

Spectral Exterior Calculus

TYRUS BERRY

Department of Mathematics, George Mason University

AND

DIMITRIOS GIANNAKIS

Courant Institute of Mathematical Sciences, New York University

Abstract

A spectral approach to building the exterior calculus in manifold learning problems is developed. The spectral approach is shown to converge to the true exterior calculus in the limit of large data. Simultaneously, the spectral approach decouples the memory requirements from the amount of data points and ambient space dimension. To achieve this, the exterior calculus is reformulated entirely in terms of the eigenvalues and eigenfunctions of the Laplacian operator on functions. The exterior derivatives of these eigenfunctions (and their wedge products) are shown to form a frame (a type of spanning set) for appropriate L^2 spaces of k -forms, as well as higher-order Sobolev spaces. Formulas are derived to express the Laplace-de Rham operators on forms in terms of the eigenfunctions and eigenvalues of the Laplacian on functions. By representing the Laplace-de Rham operators in this frame, spectral convergence results are obtained via Galerkin approximation techniques. Numerical examples demonstrate accurate recovery of eigenvalues and eigenforms of the Laplace-de Rham operator on 1-forms. The correct Betti numbers are obtained from the kernel of this operator approximated from data sampled on several orientable and non-orientable manifolds, and the eigenforms are visualized via their corresponding vector fields. These vector fields form a natural orthonormal basis for the space of square-integrable vector fields, and are ordered by a Dirichlet energy functional which measures oscillatory behavior. The spectral framework also shows promising results on a non-smooth example (the Lorenz 63 attractor), suggesting that a spectral formulation of exterior calculus may be feasible in spaces with no differentiable structure.

Contents

1. Introduction	2
2. Overview of the Spectral Exterior Calculus (SEC)	7
3. Hilbert spaces and operators in the exterior calculus	16
4. Spectral exterior calculus (SEC) on smooth manifolds	25
5. Frame and operator representations of vector fields	42
6. Galerkin method for the 1-Laplacian	48
7. Data-driven approximation	52

8. Numerical examples	59
9. Discussion	67
Appendix A. Discrete incompatibility of product rule and Leibniz rule	70
Appendix B. Computations and derivations of formulas	71
Bibliography	75

1 Introduction

The field of manifold learning has focused significant attention recently on consistently estimating the Laplacian operator on a manifold,

$$(1.1) \quad \Delta = -\operatorname{div} \operatorname{grad} = \delta d$$

(in this paper we use the positive definite Laplacian, which we also refer to as the 0-Laplacian) [7, 44, 18, 52, 12, 32, 11, 8, 43, 49, 50, 14]. Given data $\{x_i\}$ sampled from a manifold $\mathcal{M} \subset \mathbb{R}^n$, these methods build a graph with weights given by a kernel function $k(x_i, x_j)$, and then approximate the Laplacian operator with the graph Laplacian

$$(1.2) \quad L = D - K,$$

where K and D are the kernel and degree matrices associated with k , respectively. In Table 1.1, we briefly summarize the current state-of-the-art results.

The Laplacian-based approach to manifold learning is justified by the fact that the Laplace-Beltrami operator encodes all the geometric information about a Riemannian manifold. A simple demonstration of this fact arises from the product formula for the Laplacian

$$(1.3) \quad \Delta(fh) = f\Delta h + h\Delta f - 2\operatorname{grad} f \cdot \operatorname{grad} h,$$

where the dot-product above is actually the Riemannian inner product $g_x : T_x\mathcal{M} \times T_x\mathcal{M} \rightarrow \mathbb{R}$,

$$(1.4) \quad \begin{aligned} g_x(\operatorname{grad} f(x), \operatorname{grad} h(x)) &= (\operatorname{grad} f \cdot \operatorname{grad} h)(x) \\ &= \frac{1}{2}(f(x)\Delta h(x) + h(x)\Delta f(x) - \Delta(fh)(x)). \end{aligned}$$

Specifically, given any vectors $v, w \in T_x\mathcal{M}$, there must exist functions f, h with $\operatorname{grad} f(x) = v$ and $\operatorname{grad} h(x) = w$, and then the inner product

$$g_x(v, w) = g_x(\operatorname{grad} f(x), \operatorname{grad} h(x))$$

can be computed as above.

Since the geometry of a Riemannian manifold is completely determined by the Riemannian metric, the above formulas show that metric is completely recoverable from Laplacian, so learning the Laplacian is sufficient for manifold learning. Of course, this is a theoretical rather than pragmatic notion of sufficiency. If one asks certain geometric questions, such as “What is the 0-homology of the manifold?”

TABLE 1.1. Summary of results on manifold learning

-
- (1) For uniform sampling density (with respect to the Riemannian volume measure) on a compact manifold, the Gaussian kernel provides a consistent pointwise estimator of the Laplace-Beltrami operator [7].
 - (2) For nonuniform sampling density on a compact manifold, any isotropic kernel with exponential decay can be normalized to give a consistent pointwise estimator [18].
 - (3) For nonuniform sampling density on a compact manifold, any symmetric kernel with super-polynomial decay can be normalized to give a consistent pointwise estimator with respect to a geometry determined by the kernel function [12].
 - (4) The bias-variance tradeoff implies error which is exponential in the dimension of the manifold [44, 11].
 - (5) The above results can be generalized to non-compact manifolds by assuming appropriate lower bounds on the injectivity radius and either the curvature or the ratio between the intrinsic and extrinsic distances [32, 11, 14].
 - (6) For data sampled on a compact subset of $\mathbb{R}^n$, not necessarily with manifold structure, the normalized and (under additional conditions) the unnormalized graph Laplacians converge spectrally to operators on continuous functions in the infinite-data limit [52].
 - (7) For smooth compact manifolds without boundary and uniform sampling density, the graph Laplacian associated with Gaussian kernels converges spectrally to the manifold Laplacian along a decreasing sequence of kernel bandwidth parameters as the number of samples increases [9]. More recently, these results have been extended to allow spectral approximation of more general self-adjoint elliptic operators on bounded open subsets of $\mathbb{R}^n$, with specified relationships between the bandwidth and number of points, including error estimates [43, 49, 50].
 - (8) The bias and variance of the spectral estimator were computed in [14], who showed that the variance is dominated by two terms, one proportional to the eigenvalue λ (linear) and another proportional to λ^2 (quadratic), explaining why an initial part of the spectrum (close to zero) can be significantly more accurate than larger eigenvalues. For this initial part of the spectrum, the optimal bias-variance tradeoff results in a much smaller bandwidth than is optimal for pointwise estimation.
 - (9) A separate construction (closely related to kernel estimators) uses local estimators of the tangent space and orthogonal matrices that estimate the covariant derivative in order to construct an estimator of the connection Laplacian, which is closely related to the Hodge Laplacian on 1-forms [45, 27, 46].
-

TABLE 1.2. Comparison of the Spectral Exterior Calculus (SEC) introduced in this manuscript to the Discrete Exterior Calculus (DEC) and the Finite element Exterior Calculus (FEC).

Feature	DEC	FEC	SEC
Pointwise consistency	Yes	Yes	Yes
Spectral consistency	Unknown	Yes	Yes
Works on raw data	No	No	Yes
Decouples memory from data	No	N/A	Yes
Exterior Calculus structure	Partial	Partial	Partial

(i.e., the number of connected components) this can be easily answered as the dimension of the kernel (nullspace) of the Laplacian. However, if one asks for the higher homology of the manifold, or the harmonic vector fields, or the closed or exact forms, the above formulas do not suggest any practical approach. What is needed is not merely the Laplacian, but a consistent representation of the entire exterior calculus on the manifold.

In this paper, we introduce the Spectral Exterior Calculus (SEC) as a consistent representation of the exterior calculus based entirely on the eigenfunctions and eigenvalues of the Laplacian on functions. In essence, we will follow through on the above analysis and reformulate the entire exterior calculus in terms of these eigenfunctions and eigenvalues.

Discrete formulations of the exterior calculus, utilizing a finite number of sampled points on or near the manifold, have been introduced at least as early as the mid 1970s with the work of Dodziuk [22] and Dodziuk and Patodi [23]. They constructed a combinatorial Laplacian on simplicial cochains of a smooth triangulated manifold, and showed that, under refinement of the triangulation, the combinatorial Laplacian on k -cochains converges in spectrum to the Laplace-de Rham operator on k -forms (or k -Laplacian, denoted Δ_k). A key element of this construction was the use of Whitney interpolating forms [53], mapping k -cochains to k -forms on the manifold. More recently, two alternative methods of discretely representing the exterior calculus have been developed, namely the Discrete Exterior Calculus (DEC) by Hirani [33] and Desbrun et al. [21], and the Finite element Exterior Calculus (FEC) by Arnold et al. [3, 2]. Among these, the FEC includes the techniques of Dodziuk and Patodi, as well as subsequent generalizations by Baker [5] utilizing Sullivan-Whitney piecewise-polynomial forms [48], as special cases. In Table 1.2, we compare the features of the SEC to the DEC and FEC. For manifold learning applications, we focus on three requirements: consistency, applicability to raw data, and amount of data and memory required.

The first requirement is that the method should be consistent, meaning that discrete analogs of objects from the exterior calculus should converge to their continuous counterparts in the limit of large data. In this paper we will focus on the

pointwise and spectral consistency, in appropriate Hilbert spaces, of representations of vector fields and the Laplace-de Rham operators on k -forms. We chose to focus on Laplace-de Rham operators because their eigenforms form a natural ordered basis for the space of k -forms. Moreover, in the case of the eigenforms of the 1-Laplacian, the Riemannian duals form a natural basis for smooth vector fields. While the DEC formulates a discrete analog to Δ_k , currently it has not been proven to be pointwise consistent, except for Δ_0 using the cotangent formula for the special case of surfaces in $\mathbb{R}^3$. A recent preprint by Schulz and Tsogtgerel [41] shows consistency of the DEC when used to solve Poisson problems for the 1-Laplacian, but spectral convergence is not addressed. A recent thesis of Rufat [40] considers a collocation-based variant of the DEC, also called “Spectral Exterior Calculus” but abbreviated SPEX, and shows numerical examples suggesting consistency of the kernel of the 1-Laplacian. The FEC has convergence results for estimating Laplace-de Rham and related operators, including eigenvalue problems.

In Section, 6, we prove spectral convergence results for the SEC-approximated 1-Laplacian using a Galerkin technique. More generally, many of the operators encountered in exterior calculus, including the k -Laplacians, are unbounded, and the requirement of consistency must necessarily address domain issues for such operators. One of the advantageous aspects of the SEC is that the Sobolev regularity appropriate for differential operators such as k -Laplacians can be naturally enforced using the eigenvalues of the 0-Laplacian. This allows us to construct spectrally convergent Galerkin schemes using classical results from the spectral approximation theory for linear operators. This approach generalizes Galerkin approximation schemes for a class of unbounded operators on functions (generators of measure-preserving dynamical systems) [29, 30, 19] to operators acting on vector fields and higher-order k -forms.

Our second requirement is that the method should only require raw data, as the assumption of an auxiliary structure such as a simplicial complex is unrealistic for many data science applications. The FEC makes strong use of the known structure of the manifold to build their finite element constructions, which makes the FEC inappropriate for manifold learning. Indeed, it is instead targeted at solving PDEs on manifolds where the manifold structure is given explicitly. Based on this requirement we will not consider further comparison to the FEC. The DEC also makes strong use of a simplicial complex in their formulation and in the consistency results. It is conceivable that one could apply the DEC to an abstract simplicial complex based on an ε -ball or k -nearest neighbor construction, however there are no consistency results for such constructions.

Our third requirement is that the memory requirements should be decoupled from the data requirements, since data sets may be very large, rendering any method requiring memory that is even quadratic in the data impractical. In the DEC, discrete k -forms are encoded as weights on all combinations of k -neighbors of each data point. For a data set with N data points, each having ℓ neighbors, functions

would be represented as $N \times 1$ vectors, 1-forms as $N \times \ell$ matrices, and general k -forms as $N \times \ell^k$ matrices. Thus, operators such as the k -Laplacian are represented as $N\ell^k \times N\ell^k$ matrices. The SEC provides an alternative which is much more memory efficient.

It has been shown that the error in the leading eigenfunctions of the 0-Laplacian is proportional to the eigenvalues [14], which by Weyl's law grow according to $\lambda_j \propto n_j^{2/d}$ where d is the dimension of the manifold. Moreover, for larger eigenvalues and eigenfunctions the error ultimately becomes quadratic in the eigenvalue. The idea of the SEC is to formulate the exterior calculus entirely in terms of the eigenfunctions of the 0-Laplacian, approximated through graph-theoretic kernel methods, and to discretize the exterior calculus by projecting onto the first $M \ll N$ eigenfunctions. Thus, functions would be represented as $M \times 1$ vectors, 1-forms as $M \times J$ matrices, and k -forms as $M \times J^k$ matrices. As we will explain in Sections 2.3 and 4.1, J is the number of eigenfunctions required to form an embedding of the manifold in $\mathbb{R}^J$. Notice that highly redundant data sets may introduce extremely large values of N , but since M is decoupled from N this would not present a problem for the SEC. Also, for high-dimensional manifolds which require a large data set N to obtain a small number M of accurate eigenfunctions, the SEC could proceed using only these accurate eigenfunctions potentially yielding very efficient representations of higher-dimensional manifolds. Another advantageous aspect of SEC representations is that their memory cost is independent of the ambient data space dimension n (which can be very large in real-world applications). In fact, the only parts of the SEC framework with an n -dependent memory and computation cost are the initial graph-Laplacian construction and the spectral representation of the pushforward maps on vector fields, all of which depend linearly on n . In contrast, the cost of building simplicial complexes and other auxiliary constructs required by DEC and FEC approaches can be very high in large ambient space dimensions.

It is also desirable that a data-based exterior calculus should capture as much as possible of the structure of the exterior calculus from differential geometry, meaning that discrete analogs of continuous theorems should hold. While no method captures discrete analogs of all the continuous theorems, each method has some partial results. For example, the DEC beautifully captures a discrete analog of Stokes' theorem and the Leibniz rule holds exactly for closed forms, however the product rule for the Laplacian fails. In the SEC, the product rule for the Laplacian will hold exactly, however this leads inevitably to the failure of the Leibniz rule as shown in A.

Finally, even though here we do not explicitly address this issue, an important consideration in data-driven techniques is robustness to noise. The simplicial complexes employed in DEC become increasingly sensitive to noise with increasing simplex dimension. On the other hand, the noise robustness of the SEC is limited by the noise robustness of the graph Laplacian algorithm employed to approximate the eigenfunctions of the 0-Laplacian. The latter problem has been studied from

different perspectives in the literature [26, 37, 54], and it has been shown [26] that for certain classes of kernels and i.i.d. noises (including Gaussian noise of arbitrary variance) the graph Laplacian computed from noisy data converges spectrally to noise-free case in the infinite-data limit.

The central challenge of the SEC approach is obtaining the representation of the exterior calculus in the spectral basis of eigenfunctions of the 0-Laplacian. In the next section we overview how vector fields, k -forms, and the central operators of the exterior calculus can all be represented spectrally. Since the gradients of these functions do not span the set of vector fields (otherwise every vector field would be a gradient field), we instead build a *frame* (overcomplete set) [24, 36] consisting of products of Laplacian eigenfunctions and their gradient, and we represent vector fields in this frame. We proceed analogously for k -forms, using products of Laplacian eigenfunctions and k -fold wedge products of their differentials to construct frames.

The plan of this paper is as follows. We begin in Section 2 with an overview of the SEC, including the fundamental idea of our approach and tables which overview key formulas. Computation of the more complex formulas can be found in B. In Section 3, we briefly review the necessary background material and introduce our key definitions. Our central contribution to the theory of the exterior calculus is proving that our construction yields frames for L^2 and Sobolev spaces of vector fields and k -forms in Section 4. In Section 5, we discuss aspects of this frame representation for bounded vector fields, as well as associated representations of vector fields as operators on functions and the convergence properties of finite-rank approximations. Then, in Section 6, we employ this frame representation to construct a Galerkin approximation scheme for the eigenvalue problem of the 1-Laplacian, which is shown to converge spectrally. Section 7 establishes the consistency of the data-driven SEC representation of the exterior calculus. In Section 8, we present numerical results demonstrating the consistency of the SEC on a suite of numerical examples involving orientable and non-orientable smooth manifolds, as well as the fractal attractor of the Lorenz 63 system. We conclude with a summary discussion and future perspectives in Section 9. A Matlab code reproducing the results in Section 8 is included as supplementary material.

2 Overview of the Spectral Exterior Calculus (SEC)

As mentioned above, many manifold learning techniques are based on the ability to approximate the Laplacian operator on a manifold (1.1) via a graph Laplacian (1.2), defined on a graph of discrete data points sampled from the manifold. When this convergence is spectral, we may use the eigendecomposition of an appropriately constructed graph Laplacian. In Section 2.1 below we briefly summarize the Diffusion Maps approach to the construction [18]. The eigendecomposition from Diffusion Maps, or a comparable algorithm with spectral convergence guarantees (e.g., [50]), is the only input required to generate the entire SEC construction.

The SEC represents vector fields and differential forms using Laplacian eigenfunctions, and then reformulates the exterior calculus of Riemannian geometry in terms of these representations. This reformulation is described in Sections 2.2–2.5, and will be made rigorous in Sections 3 and 4. The motivation for this reformulation is that it allows us to define an exterior calculus using only the eigendecomposition of the 0-Laplacian. In other words, in a manifold learning scenario, the eigendecomposition of the graph Laplacian provides all the necessary inputs to formulas which generate the entire exterior calculus formalism. Of course, this implies a “low-pass” or truncated representation, and in Sections 6 and 7 we will prove that these truncated representations converge in the limit, as the number of eigenvectors increases.

2.1 The Diffusion Maps Algorithm for the Construction of the 0-Laplacian

Following the Diffusion Maps approach, we define a kernel matrix

$$K_{ij} = k_\varepsilon(x_i, x_j) := \exp\left(-\frac{\|x_i - x_j\|^2}{4\varepsilon}\right),$$

where $\{x_i\}_{i=1}^N \subset \mathcal{M} \subset \mathbb{R}^n$ is a data set sampled from the embedded manifold $\mathcal{M}$, under a measure with a smooth, fully supported density relative to the volume measure associated with the Riemannian metric, g , induced by the embedding. We then normalize $K = [K_{ij}]$ to a new kernel matrix $\hat{K}$ to remove the sampling bias,

$$\hat{K} = Q^{-1} K Q^{-1}, \quad Q = \text{diag}[Q_{ii}], \quad Q_{ii} = \sum_{j=1}^N K_{ij},$$

and finally normalize $\hat{K}$ into a Markov matrix P ,

$$(2.1) \quad P = D^{-1} \hat{K}, \quad D = \text{diag}[D_{ii}], \quad D_{ii} = \sum_{j=1}^N \hat{K}_{ij},$$

which approximates the heat semigroup $e^{-\varepsilon \Delta_0}$; see [18] for details on this procedure. The normalized kernel matrix has the generalized eigendecomposition

$$\hat{K} = D \Phi \Lambda \Phi^\top,$$

which can be computed by solving the fully symmetric generalized eigenvalue problem $\hat{K} \vec{\phi} = \Lambda D \vec{\phi}$. The eigenvalues satisfy $\Lambda = e^{-\varepsilon \hat{\lambda}}$ where the values $\hat{\lambda} \geq 0$ approximate the eigenvalues of the Laplacian, so we define $\hat{\lambda} = -\log(\Lambda)/\varepsilon$. An asymptotically equivalent (in the limit $N \rightarrow \infty$ and $\varepsilon \rightarrow 0$) approach is to form the graph Laplacian $L = D - \hat{K}$, and compute the generalized eigendecomposition $L = D \Phi \tilde{\Lambda} \Phi^\top$. In either approach, we approximate the eigenfunctions $\Delta_0 \phi_i = \lambda_i \phi_i$ of the Laplacian operator, and sort the columns of Φ so that the eigenvalues are increasing. The diagonal matrix D represents the Riemannian L^2 inner product on

the manifold, in the sense that if $\vec{f}_i = f(x_i)$ and $\vec{h}_i = h(x_i)$ are vector representations of (complex-valued) continuous functions, then

$$\vec{f}^\dagger \mathbf{D} \vec{h} \approx \int_{\mathcal{M}} f^*(x) h(x) d\mu(x) = \langle f, g \rangle_{L^2}$$

up to a constant proportionality factor, where $\dagger$ denotes complex-conjugate transpose, and μ is the Riemannian measure of the manifold. Thus, we can compute the generalized Fourier transform of the function f by

$$\hat{f} = \Phi^\top \mathbf{D} \vec{f}, \quad \hat{f}_j = \sum_{i=1}^N \phi_j(x_i) D_{ii} f(x_i) \approx \langle \phi_j, f \rangle_{L^2}.$$

We can then reconstruct the values of the function f on the data set by $\vec{f} = \Phi \hat{f}$, which holds exactly since $\Phi \Phi^\top = \mathbf{D}^{-1}$. If a smaller number of eigenvectors are used, then $\Phi \Phi^\top$ is not full rank, and the result is a low-pass filter.

Note that in applications (including those presented in this paper), one is frequently interested in real-valued functions and self-adjoint operators, so complex conjugation is not included in L^2 inner products as above. However, applications with complex-valued functions can also be of interest (e.g., in dynamical systems theory [30]), so in what follows we work with complex-valued functions to maintain generality.

2.2 Functions, multiplication, and the Riemannian metric

In Table 2.1, we show the basic elements of the exterior calculus and their SEC formulations. For example, complex-valued functions are represented in the SEC by their generalized Fourier transform $\hat{f}_i = \langle \phi_i, f \rangle_{L^2}$, which is justified since the Hodge theorem shows that the eigenfunctions ϕ_i form a smooth orthonormal basis for square-integrable functions on the manifold. It should be noted that, as with all L^2 expansions, f may differ from the reconstructed function $\sum_i \hat{f}_i \phi_i$ on sets of measure zero. Similarly, all frame representations of vector fields and k -forms in Table 2.1 and the ensuing discussion should be interpreted in an L^2 sense.

The two key elements of Table 2.1 are the representation of function multiplication and the Riemannian metric.

First, function multiplication will be represented by the fully symmetric three-index tensor

$$(2.2) \quad c_{ijk} = \langle \phi_i \phi_j, \phi_k \rangle_{L^2},$$

which will be a key building block of the SEC. Note that here we use the term “tensor” to represent a general multi-index object such as c_{ijk} derived from inner products of Laplacian eigenfunctions. While these objects are not geometrical tensors on the manifold, they nevertheless transform via familiar tensor laws under changes of L^2 basis preserving the Laplacian eigenspaces.

Next, the Riemannian metric is represented based on the product formula (1.3), and is given by (1.4). The power of the SEC is that we will only need to represent

TABLE 2.1. The SEC reformulation of the basic elements of the exterior calculus from Riemannian geometry.

Object	Symbolic	Spectral
Function	f	$\hat{f}_i = \langle \phi_i, f \rangle_{L^2}$
Laplacian	Δf	$\langle \phi_k, \Delta f \rangle_{L^2} = \lambda_k \hat{f}_k$
L^2 Inner Product	$\langle f, h \rangle_{L^2}$	$\sum_i \hat{f}_i^* \hat{h}_i$
Dirichlet Energy	$\langle f, \Delta f \rangle_{L^2} = \int_{\mathcal{M}} \ \text{grad } f\ ^2 d\mu$	$\sum_i \lambda_i \hat{f}_i ^2$
Multiplication	$\phi_i \phi_j$	$c_{ijk} = \langle \phi_i \phi_j, \phi_k \rangle_{L^2}$
Function Product	fh	$\sum_{ij} c_{kij} \hat{f}_i \hat{h}_j$
Riemannian Metric	$\text{grad } \phi_i \cdot \text{grad } \phi_j$	$g_{kij} \equiv \langle \text{grad } \phi_i \cdot \text{grad } \phi_j, \phi_k \rangle_{L^2}$ $= (\lambda_i + \lambda_j - \lambda_k) c_{kij} / 2$
Gradient Field	$\text{grad } f(h) = \text{grad } f^* \cdot \text{grad } h$	$\langle \phi_k, \text{grad } f(h) \rangle_{L^2} = \sum_{ij} g_{kij} \hat{f}_i \hat{h}_j$
Exterior Derivative	$df(\text{grad } h) = df^* \cdot dh$	$\sum_{ij} g_{kij} \hat{f}_i \hat{h}_j$
Vector Field (basis)	$v(f) = v^* \cdot \text{grad } f$	$\sum_j v_{ij} \hat{f}_j$
Divergence	$\text{div } v$	$\langle \phi_i, \text{div } v \rangle_{L^2} = -v_{0i}$
Frame Elements	$b_{ij}(\phi_l) = \phi_i \text{grad } \phi_j(\phi_l)$	$G_{ijkl} \equiv \langle b_{ij}(\phi_l), \phi_k \rangle_{L^2}$ $= \sum_m c_{mik} g_{mjl}$
Vector Field (frame)	$v(f) = \sum_{ij} v^{ij} b_{ij}(f)$	$\langle \phi_k, v(f) \rangle_{L^2} = \sum_{ijl} G_{ijkl} v^{ij} \hat{f}_l$
Frame Elements	$b^{ij}(v) = b^i db^j(v)$	$\langle \phi_k, b^{ij}(v) \rangle_{L^2} = \sum_{nlm} c_{kmi} G_{nlmj} v^{nl}$
1-Forms (frame)	$\omega = \sum_{ij} \omega_{ij} b^{ij}$	$\langle \phi_k, \omega(v) \rangle_{L^2} = \sum_{ij} \omega_{ij} \langle \phi_k, b^{ij}(v) \rangle_{L^2}$

the metric for gradients of (real) eigenfunctions $\text{grad } \phi_i$ and $\text{grad } \phi_j$, where we find that

(2.3)

$$g(\text{grad } \phi_i, \text{grad } \phi_j) = \frac{1}{2}(\phi_i \Delta \phi_j + \phi_j \Delta \phi_i - \Delta(\phi_i \phi_j)) = \frac{1}{2}((\lambda_i + \lambda_j) \phi_i \phi_j - \Delta(\phi_i \phi_j)).$$

We can further reduce this by writing the product $\phi_i \phi_j = \sum_k c_{ijk} \phi_k$, so that

$$g(\text{grad } \phi_i, \text{grad } \phi_j) = \frac{1}{2} \sum_k (\lambda_i + \lambda_j - \lambda_k) c_{ijk} \phi_k,$$

meaning that the k -th Fourier coefficient of the Riemannian metric is

$$(2.4) \quad g_{kij} \equiv \langle g(\text{grad } \phi_i, \text{grad } \phi_j), \phi_k \rangle_{L^2} = \frac{1}{2}(\lambda_i + \lambda_j - \lambda_k) c_{ijk}.$$

Notice that g_{kij} is symmetric in i and j but not in k . These first two simple formulas are the key to the SEC.

2.3 Vector fields

We will need two different ways of representing vector fields. The first method is called the operator representation and is based on the interpretation of a vector field as a map from smooth functions to smooth functions, defined by

$$v(f) = v^* \cdot \text{grad } f = g(v^*, \text{grad } f),$$

where v^* denotes the complex-conjugate vector field to v . Note that, as with functions, in SEC we consider vector fields, differential forms, and other tensors to be complex. This is because, ultimately, we will be concerned with spectral approximation of operators on these objects, and the complex formulation will allow us to take advantage of the full range of spectral approximation techniques for operators on Hilbert spaces over the complex numbers. Throughout, our convention will be that Riemannian inner products on complex tensors are conjugate-symmetric in their first argument, e.g., $g(fv, w) = f^*g(v, w)$ for vector fields v, w and function f . Since we have a smooth basis $\{\phi_i\}$ for functions, we can represent any vector field v in this basis by an operator with matrix elements

$$v_{ij} = \langle \phi_i, v(\phi_j) \rangle_{L^2} = \langle \phi_i, v^* \cdot \text{grad } \phi_j \rangle_{L^2} = \langle \phi_i \text{grad } \phi_j, v \rangle_{L^2_{\mathbb{X}}} = \langle b_{ij}, v \rangle_{L^2_{\mathbb{X}}},$$

where the first two inner products appearing above are the L^2 inner products on functions, the last two inner products are the Hodge inner products induced on vector fields,

$$\langle v, w \rangle_{L^2_{\mathbb{X}}} = \int_{\mathcal{M}} g(v, w) d\mu,$$

and $b_{ij} = \phi_i \text{grad } \phi_j$ are smooth vector fields. Note that the Hodge inner product defines the space of square-integrable vector fields, denoted $L^2_{\mathbb{X}}$.

The second method of representing a vector field will be as a linear combination of the vector fields b_{ij} just introduced, with coefficients v^{ij} so that

$$v = \sum_{ij} v^{ij} b_{ij}.$$

As we will show in Section 4, the vector fields $\{b_{ij}\}$ where $i = 1, \dots, \infty$ and $j = 1, \dots, J < \infty$ spans $L^2_{\mathbb{X}}$. However, instead of a basis, this set is only a *frame* for this space, i.e., a spanning set satisfying appropriate upper and lower bounds for the ℓ^2 norms of the sequence $\langle b_{ij}, v \rangle_{L^2_{\mathbb{X}}}$ for every $v \in L^2_{\mathbb{X}}$ [17]. Since $\{b_{ij}\}$ is not a basis, this representation will generally not be unique, although frame theory ensures that there is a unique choice of coefficients v^{ij} which minimizes the ℓ^2 norm [17, Lemma 5.3.6].

