Steklov Spectral Geometry for Extrinsic Shape Analysis

YU WANG, Massachusetts Institute of Technology, USA

MIRELA BEN-CHEN, Technion – Israel Institute of Technology, Israel

IOSIF POLTEROVICH, Université de Montréal, Canada

JUSTIN SOLOMON, Massachusetts Institute of Technology, USA

We propose using the *Dirichlet-to-Neumann operator* as an extrinsic alternative to the Laplacian for spectral geometry processing and shape analysis. Intrinsic approaches, usually based on the Laplace–Beltrami operator, cannot capture the spatial embedding of a shape up to rigid motion, and many previous extrinsic methods lack theoretical justification. Instead, we consider the Steklov eigenvalue problem, computing the spectrum of the Dirichlet-to-Neumann operator of a surface bounding a volume. A remarkable property of this operator is that it completely encodes volumetric geometry. We use the boundary element method (BEM) to discretize the operator, accelerated by hierarchical numerical schemes and preconditioning; this pipeline allows us to solve eigenvalue and linear problems on large-scale meshes despite the density of the Dirichlet-to-Neumann discretization. We further demonstrate that our operators naturally fit into existing frameworks for geometry processing, making a shift from intrinsic to extrinsic geometry as simple as substituting the Laplace–Beltrami operator with the Dirichlet-to-Neumann operator.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; • **Mathematics of computing** → **Partial differential equations**;

Additional Key Words and Phrases: Shape analysis, geometry processing, Steklov eigenvalue problem, Dirichlet-to-Neumann operator

ACM Reference format:

Yu Wang, Mirela Ben-Chen, Iosif Polterovich, and Justin Solomon. 2018. Steklov Spectral Geometry for Extrinsic Shape Analysis. *ACM Trans. Graph.* 38, 1, Article 7 (December 2018), 21 pages.
<https://doi.org/10.1145/3152156>

1 INTRODUCTION

Geometry processing and shape analysis tools for computer graphics typically draw from two complementary theories of geometry.

J. Solomon acknowledges funding from an MIT–Skoltech Seed Fund grant (“Boundary Element Methods for Shape Analysis”), the Skoltech–MIT Next Generation Program, the MIT Research Support Committee, an Amazon Research Award, Army Research Office grant W911NF-12-R-0011 (“Smooth Modeling of Flows on Graphs”), and National Science Foundation grant IIS-1838071 (“BIGDATA:F:Statistical and Computational Optimal Transport for Geometric Data Analysis”). I. Polterovich acknowledges the support of NSERC, FRQNT, the Canada Research Chairs program, and the Weston Visiting Professorship program of the Weizmann Institute of Science where part of this research was performed. M. Ben-Chen acknowledges funding from the European Research Council (ERC starting grant no. 714776 “OPREP”) and the Israel Science Foundation (grant no. 504/16). Y. Wang acknowledges the support of the MIT Grier Presidential Fellowship and Siebel Scholarship.

Authors’ addresses: Y. Wang, Massachusetts Institute of Technology, 32 Vassar St., Cambridge, MA, 02139; email: wangyu9@mit.edu; M. Ben-Chen, Technion – Israel Institute of Technology, Israel; email: mirela@cs.technion.ac.il; I. Polterovich, Université de Montréal, Département de mathématiques et de statistique, Montréal, QC, H3C 3J7, Canada; email: iosif@dms.umontreal.ca; J. Solomon, Massachusetts Institute of Technology, 32 Vassar St., Cambridge, MA, 02139, USA; email: jsolomon@mit.edu.

This work is licensed under a Creative Commons Attribution International 4.0 License.

Copyright 2018 held by Owner/Author
0730-0301/2018/12-ART7
<https://doi.org/10.1145/3152156>

To distinguish these, consider a closed surface embedded in three dimensions. From the viewpoint of *extrinsic geometry*, we might examine the surface as the outer boundary of a volume. This approach relies on distances and other measurements taken from the space surrounding the surface to understand its shape. Contrastingly, many techniques in differential geometry decouple extrinsic shape from *intrinsic geometry*, which is concerned with quantities like geodesic distances that can be measured without leaving the outer surface. A crowning achievement of classical differential geometry shows that certain quantities like Gaussian curvature can be measured intrinsically; this basic observation inspired theoretical exploration of purely intrinsic techniques. Computational geometry processing embraced the intrinsic perspective early on, leading to numerous applications of intrinsic computations.

Intrinsic geometry, however, is an ineffective description of shape for many applications. First, the spatial embedding information is lost. The variability in embedding can be essential. As an extreme example, consider searching a database

Fig. 1.

of origami models: from an intrinsic perspective, all origami is equivalent to the same piece of flat paper. Second, the intrinsic perspective is a counterintuitive way to describe the shapes of many real-world objects, e.g., identifying the inward and outward bumps on the cubes in Figure 1.

A naïve approach to extrinsic geometry could be to use the (x, y, z) coordinates of the embedding. This incorporates information about the embedding but is not invariant to rigid motion. Alternatively, rotation-invariant shape descriptors (e.g., built from spherical harmonic power spectra) usually involve extrinsic information about a surface as a shell rather than as the boundary of a volume.

In this article, we provide a practical and mathematically justified spectral approach to extrinsic geometry for geometry processing, via an extrinsic alternative to the intrinsic Laplace–Beltrami operator. The end result is a *surface-only* approach to *volume-aware* shape analysis, using a boundary operator that takes the interaction between nonadjacent vertices into consideration. In particular, we consider the Dirichlet-to-Neumann (DtN) operator, also known as the Poincaré–Steklov or simply the Steklov operator, and its spectrum (the *Steklov spectrum*). A physical interpretation of the DtN operator is that it maps from a temperature distribution on the surface to the resulting heat flux through the surface. As we will prove, DtN and its spectrum completely encode extrinsic geometry in the smooth case. We also show that DtN and some of its peers can be applied efficiently using the boundary element method (BEM), accelerated using preconditioning and hierarchical techniques that scale linearly in the size of a boundary triangle mesh.

We generalize our method to additional geometric primitives, notably including surfaces with open boundary. Since it is built using the Laplace equation, our discretized Dirichlet-to-Neumann operator naturally fits into existing frameworks for geometry processing and analysis. We show that in applications including compression, measurement of shape differences, and spectral distance computation, shifting from intrinsic geometry to extrinsic geometry is as simple as substituting the cotangent Laplacian with a discretized Steklov operator. This simple change leads to qualitatively different behavior for geometric algorithms.

2 RELATED WORK

2.1 Intrinsic Geometry

Intrinsic computations are ubiquitous in geometry processing, applied to problems including distance computation (Crane et al. 2013; Lipman et al. 2010), segmentation (Reuter et al. 2009), shape description (Kokkinos et al. 2012), shape retrieval (Bronstein et al. 2011), and correspondence (Ovsjanikov et al. 2012); these citations are a small sampling of the literature.

Many intrinsic algorithms use the Laplace–Beltrami operator, which can be approximated, e.g., with the celebrated cotangent formula (Pinkall and Polthier 1993) or with convergent approximations on simplicial complexes (Belkin et al. 2008) and point clouds (Belkin et al. 2009; Liang et al. 2012). Laplace–Beltrami eigenfunctions are multiscale and can characterize shape up to isometry (Aubry et al. 2011; Sun et al. 2009); truncation in this basis also provides favorable approximation properties (Aflalo et al. 2015). See the surveys by Zhang et al. (2010), Sorkine (2005), and Patané (2016) for discussion of this technique and applications.

Functional maps, proposed by Ovsjanikov et al. (2012), have inspired interest in a broader “operator-based” approach to geometry. Operator-based geometry processing now includes intrinsic techniques for shape exploration (Rustamov et al. 2013), vector field processing (Azencot et al. 2013, 2015), interpolation (Azencot et al. 2016), simulation (Azencot et al. 2014), and deformation (Boscaini et al. 2015)—among other tasks.

The intrinsic approach is desirable in many scenarios. Many real-world objects deform near isometrically, that is, without affecting intrinsic structure. Isometry invariance then creates an advantage: an object can be bent to another pose while still being considered the same intrinsically. For example, intrinsic invariance can be favorable in shape retrieval since the same object in different poses will be clustered together by design.

2.2 Extrinsic Geometry

Extrinsic techniques appear in geometry processing, usually without the multiscale and stability advantages of spectral geometry. The most common expression of discrete extrinsic geometry is in computation of mean curvature, the extrinsic counterpart of Gaussian curvature. Extrinsic methods appear in computer vision, e.g., the SHOT (Tombari et al. 2010) and D2 (Osada et al. 2002) descriptors, typically without completeness guarantees.

A few methods attempt to extend operator-based methods to incorporate extrinsic information. The volumetric heat kernel signature (Raviv et al. 2010) uses a meshing of the volume enclosed by a surface, applied to segmentation by Litman et al. (2012); Wang and

Wang (2015) present a similar construction on tetrahedral meshes. This is computationally expensive and has the property that results depend on the resolution of the interior mesh; note that methods like Patané (2015) may accelerate heat kernel evaluation but still require meshing of the interior. Hildebrandt et al. (2010) consider a family of deformation energies whose Hessian eigenmodes are sensitive to extrinsic features. Rustamov (2011) proposes extending surface-based Laplace–Beltrami eigenfunctions to the interior of a domain using generalized barycentric coordinates. This extends descriptors like the heat kernel signature (HKS) from the outer surface to the interior but is still intrinsic in the sense that when it is restricted to the outer surface, it coincides with the surface-based computation. Wu et al. (1997) exploit the correlation between the simulated charge density on a geometric domain and local concavity for shape segmentation. Wang et al. (2014) modify the Laplace–Beltrami operator to include concavity information, aiding in applications like segmentation without distinguishing between the interior and exterior of a closed surface. Similarly, Corman et al. (2017) use the Laplacian of a meshed shell around a triangulated surface, incorporating mean curvature but treating the surface as a thin shell rather than the boundary of a volume. Techniques like Chuang et al. (2009) use extrinsic calculations to approximate the Laplace–Beltrami operator; these methods aim to capture intrinsic geometry.

We use the Dirichlet-to-Neumann (DtN) operator as an extrinsic analog to the Laplace–Beltrami operator. In graphics, Gao et al. (2014) use DtN for skinning; it also appears in a recent pipeline for parameterization (Sawhney and Crane 2017). In simulation and PDEs, it appears as a Schur complement eliminating interior degrees of freedom (Bertoluzza 2003; Liu et al. 2016; Quarteroni and Valli 1999; Smith et al. 2004; Toselli and Widlund 2005).

Recent concurrent work proposes using the Dirac operator to capture extrinsic shape (Liu et al. 2017). This method is among the first to extend spectral geometry processing to include extrinsic information, but it still treats the surface as a shell rather than as the boundary of a volume. While this work opens intriguing theoretical questions regarding the informativeness of the Dirac operator (see their Section 7), the Dirichlet-to-Neumann operator we propose for shape analysis enjoys strong grounding in existing theory; this allows us to prove in Section 3.1 that the Dirichlet-to-Neumann operator fully captures extrinsic geometry up to rigid motion (Proposition 3.1). On the other hand, unlike the operators we consider in this article, their Dirac operator benefits from the efficiency of sparse linear algebra. See Section 7.7 for additional discussion and empirical comparison to this technique.

2.3 Numerical PDE

Two typical numerical methods for approximating Dirichlet-to-Neumann operators are the finite element method (FEM) and the BEM. FEM yields a sparse linear system on a discretization of the volume bounded by a surface, while BEM yields a dense linear system on just the boundary. Efficiency-wise, the added number of elements required for FEM makes the two methods similar asymptotically. We use BEM to avoid dependence on volumetric meshing.

