifacts when used to define other energies that depend on vector-field derivatives.