As we will see in the Section 2.5, there is a natural choice of basis for $L^2_{\mathbb{X}}$, and constructing this basis will be a central goal of the SEC approach, however doing so requires using the frame $\{b_{ij}\}$. To motivate this choice of frame elements, note that given a fixed point x on the manifold and a sufficient (finite) number of eigenfunctions ϕ_j , the gradients of these eigenfunctions $\text{grad } \phi_j(x)$ will span the tangent space $T_x \mathcal{M}$ (see Section 3 for details). In fact, for a d -dimensional compact

manifold, we should be able to find d eigenfunctions whose gradients form a basis for $T_x\mathcal{M}$ for a fixed x . However, in general any choice of d eigenfunctions will not span $T_x\mathcal{M}$ for every x simultaneously, meaning that the choice of eigenfunctions which span depends on x . This is easily demonstrated by the example of the sphere S^2 . On S^2 , every smooth vector field vanishes at some point x , so at that point the collection of d gradient fields will at most span a $(d-1)$ -dimensional subspace of T_xS^2 . Intuitively, given a collection of sufficiently many gradients of eigenfunctions $\{\text{grad } \phi_j\}_{j=1}^J$, and if the manifold is not too “large” (i.e., it is compact), we can span all the tangent spaces simultaneously with $J < \infty$, but of course we no longer have a basis. Given an arbitrary smooth vector field v , we can then represent v at each point $x \in \mathcal{M}$ as a linear combination of gradients of eigenfunctions,

$$v_x = \sum_{j=1}^J c_{v,j}(x) \text{grad } \phi_j(x).$$

If we can choose the coefficients $c_{v,j}(x)$ in this linear combination to be smooth functions on the manifold, then these functions can be represented in the basis $\{\phi_i\}$ of eigenfunctions, so that

$$c_{v,j}(x) = \sum_{i=0}^{\infty} \langle \phi_i, c_{v,j} \rangle_{L^2} \phi_i(x),$$

which means that we can represent the vector field v as

$$v = \sum_{j=1}^J \sum_{i=0}^{\infty} \langle \phi_i, c_{v,j} \rangle_{L^2} \phi_i \text{grad } \phi_j.$$

We now consider how to move between the operator representation and frame representation of a vector field. Substituting the frame representation into the operator representation, we find that

$$(2.5) \quad v_{ij} = \langle b_{ij}, v \rangle_{L^2_{\mathfrak{X}}} = \sum_{kl} v^{kl} \langle b_{kl}, b_{ij} \rangle_{L^2_{\mathfrak{X}}} = \sum_{kl} G_{ijkl} v^{kl},$$

where $G_{ijkl} = \langle b_{ij}, b_{kl} \rangle_{L^2_{\mathfrak{X}}}$ is the Grammian matrix of the frame elements with respect to the Hodge inner product. Thus, we see that the Hodge Grammian is the linear transformation which maps from the frame representation v^{kl} to the matrix representation v_{ij} . Crucially, the quantities G_{ijkl} can be computed in closed form from the spectral representation of the pointwise inner products in (2.3), viz.

$$(2.6) \quad G_{ijkl} = \langle b_{ij}, b_{kl} \rangle_{L^2_{\mathfrak{X}}} = \langle \phi_i \phi_j, \text{grad } \phi_k \cdot \text{grad } \phi_l \rangle_{L^2} = \frac{1}{2} \sum_{m=0}^{\infty} c_{ikm} c_{jlm} (\lambda_j + \lambda_l - \lambda_m).$$

Since the frame is overcomplete, the matrix G is necessarily rank deficient and thus there is no unique inverse transformation. However, if we also specify the minimum ℓ^2 norm then we can map from the matrix representation v_{ij} to the frame coefficients (with minimum norm) v^{kl} via the pseudo-inverse G^+ of the Hodge Grammian.

2.4 Differential forms

In order to build a formulation of the exterior calculus we need to first move from vector fields to differential k -forms. First, 0-forms are equivalent to C^∞ functions defined on the manifold, which we represent in the basis of eigenfunctions of the Laplacian ϕ_i . Since each eigenfunction can also be thought of as a 0-form, we will sometimes also denote the eigenfunctions by

$$b^i = \phi_i,$$

since the superscript notation is oftentimes used for basis elements of spaces of differential forms. Our primary focus in this paper will be 1-forms, which are duals to vector fields. That is, a 1-form takes in a vector field as its argument and returns a function. On a Riemannian manifold, we can move back and forth between vector fields and 1-forms with the $\sharp$ (sharp) and $\flat$ (flat) operators. Locally, these operators map 1-forms and vector fields, respectively, to their Riesz representatives with respect to the Riemannian inner product. In particular, if ω is a 1-form and v is a vector field, then

$$\omega(v) = g^{-1}(\omega^*, v^\flat) = g(v^*, \omega^\sharp),$$

where g^{-1} is the “inverse” metric on dual vectors. A fundamental operator on differential forms is the exterior derivative, d , which maps k -forms to $(k+1)$ -forms, so that the exterior derivative of a 0-form f is defined by the 1-form df , acting on a vector field v by

$$df(v) = v(f) = g(v^*, \text{grad } f).$$

We will sometimes use the notation d_k to explicitly exhibit the order of differential forms on which a given exterior derivative acts.

Since 1-forms are dual to vector fields, we will use a similar frame representation to that in Section 2.3, based on the eigenfunctions, $\{b^{ij} = b^i db^j\}$. As we will show in Section 4, these 1-forms span the space L_1^2 of square-integrable 1-forms. We also note that the Riemannian metric lifts to k -forms (see Section 3 for details), and takes two k -forms and returns a function. Integrating the Riemannian inner product of two k -forms,

$$\langle \psi, \omega \rangle_{L_k^2} = \int_{\mathcal{M}} g^{-1}(\psi, \omega) d\mu,$$

defines the Hodge inner product, which then defines the Hilbert space of square integrable k -forms. Finally, since $df^\sharp = \text{grad } f^*$, $\text{grad } f^\flat = df^*$, and the b_{ij} are real, we have $b_{ij}^\flat = b^{ij}$ and $(b^{ij})^\sharp = b_{ij}$, so the coefficients of a vector field in the frame representation can also be used to represent the corresponding 1-form and vice versa.

2.5 The Laplacian on forms

The Laplacian on k -forms is defined via the exterior derivative d and its Hodge dual, the codifferential δ , by

$$\Delta_k = d_{k-1}\delta_k + \delta_{k+1}d_k$$

In order to represent the eigenvalue problem for the operator Δ_1 in the frame $\{b^{ij}\}$, we need to compute the inner products

(2.7)

$$G_{ijkl} = \langle b^{ij}, b^{kl} \rangle_{L_1^2}, \quad E_{ijkl} = \langle b^{ij}, \Delta_1 b^{kl} \rangle_{L_1^2} = \langle db^{ij}, db^{kl} \rangle_{L_2^2} + \langle \delta b^{ij}, \delta b^{kl} \rangle_{L^2},$$

representing the Gramm matrix of Hodge inner products (which we will call the Hodge Grammian) and Dirichlet form matrix, respectively. We derive the expressions for these tensors in B below, and the formulas are summarized in Table 2.2. Note that both G and E can be written as symmetric matrices by numbering the frame elements. Moreover, we can easily represent the Gramm matrix with respect to the Sobolev H^1 inner product on 1-forms,

$$\langle \psi, \omega \rangle_{H_1^1} = \langle \psi, \omega \rangle_{L_1^2} + E_{1,1}(\psi, \omega), \quad E_{1,1}(\psi, \omega) = \langle d\omega, d\psi \rangle_{L_2^2} + \langle \delta\omega, \delta\psi \rangle_{L^2},$$

as

$$G^1 = G + E.$$

The importance of the Sobolev Grammian G^1 is that H_1^1 is a natural domain for weak (variational) formulations of the eigenvalue problem of the Laplacian on 1-forms.

As we will show in Section 6, the key to solving the eigenproblem of the Laplacian on 1-forms is to first express the problem in a weak sense, i.e., replace $\Delta_1 \phi = \nu \phi$ by

$$E_{1,1}(\psi, \phi) = \nu \langle \psi, \phi \rangle_{L_1^2}, \quad \forall \psi \in H_1^1,$$

which is equivalent to the minimization problem

$$\nu = \min_{\phi \in H_1^1 \setminus \{0\}} \left\{ \frac{E_{1,1}(\phi, \phi)}{\langle \phi, \phi \rangle_{L_1^2}} \right\}.$$

Intuitively, the ratio $E_{1,1}(\phi, \phi) / \langle \phi, \phi \rangle_{H_1^1}$ is a measure of “roughness”, or oscillatory behavior, of a given eigenform ϕ , much like the eigenvalues of the 0-Laplacian measure the roughness of the corresponding eigenfunctions. Thus, ordering eigenforms in order of increasing eigenvalue, as we will always do by convention, is tantamount to ordering them in order of increasing complexity that they exhibit on the manifold. As with functions, given finite amounts of data, the approximation error for eigenforms increases with the corresponding eigenvalue.

In SEC, we represent the eigenform in the frame, $\phi = \sum_{ij} \phi_{ij} b^{ij}$. The above variational problem can then in principle be written in matrix form as

$$E_{1,1}(\psi, \phi) = \nu \langle \psi, \phi \rangle_{L_1^2}, \quad \forall \psi \in H_1^1, \quad \implies \quad E\vec{\phi} = \nu G\vec{\phi}.$$

TABLE 2.2. The SEC representation of the Laplacian on 1-forms and the Dirichlet and Sobolev energy forms. Derivations can be found in B. Pairs of integers in parentheses indicate symmetries under permutations of tensor indices; e.g., (1, 2) in c_{ijkl}^n indicates that $c_{ijkl}^n = c_{jikl}^n$.

Operator	Tensor	Symmetries
Quadruple Product	$c_{ijkl}^0 = \langle \phi_i \phi_j, \phi_k \phi_l \rangle_{L^2}$ $= \sum_s c_{ijs} c_{skl}$	Fully symmetric
Product Energy	$c_{ijkl}^p = \langle \Delta^p(\phi_i \phi_j), \phi_k \phi_l \rangle_{L^2}$ $= \sum_s \lambda_s^p c_{ijs} c_{skl}$	(1,2), (3,4), (1,3), (2,4)
Hodge Grammian	$G_{ijkl} = \langle b^{ij}, b^{kl} \rangle_{L_1^2}$ $= [(\lambda_j + \lambda_l) c_{ijkl}^0 - c_{ijkl}^1] / 2$	(1,3), (2,4)
Antisymmetric	$\hat{G}_{ijkl} = \langle \hat{b}^{ij}, \hat{b}^{kl} \rangle_{L_1^2}$ $= G_{ijkl} + G_{jilk} - G_{jikl} - G_{ijlk}$	(1,3), (2,4)
Dirichlet Energy	$E_{ijkl} = E_{ijkl} = \langle b^{ij}, \Delta_1(b^{kl}) \rangle_{L_1^2}$ $= [(\lambda_i + \lambda_j + \lambda_k + \lambda_l)(c_{iljk}^1 - c_{ikjl}^1)$ $+ (\lambda_j + \lambda_l - \lambda_i - \lambda_k)c_{ijkl}^1$ $+ (c_{ijkl}^2 + c_{ikjl}^2 - c_{iljk}^2)] / 4$	(1,3), (2,4)
Antisymmetric	$\hat{E}_{ijkl} = \langle \hat{b}^{ij}, \Delta_1 \hat{b}^{kl} \rangle_{L_1^2}$ $= (\lambda_i + \lambda_j + \lambda_k + \lambda_l)(c_{iljk}^1 - c_{ikjl}^1)$ $+ (c_{ikjl}^2 - c_{iljk}^2)$	(1,3), (2,4)
Sobolev H^1 Grammian	$G_{ijkl}^1 = E_{ijkl} + G_{ijkl}$ $\hat{G}_{ijkl}^1 = \hat{E}_{ijkl} + \hat{G}_{ijkl}$	(1,3), (2,4)

However, the above eigenvector problem is not well-conditioned because $\mathbf{G}$ is not full-rank in general (since the frame is overcomplete, meaning there can be multiple representations of the same 1-form). In order to find an appropriate basis, we first diagonalize the Sobolev Grammian,

$$\mathbf{G}^1 = \mathbf{U} \mathbf{H} \mathbf{U}^\top, \quad \mathbf{H} = \text{diag}[h_{ii}],$$

and select the columns $\tilde{\mathbf{U}}$ of the orthogonal matrix $\mathbf{U}$ corresponding to the largest eigenvalues. For example, in our implementation we choose $h_{ii} > h_{11} \times 10^{-3}$. Notice that the columns of $\tilde{\mathbf{U}}$ contain the frame coefficients of unique orthogonal 1-forms. In other words, the matrix $\tilde{\mathbf{U}}$ is a choice of basis for H_1^1 represented in the frame. Thus, we can project the eigenvalue problem onto this basis by writing

$$(2.8) \quad \mathbf{L} \tilde{\mathbf{a}} = \mathbf{v} \mathbf{B} \tilde{\mathbf{a}}, \quad \mathbf{L} = \tilde{\mathbf{U}}^\top \mathbf{E} \tilde{\mathbf{U}}, \quad \mathbf{B} = \tilde{\mathbf{U}}^\top \mathbf{G} \tilde{\mathbf{U}}.$$

An eigenvector $\vec{a}$ of this generalized eigenvalue problem contains the coefficients of a frame representation for an eigenform φ of Δ_1 .

We should note that in practice we found somewhat better results using the antisymmetric elements $\hat{b}^{ij} = b^i db^j - b^j db^i$, likely due to the fact that these forms are less redundant. All of the formulas for the antisymmetric formulation of the 1-Laplacian are given in Table 2.2. The one change in the antisymmetric formulation is that in order to move from the frame representation to the operator representation, we need the additional tensor

$$\begin{aligned} H_{ijkl} &= \langle \phi_k, \hat{b}^{ij}(\text{grad } \phi_l) \rangle_{L^2} = \langle \phi_k, b^i db^j(\text{grad } \phi_l) \rangle_{L^2} - \langle \phi_k, b^j db^i(\text{grad } \phi_l) \rangle_{L^2} \\ &= G_{ijkl} - G_{jikl}. \end{aligned}$$

With this tensor, given the frame representation of a 1-form $\varphi = \sum_{ij} \hat{\phi}_{ij} \hat{b}^{ij}$, the operator representation of the corresponding vector field $v = \varphi^\sharp$ becomes $v_{kl} = \sum_{ij} H_{ijkl} \varphi_{ij}$.

In order to visualize an eigenform φ , we will visualize the corresponding vector field, $v = \varphi^\sharp$, which has the same frame coefficients as shown in Section 2.4. In particular, it follows from (2.5) that simply multiplying $\vec{a}$ by the matrix G , leads to $G\vec{a}$, which contains the operator representation of v . By reshaping $G\vec{a}$ into a $M \times M$ matrix V , we have $V_{ij} = \langle \phi_i, v(\phi_j) \rangle_{L^2}$. To visualize this vector field, we need to map it back into the original data coordinates. This “pushforward” operation on vector fields can also be represented spectrally [30, Proposition 6]. In particular, let X be the $n \times N$ matrix of N data points in $\mathbb{R}^n$. We first compute the Fourier transform of these coordinates by computing the D inner product with the $N \times M$ matrix Φ (see Section 2.1). Thus, $\hat{X} = XD\Phi$ is the $n \times M$ matrix containing the M Fourier coefficients of each of the n coordinate functions. We can now apply the vector field to each of these functions by multiplying $\hat{X}V^\top$, which now contains the Fourier coefficients of the pushforward of the coordinate functions. Finally, we can reconstruct the coordinates of the arrows by computing the inverse Fourier transform $\tilde{V} = \hat{X}V^\top\Phi^\top$, which is a $n \times N$ matrix containing the n -dimensional vectors which can be plotted at each data point. This method is used to visualize the SEC eigenforms in Section 8.

3 Hilbert spaces and operators in the exterior calculus

Consider a closed (compact and without boundary), smooth, orientable, d -dimensional manifold $\mathcal{M}$, equipped with a smooth Riemannian metric g . Without loss of generality, we assume that g is normalized such that its associated Riemannian measure, μ , satisfies $\mu(\mathcal{M}) = 1$. As stated in Section 2, we will work with vectors in the complexified tangent spaces, $T_x^{\mathbb{C}}\mathcal{M} = T_x\mathcal{M} \otimes_{\mathbb{R}} \mathbb{C}$, $x \in \mathcal{M}$, treating by convention $g(\cdot, \cdot)$ as conjugate-linear in its first argument. We denote the associated metric tensor on dual vectors by $\eta = g^{-1}$, and use the notation $v^\flat = g(v, \cdot)$, $v \in T_x^{\mathbb{C}}\mathcal{M}$, and $\alpha^\sharp = \eta(\alpha, \cdot)$, $\alpha \in T_x^{\mathbb{C}*}\mathcal{M}$, for the Riemannian duals of tangent

vectors and dual vectors. In what follows, we introduce the spaces of functions, vector fields, and differential forms that will be employed in the SEC framework.

3.1 Function spaces

Let $\Delta : C^\infty(\mathcal{M}) \rightarrow C^\infty(\mathcal{M})$ be the (positive-semidefinite) Laplace-Beltrami operator on smooth, complex-valued functions associated with the Riemannian metric g . It is a fundamental result in analysis on closed Riemannian manifolds (e.g., [47, 1, 39] that Δ extends to a unique self-adjoint operator $\bar{\Delta} : D(\bar{\Delta}) \rightarrow L^2(\mathcal{M}, \mu)$ with a dense domain $D(\bar{\Delta}) \subset L^2(\mathcal{M}, \mu)$ in the L^2 space associated with the Riemannian measure, and a pure point spectrum of eigenvalues $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ with no accumulation points, corresponding to a smooth orthonormal basis $\{\phi_j\}_{j=0}^\infty$ of eigenfunctions. By smoothness of the Laplace-Beltrami eigenfunctions, the products $\phi_i \phi_j$ lie in $L^2(\mathcal{M}, \mu)$; thus, we have

$$(3.1) \quad \phi_i \phi_j = \sum_{k=0}^{\infty} c_{ijk} \phi_k, \quad c_{ijk} = \langle \phi_k, \phi_i \phi_j \rangle_{L^2(\mathcal{M}, \mu)},$$

where the limit in the first equation is taken in the L^2 sense. As discussed in Section 2, our objective is to build a framework for tensor calculus on $\mathcal{M}$ that is defined entirely through the spectral properties of the Laplacian on functions, encoded in the eigenvalues λ_j , the corresponding eigenfunctions ϕ_j , and the coefficients c_{ijk} representing the algebraic relationships between the eigenfunctions.

We use the notation L^p , $1 \leq p \leq \infty$, to represent the standard Banach spaces of complex-valued functions on $\mathcal{M}$ associated with the Riemannian measure μ , equipped with the standard norms, $\|\cdot\|_{L^p}$. In the case $p = 2$, we use the shorthand notation $H = L^2$, and denote the corresponding Hilbert space inner product $\langle \cdot, \cdot \rangle_H$, which we take to be conjugate-linear on its first argument. We also consider Sobolev spaces of higher regularity, defined for $p \geq 0$ by

$$(3.2) \quad H^p = \left\{ \sum_{j=0}^{\infty} c_j \phi_j \in L^2 : \sum_{j=0}^{\infty} \lambda_j^p |c_j|^2 < \infty \right\}.$$

These spaces are closed with respect to the norms $\|f\|_{H^p} = \langle f, f \rangle_{H^p}^{1/2}$ associated with the inner products

$$(3.3) \quad \langle f, h \rangle_{H^p} = \sum_{q=0}^p \sum_{j=0}^{\infty} \lambda_j^q \hat{f}_j^* \hat{h}_j, \quad \hat{f}_j = \langle \phi_j, f \rangle_{L^2}, \quad \hat{h}_j = \langle \phi_j, h \rangle_{L^2}.$$

Among these, the space H^2 is precisely the domain of the self-adjoint Laplacian $\bar{\Delta}$.

We equip each H^p space with a Dirichlet form $E_p : H^p \times H^p \rightarrow \mathbb{C}$, defined as the bounded sesquilinear form

$$(3.4) \quad E_p(f, h) = \sum_{j=0}^{\infty} \lambda_j^p \hat{f}_j^* \hat{h}_j,$$

with f and h as in (3.3). This form induces a positive-semidefinite Dirichlet energy functional $E_p(f) = E_p(f, f)$. Given $f \in H^p \cap H$, the quantity $E_p(f)/\|f\|_H^2$ can be thought of as a measure of roughness of f . If f and h are smooth, $E_p(f, h)$ can be expressed in terms of the Laplace-Beltrami operator as $E_p(f, h) = \langle f, \Delta^p h \rangle_H$. Evidently, $E_p(f) = \langle f, \Delta^p f \rangle_H$ can be arbitrarily large for highly oscillatory functions.

In general, the $\{\phi_j\}$ orthonormal basis of H is not a Riesz basis of H^p , $p \geq 1$; that is, it is not the case that given any ℓ^2 sequence of expansion coefficients c_j , the vectors $f_l = \sum_{j=0}^{l-1} c_j \phi_j$ converge as $l \rightarrow \infty$ in H^p norm. This issue is manifested from the fact that the Dirichlet energies $E_p(\phi_j) = \lambda_j^p$ of the basis elements are unbounded in j , making $\{\phi_j\}$ a poorly conditioned basis of H^p for numerical calculations. On the other hand, the normalized eigenfunctions $\phi_j^{(p)}$, defined by

$$(3.5) \quad \phi_j^{(p)} = \frac{\phi_j}{\|\phi_j\|_{H^p}},$$

where $\|\phi_j^{(p)}\|_{H^p} = 1$ by construction, form orthonormal bases of the respective H^p spaces.

A related, but stronger, notion of regularity of functions on $\mathcal{M}$ to that associated with the H^p Sobolev spaces is provided by reproducing kernel Hilbert spaces (RKHSs) associated with the heat kernel on $(\mathcal{M}, g)$. In particular, let $\mathcal{H}$ be the RKHS of complex-valued functions on $\mathcal{M}$ associated with the time 1 heat kernel. We denote the inner product and norm by $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and $\|\cdot\|_{\mathcal{H}}$, respectively. A natural orthonormal basis of $\mathcal{H}$ consists of the exponentially scaled Laplace-Beltrami eigenfunctions (cf. (3.5))

$$(3.6) \quad \tilde{\phi}_j = \frac{\phi_j}{\|\phi_j\|_{\mathcal{H}}} = e^{-\lambda_j/2} \phi_j, \quad j \in \{0, 1, \dots\}.$$

After inclusion (which can be shown to be compact), $\mathcal{H}$ is a dense subspace of H consisting of all equivalence classes of functions $\sum_{j=0}^{\infty} c_j \phi_j$ satisfying the inequality $\sum_{j=0}^{\infty} e^{\lambda_j} |c_j|^2 < \infty$. In addition, $\mathcal{H}$ can be compactly embedded into every Sobolev space H^p , $p > 0$. $\mathcal{H}$ is also a dense subspace of $C^k(\mathcal{M})$ for all $k \geq 0$ [28].

In what follows, we will also be interested in spaces of bounded operators on the function spaces introduced above. Given two Banach spaces V_1 and V_2 , $\mathcal{B}(V_1, V_2)$ will denote the Banach space of bounded operators mapping V_1 to V_2 , equipped with the operator norm, $\|\cdot\|$. If V_1 and V_2 are Hilbert spaces, $\mathcal{B}_2(V_1, V_2) \subseteq \mathcal{B}(V_1, V_2)$ will denote the Hilbert space of Hilbert-Schmidt operators from V_1 to V_2 , equipped with the inner product $\langle A, B \rangle_{\text{HS}} = \text{tr}(A^* B)$ and the corresponding norm, $\|A\|_{\text{HS}} = \sqrt{\langle A, A \rangle_{\text{HS}}}$. We will also use the abbreviations $\mathcal{B}(V_1) = \mathcal{B}(V_1, V_1)$ and $\mathcal{B}_2(V_1) = \mathcal{B}_2(V_1, V_1)$.

3.2 Spaces of vector fields

We consider the space $\mathfrak{X}$ of C^∞ complex vector fields on $\mathcal{M}$ (that is, the space of derivations on the ring of smooth, complex-valued functions on $\mathcal{M}$, or, equivalently, the space of smooth sections of $T^\mathbb{C}\mathcal{M}$), where we recall that $\mathfrak{X}$ can be viewed either as a vector space over the field of complex numbers, or as a $C^\infty(\mathcal{M})$ -module. In the former case, it can be endowed with the structure of a Lie algebra with the vector field commutator, $[\cdot, \cdot] : \mathfrak{X} \times \mathfrak{X} \rightarrow \mathfrak{X}$ acting as the algebraic product. We denote the gradient and divergence operators associated with g by $\text{grad} : C^\infty(\mathcal{M}) \rightarrow \mathfrak{X}$ and $\text{div} : \mathfrak{X} \rightarrow C^\infty(\mathcal{M})$. Note that these operators are related to the positive-semidefinite Laplacian via $\Delta = -\text{div} \circ \text{grad}$. As with functions, we consider the standard Banach spaces $L_{\mathfrak{X}}^p$, $1 \leq p \leq \infty$, of vector fields associated with the norms

$$\|v\|_{L_{\mathfrak{X}}^p} = \left(\int_{\mathcal{M}} (g(v, v))^{p/2} d\mu \right)^{1/p}, \quad 1 \leq p < \infty, \quad \text{and} \quad \|v\|_{L_{\mathfrak{X}}^\infty} = \text{ess sup}_{x \in \mathcal{M}} \sqrt{g(v, v)_x}.$$

In the case $p = 2$, we set $H_{\mathfrak{X}} = L_{\mathfrak{X}}^2$ and use the notation $\langle v, w \rangle_{H_{\mathfrak{X}}} = \int_{\mathcal{M}} g(v^*, w) d\mu$ for the Hodge inner product inducing the $L_{\mathfrak{X}}^2$ norm.

With these definitions, the closure $\overline{\text{grad}} : D(\overline{\text{grad}}) \rightarrow H_{\mathfrak{X}}$ of the gradient operator has domain $D(\overline{\text{grad}}) = H^1$, and is bounded as an operator from H^1 to $H_{\mathfrak{X}}$. Similarly, the closure $\overline{\text{div}} : D(\overline{\text{div}}) \rightarrow H_{\mathfrak{X}}$ of the divergence operator has as its domain $D(\overline{\text{div}})$ the Sobolev space $H_{\mathfrak{X}, \text{div}} \subset H_{\mathfrak{X}}$, defined as the closure of $\mathfrak{X}$ with respect to the norm $\|v\|_{H_{\mathfrak{X}, \text{div}}} = \langle v, v \rangle_{H_{\mathfrak{X}, \text{div}}}^{1/2}$ induced by the inner product

$$\langle v, w \rangle_{H_{\mathfrak{X}, \text{div}}} = \langle v, w \rangle_{H_{\mathfrak{X}}} + \langle \overline{\text{div}} v, \overline{\text{div}} w \rangle_{H^p}.$$

Also, for $p \geq 0$, we introduce the Sobolev spaces

$$H_{\mathfrak{X}, \text{div}}^p = \{v \in H_{\mathfrak{X}} : \overline{\text{div}} v \in H^p\},$$

which are equipped with the inner products

$$\langle v, w \rangle_{H_{\mathfrak{X}, \text{div}}^p} = \langle v, w \rangle_{H_{\mathfrak{X}}} + \langle \overline{\text{div}} v, \overline{\text{div}} w \rangle_{H^p},$$

and the corresponding norms $\|v\|_{H_{\mathfrak{X}, \text{div}}^p} = \langle v, v \rangle_{H_{\mathfrak{X}, \text{div}}^p}^{1/2}$. As in the case of functions, we define the Dirichlet forms $E_{p, \mathfrak{X}, \text{div}} : H_{\mathfrak{X}, \text{div}}^p \times H_{\mathfrak{X}, \text{div}}^p \rightarrow \mathbb{C}$ by

$$E_{p, \mathfrak{X}, \text{div}}(u, v) = E_p(\overline{\text{div}} u, \overline{\text{div}} v).$$

The corresponding energy functionals, $E_{p, \mathfrak{X}, \text{div}}(f) = E_{p, \mathfrak{X}, \text{div}}(f, f)$ assign measures of roughness to vector fields in $H_{\mathfrak{X}, \text{div}}^p \cap H_{\mathfrak{X}}$ by $E_{p, \mathfrak{X}, \text{div}}(f) / \|f\|_{H_{\mathfrak{X}}}^2$.

An important subspace of $H_{\mathfrak{X}}$ is the closed subspace of gradient vector fields, $H_{\mathfrak{X}, \text{grad}} = \text{ran}(\overline{\text{grad}})$. This leads to the orthogonal decomposition $H_{\mathfrak{X}} = H_{\mathfrak{X}, \text{grad}} \oplus H_{\mathfrak{X}, \text{grad}}^\perp$, and it can be readily checked that any vector field in $H_{\mathfrak{X}, \text{grad}}^\perp \cap H_{\mathfrak{X}, \text{div}}$ has

vanishing divergence. A natural smooth orthonormal basis $\{u_j\}_{j=1}^\infty$ for $H_{\mathfrak{X},\text{grad}}$ is given by the normalized gradients of the Laplace-Beltrami eigenfunctions,

$$(3.7) \quad u_j = \text{grad } \phi_j / \lambda_j^{1/2}.$$

The following two lemmas characterize the behavior of vector fields as operators on functions.