Fast numerical methods can be used to accelerate BEM, including the fast multipole method (FMM) and the hierarchical matrix (H-matrix) method. FMM (Engheta et al. 1992; Rokhlin 1985)

expands the Green's function using a multipole approximation, grouping close points as a single source. The H-matrix method represents BEM operators hierarchically, reducing the cost of assembly and matrix-vector products to $O(n \log n)$ time (Börm et al. 2003; Hackbusch 1999). Either method reduces evaluation cost to $O(n \log 1/\epsilon)$ time, where ϵ is a prescribed accuracy, making it possible to solve large problems on a single workstation.

3 MATHEMATICAL PRELIMINARIES

3.1 Steklov Eigenproblem

Denote a volumetric domain as $\Omega \subseteq \mathbb{R}^3$ with boundary surface $\Gamma = \partial\Omega$. The Steklov eigenvalue problem is defined as

$$\begin{cases} \Delta\psi(x) = 0 & x \in \Omega \\ \nabla_{\mathbf{n}}\psi(x) = \lambda\psi(x) & x \in \Gamma, \end{cases}$$

where $\Delta = \partial^2/\partial x^2 + \partial^2/\partial y^2 + \partial^2/\partial z^2$ denotes the Laplacian operator and $\nabla_{\mathbf{n}}$ the normal derivative at the boundary. The spectrum of the Steklov problem is discrete and given by a sequence of eigenvalues

$$0 = \lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \dots \rightarrow \infty.$$

$\{\lambda_n\}_{n=0}^\infty$ can be interpreted as the eigenvalues of the Dirichlet-to-Neumann (DtN) operator $\mathcal{S} : H^{1/2}(\Gamma) \rightarrow H^{-1/2}(\Gamma)$ defined as $\mathcal{S}f := \nabla_{\mathbf{n}}(\mathcal{E}f)$, where H^s denotes the Sobolev space of order s (see, e.g., Steinbach (2007, Section 2.4) or Brezis (2010) for more details). $\mathcal{E}$ denotes the harmonic extension into the interior; denote $\phi_n = \psi_n|_{\Gamma} : \Gamma \rightarrow \mathbb{R}$ as the corresponding eigenfunctions. $\mathcal{S}$ is known as the Dirichlet-to-Neumann operator because it maps the boundary Dirichlet data to the Neumann data (e.g., voltage-to-current, or temperature-to-flux).

Just as the Laplace–Beltrami spectrum encodes intrinsic geometry, the spectrum of $\mathcal{S}$ encodes extrinsic information. In particular, for smooth domains in $\mathbb{R}^3$, the Steklov heat kernel admits the asymptotic expansion as $t \rightarrow 0^+$:

$$e^{-t\mathcal{S}}(x, x) = \sum_{i=0}^{\infty} e^{-t\lambda_i} \psi_i(x)^2 \sim \sum_{k=0}^{\infty} a_k(x) t^{k-2} + \sum_{l=1}^{\infty} b_l(x) t^l \log t, \quad (1)$$

where the coefficients $a_k(x)$ for $k = 0, 1, 2$ and $b_l(x)$ are local geometric invariants, in the sense that they are determined by the local geometry of Γ in a neighborhood of the point $x \in \Gamma$ (Polterovich and Sher 2015). See Duistermaat and Guillemin (1975), where such an expansion is proved for a general elliptic pseudodifferential operator.

The following expressions were obtained by Polterovich and Sher (2015):

$$a_0(x) \equiv \frac{1}{2\pi} \quad (2)$$

$$a_1(x) = \frac{H(x)}{4\pi} \quad (3)$$

$$a_2(x) = \frac{1}{16\pi} \left(H(x)^2 + \frac{K(x)}{3} \right), \quad (4)$$

where $H(x)$ and $K(x)$ are, respectively, the mean and the Gaussian curvatures of Γ . While the heat expansion for the Laplace–Beltrami

operator on the boundary captures only intrinsic geometric invariants of Γ , the first few terms of the Steklov heat expansion (Equation (1)) contain the mean curvature of the boundary, which is an extrinsic geometric quantity. We refer to the recent survey by Girouard and Polterovich (2017) for other results on the spectral geometry of $\mathcal{S}$.

Recall that the Laplace–Beltrami operator is invariant under isometries. For example, the bumpy cubes in Figure 1 have isometric boundaries, and therefore their boundary Laplacians coincide. At the same time, the boundaries of the bumpy cubes have different extrinsic geometries, and one could check that the DtN operators on these cubes are also different. Moreover, in striking contrast with the boundary Laplacian, under some assumptions the only maps preserving the DtN operator for domains in $\mathbb{R}^3$ are **rigid motions**:

PROPOSITION 3.1. *Suppose $\Omega_1, \Omega_2 \subseteq \mathbb{R}^3$ are compact domains with C^∞ boundaries Γ_1, Γ_2 and Dirichlet-to-Neumann operators $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively. Let $\alpha : \Omega_1 \rightarrow \Omega_2$ be a bijection that is C^∞ up to the boundary, and let $\tilde{\alpha} : \Gamma_1 \rightarrow \Gamma_2$ be the induced mapping between the boundaries. Suppose that the operators $\mathcal{S}_1$ and $\mathcal{S}_2$ coincide up to composition with $\tilde{\alpha}$, i.e., $\mathcal{S}_2 f = \tilde{\alpha}_* \mathcal{S}_1 \tilde{\alpha}^* f$, for any $f \in C^\infty(\Gamma_2)$, where $\tilde{\alpha}^* f = f \circ \tilde{\alpha}$, $\tilde{\alpha}_* g = g \circ \tilde{\alpha}^{-1}$ denote the pull-backs by $\tilde{\alpha}$ and $\tilde{\alpha}^{-1}$, respectively. Then α is a rigid motion.*

PROOF. The following argument was communicated to us by M. Karpukhin: by Proposition 1.3 and the discussion above it in Lee and Uhlmann (1989), the Dirichlet-to-Neumann map of a smooth surface determines the full Taylor series of the metric g at the boundary in the boundary normal coordinates. In particular, it determines the metric itself—which in our case is the first fundamental form of the boundary surface—and its first derivative, which gives us the second fundamental form (see, e.g., Kachalov et al. (2001, Section 2.1.18)). At the same time, it follows from Bonnet's theorem that the first and second fundamental forms determine a surface in $\mathbb{R}^3$ up to rigid motions. $\square$

Proving an analogous result for nonsmooth boundaries remains a significant challenge for future work.

3.2 Boundary Representation and Operators

To derive a boundary representation of the DtN operator, first consider the Laplace equation with Dirichlet boundary conditions:

$$\begin{cases} \Delta u(x) = 0 & x \in \Omega \\ u(x) = g(x) & x \in \Gamma, \end{cases} \quad (5)$$

where $g(\Gamma)$ is Dirichlet data on the boundary. A basic fact from elliptic PDE is that Equation (5) uniquely determines $u(\Omega)$ and hence its Neumann data $g_n = \frac{\partial}{\partial n} u(\Gamma)$. By definition, the Dirichlet-to-Neumann operator $\mathcal{S}$ is the map $g \mapsto g_n$.

To avoid discretizing the interior of the domain Ω , we use integral operators to bypass solving the Laplace equation (Equation (5)). Here, we introduce relevant boundary operators and notation that will be used throughout the article. We refer readers to Steinbach (2007) for a comprehensive introduction to these operators in the context of the boundary element method.

Single-Layer Potential. The single-layer potential $\mathcal{V} : H^{-1/2}(\Gamma) \rightarrow H^{1/2}(\Gamma)$ is defined via

$$[\mathcal{V}\phi](\mathbf{x}) := \int_{\Gamma} G(\mathbf{x}, \mathbf{y}) \phi(\mathbf{y}) d\Gamma(\mathbf{y}),$$

where $G(\mathbf{x}, \mathbf{y}) := \frac{1}{4\pi} \frac{1}{|\mathbf{x} - \mathbf{y}|}$ is the fundamental solution of the Laplace equation. Physically, $\mathcal{V}$ maps an input electric charge distribution ϕ to the resulting electric potential distribution.

Double-Layer Potential. The double-layer potential $\mathcal{K} : H^{1/2}(\Gamma) \rightarrow H^{1/2}(\Gamma)$ is defined via

$$[\mathcal{K}\phi](\mathbf{x}) := \int_{\Gamma} \frac{\partial G(\mathbf{x}, \mathbf{y})}{\partial n(\mathbf{y})} \phi(\mathbf{y}) d\Gamma(\mathbf{y}),$$

where the integral is understood in the sense of the Cauchy principal value (Kanwal 2013). Physically, $\mathcal{K}$ maps an input electric dipole density distribution ϕ to the resulting electric potential distribution.

Adjoint Double-Layer Potential. The adjoint double-layer potential $\mathcal{T} : H^{-1/2}(\Gamma) \rightarrow H^{-1/2}(\Gamma)$ is defined as the conormal derivative of $\mathcal{V}$:

$$[\mathcal{T}\phi](\mathbf{x}) := \int_{\Gamma} \frac{\partial G(\mathbf{x}, \mathbf{y})}{\partial n(\mathbf{x})} \phi(\mathbf{y}) d\Gamma(\mathbf{y}),$$

where the integral is understood in the sense of the Cauchy principal value. Physically, $\mathcal{T}$ maps an input electric charge density distribution ϕ to the normal derivatives of the resulting electric potential distribution.

Hypersingular Operator. The hypersingular operator $\mathcal{H} : H^{1/2}(\Gamma) \rightarrow H^{-1/2}(\Gamma)$ is defined as minus the conormal derivative of $\mathcal{K}$:

$$(\mathcal{H}\phi)(\mathbf{x}) := - \int_{\Gamma} \frac{\partial^2 G(\mathbf{x}, \mathbf{y})}{\partial n(\mathbf{x}) \partial n(\mathbf{y})} \phi(\mathbf{y}) d\Gamma(\mathbf{y}).$$

Physically, $\mathcal{H}$ maps an input electric dipole density distribution ϕ to normal derivatives of the resulting electric potential distribution.

Operator Properties. $\mathcal{V}$ and $\mathcal{H}$ are self-adjoint operators, i.e.,

$$\langle v, \mathcal{V}u \rangle_{\Gamma} \equiv \langle \mathcal{V}v, u \rangle_{\Gamma} \quad \text{and} \quad \langle v, \mathcal{H}u \rangle_{\Gamma} \equiv \langle \mathcal{H}v, u \rangle_{\Gamma},$$

where the inner product $\langle \cdot, \cdot \rangle_{\Gamma}$ is given by

$$\langle \phi, \psi \rangle_{\Gamma} := \int_{\Gamma} \phi(\mathbf{y}) \psi(\mathbf{y}) d\Gamma(\mathbf{y}).$$

$\mathcal{K}$, $\mathcal{T}$ are adjoint to each other, i.e., $\langle v, \mathcal{K}u \rangle_{\Gamma} \equiv \langle \mathcal{T}v, u \rangle_{\Gamma}$. Furthermore, $\mathcal{V}$ is positive definite, and $\mathcal{H}$ is positive semidefinite.

3.3 Calderón Projection and Dirichlet-to-Neumann Operator

The Calderón projection relates g and g_n through the relationship

$$\begin{pmatrix} g \\ g_n \end{pmatrix} = \begin{pmatrix} \frac{1}{2}\mathcal{I} - \mathcal{K} & \mathcal{V} \\ \mathcal{H} & \frac{1}{2}\mathcal{I} + \mathcal{T} \end{pmatrix} \begin{pmatrix} g \\ g_n \end{pmatrix},$$

where $\mathcal{I}$ is the identity operator (Grubb 2009). From the first row of this expression, we can derive the Dirichlet-to-Neumann operator $\mathcal{S} : g \mapsto g_n$ as the composition

$$\mathcal{S} = \mathcal{V}^{-1} \left(\frac{1}{2}\mathcal{I} + \mathcal{K} \right). \quad (6)$$

Combining both rows, however, reveals an alternative symmetric expression for the same operator:

$$\mathcal{S} = \mathcal{H} + \left(\frac{1}{2}\mathcal{I} + \mathcal{T} \right) \mathcal{V}^{-1} \left(\frac{1}{2}\mathcal{I} + \mathcal{K} \right). \quad (7)$$

Since $\mathcal{H}$ is self-adjoint and positive semidefinite, $\mathcal{V}$ is positive definite, and $(\mathcal{K}, \mathcal{T})$ is an adjoint pair, this alternative form symbolically shows that $\mathcal{S}$ is self-adjoint and positive semidefinite, with $\langle u, \mathcal{S}u \rangle_{\Gamma} = 0$ if and only if u is a constant function. In our discretization, we work with this second formula because its discretization will become symmetric and positive semidefinite by construction. Note that interior Steklov eigenfunctions are known in closed form by the representation formula

$$u(\mathbf{x}) = \int_{\Gamma} \left[G(\mathbf{x}, \mathbf{y}) g_n(\mathbf{y}) - \frac{\partial G(\mathbf{x}, \mathbf{y})}{\partial n(\mathbf{y})} g(\mathbf{y}) \right] d\Gamma(\mathbf{y})$$

for any $\mathbf{x} \in \Omega$. Here, $u(\cdot)$ is the i th eigenfunction in the interior when we let $g(\cdot)$ and $g_n(\cdot)$ be the i th eigenfunction on the surface and its normal derivative, respectively; later in Equation (15) they are denoted as $\mathbf{u}$ and $\mathbf{t}$. Our algorithm will compute both $\mathbf{u}$ and $\mathbf{t}$, so interior eigenfunctions can be easily evaluated if needed.

4 DISCRETE DIRICHLET-TO-NEUMANN OPERATOR

4.1 Weak-Form Boundary Operators

We discretize all operators discussed above. Each corresponds to a weak-form discretized operator matrix:

$$\mathcal{V} \mapsto \mathbf{V}, \quad \mathcal{K} \mapsto \mathbf{K}, \quad \mathcal{T} \mapsto \mathbf{T}, \quad \mathcal{H} \mapsto \mathbf{H}, \quad \text{and} \quad \mathcal{I} \mapsto \mathbf{M}.$$

Take the single-layer integral as an example. Assume $u(\Gamma)$ is a solution to the single-layer integral equation $\mathcal{V}u(\mathbf{x}) = f(\mathbf{x})$. Then, for any test function $v : \Gamma \rightarrow \mathbb{R}$ we have the weak form $\langle v, \mathcal{V}u \rangle_{\Gamma} = \langle v, f \rangle_{\Gamma}$. Restricting u, v, f to the piecewise-linear subspace $S_h^1(\Gamma)$ (the “hat functions” on a triangulated surface) leads to the finite-dimensional linear system $\mathbf{V}\mathbf{u} = \mathbf{M}\mathbf{f}$, where $\mathbf{V}, \mathbf{M}$ are Galerkin matrices:

$$\mathbf{V} \in \mathbb{R}^{n \times n} : (\mathbf{V})_{ij} = \langle \phi_i, \mathcal{V}\phi_j \rangle_{\Gamma} \quad (8)$$

$$\mathbf{M} \in \mathbb{R}^{n \times n} : (\mathbf{M})_{ij} = \langle \phi_i, \phi_j \rangle_{\Gamma}. \quad (9)$$

In this linear system, $\mathbf{u}, \mathbf{f}$ contain the coefficients of u and f in the piecewise-linear basis, respectively.

Similarly, we have

$$\mathbf{K} \in \mathbb{R}^{n \times n} : (\mathbf{K})_{ij} = \langle \phi_i, \mathcal{K}\phi_j \rangle_{\Gamma} \quad (10)$$

$$\mathbf{T} \in \mathbb{R}^{n \times n} : (\mathbf{T})_{ij} = \langle \phi_i, \mathcal{T}\phi_j \rangle_{\Gamma} \quad (11)$$

$$\mathbf{H} \in \mathbb{R}^{n \times n} : (\mathbf{H})_{ij} = \langle \phi_i, \mathcal{H}\phi_j \rangle_{\Gamma}. \quad (12)$$

It follows directly from properties in Section 3.2 that $\mathbf{V} = \mathbf{V}^T$, $\mathbf{H} = \mathbf{H}^T$, $\mathbf{T} = \mathbf{K}^T$, $\mathbf{V} > 0$, and $\mathbf{H} \geq 0$. We refer interested readers to Steinbach (2007, Sections 2, 3, 4, 5, 6, 7, 10, 12) for detailed discussion of these discretized operators.

Putting definitions together yields

$$\mathbf{V}_{ij} = \sum_{\substack{T_1 \in A(i) \\ T_2 \in A(j)}} \iint_{T_1 \times T_2} \frac{1}{4\pi} \frac{1}{|\mathbf{x} - \mathbf{y}|} \phi_i(\mathbf{x}) \phi_j(\mathbf{y}) d\Gamma_1(\mathbf{x}) d\Gamma_2(\mathbf{y}),$$

$$\mathbf{M}_{ij} = \sum_{T_1 \in A(i)} \int_{T_1} \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\Gamma_1(\mathbf{x}),$$

where $A(i)$ denotes the set of triangles adjacent to vertex i , and ϕ_i denotes the piecewise-linear “hat” basis function centered at vertex i . Note $\mathbf{M}_{ij}$ is the same as the mass matrix used in FEM. To

simplify notation for the remaining operators, define the following generic boundary operator $P[p(\cdot, \cdot)] \in \mathbb{R}^{n \times n}$ w.r.t. kernel $p(\mathbf{x}, \mathbf{y})$:

$$P_{ij}[p(\cdot, \cdot)] = \sum_{\substack{T_1 \in A(i) \\ T_2 \in A(j)}} \iint_{T_1 \times T_2} \frac{1}{4\pi} p(\mathbf{x}, \mathbf{y}) \phi_i(\mathbf{x}) \phi_j(\mathbf{y}) d\Gamma_1(\mathbf{x}) d\Gamma_2(\mathbf{y}). \quad (13)$$

Then, the discretized operators can be expressed as

$$\mathbf{V}_{ij} = P_{ij}[v], \quad \mathbf{K}_{ij} = P_{ij}[k], \quad \mathbf{T}_{ij} = P_{ij}[t], \quad \mathbf{H}_{ij} = P_{ij}[h],$$

where

$$\begin{aligned} v(\mathbf{x}, \mathbf{y}) &:= \frac{1}{|\mathbf{x} - \mathbf{y}|}, \\ k(\mathbf{x}, \mathbf{y}) &:= \frac{(\mathbf{x} - \mathbf{y}) \cdot \mathbf{n}(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|^3}, \quad t(\mathbf{x}, \mathbf{y}) := \frac{(\mathbf{y} - \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})}{|\mathbf{x} - \mathbf{y}|^3} = k(\mathbf{y}, \mathbf{x}), \\ h(\mathbf{x}, \mathbf{y}) &:= -\frac{\mathbf{n}(\mathbf{x}) \cdot \mathbf{n}(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|^3} - \frac{3[(\mathbf{x} - \mathbf{y}) \cdot \mathbf{n}(\mathbf{y})][(\mathbf{y} - \mathbf{x}) \cdot \mathbf{n}(\mathbf{x})]}{|\mathbf{x} - \mathbf{y}|^5}, \end{aligned}$$

and $\mathbf{n}(\mathbf{x}), \mathbf{n}(\mathbf{y})$ denote the normal directions at points $\mathbf{x}, \mathbf{y}$ on the surface, respectively.

The matrix entries are evaluated using the Gaussian quadrature method implemented in Šmigaj et al. (2015), which ensures quadrature points never coincide, even if the triangle T_1 is adjacent to or the same as the triangle T_2 (these can happen, e.g., when evaluating diagonal entries that $i=j$).

4.2 Discretized Dirichlet-to-Neumann Operator

We approximate the DtN operator by substituting the continuous operators with corresponding strong-form operators:

$$\mathcal{V} \rightarrow \mathbf{M}^{-1}\mathbf{V}, \quad \mathcal{K} \rightarrow \mathbf{M}^{-1}\mathbf{K}, \quad \mathcal{T} \rightarrow \mathbf{M}^{-1}\mathbf{T}, \quad \text{and} \quad \mathcal{I} \rightarrow \mathbf{I},$$

which leads to

$$\mathbf{S} \rightsquigarrow \mathbf{M}^{-1}[\mathbf{H} + (0.5\mathbf{M} + \mathbf{T})\mathbf{V}^{-1}(0.5\mathbf{M} + \mathbf{K})].$$

Define the (symmetric) weak form of this operator as

$$\mathbf{S} := \mathbf{H} + (0.5\mathbf{M} + \mathbf{T})\mathbf{V}^{-1}(0.5\mathbf{M} + \mathbf{K}). \quad (14)$$

5 MATRIX-FREE FORMULATION

Many applications require solving the linear system $\mathbf{S}\mathbf{u} = \mathbf{M}\mathbf{f}$ or the generalized eigenvalue problem $\mathbf{S}\mathbf{x} = \lambda\mathbf{M}\mathbf{x}$. A naïve approach is to assemble all the operators as dense matrices via Equation (14). Without acceleration, however, assembling these dense matrices would take $O(n^3)$ time, and matrix-vector products would require $O(n^2)$ operations. Furthermore, Equation (14) involves inverting the dense matrix $\mathbf{V}$, which is expensive. Instead, we apply a reformulation that avoids explicit matrix inversion.

5.1 Expansion

Define an auxiliary variable $\mathbf{t}$ as $\mathbf{t} := \mathbf{V}^{-1}(0.5\mathbf{M} + \mathbf{K})\mathbf{u}$. Notice that $\mathbf{t}$ represents the Neumann data, i.e., the normal derivative at the boundary. The linear system $\mathbf{S}\mathbf{u} = \mathbf{M}\mathbf{f}$ then can be expanded as a saddle point system

$$\begin{bmatrix} \mathbf{V} & -\mathbf{Q} \\ \mathbf{Q}^\top & \mathbf{H} \end{bmatrix} \begin{bmatrix} \mathbf{t} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{M}\mathbf{f} \end{bmatrix},$$

where $\mathbf{Q} := 0.5\mathbf{M} + \mathbf{K}$.

The same technique also applies to reformulating the eigenvalue problem $\mathbf{S}\mathbf{u} = \lambda\mathbf{M}\mathbf{u}$ as the system

$$\begin{bmatrix} \mathbf{V} & -\mathbf{Q} \\ \mathbf{Q}^\top & \mathbf{H} \end{bmatrix} \begin{bmatrix} \mathbf{t} \\ \mathbf{u} \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{0} & \\ & \mathbf{M} \end{bmatrix} \begin{bmatrix} \mathbf{t} \\ \mathbf{u} \end{bmatrix}. \quad (15)$$

The left-hand sides of these expressions do not contain matrix inverses, allowing us to use iterative linear system/eigenvalue solvers that only require matrix-vector products. See Section 5.4 for details of our solver and Section 5.6 for how we use the hierarchical techniques to *apply* $\mathbf{V}, \mathbf{Q}$ and other matrices without storing their elements.