Lemma 3.1 (vector fields as conjugate-antisymmetric operators). *To every vector field $v \in H_{\mathfrak{X},\text{div}}^1$ there corresponds a unique operator $A_v : C^\infty(\mathcal{M}) \rightarrow H$ with the property*

$$(3.8) \quad \langle f, A_v h \rangle_H = -\langle A_v^* f, h \rangle_H, \quad \forall f, h \in C^\infty(\mathcal{M}).$$

This operator is given by $A_v = v + D_v$, where $D_v : C^\infty(\mathcal{M}) \rightarrow H$ is defined as $D_v f = \text{div}(fv)$, and we also have

$$(3.9) \quad v(f) = \frac{A_v f - f A_v 1}{2}.$$

Proof. The Leibniz rule for the divergence, $\text{div}(fv) = v(f) + f \text{div } v$, and the fact that $\int_{\mathcal{M}} \text{div}(fv) d\mu$ vanishes on closed manifolds lead to

$$\langle f, v(h) \rangle_H = -\langle D_v^* f, h \rangle_H.$$

The claim in (3.8) follows from the definition of A_v and the last equation. Note that the restriction $v \in H_{\mathfrak{X},\text{div}}^1$ is important in order for D_v and A_v to be well defined. To show that A_v is unique, suppose that $u \in H_{\mathfrak{X},\text{div}}^1$, $u \neq v$, and consider $(A_u - A_v)f = 2(u - v)f + f \text{div}(u - v)$. If $\text{div } u = \text{div } v$, then $(A_u - A_v)f = 2(u - v)f$, which is non-vanishing for some $f \in C^\infty(\mathcal{M})$. On the other hand, if $\text{div } u \neq \text{div } v$, we have $(A_u - A_v)1 = \text{div}(u - v)$, which is again non-vanishing. Equation (3.9) follows from the definition of A_v and the fact that $A_v 1 = \text{div } v$. $\square$

Lemma 3.2 (vector fields as bounded operators). *Let v be a bounded vector field in $L_{\mathfrak{X}}^\infty$. Then:*

- (i) *v extends uniquely to a bounded operator $L_v \in \mathcal{B}(H^1, H)$ with operator norm $\|L_v\| \leq \|v\|_{L_{\mathfrak{X}}^\infty}$.*
- (ii) *The restriction of v to $\mathcal{H}$ is a Hilbert-Schmidt operator $\tilde{L}_v \in \mathcal{B}_2(\mathcal{H}, H)$ with norm $\|\tilde{L}_v\|_{\text{HS}} \leq C\|v\|_{L_{\mathfrak{X}}^\infty}$, where C is a constant independent of v .*

As a result, the maps $\iota : L_{\mathfrak{X}}^\infty \rightarrow \mathcal{B}(H^1, H)$ and $\iota_2 : L_{\mathfrak{X}}^\infty \rightarrow \mathcal{B}_2(\mathcal{H}, H)$ with $\iota v = L_v$ and $\iota_2 v = \tilde{L}_v$ are continuous embeddings.

Proof. (i) Consider a vector field $v \in L_{\mathfrak{X}}^\infty$, and let f be a C^∞ function. Then, we have

$$\begin{aligned} \|v(f)\|_H^2 &= \int_{\mathcal{M}} |g(v, \text{grad } f)|^2 d\mu \leq \int_{\mathcal{M}} g(v, v) g(\text{grad } f, \text{grad } f) d\mu \\ &\leq \|v\|_{L_{\mathfrak{X}}^\infty}^2 \|\text{grad } f\|_{H_{\mathfrak{X}}}^2 \leq \|v\|_{L_{\mathfrak{X}}^\infty}^2 \|f\|_{H^1}^2, \end{aligned}$$

where the first inequality in the above follows from the Cauchy-Schwarz inequality. Thus, v is a densely-defined, bounded operator from $C^\infty(\mathcal{M}) \cap H^1$ to H , and can be uniquely extended to $L_v \in \mathcal{B}(H^1, H)$ by the bounded linear transformation theorem. The fact that $\|L_v\| \leq \|v\|_{L^\infty_{\mathcal{X}}}$ follows from the inequality $\|v(f)\|_H / \|f\|_{H^1} \leq \|v\|_{L^\infty_{\mathcal{X}}}$.

(ii) Since $\mathcal{H} \subset C^\infty(\mathcal{M})$, proceeding as above we find that for any $f \in \mathcal{H}$, $\|\tilde{L}_v f\|_H \leq \|v\|_{L^\infty_{\mathcal{X}}} \|f\|_{H^1}$. Moreover, since $\mathcal{H}$ continuously embeds into H^1 , there exists a constant $\tilde{C}$, independent of f , such that $\|f\|_{H^1} \leq \tilde{C} \|f\|_{\mathcal{H}}$. This shows that $\tilde{L}_v$ lies in $\mathcal{B}(\mathcal{H}, H)$. To establish that $\tilde{L}_v$ lies in $\mathcal{B}_2(\mathcal{H}, H)$, we compute

$$\begin{aligned} \sum_{j=0}^{\infty} \langle \tilde{\phi}_j, \tilde{L}_v^* \tilde{L}_v \tilde{\phi}_j \rangle_{\mathcal{H}} &= \sum_{j=0}^{\infty} \langle \tilde{L}_v \tilde{\phi}_j, \tilde{L}_v \tilde{\phi}_j \rangle_H = \sum_{j=0}^{\infty} \|\tilde{L}_v \tilde{\phi}_j\|_H^2 \leq \|v\|_{L^\infty_{\mathcal{X}}}^2 \sum_{j=0}^{\infty} \|\tilde{\phi}_j\|_{H^1}^2 \\ &= \|v\|_{L^\infty_{\mathcal{X}}}^2 \sum_{j=0}^{\infty} e^{-\lambda_j} (1 + \lambda_j), \end{aligned}$$

where $\{\tilde{\phi}_j\}_{j=0}^{\infty}$ is the orthonormal basis of $\mathcal{H}$ from (3.6). It then follows from the Weyl estimate for Laplacian eigenvalues (see (4.17) ahead and the proofs of Theorems 4.3 and 4.4 in Section 4.2) that the quantity $C^2 = \sum_{j=0}^{\infty} e^{-\lambda_j} (1 + \lambda_j)$ is finite, and we conclude that

$$\|\tilde{L}_v\|_{\text{HS}}^2 = \text{tr}(\tilde{L}_v^* \tilde{L}_v) = \sum_{j=0}^{\infty} \langle \tilde{\phi}_j, \tilde{L}_v^* \tilde{L}_v \tilde{\phi}_j \rangle_{\mathcal{H}} \leq C^2 \|v\|_{L^\infty_{\mathcal{X}}}^2,$$

as claimed. $\square$

An implication of Lemma 3.2 is that for any $v \in L^\infty_{\mathcal{X}}$ and every sequence $f_n \in C^\infty(\mathcal{M})$, converging to f in H^1 norm, $v(f) = \lim_{n \rightarrow \infty} v(f_n)$ even though v is unbounded (and therefore discontinuous) on $C^\infty(\mathcal{M})$. It also follows from Lemma 3.2 that the operator A_v in Lemma 3.1 associated with $v \in H^1_{\mathcal{X}, \text{div}} \cap L^\infty_{\mathcal{X}}$ also extends uniquely to a bounded operator $\tilde{A}_v : H^1 \rightarrow H$.

Next, as discussed in Section 2.2, we introduce a spectral representation of pointwise Riemannian inner products between gradient vector fields. For that, we first consider the product rule for the positive-definite Laplacian on smooth functions,

$$(3.10) \quad \Delta(fh) = (\Delta f)h + f(\Delta h) - 2g(\text{grad } f^*, \text{grad } h), \quad f, h \in C^\infty(\mathcal{M}).$$

It follows by definition of the H^2 norms that the self-adjoint Laplacian $\bar{\Delta}$ is bounded as an operator from H^2 to H . As a result, given a sequence $f_j \in C^\infty(\mathcal{M})$ converging to f in H^2 norm, we have

$$(3.11) \quad \bar{\Delta}f = \bar{\Delta} \left(\lim_{j \rightarrow \infty} f_j \right) = \lim_{j \rightarrow \infty} \bar{\Delta}f_j.$$

Now, the fact that the Laplace-Beltrami eigenfunctions ϕ_j are smooth implies that given any $f \in C^\infty(\mathcal{M})$, the sequence $f_0, f_1, \dots$ with

$$f_j = \sum_{i=0}^j \hat{f}_i \phi_i, \quad \hat{f}_i = \langle \phi_i, f \rangle_H$$

is Cauchy in H^p for all $p \geq 0$, and hence (3.11) holds. As a result, we can use (3.10) in conjunction with (3.1) to obtain:

Lemma 3.3 (spectral representation of Riemannian inner products). *The Riemannian inner product between the gradient vector fields associated with two functions $f, h \in C^\infty(\mathcal{M})$ can be expressed as*

$$g(\text{grad } f, \text{grad } h) = \frac{1}{2} \sum_{j,k,l=0}^{\infty} \hat{f}_j \hat{h}_k (\lambda_j + \lambda_k - \lambda_l) c_{jkl} \phi_l, \quad \hat{f}_j = \langle \phi_j, f \rangle_H,$$

where $\hat{h}_k = \langle \phi_k, h \rangle_H$, and the sum over l in the right-hand side converges in H norm.

Note that Lemma 3.3 can be extended to $f, h \in H^1$, which follows from the fact that the map $(f, h) \mapsto g(\text{grad } f, \text{grad } h)$ is a bounded linear map with a dense domain $C^\infty(\mathcal{M}) \times C^\infty(\mathcal{M}) \subset H^1 \times H^1$.

3.3 Spaces of differential forms

We will use the symbols $\Lambda_x^k \mathcal{M}$, $\Lambda^k \mathcal{M}$, and Ω^k to represent the vector space of complex k -forms at $x \in \mathcal{M}$, the associated k -form bundle, and the space of smooth k -form fields on $\mathcal{M}$ (totally antisymmetric, k -multilinear maps on $\mathfrak{X}^k$, taking values in $C^\infty(\mathcal{M})$, or, equivalently smooth sections of $\Lambda^k \mathcal{M}$). As with vector fields, the spaces Ω^k can be viewed either as vector spaces over $\mathbb{C}$, or as $C^\infty(\mathcal{M})$ -modules. As usual, we identify Ω^0 with $C^\infty(\mathcal{M})$. We also let $\eta_x^k : \Lambda_x^k \mathcal{M} \times \Lambda_x^k \mathcal{M} \rightarrow \mathbb{R}$ be the canonical metric tensor on Λ_x^k , satisfying

$$\eta_x^k(\alpha^1 \wedge \dots \wedge \alpha^k, \beta^1 \wedge \dots \wedge \beta^k) = \det[\eta_x(\alpha^i, \beta^j)]_{ij}, \quad \forall \alpha^i, \beta^j \in T_x^{\mathbb{C}*} \mathcal{M}.$$

The metric induces a Hodge star operator $\star : \Lambda_x^k \mathcal{M} \rightarrow \Lambda_x^{d-k} \mathcal{M}$, defined uniquely through the requirement that

$$\alpha \wedge \star \beta = \eta_x^k(\alpha, \beta) \mu, \quad \forall \alpha, \beta \in \Lambda_x^k \mathcal{M}.$$

The Hodge star has the useful property

$$(3.12) \quad \star \star \alpha = (-1)^{k(\dim \mathcal{M} - k)} \alpha, \quad \forall \alpha \in \Lambda_x^k \mathcal{M}.$$

As in the case of vector fields, we introduce the Banach spaces L_k^p , $1 \leq p \leq \infty$, defined as the completion of Ω^k with respect to the norms

$$\|\alpha\|_{L_k^p} = \left(\int_{\mathcal{M}} \eta^k(\alpha, \alpha)^{p/2} d\mu \right)^{1/p}, \quad 1 \leq p \leq \infty,$$

$$\|\alpha\|_{L_k^\infty} = \operatorname{ess\,sup}_{x \in \mathcal{M}} \sqrt{\eta^k(\alpha, \alpha)_x}.$$

The case $p = 2$ is a Hilbert space, $H_k = L_k^2$, with norm $\|\cdot\|_{H_k} = \|\cdot\|_{L_k^2}$ induced from the inner product

$$\langle \alpha, \beta \rangle_{H_k} = \int_{\mathcal{M}} \alpha \wedge \star \beta = \int_{\mathcal{M}} \eta^k(\alpha, \beta) d\mu.$$

A fundamental aspect of the spaces Ω^k is that they are linked by the exterior derivative and codifferential operators, $d_k : \Omega^k \rightarrow \Omega^{k+1}$ and $\delta_k : \Omega^k \rightarrow \Omega^{k-1}$, respectively. We recall that $d_0, \dots, d_{d-1}$ are the unique linear maps with the properties:

- (1) d_0 is the differential of functions.
- (2) $d_{k+1}d_k = 0$.
- (3) The Leibniz rule,

$$(3.13) \quad d_{k+l}(\alpha \wedge \beta) = d_k \alpha \wedge \beta + (-1)^k \alpha \wedge d_l \beta,$$

holds for all $\alpha \in \Omega^k$ and $\beta \in \Omega^l$.

The codifferential operators are defined uniquely through the requirement that

$$\langle \alpha, \delta_k \beta \rangle_{H_{k-1}} = \langle d_{k-1} \alpha, \beta \rangle_{H_k};$$

i.e., δ_k is a formal adjoint of d_{k-1} . This definition of δ_k is equivalent to

$$(3.14) \quad \delta_k \alpha = (-1)^{\dim \mathcal{M} (k+1) + 1} \star d_{d-k} \star \alpha,$$

and it also implies $\delta_k \delta_{k-1} = 0$. In the case $k = 1$, the codifferential operator is related with the divergence on vector fields via $\operatorname{div} v = -\delta_1 v^\flat$, $v \in \mathfrak{X}$. Note that despite its relationship with the exterior derivative in (3.14), the codifferential does not satisfy a Leibniz rule.

Another important class of operators on differential forms are the interior product and Lie derivative associated with vector fields. Given a vector field $v \in \mathfrak{X}$, these are defined as the maps $\iota_v : \Omega^k \rightarrow \Omega^{k-1}$ and $\mathcal{L}_v : \Omega^k \rightarrow \Omega^k$, respectively, such that

$$\iota_v \alpha = \alpha(v, \cdot), \quad \mathcal{L}_v = d\iota_v + \iota_v d.$$

Both ι_v and $\mathcal{L}_v$ satisfy Leibniz rules,

$$\iota_v(\alpha \wedge \beta) = (\iota_v \alpha) \wedge \beta + (-1)^k \alpha \wedge (\iota_v \beta), \quad \mathcal{L}_v(\alpha \wedge \beta) = (\mathcal{L}_v \alpha) \wedge \beta + \alpha \wedge (\mathcal{L}_v \beta),$$

for all $\alpha \in \Omega^k$ and $\beta \in \Omega^l$. Moreover, they have the properties

$$(3.15) \quad \mathcal{L}_v d = d \mathcal{L}_v, \quad v^\flat \wedge \star \alpha = (-1)^k \star \iota_v \alpha.$$

The following lemma can be viewed as a generalization of Lemma 3.2 to spaces of differential forms.

Lemma 3.4. *For every $v \in \mathfrak{X}$, the operators $\iota_v d$, $d\iota_v$, and $\mathcal{L}_v$ extend to unique bounded operators $D_v : H_k^1 \rightarrow H_k$, $\tilde{D}_v : H_k^1 \rightarrow H_k$, and $L_v : H_k^1 \rightarrow H_k$, respectively.*

Proof. First, establish that $\|t_\nu d\omega\|_{H_k}^2 / \|\omega\|_{H_k}^2$ is bounded using local Cauchy-Schwartz inequalities as in the proof of Lemma 3.2. This, in conjunction with the bounded linear transformation theorem implies the existence and uniqueness of D_ν , as claimed. The results for $\tilde{D}_\nu$ and L_ν follow similarly. $\square$

The exterior derivative and codifferential operators lead to the Hodge Laplacian $\Delta_k : \Omega^k \rightarrow \Omega^k$ on k -forms, defined as

$$\Delta_k = \delta_{k+1} d_k + d_{k-1} \delta_k.$$

As with the Laplacian $\Delta = \Delta_0$ on functions, the Laplacian on k -forms on closed manifolds has a unique self-adjoint extension $\bar{\Delta}_k : D(\bar{\Delta}_k) \rightarrow H_k$, with a pure point spectrum of eigenvalues $0 = \lambda_{0,k} \leq \lambda_{1,k} \leq \dots$ with no accumulation points and an associated smooth orthonormal basis $\{\phi_{j,k}\}_{j=0}^\infty$ of eigenforms [47, 39].

A central result in exterior calculus on manifolds is the Hodge decomposition theorem, which states that Ω^k admits the decomposition

$$(3.16) \quad \Omega^k = \text{ran } d_{k-1} \oplus \text{ran } \delta_{k+1} \oplus \mathcal{H}_0^k, \quad \mathcal{H}_0^k = \ker \Delta_k,$$

into subspaces of closed ($\text{ran } d_{k-1}$), coclosed ($\text{ran } \delta_{k+1}$), and harmonic ($\mathcal{H}^k$) forms, all of which are invariant under Δ_k . On a compact manifold, $\mathcal{H}_0^k = \ker d_k \cap \ker \delta_k$, and the dimension of this space is finite. The Hodge decomposition in (3.16) has an L^2 extension,

$$H_k = \mathcal{H}_d^k \oplus \mathcal{H}_\delta^k \oplus \mathcal{H}_0^k, \quad \mathcal{H}_d^k = \overline{\text{ran } d_{k-1}}, \quad \mathcal{H}_\delta^k = \overline{\text{ran } \delta_{k+1}},$$

where the closed spaces $\mathcal{H}_d^k$, $\mathcal{H}_\delta^k$, and $\mathcal{H}^k$ are mutually orthogonal.

It follows directly from the definition of Δ_k that $d_k \Delta_k = \Delta_{k+1} d_k$ and $\delta_k \Delta_k = \Delta_{k-1} \delta_k$. This implies that every eigenform of Δ_k can be chosen to lie in one of the $\mathcal{H}_d^k$, $\mathcal{H}_\delta^k$, or $\mathcal{H}_0^k$ subspaces. Moreover, for every k -eigenform $\psi \in \mathcal{H}_d^k$ there exists a $(k-1)$ -eigenform $\phi \in \mathcal{H}_\delta^k$ such that $\psi = d_{k-1} \phi$, and similarly for every $\psi \in \mathcal{H}_\delta^k$ there exists as $(k+1)$ -eigenform $\omega \in \mathcal{H}_d^{k+1}$ such that $\psi = \delta_{k+1} \omega$.

Besides providing orthonormal bases for the invariant subspaces in the Hodge decomposition of H_k , the eigenfunctions of Δ_k and the corresponding eigenvalues are also useful for constructing Sobolev spaces analogous to the H^p function spaces in (3.2). Given $p \geq 0$, we define

$$H_k^p = \left\{ \sum_{j=0}^\infty c_j \phi_{j,k} \in H_k : \sum_{j=0}^\infty \lambda_{j,k}^p |c_j|^2 < \infty \right\}.$$

These spaces are Hilbert spaces with inner products

$$(3.17) \quad \langle \alpha, \beta \rangle_{H_k^p} = \sum_{q=0}^p \sum_{j=0}^\infty \lambda_{j,k}^q \hat{\alpha}_j^* \hat{\beta}_j, \quad \hat{\alpha}_j = \langle \phi_{j,k}, \alpha \rangle_{L_k^2}, \quad \hat{\beta}_j = \langle \phi_{j,k}, \beta \rangle_{L_k^2}$$

and norms $\|\alpha\|_{H_k^p} = \langle \alpha, \alpha \rangle_{H_k^p}^{1/2}$. As in the case of functions and vector fields, we equip these spaces with positive-semidefinite Dirichlet forms $E_{p,k} : H_k^p \times H_k^p \rightarrow \mathbb{C}$,

given by

$$E_{p,k} = \sum_{j=0}^{\infty} \lambda_{j,k}^p \hat{\alpha}_j^* \beta_j$$

for α, β from (3.17). Note that $E_{p,k}$ can be equivalently expressed using the exterior derivative and codifferential operators; in particular,

$$(3.18) \quad E_{1,k}(\alpha, \beta) = \langle \bar{d}_k \alpha, \bar{d}_k \beta \rangle_{H_{k+1}} + \langle \bar{\delta}_k \alpha, \bar{\delta}_k \beta \rangle_{H_{k-1}},$$

where overbars denote operator closures. Moreover, if β is smooth, $E_{p,k}(\alpha, \beta)$ can be expressed in terms of the k -Laplacian via

$$E_{p,k}(\alpha, \beta) = \langle \alpha, \Delta_k^p \beta \rangle_{H_k},$$

with an analogous expression holding if $\alpha \in \Omega^k$. The Dirichlet forms defined above induce the energy functionals $E_{p,k}(f) = E_{p,k}(f, f)$ measuring the roughness of forms in $H_k^p \cap H_k$ through $E_{p,k}(f) / \|f\|_{H_k}^2$.

For notational simplicity, henceforth we will drop the overbars from our notation for the closed differential, codifferential, and Laplacian on k -forms. We will also drop k superscripts and subscripts from η^k , d_k , δ_k , and Δ_k .

4 Spectral exterior calculus (SEC) on smooth manifolds

In this section, we introduce our representation of vector fields and forms using frames constructed from Laplace-Beltrami eigenfunctions and their derivatives. Besides rigorously satisfying the appropriate frame conditions for a number of L^2 and Sobolev spaces of interest in exterior calculus, an advantage of this representation is that it is fully spectral, and thus can also be applied in the discrete case with little modification.

4.1 Frame representation of vector fields and forms

We begin by recalling the definition of a frame of a Hilbert space [36].

Definition 4.1 (frame of a Hilbert space). Let $(V, \langle \cdot, \cdot \rangle_V)$ be a Hilbert space over the complex numbers and $u_0, u_1, \dots$ a sequence of elements $u_k \in V$. We say that the set $\{u_k\}$ is a frame if there exist positive constants C_1 and C_2 such that the following frame conditions hold for all $v \in V$:

$$(4.1) \quad C_1 \|v\|_V^2 \leq \sum_k |\langle u_k, v \rangle_V|^2 \leq C_2 \|v\|_V^2.$$

The frame $\{u_k\}$ induces a linear operator $T : V \rightarrow \ell^2$, called analysis operator, such that $Tv = \hat{v}' = (\hat{v}'_k)_k$ with $\hat{v}'_k = \langle u_k, v \rangle_V$. This operator is bounded above and below via the same constants as in (4.1); that is,

$$C_1 \|v\|_V^2 \leq \|Tv\|_{\ell^2}^2 \leq C_2 \|v\|_V^2.$$

The adjoint, $T^* : \ell^2 \rightarrow V$, is called synthesis operator.

The analysis and synthesis operators induce a positive-definite, self-adjoint, bounded operator $S : V \rightarrow V$ with bounded inverse, called frame operator, which is given by $S = T^*T$. This operator satisfies the bounds

$$C_1 \|v\|_V^2 \leq \langle v, Sv \rangle_V \leq C_2 \|v\|_V^2.$$

The fact that S has bounded inverse implies that the set $\{u'_k\}$ with $u'_k = S^{-1}u_k$ is also a frame, called dual frame. This frame has the important property

$$(4.2) \quad v = \sum_k \langle u'_k, v \rangle_V u_k = \sum_k \langle u_k, v \rangle_V u'_k, \quad \forall v \in V.$$

This means that the inner products $\langle u'_k, v \rangle_V$ between v and the dual frame elements correspond to expansion coefficients in the original frame that reconstruct v , and conversely, the coefficients $\langle u_k, v \rangle_V$ reconstruct v in the dual frame. Denoting the analysis operator associated with the dual frame by $T' : V \rightarrow \ell^2$, we have $T' = TS^{-1}$, $T'^* = S^{-1}T^*$, and (4.2) can be equivalently expressed as

$$(4.3) \quad v = T^*T'v = T'^*Tv, \quad \forall v \in V.$$

Moreover, the dual frame operator $S' : V \rightarrow V$, $S' = T'^*T'$, is equal to S^{-1} .

Another class of operators of interest in frame theory are the Gramm operators $G : \ell^2 \rightarrow \ell^2$, $G = TT^*$, and $G' : \ell^2 \rightarrow \ell^2$, $G' = T'T'^*$, associated with the frame and dual frame, respectively. While G and G' are both bounded, unlike S and S' , they are non-invertible if the frame has linearly dependent elements. Nevertheless, it follows from (4.3) that

$$(4.4) \quad Tv = GT'v, \quad T'v = G'Tv, \quad \forall v \in V,$$

which implies that G (resp. G') is invertible on the range of T' (resp. T), and its inverse is given by G' (resp. G). Denoting the canonical orthonormal basis of ℓ^2 by $\{e_k\}_{k=0}^\infty$, we have

$$G_{ij} := \langle e_i, Ge_j \rangle_{\ell^2} = \langle T^*e_i, T^*e_j \rangle_V = \langle u_i, u_j \rangle_V,$$

and similarly $G'_{ij} = \langle e_i, G'e_j \rangle_{\ell^2} = \langle u'_i, u'_j \rangle_V$. Thus, the matrix elements of the Gramm operators G and G' in the $\{e_k\}_{k=0}^\infty$ basis are equal to the pairwise inner products between the frame and dual frame elements, respectively. By boundedness of these operators, for any $\hat{f} = \sum_{j=0}^\infty \hat{f}_j e_j \in \ell^2$, we have $G\hat{f} = \sum_{i,j=0}^\infty e_i G_{ij} \hat{f}_j$ and $G'\hat{f} = \sum_{i,j=0}^\infty e_i G'_{ij} \hat{f}_j$.

Since the analysis operator T is bounded below, it follows by the closed range theorem that T and T^* have closed range, and as a result the ranges of S and G are also closed. Similarly, all of T' , T'^* , S' , and G' have closed range. An important consequence of these properties is that all of these operators have well-defined pseudoinverses. In particular, it can be shown [16] that the pseudoinverse $(T^*)^+$ of T^* is equal to the dual analysis operator, $(T^*)^+ = T'$, and similarly $(T'^*)^+ = T$. These relationships imply in turn that $G^+ = G'$ and $G'^+ = G$.

Clearly, in a separable Hilbert space, every Riesz basis is also a frame. For example, in the case $V = H = L^2(\mathcal{M}, \mu)$, a natural frame is provided by the Laplace-Beltrami eigenfunction basis $\{\phi_k\}_{k=0}^\infty$. In that case, the analysis operator T is unitary, $T^*T = S = I_H$, $TT^* = G = I_{\ell^2}$, by orthonormality of the basis. In the setting of vector fields on manifolds, a natural orthonormal set of smooth fields in $H_{\mathfrak{X}}$ is given by the normalized gradient fields u_j from (3.7). However, this only provides a basis for the space of gradient fields, $H_{\mathfrak{X}, \text{grad}}$. To construct a representation of arbitrary vector fields in $H_{\mathfrak{X}}$, we can take advantage of the $C^\infty(\mathcal{M})$ -module structure of smooth vector fields to augment this set by multiplication of gradient fields by smooth functions. Doing so will result in an overcomplete spanning set of $H_{\mathfrak{X}}$, which will turn out to meet the frame conditions in Definition 4.1. We will follow a similar approach to construct frames for the H_k spaces of differential forms, where we will also construct frames for higher-order Sobolev spaces through eigenvalue-dependent normalizations of the frame elements as in (3.5).

Remark 4.2. As alluded to in Section 2, a key property of the frames for spaces of vector fields and differential forms introduced below is that the matrix elements G_{ij} of the corresponding Gramm operators can be evaluated via closed form expressions that depend only on the eigenvalues λ_i of the Laplacian on functions and the corresponding coefficients c_{ijk} from (3.1). This allows in turn the SEC to be built entirely from the spectral properties of the Laplacian on functions.

We begin by introducing the vector fields and forms which will be employed in our frame construction and Galerkin schemes below. In the case of vector fields, we define

$$(4.5) \quad b_{ij} = \phi_i \text{grad } \phi_j, \quad \tilde{b}_{ij} = e^{-\lambda_j/2} b_{ij}, \quad b_{ij}^{(p)} = b_{ij} / \|\phi_i\|_{H^p}, \quad \tilde{b}_{ij}^{(p)} = \tilde{b}_{ij} / \|\phi_i\|_{H^p},$$

and

$$(4.6) \quad \hat{b}_{ij} = b_{ij} - b_{ji}, \quad \check{b}_{ij} = e^{-(\lambda_i + \lambda_j)/2} \hat{b}_{ij},$$

all of which are smooth vector fields in $\mathfrak{X}$. We also define the smooth forms

$$(4.7) \quad b^i = \phi_i \in \Omega^0, \\ b^{ij_1 \cdots j_k} = b^i db^{j_1} \wedge \cdots \wedge db^{j_k} \in \Omega^k, \quad \tilde{b}^{ij_1 \cdots j_k} = e^{-(\lambda_{j_1} + \cdots + \lambda_{j_k})/2} b^{ij_1 \cdots j_k}, \\ b_p^{ij_1 \cdots j_k} = b^{ij_1 \cdots j_k} / \|\phi_i\|_{H^p}, \quad \tilde{b}_p^{ij_1 \cdots j_k} = \tilde{b}^{ij_1 \cdots j_k} / \|\phi_i\|_{H^p},$$

and

$$(4.8) \quad \hat{b}^{ij_1 \cdots j_k} = b^{[ij_1 \cdots j_k]}, \quad \check{b}^{ij_1 \cdots j_k} = e^{-(\lambda_i + \lambda_{j_1} + \cdots + \lambda_{j_k})} \hat{b}^{ij_1 \cdots j_k},$$

where the square brackets $[ij_1 \cdots j_k]$ denote total antisymmetrization with respect to the enclosed indices; e.g.,

$$b^{[ij_1 j_2]} = b^{ij_1 j_2} - b^{ij_2 j_1} + b^{j_1 j_2 i} - b^{j_1 i j_2} + b^{j_2 i j_1} - b^{j_2 j_1 i}.$$

With these definitions, our main results on frames for Hilbert spaces of vector fields and forms are as follows.

Theorem 4.3 (frames for L^2 spaces of vector fields and forms). *There exists an integer $J_0 \geq d = \dim \mathcal{M}$, such that for every integer $J \geq J_0$, the sets*

$$B_{\mathcal{X}}^J = \{b_{ij} : i \in \{0, 1, \dots\}, j \in \{1, \dots, J\}\}, \quad \tilde{B}_{\mathcal{X}} = \{\tilde{b}_{ij} : i \in \{0, 1, \dots\}, j \in \{1, 2, \dots\}\},$$

are frames for $H_{\mathcal{X}}$. Moreover, the sets

$$B_k^J = \{b^{ij_1 \dots j_k} : i \in \{0, 1, \dots\}, j_1, \dots, j_k \in \{1, \dots, J\}\}, \\ \tilde{B}_k = \{\tilde{b}^{ij_1 \dots j_k} : i \in \{0, 1, \dots\}, j_1, \dots, j_k \in \{1, 2, \dots\}\},$$

with $k \in \{1, \dots, m\}$, are frames for H_k .