5.2 Symmetrization

Numerical solution to saddle-point systems like the ones in the previous subsection is a basic task in numerical PDE; see, e.g., Benzi et al. (2005). To use the conjugate gradient (CG) algorithm and other techniques requiring positive definiteness, we apply a Bramble–Pasciak transformation (1988):

$$\begin{bmatrix} \alpha\mathbf{V}\mathbf{P}_V^{-1} - \mathbf{I} & \mathbf{0} \\ -\alpha\mathbf{Q}^\top\mathbf{P}_V^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{V} & -\mathbf{Q} \\ \mathbf{Q}^\top & \mathbf{H} \end{bmatrix} \begin{bmatrix} \mathbf{t} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{M}\mathbf{f} \end{bmatrix},$$

where $\mathbf{P}_V$ is a symmetric preconditioner for the single-layer potential (our choice is defined in Section 5.3). The constant $\alpha = 1/\gamma$ is chosen to ensure the positive definiteness of the system matrix as $\gamma = 0.9\sigma_{\min}(\mathbf{P}_V^{-1}\mathbf{V}) \leq 0.9 < 1$, where $\sigma_{\min}(\cdot)$ denotes the smallest singular value. The matrix of this system is symmetric and positive semidefinite:

$$\mathbf{A} = \begin{bmatrix} \alpha\mathbf{V}\mathbf{P}_V^{-1}\mathbf{V} - \mathbf{V} & (\mathbf{I} - \alpha\mathbf{V}\mathbf{P}_V^{-1})\mathbf{Q} \\ \mathbf{Q}^\top(\mathbf{I} - \alpha\mathbf{P}_V^{-1}\mathbf{V}) & \mathbf{H} + \alpha\mathbf{Q}^\top\mathbf{P}_V^{-1}\mathbf{Q} \end{bmatrix} = \mathbf{A}^\top. \quad (16)$$

Under the Bramble–Pasciak transformation, the eigenvalue problem can be written as

$$\mathbf{A} \begin{bmatrix} \mathbf{t} \\ \mathbf{u} \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{0} & \\ & \mathbf{M} \end{bmatrix} \begin{bmatrix} \mathbf{t} \\ \mathbf{u} \end{bmatrix}, \quad \text{or more simply } \mathbf{A}\mathbf{x} = \lambda\tilde{\mathbf{M}}\mathbf{x}.$$

Although $\mathbf{A}$ and $\mathbf{M}$ are positive semidefinite, $\alpha\mathbf{A} + \beta\tilde{\mathbf{M}}$ is strictly positive definite, for any $\alpha, \beta > 0$. So it can be solved by a shifted generalized eigenvalue method.

5.3 Preconditioning

Steinbach and Wendland (1998) prove that a modified hypersingular operator is a good preconditioner for the single-layer potential and that the single-layer potential is a good preconditioner for the Dirichlet-to-Neumann operator (Equation (14)). In particular, they show that the spectral condition numbers of these two preconditioners are $O(1)$, in the sense that the condition number remains the same when the domain is upsampled. Their choices of preconditioners for $\mathbf{S}$ and $\mathbf{V}$ are

$$\begin{aligned} \mathbf{P}_V^{-1} &:= 4\mathbf{M}^{-1} \left(\mathbf{H} + \frac{\beta}{4}\mathbf{1}\mathbf{1}^\top \right) \mathbf{M}^{-1}, \\ \mathbf{P}_S^{-1} &:= \mathbf{M}^{-1}\mathbf{V}\mathbf{M}^{-1}, \end{aligned} \quad (17)$$

where $\beta = 1/(\mathbf{1}^\top\mathbf{M}\mathbf{1})^{3/2}$ is chosen to ensure scale invariance. Since $\mathbf{M}$ is sparse and diagonally dominant, Jacobi iteration could compute $\mathbf{M}^{-1}\mathbf{x}$ in $O(n)$ time. In our implementation, however, we use a diagonal lumped mass matrix that can be inverted in closed form.

For saddle-point system $\mathbf{Ax} = \lambda \tilde{\mathbf{M}}\mathbf{x}$, Bramble and Pasciak (1988) recommend the following preconditioner to Equation (16), which is spectrally equivalent to $\mathbf{A}$:

$$\mathbf{P}_A^{(1)} := \begin{bmatrix} \mathbf{V} - \gamma \mathbf{P}_V & \\ & \mathbf{P}_S \end{bmatrix}.$$

The upper left block of this matrix cannot be inverted efficiently, so it has to be combined with a modified Bramble-Pasciak CG solver. See Stoll and Wathen (2007) for a survey on other preconditioning options for the saddle-point system. In practice, we use the following simpler alternative that is easier to invert:

$$\mathbf{P}_A^{(2)} := \begin{bmatrix} \frac{1}{\alpha-1} \mathbf{P}_V & \\ & \mathbf{P}_S \end{bmatrix}.$$

$\mathbf{P}_A^{(2)}$ is a valid preconditioner as well, that is, $\mathbf{P}_A^{(2)}$ is spectrally equivalent to $\mathbf{A}$, since $\mathbf{P}_A^{(2)}$ is spectrally equivalent to $\mathbf{P}_A^{(1)}$ and $\mathbf{P}_A^{(1)}$ is spectrally equivalent to $\mathbf{A}$.¹ We find that this choice suffices for our solver; investigating the most efficient preconditioning strategy is a potential topic for future work.

5.4 Iterative Solvers

We use the preconditioned conjugate gradient (PCG) algorithm to solve linear systems of equations involving our positive definite operators. As noted above, PCG requires matrix-vector products rather than storing our operators explicitly, as would be required by Gaussian elimination.

We employ the locally optimal block preconditioned conjugate gradient (LOBPCG) solver (Knyazev 2001) to compute the spectrum efficiently. The convergence speed of this method depends on the condition number $\text{cond}(\mathbf{P}_A^{-1}\mathbf{A})$. LOBPCG allows for an initial guess of the eigenvectors. We apply a multiscale approach, computing “progressive spectra” from progressively simplified meshes obtained using (Hoppe 1996). We first compute eigenvalues and eigenfunctions on a low-resolution mesh and then upsample the low-resolution eigenfunctions as initialization for high-resolution meshes. Typically we use two to three levels of progressively simplified meshes, where at each level the mesh is simplified by a factor of $1/4$; we use the nearest neighborhood rule to upsample eigenfunctions. Note that Vaxman et al. (2010) use a related approach to evaluate the heat kernel.

5.5 Generalization

We have introduced the Dirichlet-to-Neumann operator for compact manifolds with closed boundaries. As long as we are given a normal vector field at every point, however, the integral-based definitions of inner products $\langle v, \mathcal{V}u \rangle_\Gamma$, $\langle v, \mathcal{K}u \rangle_\Gamma$, $\langle v, \mathcal{H}u \rangle_\Gamma$, and $\langle v, \mathcal{T}u \rangle_\Gamma$ as well as the resulting operator $\mathcal{S}$ remain valid even if Γ does not bound some region in $\mathbb{R}^3$. Given this observation, to

¹More precisely, by the definition of preconditioners we have:

$$\sigma_{\min}(\mathbf{P}_A^{(1)-1}\mathbf{A})\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{P}_A^{(1)}} \leq \langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{A}} \leq \sigma_{\max}(\mathbf{P}_A^{(1)-1}\mathbf{A})\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{P}_A^{(1)}}$$

$$\sigma_{\min}(\mathbf{P}_V^{-1}\mathbf{V})\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{P}_V} \leq \langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{V}} \leq \sigma_{\max}(\mathbf{P}_V^{-1}\mathbf{V})\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{P}_V}$$

where $\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{B}} := \mathbf{x}^\top \mathbf{B} \mathbf{x}$, $\forall \mathbf{B}$. These lead to $\sigma_1\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{P}_A^{(2)}} \leq \langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{A}} \leq \sigma_2\langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{P}_A^{(2)}}$,

where $\sigma_1 := \min\{1, \frac{1}{\gamma}(1-\gamma)\}$ and $\sigma_2 := \max\{1, (1-\gamma)[\frac{\sigma_{\max}(\mathbf{P}_V^{-1}\mathbf{V})}{\gamma} - 1]\}$.

Fig. 2. Generalized Steklov eigenfunctions of a hemisphere.

Fig. 3. Generalized Steklov eigenvalues of a sphere cut with a moving plane.

process oriented surfaces with boundary, we simply use the same operator defined via boundary integral in Equation (7); for related work see Bruno and Lintner (2013); Lintner and Bruno (2015) and references therein. Justifying this generalization for open surfaces theoretically is a challenging question for future work.

We experimentally verify that this natural generalization demonstrates reasonable and robust behavior. Figure 2 illustrates our “generalized Steklov” eigenfunctions computed on a hemisphere from these integral operators. Figure 3 demonstrates that the resulting spectrum blends smoothly as the shape changes from a sphere to a hemisphere by moving the cutting plane linearly.

5.6 Implementation Details

Our rephrasing of the DtN operator in terms of boundary integrals is an example of the BEM from numerical PDE. Hence, our implementation uses the BEM++ library (Šmigaj et al. 2015)

Fig. 4. The Steklov spectra of the donuts. These measurements remain stable as the donut undergoes topological change and even becomes disconnected. In the last row, the two open holes on the cut donut are closed with disks, in which case the Steklov spectrum considers the donut as two disconnected pieces.

Fig. 5. The Steklov spectrum is robust when removing random triangles progressively from the mesh, tested on the sphere and armadillo meshes.

Fig. 6. The Steklov spectrum is robust to downsampling as well as unbalanced sampling. Dragons represented with 14K (blue), 54K (red), 220K (yellow), and 117K (green) triangles have very similar Steklov eigenvalues. The dragon with 117K triangles is obtained by downsampling only the left side of the 220K dragon.

to evaluate the matrices above, whose elements are obtained using numerical quadrature. BEM++ uses hierarchical matrices (“H-matrices”) to compute matrix-vector products efficiently, approximating the integrals with hierarchical expressions that group together small contributions from far-away vertices (Börm et al. 2003). $\sigma_{\min}(\mathbf{P}_V^{-1}\mathbf{V})$ can be approximated by solving the minimum eigenvalue of the squared system, without the need for a preconditioner.

6 VALIDATION

6.1 Robustness to Topological Change

In addition to the hemisphere test in Figure 2, Figure 4 gives an example of robustness of our operator to topological change. Here, we compare the Steklov eigenvalues and eigenfunctions of a torus with zero, one, and two rings cut out. When triangles on the cut are removed, these topological changes—which drastically affect the Laplace–Beltrami operator—have little effect on the Steklov spectrum. When we close the holes with flat disc patches, however, the Steklov spectrum detects a topological change, since the donut is clearly divided into two volumetric pieces.

Figure 5 shows an example of robustness to “topological noise.” Here, we remove triangles from the mesh with uniform probability and show the resulting Steklov spectrum. Once again, our operator is remarkably stable to these changes, yielding a stable spectrum even when 50% of the triangles are removed.

6.2 Robustness to Surface Sampling

Figure 6 shows an example of robustness to downsampling and unbalanced sampling. Here we simplify the dense dragon mesh and observe that the resulting Steklov spectrum does not change significantly. Although geometric details at the belly and scales of the dragon are lost during downsampling, the eigenvalues are almost unaffected.

Since the low-order Steklov eigenvalues are robust to downsampling and in most applications we are interested in only the first ~ 50 eigenfunctions, the experiment suggests it is not necessary to use a highly detailed mesh. We find that a mesh with 7K to 15K triangles is usually sufficient to accurately approximate the top 50 eigenfunctions; in this regime, computation finishes in 4 to 10 minutes.

6.3 Robustness to Vertex Noise

Figure 7 illustrates stability of our BEM discretization when noise is added to the boundary surface. Although they are defined in Section 3.1 via a second-order differential equation, we observed that the lower Steklov eigenfunctions are particularly robust to noise, similarly to those of the Laplace–Beltrami operator. Both Laplacian and Steklov spectra scale down globally, since adding noise increases the edge lengths and thus the surface area of the mesh.

When noise is applied, the mesh is no longer watertight: there are intersecting triangles, flipped triangles, and so on caused by random vertex perturbations. For these challenging meshes, the

Fig. 7. Robustness test by adding noise to the vertex positions. Blue, red, and yellow curves correspond to the original and noisy meshes, respectively. After noise is applied, the mesh may not be exactly watertight due to self-intersections as well as many flipped triangles, and hence the smallest “Steklov eigenvalue” could be slightly positive. A decrease of eigenvalues after adding noise is expected since surface area increases. Both eigenvalues are normalized by its largest eigenvalue.

BEM++ library—not designed to handle these cases—occasionally fails because the quadrature points can coincide.

To avoid this issue, we first apply a simple repair procedure to the mesh: cut intersecting triangles into multiple nonintersecting ones, merge vertices that are closer than δ to each other, and remove duplicated triangles (if there are any). $\delta = 10^{-2}\bar{L}$ is chosen, where $\bar{L}$ is the average edge length in the original input mesh. This repair procedure leads to almost no visual difference to the mesh. In addition, for these extremely noisy tests, we find that using a regularized kernel $1/(r + \epsilon)$ instead of the original kernel $1/r$ helps to improve the solver’s robustness, where $\epsilon = 10^{-4}$ is chosen, assuming the mesh has been normalized to the unit scale. This simple repair procedure and regularization are only necessary for meshes with the largest amount of noise in Figure 7 (rightmost column), to avoid overlapping quadrature points. Generalizing the DtN operator to surfaces with self-intersections is a challenging problem and it may deserve careful numerical treatments. Our current mesh-fixing procedure is a practical workaround, leaving the theoretical question for future research.

Our Steklov solver remains robust even with a significant amount of noise. Large noise may cause partial occlusions (like the patched disks in Figure 4), which the Steklov spectrum considers as partial topological change. In this case, the smallest computed eigenvalue could be slightly positive, depending on the size of the occlusion.

As a point of contrast, no volumetric (tet-based) method can be applied before fixing the mesh to be watertight; this more complex

repair procedure is challenging and may lose geometric information. In a sense, it is not surprising that our BEM-based approach is robust to noise and triangle soup; the state-of-the-art mesh robust repair technique (Jacobson et al. 2013) in essence is also a boundary element method.

6.4 Stability Test for Volume Isometries

The Steklov spectrum is invariant to volumetric isometry, and our experiments demonstrate that it remains stable when the volume enclosed by the outer surface deforms *near* isometrically. In the first row of Figure 8, when the rectangular prisms are deformed using a near-volume-isometric bending map, the Steklov and Laplacian eigenvalues and eigenfunctions remain almost unchanged. In the second row, however, the thin box is deformed subject to a near-surface isometry but nonvolume isometry, by shearing the interior of the box. Only the Steklov spectrum captures the volumetric change as the shear increases. This example illustrates how the Steklov spectrum discriminates nuanced nonisometric volume changes, while the Laplacian spectrum fails to capture the difference.

6.5 Conditioning

Convergence of our iterative linear and eigenvalue solvers depends on the conditioning of the matrices involved. Since our operators and preconditioners are constructed using integrals that make sense in the continuum, we expect that conditioning is $O(1)$ in mesh size; i.e., it does not depend significantly on the number

(a) Both Steklov and Laplacian spectra are robust to (near)-volume-isometry.

(b) Steklov spectrum captures shearing, to which Laplacian is insensitive.

Fig. 8. Stability tests for volume isometries and nonvolume isometries.

# vertices	# faces	$\text{cond}(\mathbf{P}_V^{-1}\mathbf{V})$
258	512	3.91
1026	2048	3.96
4098	8192	4.00
16386	32768	4.00
65538	131072	3.94
262146	524288	3.95

Fig. 9. Conditioning on different meshings of the unit sphere.

Fig. 10. Relative error measured in between spherical harmonics and the eigenspaces of our Steklov operator. Relative error decreases as the mesh resolution increases. The error is measured for a fixed number of 25 iterations, computed from random initialization.

of vertices or triangles. Figure 9 verifies this relationship for the single-layer potential operator on different meshings of the unit sphere; we observe similar behavior on other meshes.

6.6 Convergence to Analytical Eigenvalues

The analytical Steklov eigenvalues and eigenfunctions are known for the sphere. We compare the numerical solution of our method on the sphere for verification purposes. Figure 10 verifies that our discretized operators are faithful to these ground-truth quantities.

(d) Plot of eigenvalues computed using different methods.

(e) Plot of absolute errors of eigenvalues (numerical solutions minus analytical solutions). For the unit sphere, the analytical solution of the Steklov eigenfunctions and eigenvalues are spherical harmonics and corresponding eigenvalues.

Fig. 11. Steklov spectrum of unit sphere, computed using BEM and FEM with different mesh resolutions. For FEM, we compute the Steklov spectrum using either fine-grained tets or coarse-grained tets, and using either a full mass matrix or a lumped mass matrix. For BEM, we experiment with full and lumped mass matrices.

6.7 Timing

For the cube with 6,146 vertices and 12,288 triangles, BEM++ takes 60 seconds to assemble all boundary element operators, one LOBPCG iteration takes 24 seconds on average, and 20 iterations are usually sufficient for random initialization. Usually fewer than 10 iterations are needed for initialization from low-resolution eigenfunctions. We observed that all the timings roughly scale linearly with size of the mesh. Our implementation is in PYTHON, and results are collected on an Alienware laptop with an Intel i7 4800MQ CPU and Linux operating system. As many graphics applications do not require a very high precision, we believe that there is much room for future work to improve the computational efficiency (see Section 8).

Fig. 12. Selected Steklov eigenfunctions on the dragon model, with comparison to Laplacian eigenfunctions. Although the Steklov eigenfunctions are computed using the boundary geometry, these eigenfunctions are aware of the geometry of the volume enclosed by the dragon surface, in particular the S-shaped bend of its body.

Fig. 13. The Steklov eigenfunctions corresponding to the smallest 11 nonzero eigenvalues, compared to the surface Laplacian eigenfunctions; this model contains 50K triangles and 25K vertices.

6.8 Comparison with FEM

An alternative approach for extrinsic spectral geometry uses the volumetric Laplacian. Spectral computations with this operator do not have a boundary integral formulation, and hence we would need a volumetric mesh of the interior of the domain. This limits applicability, since tet meshing typically requires a watertight surface as input, and the results depend on the algorithm used for generating a tet mesh. Furthermore, map-based applications such as computation of shape differences (Section 7.5) would require correspondence of not just the outer surface but the interior as well, which may not be available. As we will show in Section 7.2, we do not observe a case in which an explicit mesh of the interior leads to significantly better results for geometry processing tasks.

These high-level concerns aside, one could reasonably discretize the DtN operator using the FEM on a tet mesh. In particular, the DtN operator can be derived as the Schur complement of the positive semidefinite volumetric Laplacian matrix $\mathbf{L}^{(\text{vol})}$, i.e., the FEM discretization of the interior Laplace equation with the Neumann boundary conditions. With piecewise linear bases, $\mathbf{L}^{(\text{vol})}$ is assembled by the familiar cotangent weights (of the dihedral angle opposite to an edge); see, e.g., Jacobson (2013, Section 2.1) for an explicit formula. On a tetrahedral mesh with volumetric FEM-based Laplacian $\mathbf{L}^{(\text{vol})}$ and mass matrix $\mathbf{M}^{(\text{vol})}$, divide this matrix into blocks corresponding to boundary vertices b and interior vertices i :

$$\mathbf{L}^{(\text{vol})} = \begin{bmatrix} \mathbf{L}_{bb} & \mathbf{L}_{bi} \\ \mathbf{L}_{ib} & \mathbf{L}_{ii} \end{bmatrix}.$$

It approximates Δ in the interior and ∇_n at the boundary

$$\begin{bmatrix} \nabla_n \mathbf{u}_b \\ \Delta \mathbf{u}_i \end{bmatrix} = -\mathbf{M}^{(\text{vol})}{}^{-1} \mathbf{L}^{(\text{vol})} \mathbf{u}.$$

Setting the second row $\Delta \mathbf{u}_i = 0$ to eliminate the interior vertices as $\mathbf{u}_i = -\mathbf{L}_{ii}^{-1} \mathbf{L}_{ib} \mathbf{u}_b$, the Dirichlet energy takes the form

$$E = \mathbf{u}^T \mathbf{L}^{(\text{vol})} \mathbf{u} = \mathbf{u}_b^T \mathbf{S}^{(\text{FEM})} \mathbf{u}_b,$$

giving rise to a FEM-based Dirichlet-to-Neumann operator

$$\mathbf{S}^{(\text{FEM})} = \mathbf{L}_{bb} - \mathbf{L}_{bi} \mathbf{L}_{ii}^{-1} \mathbf{L}_{ib} \geq 0.$$

Figure 11 illustrates a drawback of using the tet-based DtN operator. Here, we compute Steklov eigenvalues on two different tetrahedral meshes of the same unit sphere using first-order FEM. This figure shows that Steklov eigenvalues depend on the choice of tetrahedral meshing of the boundary. In particular, faster sparse computations on coarse FEM meshes come at the cost of inaccurate Steklov approximation. Contrastingly, BEM faithfully approximates Steklov eigenvalues for the sphere beyond the range tested in this experiment.

We further observe that BEM with the correct full mass matrix (our method) yields the closest spectrum to ground truth by a significant margin; even BEM with a lumped diagonal mass matrix does not yield accurate results. Hence, our implementation always uses the full mass matrix in the generalized eigenvalue problem. We do use lumped mass matrices to construct preconditioners for the iterative solvers, since they can be inverted quickly; this affects only the number of iterations for the solver rather than the quality of the final output.

7 EXPERIMENTS AND APPLICATIONS

We highlight two properties of the DtN operator $\mathcal{S}$ distinguishing it from the Laplace–Beltrami operator, motivated in the theoretical discussion in Section 3:

Fig. 14. Reconstruction with increasing number of bases; color encodes relative error in position.

Fig. 15. Visualization of Laplacian and Steklov heat kernel function $h_t(x)$, respectively, in log scale. Note that the minima correspond to most negative Gaussian and mean curvatures, respectively.

- S captures the volumetric structure of the interior.
- S incorporates the mean curvature of the outer surface.

Accordingly, the algorithms we derive based on the Steklov spectrum enjoy these properties.

Many existing frameworks for intrinsic geometry processing can be extended to take extrinsic geometry into consideration by substituting the Laplacian operator with the DtN operator. There are many ways to justify this simple substitution:

- Many intrinsic geometry tasks involve the surface-based Laplace–Beltrami operator, whose weak form corresponds to a Dirichlet energy on the surface:

$$\langle \nabla_{\Gamma} u, \nabla_{\Gamma} u \rangle_{\Gamma} = \mathbf{u}^T \mathbf{L} \mathbf{u}.$$

The weak form of the DtN operator provides the closely related *volumetric* Dirichlet energy of the harmonic extension $\mathcal{E}$ of a surface-based function:

$$\langle \nabla \mathcal{E} u, \nabla \mathcal{E} u \rangle_{\Omega} = \mathbf{u}^T \mathbf{S} \mathbf{u}.$$

- Both operators' associated eigenproblems involve Rayleigh quotients with similar numerators and identical denominators (Girouard and Polterovich 2017).
- For n disconnected/closed pieces, both operators have n -dimensional null spaces.
- Laplace–Beltrami and DtN operators both can be understood as subtracting the value of a function at a point from its average over a neighborhood. The DtN operator first interpolates the function harmonically to the volume, while Laplace–Beltrami restricts it to the surface.

Fig. 16. Steklov-based scaled heat kernel signature and wave kernel signature on the cubes with inward (red) and outward (blue) bumps. In contrast, the two cubes have identical Laplacian-based signatures. Feature points are chosen as centers of the bumps.

Fig. 17. Color encodes the similarity between the signatures across the shape with signature at the marker (blue dot), which sits in the center of shoe bottom. Darker color corresponds to more similarity. *Left*: Steklov WKS is aware of regions with a large mean curvature (wrinkles on the cloth) and avoids matching them with the marker, which is on a plateau with zero mean curvature. *Right*: Steklov WKS is a more restrictive signature, while Laplacian WKS tends to confuse left and right feet.

The close analogies above hint that shifting from intrinsic geometry to extrinsic geometry can be accomplished by substituting the cotangent Laplacian operator with a discretized DtN operator.

We demonstrate this substitution in a few contexts. In particular, we experiment and study properties of Steklov-based kernel

Fig. 18. Spectral and diffusion distances computed using the Laplacian, Steklov, and volumetric Laplacian spectra; colors range from blue (zero distance) to red (large distance). The first and second rows show the top and bottom of the surface, respectively. Source point is placed at the center of the top surface (visible in the first row). The Steklov (boundary-based) and volumetric (tet-mesh-based) distances are small between the center of the top and bottom of the pancake shape, while they are far using the intrinsic Laplacian distances.