Theorem 4.4 (frames for order-1 Sobolev spaces of 1-forms). *For the same integer J as in Theorem 4.3, the sets*

$$B_{1,1}^J = \{b_1^{ij} : i \in \{0, 1, \dots\}, j \in \{1, \dots, J\}\}, \quad \tilde{B}_{1,1} = \{\tilde{b}_1^{ij} : i \in \{0, 1, \dots\}, j \in \{1, 2, \dots\}\}.$$

are frames for H_1^1 .

We will prove Theorems 4.3 and 4.4 in Sections 4.2 and 4.3, respectively. In addition, while we have not explicitly verified this, it should be possible to show via a similar approach to that in Section 4.3 that frames for H_k^p , $k, p > 1$, can be constructed using $b_p^{ij_1 \dots j_k}$ or $\tilde{b}_p^{ij_1 \dots j_k}$. It may also be possible to establish such results inductively with respect to the Sobolev order p , using the results in B. Based on these considerations, we conjecture the following:

Conjecture 4.5 (frames for Sobolev spaces of forms). *For the same integer J as in Theorem 4.3, the following sets are frames for H_k^1 , $k \in \{1, \dots, m\}$, $p \geq 0$:*

$$B_{k,p}^J = \{b_p^{ij_1 \dots j_k} : i \in \{0, 1, \dots\}, j_1, \dots, j_k \in \{1, \dots, J\}\}, \\ \tilde{B}_{k,p} = \{\tilde{b}_p^{ij_1 \dots j_k} : i \in \{0, 1, \dots\}, j_1, \dots, j_k \in \{1, 2, \dots\}\}.$$

Remark 4.6. It should also be noted that we have not established frame conditions for the antisymmetric elements, $\hat{b}_{ij}$ and $\hat{b}^{ij_1 \dots j_k}$, or their rescaled counterparts, $\check{b}_{ij}$ and $\check{b}^{ij_1 \dots j_k}$. In fact, to fully span $H_{\mathcal{X}}$ and H_k using $\hat{b}_{ij}$ and $\hat{b}^{ij_1 \dots j_k}$, respectively, one would have to use infinitely many i and j indices, leading to violations of the upper frame condition. Nevertheless, as demonstrated by the formulas derived in B and listed in Table 2.2, due to cancellation of terms by antisymmetrization, $\hat{b}_{ij}$ and $\hat{b}^{ij_1 \dots j_k}$ can sometimes lead to considerable simplification of the representation of operators of interest in exterior calculus (e.g., the 1-Laplacian). Moreover, in the applications presented in Section 8 with available analytical results for the eigenvalues and eigenforms of the 1-Laplacian (i.e., the circle and flat torus), we found that SEC formulations based on the antisymmetric elements actually exhibit a moderate performance increase over those based on the non-symmetric elements. These facts motivate further exploration of the construction of frames based on antisymmetric elements. For example, the exponentially scaled $\check{b}^{ij_1 \dots j_k}$ might provide frames for reproducing kernel Hilbert spaces associated with the heat kernel on k -forms.

For the remainder of this section, we discuss the basic properties of the vector fields and forms just defined. We begin by establishing that, while they may not form a basis, finitely many gradient vector fields of Laplace-Beltrami eigenfunctions are sufficient to generate arbitrary smooth vector fields on closed manifolds.

Lemma 4.7. *There exists an integer $J \geq \dim \mathcal{M}$ such that $\{\text{grad } \phi_1|_x, \dots, \text{grad } \phi_J|_x\}$ is a spanning set of $T_x \mathcal{M}$, and thus $T_x^\mathbb{C} \mathcal{M}$, at every $x \in \mathcal{M}$.*

Proof. The claim follows from the fact that there exists an integer J such that the map $\vec{\phi}_J : \mathcal{M} \rightarrow \mathbb{R}^J$ with $\vec{\phi}_J(x) = (\phi_1(x), \dots, \phi_J(x))$ is an embedding of $\mathcal{M}$, which is proved in [38, Theorem 4.5]. Note that B_J being an embedding implies that $J \geq \dim \mathcal{M}$. To see why it implies that $\{\text{grad } \phi_1|_x, \dots, \text{grad } \phi_J|_x\}$ is a spanning set, fix a point $x \in \mathcal{M}$, and consider the tangent vector pushforward map $\vec{\phi}_{J*}|_x : T_x \mathcal{M} \rightarrow T_{\vec{\phi}_J(x)} \mathbb{R}^J$. Since $T_{\vec{\phi}_J(x)} \mathbb{R}^J$ is canonically isomorphic to $\mathbb{R}^J$, in a coordinate chart $u : N_x \rightarrow \mathbb{R}^d$ defined on a neighborhood N_x of x , the pushforward map is represented by a $d \times J$ matrix $\Xi(x)$, $d = \dim \mathcal{M}$, with elements $\Xi_{ij}(x) = \frac{\partial \phi_j}{\partial u_i}|_x$, and because $\vec{\phi}_J$ is an embedding, that matrix has full rank, $\text{rank } \Xi(x) = d$. In this coordinate basis, the components β_{ij} of $\text{grad } \phi_j|_x = \sum_{i=1}^d \beta_{ij} \frac{\partial}{\partial u_i}|_x$ are given by $\beta_{ij} = \sum_{k=1}^d \eta_{ik} \Xi_{kj}$, where η_{ik} are the components of the dual metric $\eta_x : T^* \mathcal{M}_x \times T^* \mathcal{M}_x \rightarrow \mathbb{R}$. The η_{ik} form a $d \times d$ invertible matrix, and thus the $d \times J$ matrix with elements β_{ij} has rank d . This implies that $\text{span}\{\text{grad } \phi_1|_x, \dots, \text{grad } \phi_J|_x\} = T_x \mathcal{M}$. $\square$

Corollary 4.8. *The set $\{\text{grad } \phi_1, \dots, \text{grad } \phi_J\}$ is a generating set for $\mathfrak{X}$ viewed as a $C^\infty(\mathcal{M})$ -module. That is, for every smooth vector field $v \in \mathfrak{X}$, there exist (not necessarily unique) smooth functions $f_1, \dots, f_J$ such that $v = \sum_{j=1}^J f_j \text{grad } \phi_j$.*

Corollary 4.9. *The collection of k -form fields $d\phi_{j_1} \wedge \dots \wedge d\phi_{j_k}$ with $j_1, \dots, j_k \in \{1, \dots, J\}$ spans Λ_x^k at every $x \in \mathcal{M}$. As a result, this set is a generating set for Ω^k , which means that for every $\omega \in \Omega^k$ there exist smooth functions $f_{j_1 \dots j_k}$ such that $\omega = \sum_{j_1, \dots, j_k=1}^J f_{j_1 \dots j_k} d\phi_{j_1} \wedge \dots \wedge d\phi_{j_k}$.*

It follows from Corollary 4.8 and the fact that $\mathfrak{X}$ is dense in $H_{\mathfrak{X}}$ that for every $v \in H_{\mathfrak{X}}$ there exist functions $f_1, \dots, f_J \in H$ such that $v = \sum_{j=1}^J f_j \text{grad } \phi_j$. Expanding these functions as $f_j = \sum_{i=0}^\infty c_{ij} \phi_j$ with $c_{ij} = \langle \phi_i, f_j \rangle_H$, we conclude that every vector field $v \in H_{\mathfrak{X}}$ is expressible in the form

$$(4.9) \quad v = \sum_{i=0}^\infty \sum_{j=1}^J c_{ij} b_{ij},$$

for some (not necessarily unique) constants $c_{ij} \in \mathbb{C}$. Similarly, Corollary 4.9 implies that every k -form field in H_k can be expanded as

$$(4.10) \quad \omega = \sum_{i=0}^\infty \sum_{j_1, \dots, j_k=1}^J c_{ij_1 \dots j_k} b^{ij_1 \dots j_k}, \quad c_{ij_1 \dots j_k} \in \mathbb{C}.$$

Example 4.10. As a simple example illustrating that $B_{\mathfrak{X}}^J$ may be a spanning set with linearly dependent elements (as opposed to a basis), suppose that $(\mathcal{M}, g)$ is the circle equipped with the canonical arclength metric, normalized such that $\mu(\mathcal{M}) = 1$. Then, an orthonormal basis of H consisting of Laplace-Beltrami eigenfunctions is given by

$$\phi_i(\theta) = \begin{cases} \cos(i\theta/2), & i \text{ even}, \\ \sin((i+1)\theta/2), & i \text{ odd}, \end{cases}$$

where $\theta \in (0, 2\pi)$ is a canonical angle coordinate. Note that the coordinate basis vector field $\frac{\partial}{\partial\theta}$ extends to a globally defined harmonic vector field on $\mathcal{M}$, satisfying $\Delta \frac{\partial}{\partial\theta} = 0$. In this coordinate system, the metric and dual metric are given by $g = g_{11} d\theta \otimes d\theta$, $\eta = g'^{11} \frac{\partial}{\partial\theta} \otimes \frac{\partial}{\partial\theta}$, where $g_{11} = 1/g'^{11} = 1/(2\pi)^2$, and we have

$$\begin{aligned} \text{grad } \phi_i &= \eta(d\phi_i, \cdot) = \begin{cases} -2i\pi^2 \phi_{i-1} \frac{\partial}{\partial\theta}, & i \text{ even}, \\ 2(i+1)\pi^2 \phi_{i+1} \frac{\partial}{\partial\theta}, & i \text{ odd}, \end{cases} \\ b_{ij} &= \begin{cases} -2j\pi^2 \phi_i \phi_{j-1} \frac{\partial}{\partial\theta}, & j \text{ even}, \\ 2(j+1)\pi^2 \phi_i \phi_{j+1} \frac{\partial}{\partial\theta}, & j \text{ odd}. \end{cases} \end{aligned}$$

It therefore follows from standard trigonometric identities that for any odd $i \geq 1$,

$$\frac{b_{i+1,i} - b_{i,i+1}}{2(i+1)\pi^2} = \frac{\partial}{\partial\theta},$$

with an analogous relationship holding for i even. This shows that for $J \geq 2$, B_J contains linearly dependent elements. On the other hand, for $J = 1$, B_1 fails to be a spanning set as the harmonic vector field $\frac{\partial}{\partial\theta}$ does not lie in its span.

Remark 4.11. The circle example above might suggest that our representation of vector fields through linear combinations of elements of B_J is highly inefficient, since, after all, one could define $\xi_j = \phi_j \frac{\partial}{\partial\theta}$, and $\{\xi_j\}_{j=0}^\infty$ would be an orthonormal basis of $H_{\mathfrak{X}}$. However, such a construction implicitly makes use of a special property of the circle, namely that it is a parallelizable manifold. Equivalently, as a $C^\infty(\mathcal{M})$ -module, the space $\mathfrak{X}$ of smooth vector fields is free; that is, it contains a set $\{u_1, \dots, u_d\}$ of $d = \dim(\mathcal{M})$ nowhere-vanishing linearly independent elements. Any such set would be a basis of $\mathfrak{X}$, meaning that for every $v \in \mathfrak{X}$ there would exist unique smooth function $f_1, \dots, f_d \in C^\infty(\mathcal{M})$ such that $v = \sum_{j=1}^d f_j u_j$. In general, for non-parallelizable manifolds (e.g., the 2-sphere), $\mathfrak{X}$ does not have a basis, so any spanning set of $H_{\mathfrak{X}}$, such as B_J , that makes use of a generating set of $\mathfrak{X}$ will necessarily be overcomplete.

We continue by stating a number of useful properties of the b_{ij} fields and their antisymmetric analogs, $\widehat{b}_{ij}$. Many of these properties are also listed in Tables 2.1 and 2.2. In what follows, all equalities involving infinite sums hold in an L^2 sense.

- (1) *Relationship between antisymmetric and nonsymmetric frame elements.* Using the Leibniz rule, we compute

$$\phi_i \text{grad } \phi_j + \phi_j \text{grad } \phi_i = \text{grad}(\phi_i \phi_j) = \sum_{k=0}^{\infty} c_{ijk} \text{grad } \phi_k = \sum_{k=1}^{\infty} c_{ijk} \widehat{b}_{0k},$$

where the last equality follows from the fact that ϕ_0 is a constant equal to 1. It therefore follows that

$$(4.11) \quad b_{ij} = \frac{\phi_i \text{grad } \phi_j + \phi_j \text{grad } \phi_i}{2} + \frac{\phi_i \text{grad } \phi_j - \phi_j \text{grad } \phi_i}{2} = \sum_{k=1}^{\infty} c_{ijk} \widehat{b}_{0k} + \widehat{b}_{ij}.$$

- (2) *Riemannian inner products.* Lemma 3.3, (2.4), and (3.1) lead to the following expressions for the Riemannian inner products between the frame elements:

$$(4.12) \quad g(b_{ij}, b_{kl}) = \phi_i \phi_k g(\text{grad } \phi_j, \text{grad } \phi_l) = \sum_{n=0}^{\infty} \phi_i \phi_k \phi_n g_{njl} = \sum_{m,n,p=0}^{\infty} c_{ikm} c_{mnp} g_{njl} \phi_p.$$

Using the above, we can also compute the Riemannian inner products between the antisymmetric vector fields, i.e.,

$$(4.13) \quad g(\widehat{b}_{ij}, \widehat{b}_{kl}) = g(b_{ij}, b_{kl}) + g(b_{ji}, b_{lk}) - g(b_{ij}, b_{lk}) - g(b_{ji}, b_{kl}).$$

The Riemannian inner products between the frame elements $\widehat{b}_{ij} \in \widetilde{B}_{\mathfrak{X}}$ are given by eigenvalue-dependent rescalings of those in (4.12).

- (3) *Hodge inner products and matrix elements of the Gramm operators.* Using the spectral representation of the pointwise Riemannian inner products in (4.12) and the fact that $c_{mn0} = \delta_{mn}$, we can compute the Hodge inner products

$$G_{ijkl} = \langle b_{ij}, b_{kl} \rangle_{H_{\mathfrak{X}}} = \langle \phi_0, g(b_{ij}, b_{kl}) \rangle_H = \sum_{n=0}^{\infty} c_{ikn} g_{njl},$$

as in (2.6). In addition, we have

$$\widehat{G}_{ijkl} \equiv \langle \widehat{b}_{ij}, \widehat{b}_{kl} \rangle_{H_{\mathfrak{X}}} = G_{ijkl} + G_{jilk} - G_{ijlk} - G_{jikl}.$$

Given now any ordering $b_{i_0 j_0}, b_{i_1 j_1}, \dots$ of the frame elements in $B_{\mathfrak{X}}^J$, where $j_n \leq J$, the above can be used to compute the matrix elements of the corresponding Gramm operator $G: \ell^2 \rightarrow \ell^2$, viz.

$$G_{mn} = \langle e_m, G e_n \rangle_{\ell^2} = \langle b_{i_m j_m}, b_{i_n j_n} \rangle_{H_{\mathfrak{X}}} = G_{i_m j_m k_n l_n}.$$

The analogous expressions for the Hodge inner products and Gramm matrix elements associated with the $\widetilde{B}_{\mathfrak{X}}$ frame are given by appropriate rescalings of the G_{ijkl} .

Additional formulas for the SEC representation of vector fields are listed in Table 2.1. The next few results are for the 1-form fields in Theorems 4.3 and 4.4. They

will be employed in our proof of Theorem 4.4 and in the construction of Galerkin schemes for the Laplacian on 1-forms in Sections 4.3 and 6 ahead, respectively.

- (1) *Exterior derivative and codifferential.* It follows from the Leibniz rule for the exterior derivative in (3.13) and the definition of the codifferential in (3.14) that

$$(4.14) \quad db^{ij} = db^i \wedge db^j, \quad \delta b^{ij} = -\eta(db^i, db^j) + \lambda_j b^i b^j.$$

Similarly, we have

$$\widehat{db}^{ij} = 2db^i \wedge db^j, \quad \delta \widehat{b}^{ij} = (\lambda_j - \lambda_i) b^i b^j.$$

Observe, in particular, that if $\lambda_i = \lambda_j$ (i.e., b^i and b^j lie in the same eigenspace of the Laplacian), $\widehat{b}^{ij}$ is co-closed, $\delta \widehat{b}^{ij} = 0$. For additional details on these formulas see (B.4).

- (2) *Riemannian inner products.* Since the b^{ij} and b_{ij} are Riemannian duals to each other, we have $\eta(b^{ij}, b^{kl}) = g(b_{ij}, b_{kl})$, and the latter can be determined from (4.12). An alternative derivation of this result, directly utilizing the product rule for the Laplacian on functions, can be found in Appendix B.1. The Riemannian inner products between the antisymmetric 1-forms $\widehat{b}^{ij}$ can be computed analogously to (4.13). The Riemannian inner products between the k -form frame elements $b^{ij_1 \dots j_k}$ in B_k^J , or the antisymmetric k -forms $\widehat{b}_k^{ij_1 \dots j_k}$, can be evaluated by computing determinants of $k \times k$ matrices of Riemannian inner products between db^j 's; see Appendix B.2 for further details.
- (3) *Riemannian inner products between exterior derivatives and codifferentials of the frame elements.* In order to perform operations on the frame elements $b_1^{ij} \in B_{1,1}^J$, or $\widehat{b}_1^{ij} \in \widehat{B}_{1,1}$, with the differential operators of the exterior calculus, we need expressions for Riemannian inner products between exterior derivatives and codifferentials such as $g(db_1^{ij}, db_1^{kl})$ and $g(\delta b_1^{ij}, \delta b_1^{kl})$. Closed-form expressions for such inner products based on the Laplacian eigenvalues λ_i and the corresponding triple-product coefficients c_{ijk} can be derived using (4.14); explicit results and derivations can be found in Table 2.2 and Appnedix B.1. Analogous inner product formulas can also be derived hierarchically for the higher-order frame elements (see Appendix B.2), although currently we do not have closed-form expressions for direct evaluation of pairwise inner products between the $\delta b^{ij_1 \dots j_k}$ at arbitrary $k > 1$. The inner product relationships between $b^{ij_1 \dots j_k}$ and their exterior derivatives and codifferentials would be needed to perform operations with the frames for Sobolev spaces of k -forms, $k > 1$, in Conjecture 4.5.
- (4) *Hodge and Sobolev inner products.* These can be computed as described above for vector fields, using the additional results on Riemannian inner products between exterior derivatives and codifferentials outlined above. Specific formulas can be found in Table 2.2 and Appendix B.1. Note that

the Sobolev inner products between the frame elements for H_1^1 and the corresponding Dirichlet forms will be used in our Galerkin approximation scheme for the 1-Laplacian in Section 6 ahead.

4.2 Proof of Theorem 4.3

We will prove the theorem by establishing the upper and lower frame conditions in Definition 4.1 for $B_{\mathfrak{X}}^J$ and B_k^J , assuming that J is large-enough so that Lemma 4.7 holds. We begin from $B_{\mathfrak{X}}^J$.

Given any $v \in H_{\mathfrak{X}}$, we have

$$\begin{aligned} \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_{ij}, v \rangle_{H_{\mathfrak{X}}}|^2 &= \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle \phi_i \operatorname{grad} \phi_j, v \rangle_{H_{\mathfrak{X}}}|^2 = \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle \phi_i, v(\phi_j) \rangle_H|^2 \\ &= \sum_{j=1}^J \|v(\phi_j)\|_H^2 = \sum_{j=1}^J \|g(\operatorname{grad} \phi_j, v)\|_H^2 \\ &\leq \sum_{j=1}^J \|g(\operatorname{grad} \phi_j, v)\|_H^2 \leq \sum_{j=1}^J \left\| \sqrt{g(\operatorname{grad} \phi_j, \operatorname{grad} \phi_j)} g(v, v) \right\|_H^2 \\ &\leq \sum_{j=1}^J \|\operatorname{grad} \phi_j\|_{L_{\mathfrak{X}}^{\infty}}^2 \left\| \sqrt{g(v, v)} \right\|_H^2 = \sum_{j=1}^J \|\operatorname{grad} \phi_j\|_{L_{\mathfrak{X}}^{\infty}}^2 \|v\|_{H_{\mathfrak{X}}}^2, \end{aligned}$$

so that the upper frame condition holds with $C_2 = \sum_{j=1}^J \|\operatorname{grad} \phi_j\|_{L_{\mathfrak{X}}^{\infty}}^2$. Note that the fact that J is finite is important in the derivation of this result. Next, to verify the lower frame condition, consider the $J \times J$ Gramm matrix $\Psi(x)$, for $x \in \mathcal{M}$, with elements

$$\Psi_{ij}(x) = g(\operatorname{grad} \phi_i, \operatorname{grad} \phi_j)|_x,$$

and note that because the $\operatorname{grad} \phi_j|_x$ span $T_{x^*}\mathcal{M}$, that matrix has rank d . Therefore, writing $v = \sum_{j=1}^J f_j \operatorname{grad} \phi_j$, where the f_j are functions in H to be determined, the equation

$$g(\operatorname{grad} \phi_i, v)|_x = \sum_{j=1}^J f_j(x) \Psi_{ji}(x)$$

has a solution for μ -a.e. $x \in X$ given by

$$f_j(x) = \sum_{k=1}^J \Psi_{jk}^+(x) g(\operatorname{grad} \phi_k, v)|_x,$$

where $\Psi^+(x) = [\Psi_{jk}^+(x)]$ is the Moore-Penrose pseudoinverse of $\Psi(x)$.

Observe now that $\Psi(x)$ can be expressed using a coordinate chart as $\Psi(x) = \Xi(x)^{\top} g(x) \Xi(x)$, where the matrices $\Xi(x)$ and $g(x)$ are as in the proof of Lemma 4.7. Thus, since $\Xi(x)$ has linearly independent rows and $g(x)$ is invertible, we have

$$\Psi^+(x) = \Xi^+(x) g^{-1}(x) (\Xi^+(x))^{\top}, \quad \Xi^+(x) = \Xi(x)^{\top} (\Xi(x) \Xi(x)^{\top})^{-1},$$

where both $g^{-1}(x)$ and $(\Xi(x)\Xi(x))^{-1}$ depend smoothly on x by compactness of $\mathcal{M}$ and smoothness of $g(x)$ and $\Xi(x)$, respectively. We therefore conclude that $\Psi^+(x)$ depends smoothly on x , and thus that v admits an expansion of the form (4.9) with

$$(4.15) \quad c_{ij} = \langle \phi_i, f_j \rangle_H = \sum_{k=1}^J \langle \Psi_{jk}^+ b_{ik}, v \rangle_{H_{\mathfrak{X}}}.$$

We therefore obtain

$$\begin{aligned} \|v\|_{H_{\mathfrak{X}}}^2 &= \sum_{i=0}^{\infty} \sum_{j=1}^J \langle c_{ij} b_{ij}, v \rangle_{H_{\mathfrak{X}}} = \sum_{i=0}^{\infty} \sum_{j,k=1}^J \langle \Psi_{jk}^+ b_{ik}, v \rangle_{H_{\mathfrak{X}}}^* \langle b_{ij}, v \rangle_{H_{\mathfrak{X}}} \\ &\leq \tilde{C} \sum_{i=0}^{\infty} \sum_{j,k=1}^J |\langle b_{ik}, v \rangle_{H_{\mathfrak{X}}}| |\langle b_{ij}, v \rangle_{H_{\mathfrak{X}}}|, \end{aligned}$$

where $\tilde{C} = \max_{j,k \in \{1, \dots, J\}} \|\Psi_{jk}^+\|_{L^\infty}$. Defining now the vectors $\beta_i = (\beta_{i1}, \dots, \beta_{iJ}) \in \mathbb{R}^J$ with $\beta_{ij} = |\langle b_{ij}, v \rangle_{H_{\mathfrak{X}}}|$, it follows by equivalence of norms in finite-dimensional vector spaces that there exists a constant $\hat{C} > 0$ such that

$$\sum_{j,k=1}^J |\langle b_{ik}, v \rangle_{H_{\mathfrak{X}}}| |\langle b_{ij}, v \rangle_{H_{\mathfrak{X}}}| = \|\beta_i\|_1^2 \leq \hat{C} \|\beta_i\|_2^2,$$

where $\|\cdot\|_p$ is the canonical p -norm on $\mathbb{R}^J$. This leads to

$$\|v\|_{\mathfrak{X}}^2 \leq \tilde{C} \hat{C} \sum_{i=0}^{\infty} \|\beta_i\|_2^2 = \tilde{C} \hat{C} \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_{ij}, v \rangle_{H_{\mathfrak{X}}}|^2,$$

which proves the lower frame condition with $C_1 = 1/(\tilde{C}\hat{C})$. We have thus established that $B_{\mathfrak{X}}^J$ is a frame of $H_{\mathfrak{X}}$, as claimed.

Consider now the frame conditions for $\tilde{B}_{\mathfrak{X}}$. Using Cauchy-Schwartz inequalities as above, we can conclude that

$$\sum_{i=0}^{\infty} \sum_{j=1}^{\infty} |\langle \tilde{b}_{ij}, v \rangle_{H_{\mathfrak{X}}}|^2 \leq \sum_{j=1}^{\infty} \frac{\|\text{grad } \phi_j\|_{L_{\mathfrak{X}}^\infty}^2}{e^{\lambda_j}} \|v\|_{H_{\mathfrak{X}}}^2.$$

To bound the infinite sum in the right-hand side, we use the following estimates for the L^∞ norms of Laplace-Beltrami eigenfunctions and their gradients on smooth, closed Riemannian manifolds:

$$(4.16) \quad \|\phi_j\|_{L^\infty} \leq C \lambda_j^{(d-1)/4} \|\phi_j\|_H, \quad \|\text{grad } \phi_j\|_{L_{\mathfrak{X}}^\infty} \leq \tilde{C} \lambda_j^{1/2} \|\phi_j\|_{L^\infty}, \quad C, \tilde{C} \geq 0,$$

which hold for $j \leq 1$ and $j \geq 0$, respectively. The former is a classical result due to Hörmander [34]; the latter was proved by Shi and Xu in [42]. Combining these results with the Weyl estimate for the asymptotic distribution of Laplace-Beltrami eigenvalues as $j \rightarrow \infty$,

$$(4.17) \quad j = \hat{C} \lambda_j^{d/2} + o(\lambda_j^{(d-1)/2}),$$

we obtain

$$\|\text{grad } \phi_j\|_{L^\infty_{\mathbb{X}}} \leq \check{C} \lambda_j^{(d+1)/(2d)} \|\phi_j\|_H,$$

where $\hat{C}$ and $\check{C}$ in the last two equations are positive constants. Therefore, for any $l \geq 0$ there exists a finite constant C_l such that $e^{-\lambda_j} \|\text{grad } \phi_j\|_{L^\infty_{\mathbb{X}}}^2 \leq C_l j^{-l}$. This implies that $C_2 = \sum_{j=0}^\infty e^{-\lambda_j} \|\text{grad } \phi_j\|_{L^\infty_{\mathbb{X}}}^2$ is finite, proving the upper frame bound.

To verify the lower frame bound, start from any expansion of v in the $B_{\mathbb{X}}^J$ frame,

$$v = \sum_{i=0}^\infty \sum_{j=1}^J c_{ij} b_{ij},$$

and compute

$$\|v\|_{\mathbb{X}}^2 \leq C \sum_{i=0}^\infty \sum_{j=1}^J |\langle b_{ij}, v \rangle_{H_{\mathbb{X}}}|^2 \leq C e^{\lambda_j} \sum_{i=0}^\infty \sum_{j=1}^J |\langle \tilde{b}_{ij}, v \rangle_{H_{\mathbb{X}}}|^2 \leq C e^{\lambda_j} \sum_{i=0}^\infty \sum_{j=1}^J |\langle \tilde{b}_{ij}, v \rangle_{H_{\mathbb{X}}}|^2,$$

where C is a lower frame constant for $B_{\mathbb{X}}^J$. This shows that the lower frame condition is satisfied for $C_1 = C e^{-\lambda_j}$, and we thus conclude that $\tilde{B}_{\mathbb{X}}$ is a frame.

We now turn to the frame conditions for B_k^J and $\tilde{B}_k$. These conditions follow by similar arguments as those just made to establish the frame conditions for vector fields.

First, we introduce for convenience an ordering $l \mapsto (j_1(l), \dots, j_k(l))$ of the corresponding indices in $b^{ij_1 \dots j_k}$, where l is an integer ranging from 1 to J^k , and define $\alpha_l = d\phi^{j_1(l)} \wedge \dots \wedge d\phi^{j_k(l)}$. Then, for any $\omega \in H_k$, we have

$$\begin{aligned} \sum_{i=0}^\infty \sum_{j_1, \dots, j_k=1}^J |\langle b^{ij_1 \dots j_k}, \omega \rangle_{H_k}|^2 &= \sum_{i=0}^\infty \sum_{l=1}^{J^k} |\langle \phi_i \alpha_l, \omega \rangle_{H_k}|^2 \\ &= \sum_{i=0}^\infty \sum_{l=1}^{J^k} |\langle \phi_i, \eta(\alpha_l, \omega) \rangle_H|^2 = \sum_{l=1}^{J^k} \|\eta(\alpha_l, \omega)\|_H^2 \\ &\leq \sum_{l=1}^{J^k} \left\| \sqrt{\eta(\alpha_l, \alpha_l)} \sqrt{\eta(\omega, \omega)} \right\|_H^2 \leq \sum_{l=1}^{J^k} \|\alpha_l\|_{L_k^\infty}^2 \|\omega\|_{H_k}^2, \end{aligned}$$

establishing the upper frame condition with $C_2 = \sum_{l=1}^{J^k} \|\alpha_l\|_{L_k^\infty}^2$. To verify the lower frame condition, we use (4.10) to expand $\omega = \sum_{i=0}^\infty \sum_{l=1}^{J^k} c_{il} \phi_i \alpha_l$, where the expansion coefficients c_{il} can be chosen as (cf. (4.15))

$$c_{il} = \sum_{m=1}^{J^k} \langle \Psi_{lm}^+ \phi_i \alpha_m, \omega \rangle_{H_k},$$

and in the above $\Psi_{lm}^+(x)$ are the elements of the pseudoinverse of the $J^k \times J^k$ Gramm matrix $\Psi_{lm}(x) = \eta(\alpha_l, \alpha_m)|_x$ (these matrix elements depend smoothly on x as in the case of the corresponding Gramm matrix for vector fields). The calculation to establish the lower frame bound for $\|\omega\|_{H_k}^2 = \sum_{i=0}^\infty \sum_{l=1}^{J^k} \langle c_{il} \phi_i \alpha_l, \omega \rangle_{H_k}$ then proceeds

analogously to that in the case of vector fields, leading to the conclusion that B_k^J is an H_k -frame. Similarly, that $\tilde{B}_k$ is a frame follows analogously to the vector field case. This completes our proof of Theorem 4.3.