Fig. 19. Spectral and diffusion distances computed using the Laplacian, Steklov, and volumetric Laplacian spectra, illustrated using the same color scheme as Figure 18. In contrast to the example in Figure 18, while the two hands of the mouse are close in ambient space $\mathbb{R}^3$, the Steklov/volumetric distances between the hands are large because they are constrained to use paths through the *interior* of the volume.

signatures, distances, and shape differences, which are fundamental components of many high-level geometry processing and shape analysis algorithms.

7.1 Steklov Spectrum

Figures 12 and 13 illustrate Steklov eigenfunctions on more complex models; surface-based Laplace–Beltrami eigenfunctions are shown for comparison. Both models have more vertices than can be feasibly handled using dense matrix computations, highlighting the necessity for iterative solvers introduced in Section 5.4.

A few qualitative differences between the intrinsic and extrinsic eigenfunctions are worth pointing out. On the gargoyle model, the Steklov eigenfunctions exhibit more localization on the wings; the wings are distinctive volumetric features, whereas the outer surface blends the wings into the body/base. On the dragon model, eigenfunctions of the two operators look completely different: Steklov eigenfunctions are localized in individual folds of the dragon, while the Laplace–Beltrami eigenfunctions extend along the entire surface. This effect is reflective of the fact that intrinsically the dragon model is a long tube in which the 180° fold of the body between the front and hind legs is insignificant;

Fig. 20. Spectral and diffusion distances computed using the Laplacian, Steklov, and volumetric Laplacian spectra. For the bunny, note that the lack of surface monotonicity is the correct and expected behavior for volume distances, including the Steklov and volumetric Laplacian distances.

Fig. 21. The Steklov eigenfunctions corresponding to the smallest five nonzero Steklov eigenvalues, with comparison to surface Laplacian eigenfunctions. Note that the Steklov eigenfunctions align with the local extrema of mean curvature, suggesting its use as a segmentation cue.

volumetrically, however, the up/down bending of the dragon is a prominent geometric feature.

Figure 14 shows an experiment similar to the ones in Vallet and Lévy (2008), in which the eigenfunctions of the DtN operator are used to compress and subsequently reconstruct the xyz coordinate functions of a triangle mesh. Interestingly, fewer Steklov eigenfunctions are needed to capture key geometric features relative to the Laplace–Beltrami eigenfunctions. That is, low frequencies in DtN space appear to better capture geometric variation.

7.2 Kernel-Based Descriptors

The heat equation associated to the DtN operator is given by

$$\frac{du}{dt} = -Su, \quad (18)$$

whose solution can be written

$$u(x, t) = \int_{\Gamma} k_t(x, y) f(y) dy. \quad (19)$$

Here, $k_t(x, y)$ is the heat kernel

$$k_t(x, y) = \sum_{i=0}^{\infty} e^{-\lambda_i t} \phi_i(x) \phi_j(y), \quad (20)$$

where λ_i and ϕ_i are the i th eigenvalue and eigenfunction of $\mathcal{S}$, respectively.

If we replace DtN with the Laplace–Beltrami operator, the *heat kernel signature* (HKS) of a point $x \in \Gamma$ is defined as $h_t(x) = k_t(x, x)$ (Sun et al. 2009), the diagonal of the heat kernel. Considered as a function of t , $h_t(x)$ provides a multiscale characterization of intrinsic geometry near x ; it is a basic descriptor used in many shape matching and correspondence algorithms. The wave kernel signature (WKS) (Aubry et al. 2011), also defined using the

Fig. 22. Results of segmentation by the Steklov and Laplacian embeddings.

Fig. 23. Point-wise distortion, as well as principal distortion (top singular vector of shape difference operator), measured by the Steklov-based and Laplacian-based shape difference operators for a bending sheet of paper.

Fig. 24. Point-wise distortion measured by the Steklov-based and Laplacian-based shape difference operators. Our Steklov-based operator distinguishes bending direction of the cigar.

Laplace–Beltrami operator, outperforms the heat kernel signature for shape matching tasks using a related eigenvalue formula.

The heat and wave kernel signatures can be naturally generalized to the DtN operator by replacing the Laplace–Beltrami spectrum/eigenfunctions with Steklov spectrum/eigenfunctions. A similar construction is considered by Raviv et al. (2010), who construct a volumetric HKS via a coarse discretization of the Laplacian in the interior of a volume bounded by a watertight surface. Without this interior meshing, our DtN-based HKS and WKS capture similar extrinsic shape properties.

Figure 15 compares the Laplace–Beltrami, DtN, and volumetric HKS functions on a triangulated surface; the volumetric HKS is approximated using a tetrahedral mesh of the interior, as explained in Section 6.8. Following Sun et al. (2009), we consider the time interval $[\frac{4 \ln 10}{\lambda_{300}}, \frac{4 \ln 10}{\lambda_2}]$ in all of our examples, and Figure 15 shows

typical patterns of the signature using either a small or large time from the range. The DtN and volumetric kernel signatures capture the mean curvature in the palm of the hand, ignored by the intrinsic HKS. Similarly, Figure 16 illustrates the point signatures at the apexes of the bumps in Figure 1. Since the two cubes are isometric, Laplace–Beltrami-based signature does not discriminate the two points, while the DtN-based signature distinguishes the two models.

Figure 17 illustrates how DtN-derived descriptors can be more discriminative than their Laplace–Beltrami counterparts. Here, we mark a point on the foot of a human model and show the distance between its descriptor and those of other surface points. Fewer points have descriptors close to the descriptor of the foot, showing that the DtN WKS embedding is less ambiguous. In the second example, DtN WKS is better aware of the left-right symmetry breaking, and tends to discriminate left and right feet. This property may or may not be desirable depending on application, but it shows that DtN operators are *pose aware* and can be used to navigate databases of near-isometric models like articulated humans.

7.3 Spectral Distance

Several spectral distances are defined between points on a surface in terms of the spectrum of the Laplace–Beltrami operator. These distances enjoy certain stability and smoothness properties not satisfied by the geodesic distance and are computable using linear algebra machinery. These distances naturally generalize to the Steklov spectrum, providing spectral volume-aware distances that require computation only on the boundary.

Working in analogy to previous work (Coifman and Lafon 2006; Lipman et al. 2010), we define the Steklov diffusion distance and the bi-Steklov distance as follows:

$$d_D(x, y)^2 = \sum_{i=1}^{\infty} e^{-2t\lambda_i} (\phi_i(x) - \phi_i(y))^2$$

$$d_B(x, y)^2 = \sum_{i=1}^{\infty} \frac{1}{\lambda_i^2} (\phi_i(x) - \phi_i(y))^2.$$

These distances are aware of relationships between points that reach across the volume.

Figures 18, 19, and 20 show several examples of these distances on meshed surfaces. A few key examples point out the special properties of our distance compared to its intrinsic and/or more naïve extrinsic counterparts:

Fig. 25. PCA on collections of shape differences. The Steklov-based shape difference reveals extrinsic deformation that conformal (Laplacian-based) shape difference fails to capture. For the bumpy cube array, the Steklov-based shape difference operator successfully classifies all shapes into four clusters: inward and inward bumps, inward and outward bumps, outward and inward bumps, and outward and outward bumps.

- **FLAT DISK (FIGURE 18):** The DtN distance is small between points at the center of the top and the bottom of the disc, which are close if you cut through the interior of the disc and far along the surface.
- **MOUSE (FIGURE 19):** This example shows the opposite effect. The two hands of the mouse are close in Euclidean dis-

tance, but in Steklov distance they are far. This is because our distances are *interior aware*; unlike the disk, the shortest path between the two hands through the interior is large.

Note that computation of shortest-path distances restricted to the interior of a nonconvex triangulated surface is NP-hard (Canny

(a) 2D PCA plot of “ShapeDNAs” computed using Steklov and Laplacian eigenvalues. Shapes from the same category tend to cluster together.

(b) For both Steklov and Laplacian eigenvalues, shapes from the same category have similar eigenvalues; the Steklov spectra cluster correctly but also distinguish between shapes within the same class.

Fig. 26. The Shape “DNA” comparison

and Reif 1987). Though the Steklov family of distances is computed without interior discretization, it behaves similarly to distances computed using the volumetric Laplacian.

7.4 Volume-Aware Segmentation

Equation (2) suggests that Steklov eigenfunctions encode point-wise mean curvature, providing geometric clues for surface segmentation useful in descriptor-based algorithms (Chen et al. 2009); see Figure 21 for an example. To demonstrate, in Figure 22 we apply a naïve strategy for segmentation. For the input mesh, we first compute the Steklov spectrum embedding as

$$\left(\frac{\phi_1(\mathbf{x})}{\sqrt{\lambda_1}}, \frac{\phi_2(\mathbf{x})}{\sqrt{\lambda_2}}, \dots, \frac{\phi_k(\mathbf{x})}{\sqrt{\lambda_k}} \right),$$

where $\mathbf{x} \in \Omega$. Then, we apply the k -means clustering to this embedding with a user-specified number of clusters; to avoid local optima, we restart k -means 10 times with random initialization and keep the clustering with the lowest objective value. This simple strategy yields consistent segmentations: while the Laplacian embedding tends to segment the surface into flat patches, the Steklov embedding tends to segment the shape into volumetric parts.

7.5 Shape Differences and Variability

Rustamov et al. (2013) introduce a notion of shape differences based on an area-based inner product $b^a(u, v) = \langle u, v \rangle_\Gamma$ and a conformal inner product $b^c(u, v) = \langle \nabla_\Gamma u, \nabla_\Gamma v \rangle_\Gamma$ between surface functions u and v .

Similarly, we introduce a volume-aware inner product $b^s(u, v) := \langle \nabla \mathcal{E} u, \nabla \mathcal{E} v \rangle_\Omega$. Superficially this again could be viewed as a substitution of the Laplacian L used for intrinsic shape differences with the DtN operator S . After discretization, this new inner product induces a shape difference operator $D = S_{\mathcal{M}}^{-1} F^T S_{\mathcal{N}} F$ given a functional map taking functions on the source $\mathcal{M}$ to functions on the target $\mathcal{N}$. This shape difference operator modifies the function f such that, under the map F the Dirichlet energy $h(f, f)$ of f ’s harmonic extension will be best preserved in the sense $|b_{\mathcal{M}}^s(f, f) - b_{\mathcal{N}}^s(Ff, Ff)|$ is minimized. The action of D reflects how different the two geometric domains are. Inherited from the DtN operator, D defines a more rigid notion of shape differences that captures extrinsic shape changes not captured by intrinsic differences.

For example, Figure 23 shows a shape difference between a rectangular sheet and a bent one; Figure 24 shows a shape difference between a cigar and a bent cigar. The point-wise distortion is measured by the metric

$$(\text{distortion})_i = \frac{\mathbf{e}_i^T \mathbf{M} \mathbf{e}_i}{\mathbf{e}_i^T \mathbf{M} \mathbf{D} \mathbf{e}_i},$$

as in Rustamov et al. (2013). The Steklov-based shape difference is far less noisy and distinguishes the bending direction.

Figure 25 illustrates the effect of principal component analysis (PCA) on a collection of shape differences measured through correspondence to a base mesh; see Rustamov et al. (2013) for details of the technique. Our new difference operator captures the extrinsic changes between surfaces that are nearly identical intrinsically, most notably the “bumpy cubes” as an extreme example.

(a) Shapes to compare to [Liu et al. 2017]; the color is used to identify the shape in the following eigenvalue plots.

(c) The top 10 Steklov eigenvalues cluster bowls into two categories, according to whether the bump is inward or outward. Scaling the eigenvalues with the scaling factor $\mathcal{R}(\Omega)$ (defined in Appendix A) makes the separation more clear.

(b) The top 10 and 100 Dirac eigenvalues (which can be negative). Bowls with the same geometric details have very similar eigenvalues.

(d) The top 100 and 500 Steklov eigenvalues. The eigenvalues eventually go into four classes according to the type of geometric detail.

Fig. 27. Comparison with Liu et al. (2017).

In general, our Steklov-based shape difference matrices capture more variability in the shape collections than the Laplace–Beltrami counterparts, even when using the generalized DtN operator introduced in Section 5.5 to study differences between open surfaces (the bending sheet and falling cloth sequences).

7.6 Shape Exploration and Retrieval

In shape retrieval applications, we may wish to distinguish the same piece of geometry in different poses, e.g., for purposes of organizing a database of human scans or for processing an animated sequence of an articulated character. In these instances, the additional discriminative features provided by DtN-based computation may be desirable. As an example, Figure 26 shows an eigenspace embedding of a collection of shapes from the TOSCA dataset (Bronstein et al. 2008) using the Steklov spectrum, similar to Reuter et al. (2006).

7.7 Comparison with the Dirac Operator

We next compare our method with the Dirac operator (Liu et al. 2017). From a high level, the DtN operator is assembled using pairwise distances and relative normals, while the relative Dirac operator compares normals of nearby vertices. We highlight differences between the DtN and the Dirac operators here.

Broadly speaking, the Steklov spectrum, like the Laplacian spectrum, considers local geometric details as middle- and high-frequency geometric information; the low eigenvalues are robust (nearly invariant) to such patterns. In contrast, the low end of the

relative Dirac spectrum detects and discriminates such geometric details.

We provide an experiment in Figure 27 to illustrate this difference. In this example, we take a smooth bowl model and reflect the inward bump outward to yield an isometric shape; then, we paint noise, smooth clay, and bumps with sharp edges onto the two smooth “bowls,” to create four different types of geometric details. We compute both the Steklov spectrum and the Dirac spectrum on these eight shapes as shown in Figure 27(a). In our experiments, the Dirac spectrum is computed using the source code released by the authors, which computes the spectrum of the blended operator of the relative Dirac operator (whose weight is $\sigma = 1 - 10^{-6}$) and the Laplacian operator (whose weight is $1 - \sigma = 10^{-6}$).

The top Steklov eigenvalues in our experiment roughly cluster into two categories, according to whether the bump—a large-scale volumetric feature—points inward or outward. The separation is even more clear after we scale the Steklov eigenvalues using the isoperimetric ratio of each shape, taking into account that the scale of the shape has changed when adding geometric details. The Steklov eigenvalues gradually discriminate geometric details: the spectra of the bumpy, smooth, and noisy bowls start deviating at the 20th, 40th, and 200th eigenvalues, respectively. The principal part of the Dirac spectrum, on the other hand, discriminates the shapes into four categories, according to the type of geometric details. Note the eight shapes are nearly isometric to each other, so adjusting the weighting σ for operator blending does not help discriminate the direction of the bump.

Fig. 28. Steklov eigenfunctions 2 to 6 computed on the eight shapes. We skipped the first eigenfunction, which is trivially a constant function for all shapes. The eigenfunctions are nearly invariant to geometric detail, though the “rotation” within each eigenspace can be affected.

Fig. 29. Magnitudes of the first five (quaternion-valued) Dirac eigenfunctions. We can observe that the eigenfunction resonates most at nearly flat plateaus and aligns to the clay ridges.

The choice of using Dirac or Steklov depends on the desired behavior of the operator. The Steklov operator captures volumetric information, while the Dirac operator is surface based and accompanied by fewer guarantees. On the other hand, computing the Dirac spectrum involves sparse linear algebra, which can be cheaper than working with the dense BEM system. Geometric details, encoded at the beginning of the Dirac spectrum, are cheaper to obtain from the Dirac spectrum than from the Steklov spectrum.

8 CONCLUSION

The Dirichlet-to-Neumann operator provides an intuitive way to transfer the successes of intrinsic geometry processing to applications needing volumetric information. As we demonstrate using assorted applications in Section 7, DtN operators can substitute for Laplace–Beltrami operators *mutatis mutandis* to incorporate extrinsic geometry into existing algorithms and pipelines. The end result is an easily applied and relatively intuitive extrinsic operator backed by theoretical understanding of its behavior on smooth domains. Our experiments further demonstrate that the boundary integral formulation of the DtN operator appears to encode meaningful geometric data on surfaces with boundaries that may not enclose a volumetric region.

The DtN operator is best understood for smooth surfaces. For nonsmooth surfaces including ones with sharp structures, theoretical understanding of how the Steklov spectrum is determined by the geometry is an active topic in spectral geometry (Girouard et al. 2017; Levitin et al. 2017). In addition, numerically computing the spectrum to a high accuracy for these shapes is also a very challenging problem: the inverse single-layer potential operator often yields singularities at sharp corners and edges (if the shape contains any; see, e.g., Jackson (1998, Section 2.11)), leading to less accurate solutions; when fast convergence and high accuracy are desirable, one might want to combine our BEM scheme with more sophisticated numerical techniques like Igarashi and Honma (1996) and Lintner and Bruno (2015) to cancel out singularities and end up with a better-conditioned system. Our technique as it currently is implemented exhibits some limitations whose resolution will expand the applicability of our work. Most prominently, basic BEM algorithms can fail when triangle meshes self-overlap, e.g., if two triangles intersect. This drawback potentially can be resolved using a more complex implementation that detects and resolves self-overlap before evaluating boundary integrals, e.g., using techniques recently introduced to the graphics community by Zhou et al. (2016). Another smaller drawback of our work is the assumption that volumes enclosed by surfaces are homogeneous. In physically motivated contexts, it may be useful to incorporate anisotropy via modified Laplacian operators or by warping of the interior of the domain. This can be challenging using the boundary element method, which requires closed-form Green’s functions for the differential operators involved.

These potential extensions aside, future work on the computational side might focus on the efficiency of our technique. While asymptotically BEM matches if not surpasses the efficiency of FEM, our current BEM implementation could benefit from accelerated schemes for numerical quadrature and hierarchical matrix evaluation to bring this efficiency to practice. Additional speed could be gained by considering GPU-based implementa-

tions of BEM and eigenvalue algorithms, including the GPU-accelerated version of LOBPCG proposed in Anzt et al. (2015); their blocked matrix-vector product has some rough similarity to the hierarchical matrix method used in our BEM implementation. GPU-accelerated BEM has been considered in several recent works (Hamada 2011; Stock and Gharakhani 2010; Takahashi and Hamada 2009; Yokota et al. 2011), a developing topic in numerical analysis that will improve the practical aspects of our method as a side effect. Application-wise, this article focuses on empirical study of low-level geometric behaviors of the DtN operator and Steklov spectrum; exploring its use in high-level tasks and combining it with the existing techniques in Laplacian geometry processing are interesting directions for future work.

More broadly, our experiments with DtN operators and the Steklov eigenvalue problem reveal the value of considering extrinsic shape in geometry processing pipelines. This and other future approaches will bring a *complete* geometric characterization to the constellation of spectral and operator-based shape processing techniques.

APPENDIX

A EIGENVALUE NORMALIZATION

For shape analysis applications where scale invariance is desired, we suggest using the following isoperimetric-ratio-based scaling factor:

$$\mathcal{R}(\Omega) := \frac{\text{AREA}(\partial\Omega)}{\sqrt[3]{\text{VOL}(\Omega)}},$$

where $\text{AREA}(\partial\Omega)$ and $\text{VOL}(\Omega)$ are surface area and interior volume of the domain Ω . This scaling factor is justified by the following theorem:

THEOREM A.1 (COLBOIS ET AL. (2011, THEOREM 1.3)). *If $\Omega \in \mathbb{R}^3$ is conformally equivalent to a complete manifold with nonnegative Ricci curvature, then $\lambda_i(\Omega) \leq c \frac{i^{2/3}}{\mathcal{R}(\Omega)}, \forall i$, where c is a constant.*

Note this is not an asymptotic formula. Figure 27 shows that this scaling factor behaves reasonably for that example, canceling the effects of volume and area scaling.

Volume can be robustly evaluated, even for nonwatertight meshes, from a boundary representation as the integral

$$\begin{aligned} \text{VOL}(\Omega) &:= \frac{1}{3} \int_{\Gamma} \mathbf{x} \cdot \mathbf{n}(\mathbf{x}) \, d\Gamma(\mathbf{x}) \\ &= \frac{1}{18} \sum_{T_i} (\mathbf{x}_{i,1} + \mathbf{x}_{i,2} + \mathbf{x}_{i,3}) \cdot ((\mathbf{x}_{i,3} - \mathbf{x}_{i,2}) \times (\mathbf{x}_{i,2} - \mathbf{x}_{i,1})), \end{aligned}$$

where $\mathbf{x}_{i,\{1,2,3\}}$ are positions of the three vertices in the triangle T_i .

ACKNOWLEDGMENTS

The authors thank Mikhail Karpukhin for communicating the proof of Proposition 3.1, as well as Chris Wojtan and David Hahn for early conversations about BEM. Meshes are courtesy of the AIM@Shape Repository, the TOSCA project, the Stanford Computer Graphics Laboratory, MIT, and the Princeton Segmentation Benchmark.