4.3 Proof of Theorem 4.4

Let ω be an arbitrary 1-form field in H_1^1 . We begin by stating two auxiliary results on the inner products between the exterior derivative (codifferential) of ω and the exterior derivative (codifferential) of the frame elements b_1^{ij} and $\tilde{b}_1^{ij}$.

Lemma 4.12. (i) *There exist constants U_J and V_J , independent of ω , such that*

$$\sum_{i=0}^{\infty} \sum_{j=1}^J |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2 \leq U_J \|\omega\|_{H_1^1}^2, \quad \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle \delta b_1^{ij}, \delta\omega \rangle_H|^2 \leq V_J \|\omega\|_{H_1^1}^2.$$

Moreover, there exist a positive real number $\hat{C}$ and a positive integer $q \geq 0$, both independent of J , such that U_J and V_J are both bounded above by $\hat{C}J^q$.

(ii) *There exist finite constants $\tilde{U}$ and $\tilde{V}$, independent of ω , such that*

$$\sum_{i=0}^{\infty} \sum_{j=1}^{\infty} |\langle \tilde{db}_1^{ij}, d\omega \rangle_{H_2}|^2 \leq \tilde{U} \|\omega\|_{H_1^1}^2, \quad \sum_{i=0}^{\infty} \sum_{j=1}^{\infty} |\langle \tilde{\delta b}_1^{ij}, \delta\omega \rangle_H|^2 \leq \tilde{V} \|\omega\|_{H_1^1}^2.$$

A proof of this lemma will be given below. Assuming, for now, that it is valid, it leads to the following corollary:

Corollary 4.13. *The frame elements b_1^{ij} and $\tilde{b}_1^{ij}$ satisfy the bounds*

$$\sum_{i=0}^{\infty} \sum_{j=1}^J |E_{1,1}(b_1^{ij}, \omega)|^2 \leq C \|\omega\|_{H_1^1}^2, \quad \sum_{i=0}^{\infty} \sum_{j=1}^{\infty} |E_{1,1}(\tilde{b}_1^{ij}, \omega)|^2 \leq \tilde{C} \|\omega\|_{H_1^1}^2,$$

where $E_{1,1}$ is the Dirichlet form from (3.18), and C and $\tilde{C}$ are constants independent of ω .

Proof. We first verify the claim for b_1^{ij} . By Lemma 4.12, the sequences $(i, j) \mapsto |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2$ and $(i, j) \mapsto |\langle \delta b_1^{ij}, \delta\omega \rangle_H|^2$ are in ℓ^2 . Therefore, using the Cauchy-Schwartz inequality for ℓ^2 and Lemma 4.12(i), we obtain,

$$\begin{aligned} \sum_{i=0}^{\infty} \sum_{j=1}^J |E_{1,1}(b_1^{ij}, \omega)|^2 &= \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle db_1^{ij}, d\omega \rangle_{H_2} + \langle \delta b_1^{ij}, \delta\omega \rangle_H|^2 \\ &\leq \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2 + \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle \delta b_1^{ij}, \delta\omega \rangle_H|^2 \\ &\quad + 2 \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle db_1^{ij}, d\omega \rangle_{H_2} \langle \delta b_1^{ij}, \delta\omega \rangle_H| \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2 + \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle \delta b_1^{ij}, \delta \omega \rangle_H|^2 \\
&\quad + 2 \left(\sum_{i=0}^{\infty} \sum_{j=1}^J |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2 \right)^{1/2} \\
&\quad \times \left(\sum_{k=0}^{\infty} \sum_{l=1}^J |\langle \delta b_1^{kl}, \delta \omega \rangle_H|^2 \right)^{1/2} \\
&\leq \left(U_J^{1/2} + V_J^{1/2} \right)^2 \|\omega\|_{H_1^1}^2,
\end{aligned}$$

and the first claim of the corollary follows with $C = \left(U_J^{1/2} + V_J^{1/2} \right)^2$. To verify the second claim, we proceed as above using Lemma 4.12(ii) to derive the bound

$$\sum_{i=0}^{\infty} \sum_{j=1}^J |E_{1,1}(\tilde{b}_1^{ij}, \omega)|^2 \leq \tilde{C} \|\omega\|_{H_1^1}^2, \quad \tilde{C} = \left(\tilde{U}^{1/2} + \tilde{V}^{1/2} \right)^2. \quad \square$$

We now return to the proof of Theorem 4.4. By Corollary 4.13, the upper frame condition for $B_{1,1}^J$ follows from

$$\begin{aligned}
\sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1^1}|^2 &= \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1} + E_{1,1}(b_1^{ij}, \omega)|^2 \\
&\leq \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1}|^2 + \sum_{i=0}^{\infty} \sum_{j=1}^J |E_{1,1}(b_1^{ij}, \omega)|^2 \\
&\quad + 2 \left(\sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1}|^2 \right)^{1/2} \left(\sum_{k=0}^{\infty} \sum_{l=1}^J |E_{1,1}(b_1^{kl}, \omega)|^2 \right)^{1/2} \\
&= \sum_{i=0}^{\infty} \sum_{j=1}^J \frac{|\langle b_1^{ij}, \omega \rangle_{H_1}|^2}{\|b^i\|_{H^1}^2} + \sum_{i=0}^{\infty} \sum_{j=1}^J |E_{1,1}(b_1^{ij}, \omega)|^2 \\
&\quad + 2 \left(\sum_{i=0}^{\infty} \sum_{j=1}^J \frac{|\langle b_1^{ij}, \omega \rangle_{H_1}|^2}{\|b^i\|_{H^1}^2} \right)^{1/2} \left(\sum_{k=0}^{\infty} \sum_{l=1}^J |E_{1,1}(b_1^{kl}, \omega)|^2 \right)^{1/2} \\
&\leq \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1}|^2 + \sum_{i=0}^{\infty} \sum_{j=1}^J |E_{1,1}(b_1^{ij}, \omega)|^2 \\
&\quad + 2 \left(\sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1^1}|^2 \right)^{1/2} \left(\sum_{k=0}^{\infty} \sum_{l=1}^J |E_{1,1}(b_1^{kl}, \omega)|^2 \right)^{1/2} \\
&\leq \left(C_2^{1/2} + C^{1/2} \right)^2 \|\omega\|_{H_1^1}^2,
\end{aligned}$$

where C_2 is an upper frame constant for B_1^J and C the constant in Corollary 4.13. The upper frame condition for $\tilde{B}_{1,1}^J$ follows similarly.

What remains in the proof of the upper frame conditions in Theorem 4.4 is to verify Lemma 4.12. For that, first observe that

$$\begin{aligned} \langle db^{ij}, dv \rangle_{H_2} &= \langle db^i \wedge db^j, dv \rangle_{H_2} = \int_{\mathcal{M}} db^i \wedge db^j \wedge \star dv \\ &= \int_{\mathcal{M}} db^i \wedge \star \star (db^j \wedge \star dv) = \langle db^i, \star (db^j \wedge \star dv) \rangle_{H_1}, \end{aligned}$$

where we have used (4.14) and (3.12) in the first and third equalities, respectively. By the Hodge decomposition theorem, there exists a unique function $f_j \in H^1$, a unique 2-form $\alpha_j \in H_2^1$, and a unique harmonic 1-form $\chi_j \in \mathcal{H}_0^1$ such that

$$\star (db^j \wedge \star d\omega) = df_j + \delta \alpha_j + \chi_j;$$

as a result,

$$\begin{aligned} \sum_{i=0}^{\infty} |\langle db_1^i, \star (db^j \wedge \star d\omega) \rangle_{H_1}|^2 &= \sum_{i=0}^{\infty} |\langle db_1^i, df_j + \delta \alpha_j + \chi_j \rangle_{H_1}|^2 = \sum_{i=0}^{\infty} |\langle \delta db_1^i, f_j \rangle_H|^2 \\ &= \sum_{i=0}^{\infty} \lambda_i |\langle b^i, f_j \rangle_H|^2 = E_1(f_j). \end{aligned}$$

We therefore have,

$$\begin{aligned} \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2 &= \sum_{i=0}^{\infty} |\langle db_1^i, \star (db^j \wedge \star d\omega) \rangle_{H_1}|^2 \\ &= \sum_{j=1}^J E_1(f_j) = \sum_{j=1}^J \langle df_j, df_j \rangle_{H_1} \\ &\leq \sum_{j=1}^J \|\star (db^j \wedge \star d\omega)\|_{H_1}^2 = \sum_{j=1}^J \|db^j \wedge \star d\omega\|_{H_{d-1}}^2 \\ &= \sum_{j=1}^J \|\star \iota_{\text{grad } b^j} \star \star d\omega\|_{H_{d-1}}^2 = \sum_{j=1}^J \|\iota_{\text{grad } b^j} d\omega\|_{H_1}^2, \end{aligned}$$

where we have used (3.15) to obtain the first equality in the second line. It then follows from Lemma 3.4 that

$$\sum_{j=1}^J \sum_{i=0}^{\infty} |\langle db_1^{ij}, d\omega \rangle_{H_2}|^2 = \sum_{j=1}^J \|D_{\text{grad } b^j} \omega\|_{H_1}^2 \leq U_J \|\omega\|_{H_1^1}^2, \quad U_J = \sum_{j=1}^J \|D_{\text{grad } b^j}\|^2,$$

as claimed in part (i) of the lemma. It can further be shown via local Cauchy-Schwartz inequalities as in the proof of Lemma 3.4 that the operator norms $\|D_{\text{grad } b^j}\| = \|D_{\text{grad } \phi_j}\|$ can be bounded above by $\tilde{C} \|\phi_j\|_{L^\infty}^{\tilde{q}}$ for some positive constants $\tilde{C}$ and $\tilde{q}$ that do not depend on j . This, in conjunction with the Hörmander bound in (4.16)

implies that there exists a positive constant C_U and a positive integer q_U , both independent of J , such that

$$(4.18) \quad U_J \leq C_U J^{q_U}.$$

Moving on to the second claim of Lemma 4.12(i), consider

$$\sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \delta b_1^{ij}, \delta \omega \rangle_H|^2 = \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle -\eta(db_1^i, db^j) + \lambda_j b_1^i b^j, \delta \omega \rangle_H|^2,$$

where we have used the expression for δb^{ij} in (4.14). Expanding

$$\delta \omega db^j = df_j + \delta \alpha_j + \chi_j,$$

where $f_j \in H^1$, $\alpha_j \in H_2^1$, and $\chi_j \in \mathcal{H}_0^1$ are unique, we compute,

$$\begin{aligned} \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \eta(db_1^i, db^j), \delta \omega \rangle_H|^2 &= \sum_{j=1}^J \sum_{i=0}^{\infty} \left| \int_{\mathcal{M}} db_1^i \wedge \star db^j \delta \omega \right|^2 \\ &= \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle db_1^i, \delta \omega db^j \rangle_{H_1}|^2 = \sum_{j=1}^J E_1(f_j) \\ &\leq \sum_{j=1}^J \|\delta \omega db^j\|_{H_1}^2 = \sum_{j=1}^J \left\| \sqrt{\eta(db^j, db^j)} \delta \omega \right\|_H^2 \\ (4.19) \quad &\leq \sum_{j=1}^J \|\text{grad } \phi_j\|_{L_x^\infty}^2 \|\delta \omega\|_H^2 \leq \sum_{j=1}^J \|\text{grad } \phi_j\|_{L_x^\infty}^2 \|\omega\|_{H_1^1}^2. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \lambda_j b_1^i b^j, \delta \omega \rangle_H|^2 &= \sum_{j=1}^J \sum_{i=0}^{\infty} \frac{\lambda_j^2}{\|b^i\|_{H^1}^2} |\langle b^i b^j, \delta \omega \rangle_H|^2 \\ &\leq \sum_{j=1}^J \sum_{i=0}^{\infty} \lambda_j^2 |\langle b^i, b^j \delta \omega \rangle_H|^2 = \sum_{j=1}^J \lambda_j^2 \|b^j \delta \omega\|_H^2 \\ (4.20) \quad &\leq \sum_{j=1}^J \lambda_j^2 \|b^j\|_{L^\infty}^2 \|\omega\|_{H_1^1}^2. \end{aligned}$$

Equations (4.19) and (4.20) imply that the sequences $(i, j) \mapsto |\langle \eta(db_1^i, db^j), \delta \omega \rangle_H|$ and $(i, j) \mapsto |\langle \lambda_j b_1^i b^j, \delta \omega \rangle_H|$ are both in ℓ^2 . Therefore, by the Cauchy-Schwarz inequality for that space we can conclude that

$$\begin{aligned} \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \delta b_1^{ij}, \delta \omega \rangle_H|^2 &= \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle -\eta(db_1^i, db^j) + \lambda_j b_1^i b^j, \delta \omega \rangle_H|^2 \\ &\leq \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \eta(db_1^i, db^j), \delta \omega \rangle_H|^2 + \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \lambda_j b_1^i b^j, \delta \omega \rangle_H|^2 \end{aligned}$$

$$\begin{aligned}
& + 2 \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \eta(db_1^i, db^j), \delta\omega \rangle_H \langle \lambda_j b_1^i b^j, \delta\omega \rangle_H| \\
& \leq \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \eta(db_1^i, db^j), \delta\omega \rangle_H|^2 + \sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \lambda_j b_1^i b^j, \delta\omega \rangle_H|^2 \\
& \quad + 2 \left(\sum_{j=1}^J \sum_{i=0}^{\infty} |\langle \eta(db_1^i, db^j), \delta\omega \rangle_H|^2 \right)^{1/2} \\
& \quad \times \left(\sum_{l=1}^J \sum_{k=0}^{\infty} |\langle \lambda_l b_1^k b^l, \delta\omega \rangle_H|^2 \right)^{1/2} \\
& \leq V_J \|\omega\|_{H_1^1},
\end{aligned}$$

where

$$V_J^{1/2} = \left(\sum_{j=1}^J \|\text{grad } \phi_j\|_{L_{\mathfrak{X}}^{\infty}}^2 \right)^{1/2} + \left(\sum_{j=0}^J \lambda_j^2 \|\phi_j\|_{L^{\infty}}^2 \right)^{1/2}.$$

This establishes the existence of the ω -independent constants V_J claimed in the lemma. Invoking L^{∞} and Weyl bounds as in the case of U_J , we can also deduce that there exist a real number C_V and a positive integer q_V such that

$$(4.21) \quad V_J \leq C_V J^{q_V}.$$

Combining (4.18) and (4.21) leads to $U_J \leq \hat{C} J^q$ and $V_J \leq \hat{C} J^q$ with $\hat{C} = \max\{C_U, C_V\}$ and $q = \max\{q_U, q_V\}$. This completes our proof of Lemma 4.12(i) and thus the upper frame condition for $B_{1,1}^J$.

To prove Lemma 4.12(ii), we proceed as above to establish that

$$\sum_{j=1}^k \sum_{i=0}^{\infty} |\langle \tilde{db}_1^{ij}, d\omega \rangle_{H_2}|^2 \leq \tilde{U}_k \|\omega\|_{H_1^1}^2, \quad \tilde{U}_k = \sum_{j=1}^k e^{-\lambda_j} \|D_{\text{grad } b^j}\|^2,$$

and

$$\begin{aligned}
& \sum_{j=1}^k \sum_{i=0}^{\infty} |\langle \tilde{\delta b}_1^{ij}, \delta\omega \rangle_H|^2 \leq \tilde{V}_k \|\omega\|_{H_1^1}^2, \\
& \tilde{V}_k^{1/2} = \left(\sum_{j=1}^k e^{-\lambda_j} \|\text{grad } \phi_j\|_{L_{\mathfrak{X}}^{\infty}}^2 \right)^{1/2} + \left(\sum_{j=0}^k e^{-\lambda_j} \lambda_j^2 \|\phi_j\|_{L^{\infty}}^2 \right)^{1/2}.
\end{aligned}$$

The L^{∞} and Weyl estimates in (4.16) and (4.17), respectively, then again imply that $\tilde{U} = \lim_{k \rightarrow \infty} \tilde{U}_k$ and $\tilde{V} = \lim_{k \rightarrow \infty} \tilde{V}_k$ is finite, proving Lemma 4.12(ii), and completing our proof of the upper frame condition for $\tilde{B}_{1,1}$.

Next, to verify the lower frame conditions, we express ω as a linear combination $\omega = \sum_{j=1}^J f_j d\phi_j$, where $f_1, \dots, f_J$ are H^1 functions satisfying

$$f_j(x) = \sum_{k=1}^J \Psi_{jk}^+(x) \eta(d\phi_k, \omega)$$

for μ -a.e. $x \in \mathcal{M}$, and $\Psi_{jk}^+(x)$ are the elements of the Moore-Penrose pseudoinverse of the $J \times J$ Gramm matrix

$$\Psi_{jk}(x) = \eta(d\phi_i, d\phi_j) = g(\text{grad } \phi_i, \text{grad } \phi_j)$$

from Section 4.2. Note that the existence of such an expansion for ω follows from the fact that $\{d\phi_1, \dots, d\phi_J\}$ is a generating set of the space of smooth 1-forms Ω^1 , and the latter is dense in H^1 . The f_j functions can be expanded in the $\{\phi_i^{(1)}\}_{i=0}^\infty$ basis of H^1 from (3.5), viz.

$$f_j = \sum_{i=0}^\infty c_{ij} \phi_i^{(1)}, \quad c_{ij} = \langle \phi_i^{(1)}, f_j \rangle_{H^1} = \sum_{k=1}^J \langle \phi_i^{(1)}, \Psi_{jk}^+ \eta(d\phi_k, \omega) \rangle_{H^1}.$$

Therefore, setting $b^i = \phi_i$ and $b_1^i = \phi_i^{(1)}$ per our notational convention for frame elements, we obtain

$$\begin{aligned} \|\omega\|_{H^1}^2 &= \langle \omega, \omega \rangle_{H^1} = \sum_{i=0}^\infty \sum_{j,k=1}^J \langle b_1^i, \Psi_{jk}^+ \eta(db^k, \omega) \rangle_{H^1} \langle b_1^{ij}, \omega \rangle_{H^1} \\ &\leq \sum_{i=0}^\infty \sum_{j,k=1}^J |\langle b_1^i, \Psi_{jk}^+ \eta(db^k, \omega) \rangle_{H^1} \langle b_1^{ij}, \omega \rangle_{H^1}| \\ (4.22) \quad &\leq \left(\sum_{i=0}^\infty \sum_{j,k=1}^J |\langle b_1^i, \Psi_{jk}^+ \eta(db^k, \omega) \rangle_{H^1}|^2 \right)^{1/2} \left(\sum_{n=0}^\infty \sum_{p=1}^J |\langle b_1^{np}, \omega \rangle_{H^1}|^2 \right)^{1/2}. \end{aligned}$$

Note that to arrive at the inequality in the last line we used the ℓ^2 Cauchy-Schwartz inequality on the sequences $(i, j) \mapsto \langle b_1^i, \Psi_{jk}^+ \eta(db^k, \omega) \rangle_{H^1}$ and $(i, j) \mapsto \langle b_1^{ij}, \omega \rangle_{H^1}$, both of which can be verified to indeed lie in that space. We now proceed to bound the first term in the last line.

First, since $\{b_1^i\}_{i=0}^\infty$ is an orthonormal basis of H^1 , we have

$$\sum_{i=0}^\infty \sum_{j,k=1}^J |\langle b_1^i, \Psi_{jk}^+ \eta(db^k, \omega) \rangle_{H^1}|^2 = \sum_{j,k=1}^J \|\Psi_{jk}^+ \eta(db^k, \omega)\|_{H^1}^2.$$

Moreover, observe that for any $f \in H^1$ and $h \in C^\infty(\mathcal{M})$, $\|hf\|_{H^1}$ can be bounded above by $\tilde{C}\|f\|_{H^1}$, where $\tilde{C}^2$ is a polynomial function of $\|h\|_{L^\infty}$ and $\|\text{grad } h\|_{L_x^\infty}$. This implies that there exists a constant C such that

$$\sum_{j,k=1}^J \|\Psi_{jk}^+ \eta(db^k, \omega)\|_{H^1}^2 \leq C \sum_{k=1}^J \|\eta(db_k, \omega)\|_{H^1}^2$$

$$(4.23) \quad = C \sum_{k=1}^J \left(\|\eta(db^k, \omega)\|_H^2 + \|d\eta(db^k, \omega)\|_{H_1}^2 \right).$$

In the above, the term $\|\eta(db^k, \omega)\|_H^2$ can be bounded above by $\|\text{grad } b^k\|_{L^\infty}^2 \|\omega\|_{H_1}^2$ using local Cauchy-Schwartz inequalities and the fact that $\|\cdot\|_{H_1} \leq \|\cdot\|_{H_1^1}$. To bound $\|d\eta(db^k, \omega)\|_{H_1}$, we use (3.15) to write down

$$d\eta(db^k, \omega) = d\iota_{\text{grad } b^k} \omega.$$

It follows from Lemma 3.4 that there exists a constant $\hat{C}$ such that

$$(4.24) \quad \|d\eta(db^k, \omega)\| \leq \hat{C} \|\omega\|_{H_1^1}.$$

Combining (4.22)–(4.24), we conclude that there exists a constant $\bar{C}$ such that

$$\|\omega\|_{H_1^1}^2 \leq \bar{C} \|\omega\|_{H_1^1} \left(\sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1^1}|^2 \right)^{1/2},$$

and thus the lower frame condition for $B_{1,1}^J$ holds with $C_1 = 1/\bar{C}^2$. To verify the lower frame condition for $\tilde{B}_{1,1}$, we use the results just established to compute

$$\begin{aligned} \|\omega\|_{H_1^1}^2 &\leq C_1 \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle b_1^{ij}, \omega \rangle_{H_1^1}|^2 \leq C_1 e^{\lambda_J} \sum_{i=0}^{\infty} \sum_{j=1}^J |\langle \tilde{b}_1^{ij}, \omega \rangle_{H_1^1}|^2 \\ &\leq C_1 e^{\lambda_J} \sum_{i=0}^{\infty} \sum_{j=1}^{\infty} |\langle \tilde{b}_1^{ij}, \omega \rangle_{H_1^1}|^2. \end{aligned}$$

This proves the lower frame condition for $\tilde{B}_{1,1}$, and completes our proof of Theorem 4.4.

5 Frame and operator representations of vector fields

As described in Section 2.3, the SEC is based on alternative representations of vector fields with respect to frames, or as operators on functions. In this section, we make these notions precise, and further examine the convergence properties of finite-rank analogs of these representations.

Unless otherwise stated, throughout this section, $T : H_{\mathfrak{X}} \rightarrow \ell^2$, $T^* : \ell^2 \rightarrow H_{\mathfrak{X}}$, $S = T^*T : H_{\mathfrak{X}} \rightarrow H_{\mathfrak{X}}$, and $G = TT^* : \ell^2 \rightarrow \ell^2$ will be the analysis, synthesis, frame, and Gramm operators, respectively, associated with one of the frames for the $H_{\mathfrak{X}} = L_{\mathfrak{X}}^2$ space of vector fields from Theorem 4.3. We also let $T' : H_{\mathfrak{X}} \rightarrow \ell^2$, $T'^* : \ell^2 \rightarrow H_{\mathfrak{X}}$, $S' = T'^*T' : H_{\mathfrak{X}} \rightarrow H_{\mathfrak{X}}$, and $G' = T'T'^* : \ell^2 \rightarrow \ell^2$ be the corresponding operators for the dual frame. For notational simplicity, we use the symbols α_k and α'_k with $k \in \{1, 2, \dots\}$ to represent the frame and dual frame elements, respectively. For example, in the case of the $B_{\mathfrak{X}}^J$ frames from Theorem 4.3, we set $\alpha_k = b_{p_k q_k}$, where $k \mapsto (p_k, q_k)$ is any ordering of the (i, j) indices in b_{ij} with $k \in \{0, 1, \dots\}$. A convenient choice of such ordering is a lexicographical ordering,

i.e., $(p_0, q_0) = (0, 1), \dots, (p_{J-1}, q_{J-1}) = (0, J), (p_J, q_J) = (1, 0), \dots$. We use the notation $(i, j) \mapsto r_{ij}$ to represent the inverse of this ordering map. Note that for any choice of frame from Theorem 4.3 we have $\alpha_k = s_k \phi_{p_k} \text{grad } \phi_{q_k}$, where $s_k = 1$ in the case of the $B_{\mathcal{X}}^J$ frames and $s_k = e^{-\lambda_{q_k}/2}$ in the case of the $\tilde{B}_{\mathcal{X}}$ frame. We also let $\{e_k\}_{k=0}^\infty$ be the canonical orthonormal basis of ℓ^2 , and $\pi_l : \ell^2 \rightarrow \ell^2$ the orthogonal projection operators with range $\tilde{W}_l = \text{span}\{e_0, \dots, e_{l-1}\}$. Moreover, the set $\{e_{ij}\}_{i,j=0}^\infty$, $e_{ij} \in \mathcal{B}_2(\ell^2)$, will be the canonical orthonormal basis of $\mathcal{B}_2(\ell^2)$ with $e_{ij} = e_i \langle e_j, \cdot \rangle_{\ell^2}$.

In addition to the various frame operators acting on vector fields, we will consider the unitary Fourier operators $U : H \rightarrow \ell^2$, $U^{(1)} : H^1 \rightarrow \ell^2$, and $\tilde{U} : \mathcal{H} \rightarrow \ell^2$ associated with the $\{\phi_k\}_{k=0}^\infty$, $\{\phi_k^{(1)}\}_{k=0}^\infty$, and $\{\tilde{\phi}_k\}_{k=0}^\infty$ orthonormal bases of H , H^1 , and $\mathcal{H}$, respectively, where $Uf = (\langle \phi_k, f \rangle_H)_k$, $U^{(1)}f = (\langle \phi_k^{(1)}, f \rangle_{H^1})_k$, and $\tilde{U}f = (\langle \tilde{\phi}_k, f \rangle_{\mathcal{H}})_k$. As noted in Section 4.1, U , $U^{(1)}$, and $\tilde{U}$ are special cases of analysis operators. Together, U and $U^{(1)*}$ induce the linear isometry $V : \mathcal{B}(H^1, H) \rightarrow \mathcal{B}(\ell^2)$ with $VA = UAU^{(1)*}$, while U and $\tilde{U}$ induce the unitary map $\tilde{V} : \mathcal{B}_2(\mathcal{H}, H) \rightarrow \mathcal{B}_2(\ell^2)$ with $\tilde{V}A = U\tilde{U}^*$.

5.1 SEC representations of vector fields and their correspondence

Let v be an arbitrary bounded vector field in $H_{\mathcal{X}} \cap L_{\mathcal{X}}^\infty$. The SEC is based on the following three representations of v :

- (1) *Frame representation*, given by the sequence $\hat{v} = T'v \in \ell^2$, such that $v = \sum_{k=0}^\infty \hat{v}_k \alpha_k$.
- (2) *Dual frame representation*, given by the sequence $\hat{v}' = Tv \in \ell^2$, such that $v = \sum_{k=0}^\infty \hat{v}'_k \alpha'_k$.
- (3) *Operator representation*, given by the bounded operator $L = Wv \in \mathcal{B}(\ell^2)$, $W = \iota \circ V$, or the Hilbert-Schmidt operator $\tilde{L} = \tilde{W}v \in \mathcal{B}_2(\ell^2)$, $\tilde{W} = \iota_2 \circ \tilde{V}$, where $\iota : L_{\mathcal{X}}^\infty \rightarrow \mathcal{B}(H^1, H)$ and $\iota_2 : L_{\mathcal{X}}^\infty \rightarrow \mathcal{B}_2(\mathcal{H}, H)$ are the embeddings from Lemma 3.2. When we wish to distinguish between L and $\tilde{L}$ we will refer to the former as the *bounded operator representation* and the latter as the *Hilbert-Schmidt operator representation* of v .

Among these, the frame and dual frame representations only make use of the inner-product-space structure of $H_{\mathcal{X}} \cap L_{\mathcal{X}}^\infty$. The operator representations make use of the relationship between $L_{\mathcal{X}}^\infty$ and bounded or Hilbert-Schmidt operators on functions, which is special to vector fields.