REFERENCES

- Yonathan Aflalo, Haim Brezis, and Ron Kimmel. 2015. On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences* 8, 2 (2015), 1141–1160.
- Hartwig Anzt, Stanimire Tomov, and Jack Dongarra. 2015. Accelerating the LOBPCG method on GPUs using a blocked sparse matrix vector product. In *Proc. Symposium on High Performance Computing*. 75–82.
- Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proceedings of the International Conference on Computer Vision (ICCV'11)*. 1626–1633.
- Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Maks Ovsjanikov. 2013. An operator approach to tangent vector field processing. In *Computer Graphics Forum*, Vol. 32. 73–82.
- Omri Azencot, Maks Ovsjanikov, Frédéric Chazal, and Mirela Ben-Chen. 2015. Discrete derivatives of vector fields on surfaces—An operator approach. *Transactions on Graphics* 34, 3 (2015), 29.
- Omri Azencot, Orestis Vantzos, and Mirela Ben-Chen. 2016. Advection-based function matching on surfaces. In *Computer Graphics Forum*, Vol. 35. 55–64.
- Omri Azencot, Steffen Weißmann, Maks Ovsjanikov, Max Wardetzky, and Mirela Ben-Chen. 2014. Functional fluids on surfaces. In *Computer Graphics Forum*, Vol. 33. 237–246.
- Mikhail Belkin, Jian Sun, and Yusu Wang. 2008. Discrete Laplace operator on meshed surfaces. In *Proceedings of the Symposium on Discrete Algorithms (SODA'08)*. 278–287.
- Mikhail Belkin, Jian Sun, and Yusu Wang. 2009. Constructing Laplace operator from point clouds in $\mathbb{R}^d$. In *Proceedings of the Symposium on Discrete Algorithms (SODA'09)*. 1031–1040.
- Michele Benzi, Gene H. Golub, and Jörg Liesen. 2005. Numerical solution of saddle point problems. *Acta Numerica* 14 (2005), 1–137.
- Silvia Bertoluzza. 2003. Non conforming domain decomposition: The Steklov–Poincaré operator point of view. In *Domain Decomposition Methods in Science and Engineering*. 15–26.
- Steffen Börm, Lars Grasedyck, and Wolfgang Hackbusch. 2003. *Hierarchical Matrices*. Technical Report. Max Planck Institute.
- Davide Boscaini, Davide Eynard, Drosos Kourounis, and Michael M. Bronstein. 2015. Shape-from-operator: Recovering shapes from intrinsic operators. In *Computer Graphics Forum*, Vol. 34. 265–274.
- James H. Bramble and Joseph E. Pasciak. 1988. A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems. *Mathematics of Computation* 50, 181 (1988), 1–17.
- Haim Brezis. 2010. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer Science & Business Media.
- Alexander M. Bronstein, Michael M. Bronstein, Leonidas J. Guibas, and Maks Ovsjanikov. 2011. Shape Google: Geometric words and expressions for invariant shape retrieval. *Transactions on Graphics* 30, 1 (2011), 1.
- Alexander M. Bronstein, Michael M. Bronstein, and Ron Kimmel. 2008. *Numerical Geometry of Non-rigid Shapes*. Springer Science & Business Media.
- Oscar P. Bruno and Stéphane K. Lintner. 2013. A high-order integral solver for scalar problems of diffraction by screens and apertures in three-dimensional space. *Journal of Computational Physics* 252 (2013), 250–274.
- John Canny and John Reif. 1987. New lower bound techniques for robot motion planning problems. In *Annual Symposium on Foundations of Computer Science*. 49–60.
- Xiaobai Chen, Aleksey Golovinskiy, and Thomas Funkhouser. 2009. A benchmark for 3D mesh segmentation. In *Transactions on Graphics*, Vol. 28. ACM, 73.
- Ming Chuang, Linjie Luo, Benedict J. Brown, Szymon Rusinkiewicz, and Michael Kazhdan. 2009. Estimating the laplace–Beltrami operator by restricting 3D functions. In *Computer Graphics Forum*, Vol. 28. 1475–1484.
- Ronald R. Coifman and Stéphane Lafon. 2006. Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 1 (2006), 5–30.
- Bruno Colbois, Ahmad El Soufi, and Alexandre Girouard. 2011. Isoperimetric control of the Steklov spectrum. *Journal of Functional Analysis* 261, 5 (2011), 1384–1399.
- Etienne Corman, Justin Solomon, Mirela Ben-Chen, Leonidas Guibas, and Maks Ovsjanikov. 2017. Functional characterization of intrinsic and extrinsic geometry. *Transactions on Graphics* 36, 2 (March 2017), 14:1–14:17.
- Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *Transactions on Graphics* 32, 5 (2013), 152.
- Johannes J. Duistermaat and Victor W. Guillemin. 1975. The spectrum of positive elliptic operators and periodic bicharacteristics. *Inventiones Mathematicae* 29 (1975), 39–80.
- Nader Engheta, William D. Murphy, Vladimir Rokhlin, and Marius S. Vassiliou. 1992. The fast multipole method (FMM) for electromagnetic scattering problems. *Transactions on Antennas and Propagation* 40, 6 (1992), 634–641.
- Ming Gao, Nathan Mitchell, and Eftychios Sifakis. 2014. Steklov–Poincaré skinning. In *Proc. Symposium on Computer Animation*. 139–148.
- Alexandre Girouard, Jean Lagacé, Iosif Polterovich, and Alessandro Savo. 2017. The Steklov spectrum of cuboids. *arXiv preprint arXiv:1711.03075* (2017).
- Alexandre Girouard and Iosif Polterovich. 2017. Spectral geometry of the Steklov problem. *Journal of Spectral Theory* 7, 2 (2017), 321–359.
- Gerd Grubb. 2009. Pseudodifferential methods for boundary value problems. *Distributions and operators*. Springer Science and Business Media, 305–333.
- Wolfgang Hackbusch. 1999. A sparse matrix arithmetic based on $\mathcal{H}$ -matrices. Part I: Introduction to $\mathcal{H}$ -matrices. *Computing* 62, 2 (1999), 89–108.
- Shoji Hamada. 2011. GPU-accelerated indirect boundary element method for voxel model analyses with fast multipole method. *Computer Physics Communications* 182, 5 (2011), 1162–1168.
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2010. Eigenmodes of surface energies for shape analysis. In *Proceedings of Geometric Modeling and Processing (GMP'10)*. Springer, 296–314.
- Hugues Hoppe. 1996. Progressive meshes. In *Proc. SIGGRAPH*. 99–108.
- Hajime Igarashi and Toshihisa Honma. 1996. A boundary element method for potential fields with corner singularities. *Applied Mathematical Modelling* 20, 11 (1996), 847–852.
- John David Jackson. 1998. Classical electrodynamics. In *Classical Electrodynamics*, 3rd ed., John David Jackson (Ed.). Wiley-VCH, 832.
- Alec Jacobson. 2013. *Algorithms and Interfaces for Real-time Deformation of 2d and 3d Shapes*. Ph.D. Dissertation. ETH Zurich.
- Alec Jacobson, Ladislav Kavan, and Olga Sorkine-Hornung. 2013. Robust inside-outside segmentation using generalized winding numbers. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 33.
- Alexander Kachalov, Yaroslav Kurylev, and Matti Lassas. 2001. *Inverse Boundary Spectral Problems*. CRC Press.
- Ram P. Kanwal. 2013. *Linear Integral Equations*. Springer.
- Andrew V. Knyazev. 2001. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM Journal on Scientific Computing* 23, 2 (2001), 517–541.
- Iasonas Kokkinos, Michael M. Bronstein, Roei Litman, and Alex M. Bronstein. 2012. Intrinsic shape context descriptors for deformable shapes. In *Proceedings of the Computer Vision and Pattern Recognition (CVPR'12)*. 159–166.
- John M. Lee and Gunther Uhlmann. 1989. Determining anisotropic real-analytic conductivities by boundary measurements. *Communications on Pure and Applied Mathematics* 42, 8 (1989), 1097–1112.
- Michael Levitin, Leonid Parnowski, Iosif Polterovich, and David A. Sher. 2017. Sloshing, Steklov and corners I: Asymptotics of sloshing eigenvalues. *arXiv Preprint arXiv:1709.01891* (2017).
- Jian Liang, Rongjie Lai, Tsz Wai Wong, and Hongkai Zhao. 2012. Geometric understanding of point clouds using Laplace–Beltrami operator. In *Proceedings of the Computer Vision and Pattern Recognition (CVPR'12)*. 214–221.
- Stéphane K. Lintner and Oscar P. Bruno. 2015. A generalized Calderón formula for open-arc diffraction problems: Theoretical considerations. *Proceedings of the Royal Society of Edinburgh Section A: Mathematics* 145, 2 (2015), 331–364.
- Yaron Lipman, Raif M. Rustamov, and Thomas A. Funkhouser. 2010. Biharmonic distance. *Transactions on Graphics* 29, 3 (2010), 27.
- Roei Litman, Alexander M. Bronstein, and Michael M. Bronstein. 2012. Stable volumetric features in deformable shapes. *Computers & Graphics* 36, 5 (2012), 569–576.
- Haixiang Liu, Nathan Mitchell, Mridul Aanjaneya, and Eftychios Sifakis. 2016. A scalable Schur-complement fluids solver for heterogeneous compute platforms. *Transactions on Graphics* 35, 6 (2016), 201.
- Hsueh-Ti Derek Liu, Alec Jacobson, and Keenan Crane. 2017. A Dirac operator for extrinsic shape analysis. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 139–149.
- Robert Osada, Thomas Funkhouser, Bernard Chazelle, and David Dobkin. 2002. Shape distributions. *Transactions on Graphics* 21, 4 (2002), 807–832.
- Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional maps: A flexible representation of maps between shapes. *Transactions on Graphics* 31, 4 (2012), 30.
- Giuseppe Patané. 2015. Volumetric heat kernel: Padé–Chebyshev approximation, convergence, and computation. *Computers & Graphics* 46 (2015), 64–71.
- Giuseppe Patané. 2016. Laplacian spectral kernels and distances for geometry processing and shape analysis. In *Computer Graphics Forum*, Vol. 35. 599–624.
- Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1 (1993), 15–36.
- Iosif Polterovich and David A. Sher. 2015. Heat invariants of the Steklov problem. *Journal of Geometric Analysis* 25, 2 (2015), 924–950.
- Alfio Quarteroni and Alberto Valli. 1999. *Domain Decomposition Methods for Partial Differential Equations*. Oxford.
- Dan Raviv, Michael M. Bronstein, Alexander M. Bronstein, and Ron Kimmel. 2010. Volumetric heat kernel signatures. In *Proceedings of the Workshop on 3D Object Retrieval (3DOR'10)*. 39–44.
- Martin Reuter, Silvia Biasotti, Daniela Giorgi, Giuseppe Patané, and Michela Spagnuolo. 2009. Discrete Laplace–Beltrami operators for shape analysis and segmentation. *Computers & Graphics* 33, 3 (2009), 381–390.
- 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.

- Vladimir Rokhlin. 1985. Rapid solution of integral equations of classical potential theory. *Journal of Computational Physics* 60, 2 (1985), 187–207.
- Raif M. Rustamov. 2011. Interpolated eigenfunctions for volumetric shape processing. *Visual Computer* 27, 11 (2011), 951–961.
- Raif M. Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Leonidas Guibas. 2013. Map-based exploration of intrinsic shape differences and variability. *Transactions on Graphics* 32, 4 (2013), 72.
- Rohan Sawhney and Keenan Crane. 2017. Boundary first flattening. *ACM Transactions on Graphics (TOG)* 37, 1 (2017), 5.
- Wojciech Śmigaj, Timo Betcke, Simon Arridge, Joel Phillips, and Martin Schweiger. 2015. Solving boundary integral problems with BEM++. *Transactions on Mathematical Software* 41, 2 (2015), 6.
- Barry Smith, Petter Bjørstad, and William Gropp. 2004. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge.
- Olga Sorkine. 2005. Laplacian mesh processing. In *Eurographics State of the Art Report*.
- Olaf Steinbach. 2007. *Numerical Approximation Methods for Elliptic Boundary Value Problems: Finite and Boundary Elements*. Springer.
- Olaf Steinbach and Wolfgang L. Wendland. 1998. The construction of some efficient preconditioners in the boundary element method. *Advances in Computational Mathematics* 9, 1–2 (1998), 191–216.
- Mark J. Stock and Adrin Gharakhani. 2010. A GPU-accelerated boundary element method and vortex particle method. In *AIAA 40th Fluid Dynamics Conference and Exhibit*, Vol. 1.
- Martin Stoll and A. J. Wathen. 2007. *The Bramble–Pasciak Preconditioner for Saddle Point Problems*. Technical Report. University of Oxford.
- 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.
- Toru Takahashi and Tsuyoshi Hamada. 2009. GPU-accelerated boundary element method for Helmholtz equation in three dimensions. *International Journal for Numerical Methods in Engineering* 80, 10 (2009), 1295–1321.
- Federico Tombari, Samuele Salti, and Luigi Di Stefano. 2010. Unique signatures of histograms for local surface description. In *Proceedings of the European Conference on Computer Vision (ECCV’10)*. 356–369.
- Andrea Toselli and Olof B. Widlund. 2005. *Domain Decomposition Methods: Algorithms and Theory*. Vol. 34. Springer.
- Bruno Vallet and Bruno Lévy. 2008. Spectral geometry processing with manifold harmonics. In *Computer Graphics Forum*, Vol. 27. 251–260.
- Amir Vaxman, Mirela Ben-Chen, and Craig Gotsman. 2010. A multi-resolution approach to heat kernels on discrete surfaces. *Transactions on Graphics* 29, 4 (2010), 121.
- Gang Wang and Yalin Wang. 2015. Multi-scale heat kernel based volumetric morphology signature. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 751–759.
- 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.
- Kenong Wu and Martin D. Levine. 1997. 3D part segmentation using simulated electrical charge distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 11 (1997), 1223–1235.
- Rio Yokota, Jaydeep P. Bardhan, Matthew G. Knepley, Lorena A. Barba, and Tsuyoshi Hamada. 2011. Biomolecular electrostatics using a fast multipole BEM on up to 512 GPUs and a billion unknowns. *Computer Physics Communications* 182, 6 (2011), 1272–1283.
- Hao Zhang, Oliver Van Kaick, and Ramsay Dyer. 2010. Spectral mesh processing. In *Computer Graphics Forum*, Vol. 29. 1865–1894.
- Qingnan Zhou, Eitan Grinspun, Denis Zorin, and Alec Jacobson. 2016. Mesh arrangements for solid geometry. *Transactions on Graphics* 35, 4 (July 2016), 39:1–39:15.

Received July 2017; revised July 2018; accepted October 2018