As one might expect, the need to pass between these representations arises in a number of cases. On the one hand, many of the numerical procedures involving vector fields that one can envision being formulated via SEC produce output in the frame representation (that is, as linear combinations of frame elements), and in order to act with these vector fields on functions an operator representation is needed. For instance, the Galerkin approximation scheme for the eigenforms of the 1-Laplacian in Section 6 yields approximate eigenforms as linear combinations




FIGURE 5.1. Commutative diagram illustrating the frame, dual frame, bounded operator, and Hilbert-Schmidt operator representations of bounded vector fields in $L^\infty_{\mathcal{X}}$. $\mathcal{T} \subset \ell^2$ and $\mathcal{T}' \subset \ell^2$ are the ranges of the analysis and dual analysis operators, T and T' , respectively, restricted to $L^\infty_{\mathcal{X}}$. $\mathcal{W} \subset \mathcal{B}(\ell^2)$ and $\tilde{\mathcal{W}} \subset \mathcal{B}_2(\ell^2)$ are the ranges of the operators W and $\tilde{W}$ yielding the bounded operator and Hilbert-Schmidt operator representations vector fields in $L^\infty_{\mathcal{X}}$, respectively. The operators R and $\tilde{R}$ carry out the transformation from the bounded and Hilbert-Schmidt operator representations, respectively, to the dual frame representation. The Gramm and dual Gramm operators, G and G' , respectively, map between the frame and dual frame representations.

of 1-form frame elements, and in order to visualize these forms we compute the pushforwards of the their vector field duals into data space. The pushforward operation requires evaluation of the action of these vector fields on the embedding map of the manifold, which we carry out using the operator representation. Conversely, one may be given v in the operator representation (e.g., from data sampled along integral curves of the flow generated by v [30]), and then seek a frame representation for denoising and/or further use in a numerical procedure.

The correspondence between the frame and operator representations of vector fields in SEC is illustrated with a commutative diagram in Fig. 5.1. For the remainder of this section, we discuss aspects of this correspondence in the infinite-dimensional setting. Then, in Section 5.2 we examine finite-rank representations and their convergence properties.

We first consider how to pass from the frame representation $\hat{v} = T'v \in \ell^2$ of a bounded vector field v to the corresponding operator representation $L = Wv \in \mathcal{B}(\ell^2)$. Let $\mathcal{T}' = T'L^\infty_{\mathcal{X}} \subseteq \ell^2$ be the range of the dual analysis operator restricted to $L^\infty_{\mathcal{X}}$. Then, the inverse of $T'|_{L^\infty_{\mathcal{X}}}$ is given by $T^*|_{\mathcal{T}'}$, and we have

$$(5.1) \quad L = WT^*\hat{v}.$$

Since every bounded operator $A \in \mathcal{B}(\ell^2)$ is uniquely characterized by the coefficients $A_{ij} = \langle e_i, A e_j \rangle_{\ell^2}$ (its “matrix elements”), it suffices to compute

$$(5.2) \quad \begin{aligned} L_{ij} &= \langle e_i, L e_j \rangle_{\ell^2} = \langle e_i, U \iota(v) U^{(1)*} e_j \rangle_{\ell^2} = \langle U^* e_i, \iota(v) U^{(1)*} e_j \rangle_H \\ &= \langle \phi_i, \iota(v) \phi_j^{(1)} \rangle_H = \frac{1}{\sqrt{1 + \lambda_j}} \langle \phi_i, v(\phi_j) \rangle_H, \end{aligned}$$

given the sequence $\hat{v} = T'v \in \ell^2$. Substituting $v = T^* T'v = \sum_{m=0}^{\infty} \hat{v}_m \alpha_m$, we obtain

$$L_{ij} = \frac{1}{\sqrt{1 + \lambda_j}} \sum_{m=1}^{\infty} \hat{v}_m \langle \phi_i, \alpha_m \phi_j \rangle_H,$$

and since, for the frame elements in Theorem 4.3, $\alpha_m = s_m \phi_{p_m} \text{grad } \phi_{q_m}$, it follows that

$$(5.3) \quad L_{ij} = \sum_{m=0}^{\infty} \frac{s_m}{\sqrt{1 + \lambda_j}} \langle \phi_{p_m}, \text{grad } \phi_{q_m}(\phi_j) \rangle_{H_{\mathcal{X}}} = \sum_{m=0}^{\infty} \frac{s_m}{\sqrt{1 + \lambda_j}} G_{ijp_m q_m} \hat{v}_m,$$

where $G_{ijp_m q_m}$ are the Hodge inner products from (2.6). The above expression fully characterizes the operator WT^* in (5.1).

Next, we consider how to pass from the operator representation to the *dual* frame representation. Defining $\mathcal{W} = WL_{\mathcal{X}}^{\infty}$, this procedure amounts to computing the ℓ^2 sequence $\hat{v}' = Tv$ given $L \in \mathcal{W}$, where v is the unique $L_{\mathcal{X}}^{\infty}$ vector field with the operator representation $L = Wv$. In this case, the matrix elements $L_{ij} = \langle e_i, L e_j \rangle_{\ell^2}$ are known, and we have

$$(5.4) \quad \hat{v}'_k = \langle \alpha_k, v \rangle_{H_{\mathcal{X}}} = s_k \langle \phi_{p_k}, \text{grad } \phi_{q_k}, v \rangle_{H_{\mathcal{X}}} = s_k \langle \phi_{p_k}, v(\phi_{q_k}) \rangle_H = L_{p_k q_k}.$$

The above expression defines a linear map $R : \mathcal{W} \rightarrow \mathcal{T} = TL_{\mathcal{X}}^{\infty}$, that carries out the transformation from the operator representation to the dual frame representation. In particular, it follows from (5.4) that the components $\hat{v}'_k$ of the dual frame representation are equal to a subset of the matrix elements L_{ij} of the operator representation, as appropriate for the frame from Theorem 4.3 used. In the special case of the $\tilde{B}_{\mathcal{X}}$ frame, where all combinations of $\phi_i \text{grad } \phi_j$ are used as frame elements, the mapping $k \mapsto (p_k, q_k) \in \mathbb{N}^2$ is bijective, and the components of the dual frame representation are in one-to-one correspondence with the operator frame elements.

What remains to complete the portion of the commutative diagram in Fig. 5.1 involving the bounded operator representation is a map from the dual frame representation $\hat{v}' = Tv \in \mathcal{T}'$ to the frame representation $\hat{v} = T'v \in \mathcal{T}$ of $v \in L_{\mathcal{X}}^{\infty}$. This is accomplished by means of the dual Gramm operator G' (see (4.4)), i.e.,

$$(5.5) \quad \hat{v} = T'v = T' T'^* \hat{v}' = G' \hat{v}' = G^+ \hat{v}'.$$

Note that, unlike WT^* and R , we do not have a closed-form expression for the action of G' on sequences. Nevertheless, in Section 5.2 we will see that this operator can be approximated by a strongly convergent sequence of finite-rank operators associated with finite collections of frame elements, whose action can be explicitly evaluated.

We now turn attention to the task of passing between the frame and Hilbert-Schmidt operator representations. In this case, given $v \in L_{\mathfrak{X}}^\infty$, we have

$$\tilde{L} = \tilde{W}v = \tilde{W}T^*\hat{v} \in \mathcal{B}_2(\ell^2),$$

and we can expand $\tilde{L}$ in the Hilbert-Schmidt operator basis $\{e_{ij}\}$, viz.

$$\tilde{L} = \sum_{i,j=0}^{\infty} \tilde{L}_{ij} e_{ij}.$$

By construction of the e_{ij} , the expansion coefficients $\tilde{L}_{ij}$ are equal to the matrix elements of $\tilde{L}$ in the $\{e_i\}$ basis of ℓ^2 ; that is (cf. (5.2)),

$$\begin{aligned} \tilde{L}_{ij} &= \langle e_{ij}, \tilde{L} \rangle_{\text{HS}} = \langle e_i, \tilde{L}e_j \rangle_{\ell^2} = \langle e_i, U\iota_2(v)\tilde{U}^*e_j \rangle_{\ell^2} \\ &= \langle U^*e_i, \iota_2(v)\tilde{U}^*e_j \rangle_H = \langle \phi_i, \iota_2(v)\tilde{\phi}_j \rangle_H = e^{-\lambda_j/2} \langle \phi_i, v(\phi_j) \rangle_H. \end{aligned}$$

Using this result and proceeding as in (5.3), we find

$$L_{ij} = \sum_{n=0}^{\infty} e^{-\lambda_j/2} s_n G_{ijp_nq_n} \hat{v}_n,$$

which fully characterizes the operator $\tilde{W}T^*$. Observe now that in the special case of the $\tilde{B}_{\mathfrak{X}}$ frame, we have

$$e^{-\lambda_j/2} s_n G_{ijp_nq_n} = \frac{\langle b_{ij}, b_{p_nq_n} \rangle_{H_{\mathfrak{X}}}}{e^{(\lambda_j + \lambda_{q_n})/2}} = \langle \tilde{b}_{ij}, \tilde{b}_{p_nq_n} \rangle_{H_{\mathfrak{X}}} = \langle \alpha_{r_{ij}}, \alpha_n \rangle_{H_{\mathfrak{X}}} = G_{r_{ij}n}.$$

That is, for this frame, the matrix elements of the operator $\tilde{W}T^*$ are in one-to-one correspondence with the matrix elements of the Gramm operator.

The remaining transformations to pass from the Hilbert-Schmidt operator representation to the dual frame representation, and from the dual frame representation to the frame representation are completely analogous to those in (5.4) and (5.5), respectively, so we do not discuss them further here.

5.2 Finite-dimensional representations and their convergence properties

We now consider how to construct finite-rank analogs of the representations of vector fields introduced in Section 5.1. We begin by introducing the finite-rank (hence, compact) analysis and synthesis operators, $T_l = \pi_l T$ and $T_l^* = T^* \pi_l$, respectively, where

$$T_l v = (\langle \alpha_0, v \rangle_{H_{\mathfrak{X}}}, \dots, \langle \alpha_{l-1}, v \rangle_{H_{\mathfrak{X}}}, 0, 0, \dots), \quad T_l^*(c) = \sum_{k=0}^{l-1} c_k \alpha_k,$$

for $v \in H_{\mathfrak{X}}$ and $c = (c_0, c_1, \dots) \in \ell^2$. We also define finite-rank analysis and synthesis operators for the dual frame, $T'_l = \pi_l T'$ and $T'^*_l = T'^* \pi_l$, and the finite-rank Gramm operators $G_l = T_l T_l^*$, $G'_l = T'_l T'^*_l$. As in the case of the full frame and its dual (see Section 4.1), we can relate the finite-rank operators associated with the

frame and dual frames by pseudoinverses. In particular, since π_l is an orthogonal projection, we have $\pi_l^+ = \pi_l$, and therefore

$$(T_l^*)^+ = (T^* \pi_l)^+ = \pi_l^+ (T^*)^+ = \pi_l T' = T_l'.$$

Similarly, we have $(T_l'^*)^+ = T_l$ and $G_l^+ = G_l'$. It is important to note that by virtue of its relationship with G_l^+ , it is possible to compute the nonzero matrix elements $\langle e_i, G_l' e_j \rangle_{\ell^2}$ of G_l' numerically by computing the pseudoinverse of the $l \times l$ Gramm matrix $\mathbf{G}_l = [\langle e_i, G_l e_j \rangle]_{i,j=1}^l$, whose elements are known analytically. Thus, we can consider G_l' to be available to us in applications, albeit not in closed form.

In addition to finite-rank operators for the frame and dual frame, we will need finite-rank Fourier operators $U_l = \pi_l U$, $U_l^{(1)} = \pi_l U^{(1)}$, and $\tilde{U}_l = \pi_l \tilde{U}$ for the $\{\phi_j\}$, $\{\phi_j^{(1)}\}$, and $\{\tilde{\phi}_j\}$ bases of H , H^1 , and $\mathcal{H}$, respectively. These operators lead to finite-rank analogs $V_l : \mathcal{B}(H^1, H) \rightarrow \mathcal{B}(\ell^2)$ and $\tilde{V}_l : \mathcal{B}_2(\mathcal{H}, H) \rightarrow \mathcal{B}_2(\ell^2)$ of V and $\tilde{V}$, respectively, such that $V_l A = U_l A U_l^{(1)}$ and $\tilde{V}_l = U_l A \tilde{U}_l$. In addition, $W_l : L_{\mathfrak{X}}^\infty \rightarrow \mathcal{B}(\ell^2)$ and $\tilde{W}_l : L_{\mathfrak{X}}^\infty \rightarrow \mathcal{B}_2(\ell^2)$ with $W_l = \iota \circ V_l$ and $\tilde{W}_l = \iota_2 \circ \tilde{V}_l$ are finite-rank analogs of W and $\tilde{W}$, respectively. As with G_l and G_l' , the operators W_l and $\tilde{W}_l$ can be represented by $l \times l$ matrices.

With the above definitions, we can construct finite-rank analogs of the frame, dual frame, and operator representations for vector fields introduced in Section 5.1. In particular, $\hat{v}_l = T_l' v$, $\hat{v}_l' = T_l v$, $L_l = W_l v$, and $\tilde{L}_l = \tilde{W}_l v$ are respectively finite-rank frame, dual frame, bounded operator, and Hilbert-Schmidt operator representations of a vector field $v \in L_{\mathfrak{X}}^\infty$. To examine the convergence properties of these representations, note that since $\{e_k\}$ is a basis, the projection operators π_l converge strongly to the identity as $l \rightarrow \infty$ (i.e., $\pi_l c \rightarrow c$ for any $c \in \ell^2$). This implies that $T_l \rightarrow T$, $T_l' \rightarrow T'$, $U_l \rightarrow U$, $U_l^{(1)} \rightarrow U^{(1)}$, $\tilde{U}_l \rightarrow \tilde{U}$, and as a result $G_l \rightarrow G$, $G_l' \rightarrow G'$, $W_l \rightarrow W$, $\tilde{W}_l \rightarrow \tilde{W}$, where all limits are taken in the respective strong operator topologies. Therefore, $\hat{v}_l$ and $\hat{v}_l'$ converge to $\hat{v} = T v$ and $\hat{v}' = T' v$, respectively, in ℓ^2 norm, and L_l and $\tilde{L}_l$ converge to $L = W v$ and $\tilde{L} = \tilde{W} v$, respectively, strongly. In fact, since $\tilde{L}_l = \sum_{i,j=0}^{l-1} \tilde{L}_{ij} e_{ij}$, $\tilde{L}_{ij} = \langle e_i, \tilde{L} e_j \rangle_{\ell^2}$, it follows that $\tilde{L}_l$ converges to $\tilde{L} = \sum_{i,j=0}^\infty \tilde{L}_{ij} e_{ij}$ in Hilbert-Schmidt operator norm. This implies convergence in $\mathcal{B}(\mathcal{H}, H)$ operator norm, which implies in turn the strong convergence just stated. In effect, by restricting vector fields to act on the RKHS $\mathcal{H}$ containing functions of higher regularity than H^1 , the Hilbert-Schmidt operator representation allows for a stronger mode of convergence than the bounded operator representation.

Formulas for passing between the finite-rank frame, dual frame, and operator representations can be constructed analogously to those described in Section 5.1. As an application of the operator representation of vector fields, which was already mentioned in Sections 2.5 and 5.1, and will be employed in Section 8, we note that the pushforward map $F_* : \mathfrak{X} \rightarrow \Gamma \mathbb{R}^n$ on vector fields associated with an embedding $F : \mathcal{M} \rightarrow \mathbb{R}^n$ (here, $\Gamma \mathbb{R}^n \simeq \mathbb{R}^n$ is the space of smooth vector fields on $\mathbb{R}^n$) is given

by $F_*v = v(F)$, where $v \in \mathfrak{X}$ acts on F componentwise. As a result, we can consistently approximate F_*v by $L_l F$, where $L_l = W_l v$ is the operator representation of v . The result of this operation is an “arrow plot”, consisting of approximate tangent vectors to the image $F(\mathcal{M}) \subset \mathbb{R}^n$ of the manifold under the embedding.

6 Galerkin method for the 1-Laplacian

We now apply the framework developed in Section 4 to construct a Galerkin approximation scheme for the eigenvalues and eigenforms of the 1-Laplacian, Δ_1 . Eigenvalue problems for other differential operators of interest in exterior calculus can be formulated analogously. For notational simplicity, in this section we will use the symbols (v_k, φ_k) to denote a general eigenvalue-eigenfunction pair of Δ_1 , as opposed to the multi-index notation $(\lambda_{j,1}, \phi_{j,1})$ from Section 3.3.

6.1 Variational eigenvalue problem for the 1-Laplacian and its Galerkin approximation

We begin by stating the eigenvalue problem for the 1-Laplacian in strong form. This amounts to finding $\varphi_k \in \Omega^1$ and $v_k \in \mathbb{C}$ such that

$$(6.1) \quad \Delta_1 \varphi_k = v_k \varphi_k.$$

As is well known [47, 39], for the class of smooth closed Riemannian manifolds studied here, Δ_1 has a unique self-adjoint extension $\bar{\Delta}_1 : D(\bar{\Delta}_1) \rightarrow H_1$, with a dense domain $D(\bar{\Delta}_1) \simeq H_1^2 \subset H_1$ and a compact resolvent. As a result, we can obtain weak solutions to (6.1) by passing to a variational formulation, with an associated well posed Galerkin approximation scheme [4]. To construct this variational eigenvalue problem, we introduce the sesquilinear forms $L_\theta : H_1^1 \times H_1^1 \rightarrow \mathbb{C}$, $\theta > 0$, and $B : H_1^1 \times H_1^1 \rightarrow \mathbb{C}$, defined by

$$L_\theta(\psi, \omega) = E_{1,1}(\psi, \omega) + \theta \langle \psi, \omega \rangle_{H_1}, \quad B(\psi, \omega) = \langle \psi, \omega \rangle_{H_1},$$

where $E_{1,1}$ is the Dirichlet form on H_1 from (3.18). Note that $E_{1,1}(\psi, \omega)$ can be formally obtained by performing integration by parts on the expression $\langle \psi, \Delta_1 \omega \rangle_{H_1}$, taking ψ and ω to be smooth 1-form fields. The term in $L_\theta(\psi, \omega)$ proportional to θ is a regularization term, ensuring that L_θ has a coercivity property important to the well-posedness of our Galerkin scheme. Specifically, we seek to solve the following variational eigenvalue problem:

Definition 6.1 (eigenvalue problem for the 1-Laplacian, weak form). We say that $v_k \in \mathbb{C}$ and $\varphi_k \in H_1^1$, solve the weak eigenvalue problem for L_θ if the equality

$$L_\theta(\psi, \varphi_k) = v_k B(\psi, \varphi_k)$$

holds for all $\psi \in H_1^1$.

We refer to the solutions (v_k, φ_k) of the problem in Definition 6.1 as weak eigenvalues and eigenfunctions of the 1-Laplacian, respectively. Clearly, every classical eigenvalue (eigenfunction) from (6.1) is also a weak eigenvalue (eigenfunction).

We now discuss the well-posedness of the weak eigenvalue just formulated, and establish its Galerkin approximation using our frames for H_1^1 . A data-driven analog of this Galerkin method, utilizing frame elements of H_1^1 approximated from data, will be presented in Section 7.

Lemma 6.2. *The sesquilinear forms L_θ and B obey the bounds*

$$\begin{aligned} |L_\theta(\psi, \omega)| &\leq (1 + \theta) \|\psi\|_{H_1^1} \|\omega\|_{H_1^1}, \quad L_\theta(\omega, \omega) \geq \min\{\theta, 1\} \|\omega\|_{H_1^1}^2, \\ |B(\psi, \omega)| &\leq \|\psi\|_{H_1^1} \|\omega\|_{H_1^1}, \end{aligned}$$

for all $\psi, \omega \in H_1^1$.

Proof. The upper bounds on $|L_\theta(\psi, \omega)|$ and $|B(\psi, \omega)|$ follow directly from application of Cauchy-Schwartz inequalities. To verify the lower (coercivity) bound on $L_\theta(\omega, \omega)$, note that if $0 < \theta \leq 1$, we have

$$L_\theta(\omega, \omega) = \theta(\theta^{-1}E_{1,1}(\omega, \omega) + \|\omega\|_{H_1^1}^2) \geq \theta(E_{1,1}(\omega, \omega) + \|\omega\|_{H_1^1}^2) = \theta\|\omega\|_{H_1^1}^2.$$

The claim for $\theta > 1$ can be verified similarly. $\square$

According to classical results in spectral approximation theory [4], Lemma 6.2 implies that there exists a compact operator $A_\theta : H_1^1 \rightarrow H_1^1$ such that

$$(6.2) \quad L_\theta(\psi, A_\theta \omega) = B(\psi, \omega), \quad \forall \psi, \omega \in H_1^1.$$

This implies in turn that (v_k, φ_k) , $v_k \neq 0$, is a weak eigenvalue-eigenvector pair if and only if

$$A_\theta \varphi_k = v_k^{-1} \varphi_k.$$

Due to the above, A_θ can be thought of as a “solution operator” for the variational eigenvalue problem in Definition 6.1. In particular, the properties of spectral approximations to the solutions of that problem can be analyzed in terms of approximations of the eigenvalue problem for A_θ . This approach leads to a Galerkin approximation scheme, as follows.

Let $\Pi_1, \Pi_2, \dots$ be a family of finite-rank projection operators on H_1^1 , converging pointwise to the identity; that is, $\Pi_l^2 = \Pi_l$ and $\lim_{l \rightarrow \infty} \Pi_l \omega = \omega$ for every $\omega \in H_1^1$. Let also W_l be the closed subspaces of H_1^1 defined as $W_l = \text{ran } \Pi_l$. These spaces, which will be constructed explicitly below, will be our Galerkin approximation spaces. In particular, we will solve:

Definition 6.3 (eigenvalue problem for the 1-Laplacian, Galerkin approximation). We say that $v_{k,l} \in \mathbb{C}$ and $\varphi_{k,l} \in W_l$, solve the Galerkin truncated eigenvalue problem for L_θ if the equality

$$L_\theta(\psi, \varphi_{k,l}) = v_{k,l} B(\psi, \varphi_{k,l})$$

holds for all $\psi \in W_l$.

Given a basis $\{v_{0,l}, \dots, v_{w_l-1,l}\}$ of W_l , where $w_l = \dim W_l$, the eigenvalue problem in Definition 6.3 is equivalent to a generalized matrix eigenvalue problem. That is, $(v_{k,l}, \varphi_{k,l})$ with $\varphi_{k,l} = \sum_{j=0}^{w_l-1} c_j v_{j,l}$ is a solution if and only if

$$(6.3) \quad L_l \vec{a}_{k,l} = v_{k,l} B_l \vec{a}_{k,l},$$

where L_l and B_l are the $w_l \times w_l$ matrices with elements $[L_l]_{ij} = L_\theta(v_{i,l}, v_{j,l})$ and $[B_l]_{ij} = B(v_{i,l}, v_{j,l})$, respectively, and $\vec{a}_{k,l} = (a_0, \dots, a_{w_l-1})^\top \in \mathbb{R}^{w_l}$ is the w_l -dimensional column vector containing the expansion coefficients of $\varphi_{k,l}$ in the $\{v_{j,l}\}$ basis of W_l . In addition, one can verify that $(v_{k,l}, \varphi_{k,l})$ solves the eigenvalue problem in Definition 6.3 if and only if

$$A_{\theta,l} \varphi_{k,l} = v_{k,l}^{-1} \varphi_{k,l},$$

where $A_{\theta,l} : H_1^1 \rightarrow H_1^1$ is the finite-rank operator given by $A_{\theta,l} = \Pi_l A_\theta$. Note that this operator satisfies (cf. (6.2))

$$L_\theta(\psi, A_{\theta,l} \omega) = B(\psi, \omega), \quad \forall \psi \in W_l, \quad \forall \omega \in H_1^1.$$

Now, because A_θ is compact, the fact that Π_l converges pointwise to the identity implies that $A_{\theta,l}$ converges to A_θ in norm. This implies in turn that for every eigenvalue v_k of A_θ (which is nonzero by coercivity of L_θ , and thus isolated and with finite geometric multiplicity by compactness of A_θ) there exists a sequence $v_{k,l}$ of eigenvalues of $A_{l,\theta}$ converging as $l \rightarrow \infty$ to v_k . Moreover, for every eigenfunction φ_k in the eigenspace of A_θ at eigenvalue v_k , there exists a sequence $\varphi_{k,l}$ of eigenfunctions of $A_{\theta,l}$ at eigenvalue $v_{k,l}$ converging in H_1^1 norm to φ_k . This establishes convergence of the solutions of the Galerkin scheme in Definition 6.3 to those of the eigenvalue problem in Definition 6.1.

6.2 Construction of the Galerkin approximation spaces

What remains is to construct the projection operators Π_l and the associated subspaces W_l . Here, we will construct these operators making use of the result established in Theorem 4.4 that $\{b_1^{ij}\}$ with $i \in \{0, 1, \dots\}$ and $j \in \{1, \dots, J\}$ is a frame of H_1^1 . As in Section 5, for notational simplicity, we set $\alpha_k = b_1^{p_k q_k}$, where $k \mapsto (p_k, q_k)$ is any ordering of the (i, j) indices in b_1^{ij} with $k \in \{0, 1, \dots\}$. We also let $T : H_1^1 \rightarrow \ell^2$, $T^* : \ell^2 \rightarrow H_1^1$, and $S = T^* T : \ell^2 \rightarrow \ell^2$ be the corresponding analysis, synthesis, and frame operators, respectively. We also consider the finite-rank operators $T_l = \pi_l T$, $T_l^* = T^* \pi_l$, $S_l = T_l^* T_l$, and $G_l = T_l T_l^*$, associated with the canonical orthogonal projection operators $\pi_l : \ell^2 \rightarrow \ell^2$. All of these operators converge strongly to their infinite-rank counterparts as $l \rightarrow \infty$; see Section 5.2.

By construction, at each l , the operator G_l is a positive-semidefinite, self-adjoint, compact operator on ℓ^2 . As a result, there exists an orthonormal basis $\{u_{k,l}\}_{k=1}^\infty$ of ℓ^2 consisting of eigenvectors of G_l . We denote the corresponding eigenvalues by $\eta_{k,l}$, and order the eigenpairs $(\eta_{k,l}, u_{k,l})$ in order of decreasing $\eta_{k,l}$.

We let w_l be the number of nonzero eigenvalues, where $w_l = \text{rank } T_l \leq l$. The following lemma, whose proof is left to the reader, summarizes certain properties of the $(\eta_{k,l}, u_{k,l})$ eigenpairs.

Lemma 6.4. *The eigenpairs $(\eta_{k,l}, u_{k,l})$ have the following properties:*

- (i) $\eta_{k,l}$ is bounded above by the operator norm of G .
- (ii) $\eta_{k,l}$ is a nonzero eigenvalue of G_l if and only if it is a nonzero eigenvalue of S_l . Moreover, the vectors

$$v_{k,l} = \frac{1}{\eta_{k,l}^{1/2}} T^* u_{k,l}, \quad 0 \leq k \leq w_l - 1,$$

are orthonormal eigenvectors of S_l corresponding to the same eigenvalues as $u_{k,l}$.

By Lemma (6.4)(ii), we can approach the problem of constructing orthonormal sets in H_1^1 , consisting of eigenvectors of S_l corresponding to nonzero eigenvalues, through the eigenvalue problem of G_l . This is advantageous, since the eigenvectors of G_l corresponding to nonzero eigenvalues can be computed via the eigenvalue problem of the $l \times l$ Grammian matrix for H_1^1 introduced in Section 2.5, which we denote here by G_l^1 to make its dependence on $l = MJ$ explicit. In particular, observe that given any $c = (c_0, c_1, \dots) \in \ell^2$, we have $d = G_l c = (d_0, \dots, d_{l-1}, 0, 0, \dots)$, where the first l elements of d are given by $d_j = \sum_{i=0}^{l-1} G_{ij} c_j$, with

$$G_{ij} = \langle e_i, G e_j \rangle_{\ell^2} = \langle T^* e_i, T^* e_j \rangle_{H_1^1} = \langle \alpha_i, \alpha_j \rangle_{H_1^1},$$

and the inner products $\langle \alpha_i, \alpha_j \rangle_{H_1^1}$ can be computed in closed form via the formulas in Table 2.2. The above implies that the nonzero eigenvalues $\eta_{k,l}$ of G_l are equal to the nonzero eigenvalues of G_l^1 , and the corresponding eigenvectors $\vec{u}_{k,l} = (u_{0,k,l}, \dots, u_{l-1,k,l})^\top$ of that matrix yield $u_{k,l} = (u_{0,k,l}, \dots, u_{l-1,k,l}, 0, 0, \dots)$. We thus obtain

$$v_{k,l} = \frac{1}{\eta_{k,l}^{1/2}} \sum_{j=0}^{l-1} u_{j,k,l} \alpha_j.$$

Based on these considerations, we define our Galerkin approximation spaces as

$$(6.4) \quad W_l = \text{span}\{v_{1,l}, \dots, v_{w_l,l}\},$$

and the projection operators $\Pi_l : H_1^1 \rightarrow H_1^1$ as orthogonal projectors onto those subspaces.

Lemma 6.5. *The sequence Π_l of projection operators converges pointwise to the identity; that is, for any $\omega \in H_1^1$, $\lim_{l \rightarrow \infty} \Pi_l \omega = \omega$.*

Proof. Since the frame operator S has a bounded inverse, it suffices to show that $S \Pi_l \omega$ converges to $S \omega$ for any $\omega \in H_1^1$. To verify this, observe first that

$$T \Pi_l \omega = \sum_{k=0}^{w_l-1} \langle v_{k,l}, \omega \rangle_{H_1^1} T v_{k,l} = \sum_{k=0}^{w_l-1} \langle u_{k,l}, T \omega \rangle_{\ell^2} u_{k,l} = \pi_l \sum_{k=0}^{w_l-1} \langle u_{k,l}, T \omega \rangle_{\ell^2} u_{k,l} =$$

$$= \pi_l \sum_{k=0}^{\infty} \langle u_{k,l}, T\omega \rangle_{\ell^2} u_{k,l} = \pi_l T\omega.$$

We therefore have

$$S\Pi_l\omega = T^*T\Pi_l\omega = T^*\pi_l T\omega = S_l\omega,$$

which converges to $S\omega$ by the pointwise convergence of S_l to S . $\square$

Lemma 6.5 implies that with the choice of approximation spaces in (6.4), the Galerkin scheme in Definition 6.3 converges. Moreover, all of the matrix elements of the associated sesquilinear forms can be evaluated using Lemma 6.4 in conjunction with the formulas listed in Section 4.1 and B. Explicitly, the $w_l \times w_l$ matrices appearing in the generalized eigenvalue problem in (6.3) are given by

$$(6.5) \quad L_l = H_l^{-1/2} U_l^\dagger (E_l - \theta G_l^1) U_l H_l^{-1/2}, \quad B_l = H_l^{-1/2} U_l^\dagger G_l U_l H_l^{-1/2},$$

where H_l is a $w_l \times w_l$ diagonal matrix with $[H_l]_{ii} = \eta_{i,l}$, U_l an $l \times w_l$ matrix with $[U_l]_{ij} = u_{i,j,l}$, E_l an $l \times l$ matrix with $[E_l]_{ij} = E_{1,1}(\alpha_i, \alpha_j)$ determined from Table 2.2, and G_l the $l \times l$ Grammian matrix for H_1 with $[G_l]_{ij} = \langle \alpha_i, \alpha_j \rangle_{H_1}$ determined from Table 2.2. Note that the matrices L_l and B_l in (6.5) differ from the corresponding matrices appearing in the generalized eigenvalue problem in (2.8) in that they include H_l - and θ -dependent terms, which do not appear in (2.8). Due to the absence of these terms, (2.8) represents the problem in Definition 6.3 in a basis of W_l that exhibits unbounded growth of H_1^1 norm with l (controlled by the H_l terms in (6.5)), and is also not compatible with the coercivity condition in Lemma 6.2 (enforced by the θ -dependent terms). While both of these issues could potentially affect the numerical conditioning of (2.8), especially at large spectral orders l , in the examples studied in Section 8 we found that (2.8) and (6.5) perform comparably.

7 Data-driven approximation

All of the schemes in Sections 4 and 6 can be implemented given knowledge of the eigenvalues and eigenfunctions of the Laplacian on functions. The problem of approximating these objects from finite sets of points in a purely data-driven manner (that is, without requiring explicit knowledge of the manifold $\mathcal{M}$ and/or its embedding in data space) has been studied extensively in recent years [7, 18, 44, 52, 11, 12, 14, 8, 43, 49, 50, 32], leading to the development of approximation techniques with well-established pointwise and spectral convergence guarantees. In this section, we summarize the main properties of one such techniques, namely the diffusion maps algorithm [18], and describe the analogs of the methods of Sections 4 and 6 in a data-driven, discrete setting. It should be noted that, while generally expected on the basis of results for related techniques [8, 43, 49, 50], to our knowledge, the spectral convergence result for diffusion maps in Theorem 7.3 below has not been stated elsewhere in the literature, so we have included here a self-contained proof for completeness.

7.1 Assumptions for data-driven approximation

We consider that the Riemannian manifold $(\mathcal{M}, g)$ is embedded in n -dimensional Euclidean space by means of a smooth, isometric embedding $F : \mathcal{M} \rightarrow \mathbb{R}^n$. Using the notation $y \cdot z = \sum_{i=1}^n y^i z^i$ to represent the canonical Euclidean inner product between two vectors $y = (y^1, \dots, y^n)$ and $z = (z^1, \dots, z^n)$ in $\mathbb{R}^n$, we thus have $g_x(u, v) = F_{*,x}u \cdot F_{*,x}v$ for any point $x \in \mathcal{M}$ and tangent vectors $u, v \in T_x\mathcal{M}$, where $F_{*,x} : T_x\mathcal{M} \rightarrow T_{F(x)}\mathbb{R}^n$ is the pushforward map on tangent vectors associated with F , and we have used the canonical isomorphism $T_{F(x)}\mathbb{R}^n \simeq \mathbb{R}^n$. If the embedding F is not isometric, then the method described below can be modified via the techniques developed in [14] to yield approximations of Laplacian eigenvalues and eigenfunctions with respect to any (known) metric g .

We also assume that we have access to a dataset consisting of N samples $y_1, \dots, y_N$ in $\mathbb{R}^n$ with $y_j = F(x_j)$ taken on a sequence of distinct points $x_1, x_2, \dots$ in $\mathcal{M}$, which is equidistributed with respect to a smooth sampling measure σ supported on $\mathcal{M}$. By that, we mean that given any continuous function $f : \mathcal{M} \rightarrow \mathbb{C}$, the result

$$(7.1) \quad \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N f(x_j) = \int_{\mathcal{M}} f d\sigma$$

holds, and moreover σ has a smooth density $\rho = d\sigma/d\mu$ with respect to the Riemannian measure μ , bounded away from zero. Such an equidistributed sequence can be provided, e.g., by i.i.d. points on $\mathcal{M}$ (as is commonly assumed in machine learning applications), or by an orbit of an ergodic dynamical system (in which case, the x_j are not independent). The requirement in (7.1) is equivalent to assuming that the sequence of sampling measures $\sigma_N = N^{-1} \sum_{n=0}^{N-1} \delta_{x_j}$ weak-converges to σ ; that is,

$$\lim_{N \rightarrow \infty} \int_{\mathcal{M}} f d\sigma_N = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{j=1}^N f(x_j) = \int_{\mathcal{M}} f d\sigma, \quad \forall f \in C(\mathcal{M}).$$

In this data-driven setting, we will be working with the N -dimensional Hilbert space $L^2(\mathcal{M}, \sigma_N)$ associated with the discrete sampling measure σ_N , equipped with the inner product

$$\langle f, h \rangle_{L^2(\mathcal{M}, \sigma_N)} = \int_{\mathcal{M}} f^* h d\sigma_N = \frac{1}{N} \sum_{j=1}^N f^*(x_j) h(x_j).$$

Note that $L^2(\mathcal{M}, \sigma_N)$ consists of equivalence classes of functions on $\mathcal{M}$ which are equal up to sets of zero σ_N measure; that is, two functions $f : \mathcal{M} \rightarrow \mathbb{C}$ and $h : \mathcal{M} \rightarrow \mathbb{C}$ satisfying $f(x_j) = h(x_j)$ for all $j \in \{1, \dots, N\}$, but taking arbitrarily different values at other points, lie in the same $L^2(\mathcal{M}, \sigma_N)$ equivalence class. Because the points $x_1, \dots, x_N$ are all distinct, $L^2(\mathcal{M}, \sigma_N)$ is isomorphic as a Hilbert space to $\mathbb{C}^N$ equipped with the normalized dot product $f \cdot g / N$, but here we prefer to work with $L^2(\mathcal{M}, \sigma_N)$ to emphasize the fact that our data-driven approximation spaces

contain equivalence classes of functions on the same underlying manifold as the equivalence classes comprising $L^2(\mathcal{M}, \mu)$.

7.2 Kernel method for the eigenvalues and eigenfunctions of the Laplacian on functions

Following the approach introduced in the diffusion maps algorithm [18], and further generalized in [12], we compute data-driven approximations of the eigenvalues and eigenfunctions of the Laplacian on functions through the eigenvalues and eigenfunctions of a kernel integral operator approximating, in a suitable sense, the heat operator $e^{-\tau\bar{\Delta}}$, $\tau \geq 0$, on $L^2(\mathcal{M}, \mu)$, where $\bar{\Delta}$ is the self-adjoint Laplacian (see Section 3.1). This kernel integral operator is constructed from a smooth, exponentially decaying kernel $k_\varepsilon : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}_+$, bounded away from zero. Here, as a concrete example, we work with a Gaussian kernel,

$$k_\varepsilon(x, x') = \exp\left(-\frac{\|F(x) - F(x')\|_{\mathbb{R}^n}^2}{\varepsilon}\right),$$

where ε is a positive bandwidth parameter. Approximation techniques based on other classes of kernels, including kernels with variable bandwidth functions [11, 14], have equivalent asymptotic properties while generally achieving higher performance in terms of approximation accuracy and noise robustness, particularly in applications with large variations in the sampling density ρ .

Having specified an appropriate kernel, we introduce the associated kernel integral operators $\hat{K}_\varepsilon : L^2(\mathcal{M}, \mu) \rightarrow C(\mathcal{M})$ and $\hat{K}_{\varepsilon,N} : L^2(\mathcal{M}, \sigma_N) \rightarrow C(\mathcal{M})$, where

$$\hat{K}_\varepsilon f = \int_{\mathcal{M}} k_\varepsilon(\cdot, x) f(x) \rho(x) d\mu(x), \quad \hat{K}_{\varepsilon,N} h = \int_{\mathcal{M}} k_\varepsilon(\cdot, x) h(x) d\sigma_N(x).$$

Composing $\hat{K}_\varepsilon$ with the canonical inclusion operator $\iota : C(\mathcal{M}) \rightarrow L^2(\mathcal{M}, \mu)$, we also define $\tilde{K}_\varepsilon : L^2(\mathcal{M}, \mu) \rightarrow L^2(\mathcal{M}, \mu)$ and $K_\varepsilon : C(\mathcal{M}) \rightarrow C(\mathcal{M})$, where $\tilde{K}_\varepsilon = \iota \hat{K}_\varepsilon$ and $K_\varepsilon = \hat{K}_\varepsilon \iota$. Similarly, we define $\tilde{K}_{\varepsilon,N} : L^2(\mathcal{M}, \sigma_N) \rightarrow L^2(\mathcal{M}, \sigma_N)$ and $K_{\varepsilon,N} : C(\mathcal{M}) \rightarrow C(\mathcal{M})$, where $\tilde{K}_\varepsilon = \iota_N \hat{K}_{\varepsilon,N}$, $K_{\varepsilon,N} = \hat{K}_{\varepsilon,N} \iota_N$, and $\iota_N : C(\mathcal{M}) \rightarrow L^2(\mathcal{M}, \sigma_N)$ is the canonical restriction operator from $C(\mathcal{M})$ to $L^2(\mathcal{M}, \sigma_N)$.

Proposition 7.1. *The operators $\tilde{K}_\varepsilon$, K_ε , $\tilde{K}_{\varepsilon,N}$, and $K_{\varepsilon,N}$ have the following properties.*

(i) *They are all compact.*

(ii) *As $N \rightarrow \infty$, $K_{\varepsilon,N}$ converges pointwise to K_ε ; that is, for any $f \in C(\mathcal{M})$, we have $\lim_{N \rightarrow \infty} K_{\varepsilon,N} f = K_\varepsilon f$ in uniform norm.*

Proof. (i) That $\tilde{K}_{\varepsilon,N}$ and $K_{\varepsilon,N}$ are compact follows immediately from the fact that they have finite rank. The compactness of $\tilde{K}_\varepsilon$ follows from the facts that k_ε is a Hilbert-Schmidt kernel on $L^2(\mathcal{M} \times \mathcal{M}, \sigma \times \sigma)$ (i.e., $\int_{\mathcal{M}} \int_{\mathcal{M}} k_\varepsilon(x, x') d\sigma(x) d\sigma(x') < \infty$), and $L^2(\mathcal{M}, \mu)$ and $L^2(\mathcal{M}, \sigma)$ are isomorphic Hilbert spaces (by smoothness of ρ and compactness of $\mathcal{M}$). The compactness of K_ε can be verified using the Arzelà-Ascoli theorem in conjunction with the continuity of k_ε ; see, e.g., [19].

(ii) The claim is a direct consequence of (7.1) and the fact that $\int_{\mathcal{M}} f \rho d\mu = \int_{\mathcal{M}} f d\sigma$ for any $f \in L^2(\mathcal{M}, \mu)$. $\square$

Proposition 7.1(ii) shows that, on $C(\mathcal{M})$, we can approximate kernel integral operators with respect to the Riemannian measure by kernel integral operators with respect to the sampling measure. However, the pointwise convergence established there does not, in general, imply spectral convergence for these operators. Moreover, in applications we work in the finite-dimensional Hilbert space $L^2(\mathcal{M}, \sigma_N)$, as opposed to the infinite-dimensional Banach space $C(\mathcal{M})$, which necessitates establishing connections between the spectral properties of $\tilde{K}_{\varepsilon,N}$ and $K_{\varepsilon,N}$. Another issue that must be addressed is that of approximating the heat operator $e^{-\tau\tilde{\Delta}}$ by a suitable modification of $\tilde{K}_{\varepsilon}$.

Following [18, 12], we proceed by normalizing k_{ε} to construct a smooth ergodic Markov kernel $p_{\varepsilon} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}_+$, satisfying $\int_{\mathcal{M}} p_{\varepsilon}(x, \cdot) d\sigma = 1$ for all $x \in \mathcal{M}$. For that, we introduce the normalization functions $r_{\varepsilon} = \hat{K}_{\varepsilon}1$ and $l_{\varepsilon} = \hat{K}_{\varepsilon}(1/r_{\varepsilon})$, which are both smooth, positive, and bounded away from zero, and define

$$p_{\varepsilon}(x, x') = \frac{k_{\varepsilon}(x, x')}{l_{\varepsilon}(x)r_{\varepsilon}(x')}.$$

The Markov property of p_{ε} then follows by construction, and its ergodicity follows from the fact that it is bounded below. In the data-driven case, we define

$$(7.2) \quad r_{\varepsilon,N} = \hat{K}_{\varepsilon,N}1, \quad l_{\varepsilon,N} = \hat{K}_{\varepsilon,N}(1/r_{\varepsilon,N}), \quad p_{\varepsilon,N}(x, x') = \frac{k_{\varepsilon}(x, x')}{l_{\varepsilon,N}(x)r_{\varepsilon,N}(x')},$$

and $p_{\varepsilon,N}$ is a smooth Markov kernel satisfying $\int_{\mathcal{M}} p_{\varepsilon,N}(\cdot, x) d\sigma_N(x) = 1$. As in the case of the kernel k_{ε} , we define the kernel integral operators $\hat{P}_{\varepsilon} : L^2(\mathcal{M}, \mu) \rightarrow C(\mathcal{M})$ and $\hat{P}_{\varepsilon,N} : L^2(\mathcal{M}, \sigma_N) \rightarrow C(\mathcal{M})$ via

$$\hat{P}_{\varepsilon}f = \int_{\mathcal{M}} p_{\varepsilon}(\cdot, x)f(x)\rho(x) d\mu(x), \quad \hat{P}_{\varepsilon,N}h = \int_{\mathcal{M}} p_{\varepsilon,N}(\cdot, x)h(x) d\sigma_N(x),$$

and also introduce the operators $\tilde{P}_{\varepsilon} : L^2(\mathcal{M}, \mu) \rightarrow L^2(\mathcal{M}, \mu)$, $P_{\varepsilon} : C(\mathcal{M}) \rightarrow C(\mathcal{M})$, $\tilde{P}_{\varepsilon,N} : L^2(\mathcal{M}, \sigma_N) \rightarrow L^2(\mathcal{M}, \sigma_N)$, and $P_{\varepsilon,N} : C(\mathcal{M}) \rightarrow C(\mathcal{M})$, where $\tilde{P}_{\varepsilon} = \iota\hat{P}_{\varepsilon}$, $P_{\varepsilon} = \hat{P}_{\varepsilon}\iota$, $\tilde{P}_{\varepsilon,N} = \iota_N\hat{P}_{\varepsilon,N}$, and $P_{\varepsilon,N} = \hat{P}_{\varepsilon,N}\iota_N$. These operators have the analogous properties to those stated in Proposition 7.1. Among them, $\tilde{P}_{\varepsilon,N}$ is represented by the Markov matrix $\mathbf{P}$ from (2.1) (note that, due to cancellation of terms, the normalization of $k_{\varepsilon}(x, x')$ by $r_{\varepsilon,N}$ and $l_{\varepsilon,N}$ to construct $p_{\varepsilon,N}$ is equivalent to the normalization procedure used to construct $\mathbf{P}$; see [12] for details). The following theorem summarizes how the nonzero eigenvalues and corresponding eigenvectors of $\tilde{P}_{\varepsilon}$ can be approximated by the corresponding eigenvalues and eigenvectors of $\tilde{P}_{\varepsilon,N}$, which are accessible from data as described in Section 2.1.

Theorem 7.2. *The following hold:*

(i) $\Lambda_{j,\varepsilon}$ is a nonzero eigenvalue of $\tilde{P}_{\varepsilon}$ if and only if it is a nonzero eigenvalue of P_{ε} . Similarly, $\Lambda_{j,\varepsilon,N}$ is a nonzero eigenvalue of $\tilde{P}_{\varepsilon,N}$ if and only if it is a nonzero

eigenvalue of $P_{\varepsilon,N}$. Moreover, the $\Lambda_{j,\varepsilon,N}$ and $\Lambda_{j,\varepsilon,N}$ are all real, and thus lie in the interval $[-1, 1]$ by the Markov property of p_ε and $p_{\varepsilon,N}$.

(ii) If $\tilde{\phi}_{j,\varepsilon} \in L^2(\mathcal{M}, \mu)$ and $\vec{\phi}_{j,\varepsilon,N} \in L^2(\mathcal{M}, \sigma_N)$ are eigenvectors of P_ε and $P_{\varepsilon,N}$ at nonzero eigenvalues $\Lambda_{j,\varepsilon}$ and $\Lambda_{j,\varepsilon,N}$, respectively, then the smooth functions

$$\phi_{j,\varepsilon} = \frac{1}{\Lambda_{j,\varepsilon}} \hat{P}_\varepsilon \tilde{\phi}_{j,\varepsilon}, \quad \phi_{j,\varepsilon,N} = \frac{1}{\Lambda_{j,\varepsilon,N}} \hat{P}_{\varepsilon,N} \vec{\phi}_{j,\varepsilon,N}$$

are eigenvectors of P_ε and $P_{\varepsilon,N}$, respectively, at the same eigenvalues.

(iii) For every nonzero eigenvalue $\Lambda_{j,\varepsilon}$ of P_ε , the sequence of eigenvalues $\Lambda_{j,\varepsilon,N}$ of $P_{\varepsilon,N}$ converges as $N \rightarrow \infty$ to $\Lambda_{j,\varepsilon}$. Moreover, if $\phi_{j,\varepsilon}$ is an eigenvector of P_ε corresponding to eigenvalue $\Lambda_{j,\varepsilon}$, then there exists a sequence of eigenvectors $\phi_{j,\varepsilon,N}$ of $P_{\varepsilon,N}$ at eigenvalues $\Lambda_{j,\varepsilon,N}$, converging as $N \rightarrow \infty$ to $\phi_{j,\varepsilon}$ in uniform norm.

Proof. (i,ii) The claims on the relationships between the eigenvalues and eigenvectors of $\tilde{P}_\varepsilon$ and P_ε (and those of $\tilde{P}_{\varepsilon,N}$ and $P_{\varepsilon,N}$) can be verified from the definition of these operators. In addition, it can be verified that $\tilde{P}_\varepsilon$ and $\tilde{P}_{\varepsilon,N}$ are related to self-adjoint operators by similarity transformations, which implies that their eigenvalues are real.

(iii) The convergence of the eigenvalues follows by showing that the operators $P_{\varepsilon,N}$ converge compactly to P_ε (a stronger notion of convergence than pointwise convergence, but weaker than convergence in operator norm); see [52] for additional details. A proof of the convergence of the eigenvectors can be found in [19]. We also note that [52] establishes pointwise convergence of projection operators onto the corresponding eigenspaces. $\square$

That we can approximate eigenvalues and eigenvectors of $\tilde{P}_\varepsilon$ through eigenvalues and eigenvectors of $\tilde{P}_{\varepsilon,N}$ is important since the eigenvalue problem for the latter operator is equivalent to the numerically solvable $N \times N$ matrix eigenvalue problem for P . Note that by ergodicity and the Markov property, the eigenvalues $\Lambda_{j,\varepsilon}$ and $\Lambda_{j,\varepsilon,N}$ can be ordered as $1 = \Lambda_{0,\varepsilon} > \Lambda_{1,\varepsilon} \geq \Lambda_{2,\varepsilon} \geq \dots$ and $1 = \Lambda_{0,\varepsilon,N} > \Lambda_{1,\varepsilon,N} \geq \Lambda_{2,\varepsilon,N} \geq \dots \Lambda_{N-1,\varepsilon,N}$, respectively. We will adopt these orderings for the remainder of the paper. What remains is to establish how the eigenvalues and eigenvectors of $\tilde{P}_\varepsilon$ approximate in turn eigenvalues and eigenvectors of the Laplacian.

Theorem 7.3. *For every $j \in \mathbb{N}_0$ and as $\varepsilon \rightarrow 0^+$, the quantities $\lambda_{j,\varepsilon} = (1 - \Lambda_{j,\varepsilon})/\varepsilon$ converge to the Laplace-Beltrami eigenvalue λ_j . Moreover, for each Laplace-Beltrami eigenfunction ϕ_j corresponding to λ_j , there exist eigenfunctions $\tilde{\phi}_{j,\varepsilon}$ of $\tilde{P}_\varepsilon$ corresponding to eigenvalue $\lambda_{j,\varepsilon}$ converging, as $\varepsilon \rightarrow 0^+$, to ϕ_j in $L^2(\mathcal{M}, \mu)$ norm.*

Proof. By [18, Proposition 10], as $\varepsilon \rightarrow 0^+$, the family of operators $\tilde{\Delta}_\varepsilon := (I - \tilde{P}_\varepsilon)/\varepsilon$ converges to $\tilde{\Delta}$, pointwise on $C^\infty(\mathcal{M}) \cap L^2(\mathcal{M}, \mu)$. Moreover, as can be directly verified from the normalization procedure in (7.2), $\tilde{P}_\varepsilon$ is related to a self-adjoint operator $\bar{P}_\varepsilon : L^2(\mathcal{M}, \mu) \rightarrow L^2(\mathcal{M}, \mu)$ by a similarity transformation by a bounded multiplication operator with a bounded inverse; specifically, $\tilde{P}_\varepsilon = D_\varepsilon^{1/2} \bar{P}_\varepsilon D_\varepsilon^{-1/2}$,

where $D_\varepsilon : L^2(\mathcal{M}, \mu) \rightarrow L^2(\mathcal{M}, \mu)$ is the multiplication operator by the smooth, strictly positive function $d_\varepsilon = r_\varepsilon / (l_\varepsilon \rho)$. It can also be shown through small- ε expansions that, as $\varepsilon \rightarrow 0^+$, d_ε converges to a constant in H^2 norm. By virtue of the above, $\bar{\Delta}_\varepsilon = (I - \bar{P}_\varepsilon) / \varepsilon$ is a family of self-adjoint, compact operators, with the same eigenvalues $\lambda_{j,\varepsilon}$ as $\tilde{\Delta}_\varepsilon$, converging pointwise to $\bar{\Delta}$ on $C^\infty(\mathcal{M}) \cap L^2(\mathcal{M}, \mu)$. Now, observe that $C^\infty(\mathcal{M})$ is a *core* for $\bar{\Delta}$ (i.e., $\bar{\Delta}$ is equal to the closure of its restriction on $C^\infty(\mathcal{M})$), which follows immediately from the fact that, on manifolds without boundary, the Laplacian Δ on $C^\infty(\mathcal{M})$ functions is essentially self-adjoint). By results from spectral approximation theory of self-adjoint operators [20, Proposition 10.1.18], this actually implies that the eigenvalues $\lambda_{j,\varepsilon}$ converge to λ_j , as claimed. Moreover, by related results [25, Chapter X.7, Corollary 3], the orthogonal projections to the eigenspaces of $\bar{\Delta}_\varepsilon$ corresponding to $\lambda_{j,\varepsilon}$ converge pointwise to the projectors onto the eigenspaces of $\bar{\Delta}$ at eigenvalue λ_j . The latter, in conjunction with the fact that D_ε converges to a multiplication operator by a constant function, leads to the claim on the convergence of the eigenfunctions of $\tilde{P}_\varepsilon$ to Laplace-Beltrami eigenfunctions as $\varepsilon \rightarrow 0^+$. $\square$

In summary, we can conclude from Theorems 7.2 and 7.3 that the quantities

$$(7.3) \quad \lambda_{j,\varepsilon} = \frac{1 - \Lambda_{j,\varepsilon}}{\varepsilon}, \quad \lambda_{j,\varepsilon,N} = \frac{1 - \Lambda_{j,\varepsilon,N}}{\varepsilon},$$

converge to the eigenvalues of the Laplacian, i.e.,

$$\lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \lambda_{j,\varepsilon,N} = \lim_{\varepsilon \rightarrow 0^+} \lambda_{j,\varepsilon} = \lambda_j.$$

Moreover, for any eigenfunction $\phi_j \in C^\infty(\mathcal{M})$ of Δ at eigenvalue λ_j , there exist eigenfunctions $\tilde{\phi}_{j,\varepsilon}$ and $\tilde{\phi}_{j,\varepsilon,N}$ of P_ε and $P_{\varepsilon,N}$, respectively, such that the smooth functions

$$(7.4) \quad \phi_{j,\varepsilon} = \frac{1}{\Lambda_{j,\varepsilon}} \hat{P}_\varepsilon \tilde{\phi}_{j,\varepsilon}, \quad \phi_{j,\varepsilon,N} = \frac{1}{\Lambda_{j,\varepsilon,N}} \hat{P}_{\varepsilon,N} \tilde{\phi}_{j,\varepsilon,N}$$

satisfy

$$\lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \phi_{j,\varepsilon,N} = \lim_{\varepsilon \rightarrow 0^+} \phi_{j,\varepsilon} = \phi_j,$$

where the limits are taken with respect to uniform norm.

Remark 7.4. In addition to the convergence results stated above as iterated ($\varepsilon \rightarrow 0^+$ after $N \rightarrow \infty$) limits, in applications it is clearly important to have convergence results for limits where $N \rightarrow \infty$ and $\varepsilon \rightarrow 0^+$ simultaneously. In particular, note that at fixed $N \in \mathbb{N}$, the eigenvalues $\lambda_{j,\varepsilon,N}$ degenerate as $\varepsilon \rightarrow 0^+$, and fail to provide a good approximation of λ_j . This necessitates taking $N \rightarrow \infty$ limits along a sequence $\varepsilon(N)$ decreasing towards zero at a sufficiently slow rate. In the literature, this problem has mainly been studied in the context of i.i.d. samples [8, 43, 49, 50, 14]. For example, [50] shows that in dimension $d = \dim \mathcal{M} \geq 3$ it suffices to take $\varepsilon(N)$ such that $\lim_{N \rightarrow \infty} [(\log N)^{1/d} / (N^{1/d} \varepsilon(N))] = 0$. Here, we have opted to state spectral convergence results in terms of iterated limits, as they are valid

for arbitrary sampling scenarios satisfying the weak convergence property in (7.1). As previously noted, a common scenario with non-i.i.d. sampling is that of time-ordered data taken along orbits of ergodic dynamical systems.

7.3 Data-driven frame elements and approximation of sesquilinear forms

Using the approximate eigenvalues and eigenfunctions from (7.3) and (7.4), we can construct data-driven analogs of the various basis and frame elements for functions, vector fields, and forms introduced in Sections 3 and 4. For example,

$$\phi_{j,\varepsilon,N}^{(p)} = \frac{\phi_{j,\varepsilon,N}}{\lambda_{j,\varepsilon,N}^{p/2}}, \quad b_{ij,\varepsilon,N}^{(p)} = \frac{\phi_{i,\varepsilon,N} \operatorname{grad} \phi_{j,\varepsilon,N}}{\lambda_{i,\varepsilon,N}^{p/2}}, \quad b_{p,\varepsilon,N}^{ij_1 \cdots j_k} = \frac{d\phi_{j_1,\varepsilon,N} \wedge \cdots \wedge d\phi_{j_k,\varepsilon,N}}{\lambda_{i,\varepsilon,N}^{p/2}}$$

are data-driven analogs of the basis functions $\phi_j^{(p)}$ in (3.5), the frame elements $b_{ij}^{(p)}$ in (4.6) for vector fields, and the frame elements $b_p^{ij_1 \cdots j_k}$ in (4.7) for forms, respectively. All of these objects are “concrete”, i.e., they take values pointwise on $\mathcal{M}$ as opposed to being defined up to null sets. Moreover, their pointwise evaluation in practice relies on the ability to compute derivatives of the kernel k_ε .

In SEC, however, it oftentimes suffices to consider quantities that can be computed using only the “weak” counterparts $\tilde{\phi}_{j,\varepsilon,N}$ of $\phi_{j,\varepsilon,N}$ lying in $L^2(\mathcal{M}, \sigma_N)$. As a concrete example, consider the Galerkin method for the 1-Laplacian in Definition 6.3. To construct a data-driven analog of this scheme we compute the following quantities, using the shorthand notation $\alpha_{k,\varepsilon,N} = b_{1,\varepsilon,N}^{ikjk}$ as in Section 6:

- (1) Triple products $c_{ijk,\varepsilon,N} = \langle \tilde{\phi}_{i,\varepsilon,N}, \tilde{\phi}_{j,\varepsilon,N} \tilde{\phi}_{k,\varepsilon,N} \rangle_{L^2(\mathcal{M}, \sigma_N)}$.
- (2) Approximations $G_{ij,\varepsilon,N}$ and $E_{ij,\varepsilon,N}$ of the H_1 inner products $\langle \alpha_{i,\varepsilon,N}, \alpha_{j,\varepsilon,N} \rangle_{H_1}$ and Dirichlet energies $E_{1,1}(\alpha_{i,\varepsilon,N}, \alpha_{j,\varepsilon,N})$, computed via the formulas in Table 2.2, with the eigenvalues λ_i and triple products c_{ijk} replaced by $\lambda_{i,\varepsilon,N}$ and $c_{ijk,\varepsilon,N}$, respectively.
- (3) Approximations $u_{k,l,\varepsilon,N} \in \ell^2$ of the eigenvectors of the operator G_l , computed by solving the eigenvalue problem of the $l \times l$ matrix $\mathbf{G}_{l,\varepsilon,N}^1$ with elements $[\mathbf{G}_{l,\varepsilon,N}^1]_{ij} = G_{ij,\varepsilon,N} + E_{ij,\varepsilon,N}$.
- (4) Approximation of the $w_l \times w_l$ matrices $\mathbf{L}_l$ and $\mathbf{B}_l$ by $w_{l,\varepsilon,N} \times w_{l,\varepsilon,N}$ matrices $\mathbf{L}_{l,\varepsilon,N}$ and $\mathbf{B}_{l,\varepsilon,N}$, respectively, where $w_{l,\varepsilon,N}$ is the number of nonzero eigenvalues of $\mathbf{G}_{l,\varepsilon,N}^1$, and $\mathbf{L}_{l,\varepsilon,N}$ and $\mathbf{B}_{l,\varepsilon,N}$ are computed via (6.5), using the results of Steps 2 and 3 above as appropriate.

By Theorems 7.2 and 7.3,

$$\lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \mathbf{L}_{l,\varepsilon,N} = \mathbf{L}_l, \quad \lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \mathbf{B}_{l,\varepsilon,N} = \mathbf{B}_l,$$

in any matrix norm. Thus, for any solution $(v_{k,l}, \vec{a}_{k,l})$ of the generalized eigenvalue problem in (6.3) there exist solutions $(v_{k,l,\varepsilon,N}, \vec{a}_{k,l,\varepsilon,N})$ of the generalized eigenvalue problems

$$\mathbf{L}_{l,\varepsilon,N} \vec{a}_{k,l,\varepsilon,N} = v_{k,l,\varepsilon,N} \mathbf{B}_{l,\varepsilon,N} \vec{a}_{k,l,\varepsilon,N},$$

such that $\lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} v_{k,l,\varepsilon,N} = v_{k,l}$ and $\lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \vec{a}_{k,l,\varepsilon,N} = \vec{a}_{k,l}$. The latter, in conjunction with the convergence of the Laplacian eigenfunctions in Theorem 7.2(iii), implies that the reconstructed 1-forms

$$\varphi_{k,l,\varepsilon,N} = \sum_{j=0}^{w_{l,\varepsilon,N}-1} [\vec{a}_{k,l,\varepsilon,N}]_j \alpha_{j,\varepsilon,N}$$

converge, in H_1^1 norm, to the reconstructed form $\varphi_{k,l}$ associated with $\vec{a}_{k,l}$. By convergence of the Galerkin scheme in Definition 6.3, we therefore conclude that

$$\lim_{l \rightarrow \infty} \lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} v_{k,l,\varepsilon,N} = v_k, \quad \lim_{l \rightarrow \infty} \lim_{\varepsilon \rightarrow 0^+} \lim_{N \rightarrow \infty} \varphi_{k,l,\varepsilon,N} = \varphi_{k,l},$$

where (v_k, φ_k) is a weak eigenvalue-eigenvector pair of the 1-Laplacian, solving the variational eigenvalue problem in Definition 6.1.

8 Numerical examples

In this section, we apply the SEC to several smooth manifolds and a fractal set to verify and demonstrate the utility of our approach. In each example, we constructed the 1-Laplacian and its eigenvalue and eigenforms using the same procedure, which we describe here. First, we applied the diffusion maps algorithm described in Section 2.1 to the data in order to estimate the eigenvalues and eigenfunctions of the 0-Laplacian. Setting $M = 20$ eigenfunctions, we used 100 eigenfunctions to compute the $M \times M \times M$ tensor c which is the Fourier representation of function multiplication (see Table 2.1 in Section 2.2). Using the formulas in Table 2.2 in Section 2.5, we constructed the $M^2 \times M^2$ energy matrix $\hat{E}$ and Hodge Grammian $\hat{G}$ for the anti-symmetric formulation of the SEC. Following the procedure in Section 2.5, we then projected the eigenvalue problem onto the appropriate Sobolev H_1^1 basis, and computed the eigenvalues and eigenforms of the 1-Laplacian. Finally, we visualized the vector fields corresponding to the eigenforms by computing their operator representation, and pushing forward these vector fields into the original data space as described in Sections 2.5 and 5.2.

In the online supplementary material we have included Matlab code which implements the SEC 1-Laplacian construction along with a Diffusion Maps implementation. We also include code that generates all the data sets shown below and a simple “DEMO.m” file to replicate our results.

8.1 Validation of the SEC 1-Laplacian spectra

In this section we apply the SEC based construction of the 1-Laplacian to two examples where the spectrum of the 1-Laplacian can easily be worked out analytically.

We first consider the circle S^1 , where there is a nowhere-vanishing, harmonic 1-form $d\theta$ associated with a canonical angle coordinate θ . Because S^1 is one-dimensional and $d\theta$ is nowhere-vanishing, every smooth 1-form can be represented as $f d\theta$, with f a smooth function. Moreover, $d(f d\theta) = df \wedge d\theta = 0$, so that




FIGURE 8.1. Top Row, Left: Eigenvalues of the Laplacian on 1-forms estimated using the SEC (red, dashed) compared to the analytical spectrum (gray, solid) for a unit circle. Middle: The vector field representation of the 1-form corresponding to the smallest eigenvalue ($v_1 \approx 2.46 \times 10^{-6}$) clearly corresponds to the harmonic form $d\theta$. Right: The eigenform of the second smallest eigenvalue ($v_2 \approx 1.06$) is a linear combination of $\sin(\theta)d\theta$ and $\cos(\theta)d\theta$, which both have eigenvalue 1. Bottom Row: Computation is repeated using 500 points randomly sampled from the uniform distribution on the circle. Notice the agreement between the initial part of the spectrum and shape of the smoothest eigenforms. However, the spectrum diverges from the true spectrum faster than for the uniform grid of data due to variance of the estimators.

$\Delta_1(fd\theta) = \delta d(fd\theta) + d\delta(fd\theta) = d\delta(fd\theta)$, and $\delta d\theta = \Delta\theta = 0$ on S^1 . Therefore,

$$\begin{aligned} \Delta_1(fd\theta) &= d\delta(fd\theta) = d(-g(df, d\theta) + f\delta d\theta) = -d(g(df, d\theta)) \\ &= -d(d\theta(\nabla f)) = -d\left(\frac{df}{d\theta}\right) = -\frac{d^2f}{d\theta^2}d\theta, \end{aligned}$$

meaning that $\Delta_1(fd\theta) = \Delta(f)d\theta$. Thus, the eigenforms of the 1-Laplacian on S^1 are simply $\sin(k\theta)d\theta$ and $\cos(k\theta)d\theta$, $k \in \mathbb{N}_0$, and the eigenvalues are the same as those of Δ which are simply $\{0, 1, 1, 4, 4, \dots, k^2, k^2, \dots\}$.

We apply the SEC by generating 101 uniformly spaced data points on the unit circle in $\mathbb{R}^2$, and using the method outlined at the beginning of the section. In the leftmost plot in Fig. 8.1, we compare the analytic eigenvalues of Δ_1 (gray, solid) to the eigenvalues estimated by the SEC (red, dashed). We also show the first two eigenforms of the 1-Laplacian in Fig. 8.1, and note that the first eigenform clearly approximates the harmonic form $d\theta$. Notice that a closed integral curve of a harmonic form should correspond to a unique representative of a nontrivial




FIGURE 8.2. Top, Left: Eigenvalues of the Laplacian on 1-forms estimated using the SEC (red, dashed) compared to the analytical spectrum (gray, solid) for a flat torus. Top, Middle/Right: The vector field representation of the 1-form corresponding to the smallest eigenvalue ($v_1 \approx 1.20 \times 10^{-4}$) shown in the embedding coordinates $(\cos(\theta), \sin(\theta))$ (middle) and $(\cos(\phi), \sin(\phi))$ (right); this eigenform approximates $d\theta - d\phi$ (for clarity we only draw arrows every fifth data point). Bottom, Middle/Right: The eigenform of the second smallest eigenvalue ($v_2 \approx 2.43 \times 10^{-4}$) shown in the same coordinates as the first eigenform; this eigenform approximates $-d\theta - d\phi$. Notice that the two eigenforms together form a basis for the harmonic 1-forms which are the span of $d\theta$ and $d\phi$.

1-homology class, as will be further demonstrated below. In order to demonstrate the effect of random sampling on the manifold (as opposed to the uniform grid) we repeated this experiment using 500 independent uniformly distributed random points on S^1 . More points were required in this example due to the increased variance in the eigenvalues and eigenfunctions produced by the diffusion maps algorithm. The increased error from diffusion maps causes increased error in the SEC results, as shown in the eigenvalues of the 1-Laplacian. However, the SEC still obtained good approximations of the eigenforms at the level of having the coarse-grained structure (note the similarity of the first two eigenforms to those produced using the 101 evenly spaced data points). This demonstrates both the sensitivity of the SEC regarding eigenvalue precision, as well as the robustness of the SEC with respect to coarse-grained structure, which we believe is very desirable for applications.

We next consider the 1-Laplacian on the flat torus which has two nontrivial 1-homology classes. Every smooth 1-form on the flat torus can be written as $f d\theta + h d\phi$, where θ, ϕ are canonical angle coordinates and f, h smooth functions.

Considering $f d\theta$, we compute

$$\begin{aligned}\delta d(f d\theta) &= \delta(df \wedge d\theta) = \delta\left(\frac{\partial f}{\partial \phi} d\phi \wedge d\theta\right) = \delta\left(-\frac{\partial f}{\partial \phi} d\theta \wedge d\phi\right) \\ &= -\star d\star\left(-\frac{\partial f}{\partial \phi} d\theta \wedge d\phi\right) = \star d\left(\frac{\partial f}{\partial \phi}\right) = \star\left(\frac{\partial^2 f}{\partial \phi^2} d\phi + \frac{\partial^2 f}{\partial \theta \partial \phi} d\theta\right) \\ &= -\frac{\partial^2 f}{\partial \phi^2} d\theta + \frac{\partial^2 f}{\partial \theta \partial \phi} d\phi,\end{aligned}$$

where we choose the ordering $d\theta \wedge d\phi = -d\phi \wedge d\theta$ so that $\star d\theta = d\phi$ and $\star d\phi = -d\theta$. Next,

$$\begin{aligned}\delta d(f d\theta) &= -d(\star d\star(f d\theta)) = -d\star d(f d\phi) = -d\star(df \wedge d\phi) \\ &= -d\star\left(\frac{\partial f}{\partial \theta} d\theta \wedge d\phi\right) = -d\left(\frac{\partial f}{\partial \theta}\right) = -\frac{\partial^2 f}{\partial \theta^2} d\theta - \frac{\partial^2 f}{\partial \theta \partial \phi} d\phi,\end{aligned}$$

so that

$$\Delta_1(f d\theta) = \delta d(f d\theta) + d\delta(f d\theta) = \left(-\frac{\partial^2 f}{\partial \phi^2} - \frac{\partial^2 f}{\partial \theta^2}\right) d\theta = \Delta_0(f) d\theta,$$

and similarly $\Delta_1(h d\phi) = \Delta_0(h) d\phi$. Thus, for an eigenform with eigenvalue ν we have

$$\nu(f d\theta + h d\phi) = \Delta_1(f d\theta + h d\phi) = \Delta_0(f) d\theta + \Delta_0(h) d\phi$$

which implies that both f and h must be eigenfunctions of the 0-Laplacian with eigenvalue ν . Thus, the non-zero eigenvalues of the 1-Laplacian are the same as those of the 0-Laplacian up to multiplicity. Due to the two harmonic forms $d\theta$ and $d\phi$, each eigenvalue of the 1-Laplacian has double the multiplicity of the same eigenvalue for the 0-Laplacian.

To verify the SEC on this example, we generate 10,000 uniformly spaced data points on a flat torus in $\mathbb{R}^4$ with the embedding $(\cos \theta, \sin \theta, \cos \phi, \sin \phi)^\top$. In the top left of Fig. 8.2, we plot the analytic spectrum of the 1-Laplacian (gray, solid curve) along with the SEC approximation of the spectrum (red, dashed curve). We also show the vector fields corresponding to the first two SEC eigenforms. The latter approximate $d\theta - d\phi$ and $-d\theta - d\phi$, which span the space of harmonic forms which is the span of $d\theta$ and $d\phi$.

Notice that this example requires 100 times more data than the circle (data required grows exponentially in the intrinsic dimension) for the diffusion maps algorithm to yield the same accuracy for the 0-Laplacian eigenfunctions and eigenvalues. This also means that the diffusion maps algorithm takes significantly longer to run on this larger data set. Crucially, the SEC still only uses $M = 20$ (and 100 eigenfunctions to compute the c tensor) so the matrices in the SEC were the same size for the torus example as for the circle example. Thus, following the initial diffusion maps step, the SEC algorithm runtime is the same in the torus example as for the circle example. This demonstrates how the SEC formulation allows us to

TABLE 8.1. First eight eigenvalues of the Laplacian on 1-forms estimated using the SEC for a Möbius band, torus (standard embedding in $\mathbb{R}^3$), a genus two surface (two-holed torus), a two-dimensional sphere in $\mathbb{R}^3$, and the Lorenz 63 attractor (L63). Notice that the eigenvalues close to zero represent harmonic forms, and the number of the eigenvalues close to zero matches the first Betti numbers for the manifolds which are 1, 2, 4, and 0 respectively. While the Lorenz 63 attractor is not a manifold, a coarse approximation as a manifold would suggest a Betti number of 2 due to the two holes.

Möbius	Torus	Genus 2	Sphere	L63
0.0242	0.0040	0.0021	1.9349	0.0011
1.0415	0.0093	0.0026	1.9521	0.0017
1.0449	0.2574	0.0026	1.9781	0.0030
3.8684	0.2575	0.0041	1.9817	0.0072
3.8948	0.2575	0.0893	2.0042	0.0105
8.0352	0.2587	0.0901	2.0172	0.0109
8.1018	0.8061	0.2151	5.8001	0.0205
8.9369	0.8067	0.2175	5.8142	0.0262




FIGURE 8.3. First two eigenforms on the Möbius band. The first represents the one 1-homology class.

decouple the representation of differential forms and their associated Laplacian operators from the amount of data. In particular, larger data sets are only needed in the initial diffusion maps step to obtain the best possible estimate of the 0-Laplacian eigenfunctions and eigenvalues. Since many of these eigenfunctions and eigenvalues are known to be poor estimates, we can set M much less than the number of data points to obtain high quality 1-Laplacian representations.

8.2 Topological features and eigenforms via the SEC

We now apply the SEC on several surfaces embedded in $\mathbb{R}^3$ in order to demonstrate its connection to the manifold homology. We also demonstrate the algorithm




FIGURE 8.4. First two eigenforms on the sphere S^2 . Notice that the SEC constructs valid smooth vector fields which must vanish on the sphere.




FIGURE 8.5. First four eigenforms on the standard embedding of the torus in $\mathbb{R}^3$. Notice that the first two represent the two 1-homology classes.

on a data set which is not sampled from a manifold to show that the SEC has potential applications even when the assumption of an underlying manifold does not hold. While the results on the circle and flat torus above used large uniformly distributed data sets to validate the algorithm, in this case we work with much smaller




FIGURE 8.6. First four eigenforms on the genus two surface, representing the four 1-homology classes.

data sets that may not be uniformly distributed with respect to the volume form. Details on the sampling procedure employed in each example can be found in the code in the online supplementary material.

We first consider four surfaces, namely a Möbius band, sphere, torus, and a genus-2 surface. Our goal for these examples is to show that the SEC correctly captures the coarse topological features of these smaller nonuniform data sets. The 1-homology of these manifolds corresponds to the kernel of the 1-Laplacian so that the first Betti number should be equal to the multiplicity of the zero eigenvalue of Δ_1 . This is shown in Table 8.1 if we consider the eigenvalues closest to zero to represent 1-homology. We can also make this connection by visualizing the eigenforms via their corresponding vector fields as shown in Figs. 8.3–8.6. Following the vector field corresponding to a harmonic form (having eigenvalue zero) should generate a closed curve, which is a representative of a unique homology class. We also show in Figs. 8.3–8.6 that the SEC approximations are smooth forms, demonstrated both by the smooth changes in the arrows, as well as the fact that the corresponding vector fields each vanish at some point on the sphere (since there are no smooth non-vanishing vector fields on a sphere).

We should note that the Möbius band is a manifold with boundary, which violates the assumption of manifold without boundary used in our theoretical derivations. Nevertheless, the algorithm produces reasonable results which suggest the theory may be able to be extended to manifolds with boundary. A potential theoretical barrier is the requirement that the gradients of the eigenfunctions must

span every tangent space (see Lemma 4.7), yet the estimator of the Laplacian from diffusion maps produces Neumann eigenfunctions [18], which will all have gradient orthogonal to the normal direction. While this failure to span the tangent spaces may preclude the corresponding 1-forms b^{ij} and $\tilde{b}^{ij}$ from providing frames for the L^2 space H_1 , it is still possible that they provide frames for higher-order Sobolev spaces associated with Neumann boundary conditions. In particular, it is possible that an analog of Theorem 4.4 holds for the Neumann $H_{1,1}$ Sobolev space, so that the corresponding Galerkin scheme from Definition 6.3 approximates the spectrum of the Neumann Laplacian on 1-forms. While a rigorous study of the properties of SEC in the Neumann setting lies outside the scope of this work, the numerical results for the Möbius band described above are consistent with this behavior. If approximation of the full H_1 space, or Sobolev spaces associated with other boundary conditions is desired, a possible remedy would be to employ the normal direction estimator and/or the distance to boundary estimator which were recently developed in [13]. These estimators could be used to ensure that the full tangent space is spanned on the boundary (and near the boundary, which may be an issue for finite data sets).

As a final example, we demonstrate the SEC on data sampled from the chaotic attractor of the Lorenz 63 dynamical system [35] (we refer to this set as “L63”), a fractal set having no differential structure. While there is no exterior calculus defined on L63, it is a well-defined compact subset on $\mathbb{R}^3$ [51] (with an induced metric topology), and exhibits certain coarse-grained topological features, most notably a hole in each of the two lobes of the attractor as shown in Fig. 8.7. Moreover, the diffusion maps and SEC algorithms can easily be applied to data sampled from this set (or indeed any data set in a metric space). The SEC spectrum for the 1-Laplacian is shown in Fig. 8.1, and the first two eigenvalues are very close to zero, while the corresponding eigenforms, shown in Fig. 8.7 seem to capture these two coarse topological features.

In summary, the examples in this section show that the SEC can generate a collection of 1-forms (equivalently, vector fields) which can be used as a basis for vector fields defined on the data set. These vector fields are ordered by smoothness based on their corresponding eigenvalue (Dirichlet energy; see Section 2.5), and higher eigenvalues correspond to more oscillatory vector fields. Moreover, our theory shows that in the limit of large data and large M (number of eigenfunctions of Δ_0 used), the SEC basis corresponds to the natural basis for square integrable 1-forms. However, even for small nonuniform data sets the SEC reflects coarse topological features of the underlying continuous space, which indicates that the SEC-derived approximations could be useful even outside of the large data limit.




FIGURE 8.7. First two eigenforms of the SEC 1-Laplacian on 10000 data points sampled from the Lorenz 63 attractor. Notice that closed integral curves of these vector fields correspond to independent representatives of the two “holes” in the attractor.

9 Discussion

In this paper, we have developed a spectral framework, called spectral exterior calculus (SEC), for the exterior calculus on Riemannian manifolds. A central underpinning of this approach is a family of frames for L^2 and higher-order Sobolev spaces of vector fields and differential forms, built entirely from the eigenvalues and eigenfunctions of the Laplacian on functions. By virtue of this construction, our framework lends itself well to data-driven approximation of the objects of interest in exterior calculus, such as vector fields and differential forms, as well as operators acting on these objects (e.g., the Laplacian on forms), requiring no additional information beyond point-clouds of data. In particular, SEC extends the applicability of the extensive array of graph-theoretic techniques for pointwise and spectral approximation of the Laplacian on functions [7, 44, 18, 52, 12, 32, 11, 8, 43, 49, 50, 14] to learning problems involving higher-order objects, with rigorous convergence guarantees in the asymptotic limit of large data. Crucially, by relying solely on approximations of the eigenvalues and eigenfunctions of the Laplacian on functions, SEC decouples the computational complexity of approximation of vector fields, forms, and related operators from the number of samples and ambient data space dimension.

Another key aspect of SEC is its focus on L^2 -convergent, as opposed to pointwise-convergent approximations, as the latter typically require additional structures such as simplicial complexes that are difficult to estimate from data alone. Here, we have shown that our frames for H^1 Sobolev spaces of 1-forms lead naturally to Galerkin approximation schemes for the eigenvalue problem of the Laplacian on 1-forms, which are provably well-posed by classical results on variational formulations of elliptic eigenvalue problems [4]. These techniques extend previously developed

data-driven Galerkin approximation techniques for unbounded operators on functions [29, 30, 19] to the setting of 1-forms. These Galerkin methods for the Laplacian on one-forms, and more broadly approximation techniques for unbounded operators in exterior calculus, crucially depend on the availability of well-conditioned approximation spaces for H^p spaces of higher regularity than L^2 . In the framework of SEC, such spaces are naturally constructed through singular value decomposition of sparse Grammian matrices associated with H^p frame elements, which are computable via closed-form expressions involving the eigenvalues of the Laplacian on functions, and inner-product relationships between products of the corresponding eigenfunctions. Once computed, the Galerkin-approximated Laplacian eigenforms can be visualized through their dual vector fields, reconstructed via a data-driven spectral approximation of the pushforward map. In addition to frame representations, SEC provides data-driven representations of vector fields as operators on functions, which is useful for tasks such as dynamical systems forecasting [10, 30], among other applications.

We have demonstrated the efficacy of SEC approximations to the eigenvalues and eigenforms of the 1-Laplacian, and their associated vector field duals, in a suite of examples involving orientable (circle, flat and curved 2-tori, 2-sphere, genus-2 surface) and non-orientable (Möbius band) smooth manifolds, as well as a fractal set having no differentiable structure (the L63 attractor). In the circle and flat-torus examples, where analytical expressions for the 1-Laplacian eigenvalues and eigenforms are available, we found that SEC accurately approximates the leading 50 to 100 eigenvalues using a modest number (20) of 0-Laplacian eigenfunctions to build the SEC frames, and a moderate dimension (40 for circle and 107 for the flat torus) of the corresponding Galerkin approximation spaces. In the curved torus, sphere, and Möbius band examples (where analytical expressions for the eigenvalues and eigenforms of the 1-Laplacian are not readily available), we demonstrated that the SEC results are consistent with the 1-homology of these manifolds. In particular, the number of eigenvalues of the SEC-approximated 1-Laplacian numerically close to zero was found to be equal to the first Betti number of the manifold under study, and the corresponding eigenforms were found to generate closed curves (through the integral curves of their vector field duals), representing 1-homology classes. It should be noted that due to the presence of a boundary, the Möbius example lies outside the theoretical domain of applicability of the SEC formulation developed here, but the numerical results were found to be qualitatively consistent with the application of Neumann boundary conditions, implicitly enforced through diffusion maps. Moreover, even though our theory was developed in the smooth-manifold setting, the SEC-derived eigenforms for the L63 attractor were also found to be qualitatively consistent with the 1-homology of a coarse graining of the attractor. In this example, we obtained two eigenforms at small corresponding eigenvalue, generating closed curves around each of the “holes” in the lobes of the attractor (Fig. 8.7), which is what one would intuitively expect for the L63 topology.

There are several avenues of future research stemming from this work. First, it would be of interest to develop SEC approximation schemes for other operators of interest in exterior calculus, such as the Hodge star operator and the Lie derivative. Among other applications, such approximation schemes are likely to be of interest in dynamical systems modeling. For instance, for a dynamical system on a manifold generated by a vector field v , there is an associated dynamical system acting on the tangent bundle, which can be thought of as a pointwise linearization of the system. In this setting, the Lie derivative $\mathcal{L}_v$ on vector fields generates the action of this system on sections of the tangent bundle, and is known to have useful spectral properties characterizing growth rates of perturbations [15]. Approximating these spectra from data would thus provide a computationally efficient empirical method for characterizing the principal modes of instability in datasets generated by dynamical systems; a topic of considerable current interest that has spurred the development of powerful geometrical approaches such as the theory of covariant Lyapunov vectors [31].

Other research directions would involve extending SEC to manifolds with boundary, or non-smooth topological spaces. In particular, our results for the L63 attractor suggest that a spectral formulation of exterior calculus, possibly with a very similar structure to the one presented here, may be possible provided one can construct an appropriate diffusion operator generalizing the notion of the Laplacian on Riemannian manifolds. To that end, it is worthwhile noting that the Markov integral operators employed here to approximate the heat semigroup and its associated Laplacian are, in fact, well defined as operators on C^0 and L^2 function spaces on Borel measure spaces without differentiable structure. Moreover, recent work [19] has shown that, at least in some scenarios, such operators can be consistently approximated from data under natural assumptions (e.g., data sampled from a dynamical system possessing a physical ergodic invariant measure, such as the L63 system studied here). Yet, a key step relevant to the SEC framework, namely the consistent approximation of the Laplacian from the heat semigroup, achieved by taking $\epsilon \rightarrow 0$ kernel bandwidth limits, has, to our knowledge, no known generalization to non-differentiable spaces. In such a setting, the challenge would be to approximate the generator of an appropriate C^0 diffusion semigroup based on kernel integral operators approximating that semigroup (akin to the operators $\tilde{P}_\epsilon$ in Section 7.2 approximating the heat semigroup). Such a diffusion operator would lead to an analog of the product rule for the 0-Laplacian on smooth functions through its associated symmetric bilinear form [6], providing one of the necessary ingredients to building the SEC framework. However, at present, the construction of this operator and its approximation remain elusive.

Finally, while the present SEC formulation has focused heavily on H^p Sobolev spaces and their associated notions of regularity based on Dirichlet energy, it should be noted that there exist analogs of the RKHSs associated with the heat kernel on functions in the setting of differential forms. In addition to providing well-defined notions of regularity through the corresponding RKHS norms, the fact that RKHS

spaces have continuous pointwise evaluation functionals allows for pointwise- (as opposed to L^2 -) convergent approximations when working in such spaces, with well-known applications in smoothing and interpolation. These properties motivate the construction of frames for RKHSs of forms, and their use in Galerkin approximation schemes for operators on forms, analogous to the schemes presented here utilizing Sobolev spaces.

Appendix A: Discrete incompatibility of product rule and Leibniz rule

Let f , h , and b be arbitrary smooth, real-valued functions on a Riemannian manifold. The product rule for the Laplacian operator can be written as

$$\Delta(fh) = f\Delta h + h\Delta f - 2\operatorname{grad} f(h),$$

where $\operatorname{grad} f(h) = \operatorname{grad} f \cdot \operatorname{grad} h$. The Leibniz rule for vector fields states that

$$v(fh) = fv(h) + hv(f),$$

and since $\operatorname{grad} f$ is a vector field, we can easily derive a triple product rule for the Laplacian

$$\begin{aligned} \Delta(fhb) &= f\Delta(hb) + hb\Delta f - 2\operatorname{grad} f(hb) \\ &= f\Delta(hb) + hb\Delta f - 2h\operatorname{grad} f(b) - 2b\operatorname{grad} f(h) \\ &= f\Delta(hb) + hb\Delta f - hf\Delta b - hb\Delta f + h\Delta(fb) - bf\Delta h - bh\Delta f + b\Delta(hf). \end{aligned}$$

Rearranging the triple product rule, we have

$$0 = \Delta(fhb) - f\Delta(hb) - h\Delta(fb) - b\Delta(fh) + fh\Delta b + fb\Delta h + hb\Delta f.$$

Now assume that L is a discrete Laplacian operator (meaning a matrix), without any assumption on the form of the representation of vector fields. If we assume that gradient fields are represented in such a way that both the product rule for the Laplacian and the Leibniz rule hold, then the same derivation will produce the triple product rule for the discrete Laplacian. However, if the triple product rule holds on the standard basis vectors $\{e_i\}$, then we find that

$$0 = L(e_i e_j e_k) - e_i L(e_j e_k) - e_j L(e_i e_k) - e_k L(e_i e_j) + e_i e_j L e_k + e_i e_k L e_j + e_j e_k L e_i,$$

where vector-vector products are componentwise. Setting $i = j = k$, we find

$$0 = L(e_i) - 3e_i L(e_i e_i) + 3e_i e_i L e_i = L e_i$$

and since $L e_i = 0$ for all i , we conclude that the only matrix which satisfies the triple product rule is the zero matrix. This shows that no discrete representation of vector fields can simultaneously satisfy both the Laplacian product rule and the Leibniz rule.

Appendix B: Computations and derivations of formulas

In this appendix, we derive the formulas for the inner products and Dirichlet energies of the SEC frame elements for 1-forms listed in Table 2.2 (B.1). In addition, we discuss how these formulas can be extended to higher-order k -forms (B.2). Since the SEC frame elements are real, in these derivations we will consider that all Hilbert spaces of functions, vector fields, and k -forms are over the real numbers, i.e., there will be no complex conjugation in the corresponding inner products.

B.1 Inner products and Dirichlet forms for 1-forms

Recall that the Hodge inner product 1-forms ω, ν is defined by

$$\langle \omega, \nu \rangle_{H_1} = \int_{\mathcal{M}} \omega \wedge \star \nu = \langle 1, \eta(\omega, \nu) \rangle_H = \langle 1, \omega^\sharp \cdot \nu^\sharp \rangle_H,$$

where we abbreviate the Riemannian inner product for vector fields w, v by $w \cdot v \equiv g(w, v)$. Similarly for 1-forms ω, ν we will abbreviate the Riemannian inner product by $\omega \cdot \nu \equiv \eta(\omega, \nu) = g(\omega^\sharp, \nu^\sharp)$. In particular, if $\omega = f dh$ and $\nu = \alpha d\beta$, then the Hodge inner product can be further simplified as

$$\langle f dh, \alpha d\beta \rangle_{H_1} = \langle dh \cdot d\beta, f \alpha \rangle_H = \frac{1}{2} \langle h \Delta \beta + \beta \Delta h - \Delta(h\beta), f \alpha \rangle_H,$$

and substituting eigenfunctions b^i of Δ into this formula, we define

$$\begin{aligned} G_{ijkl} &\equiv \langle b^i db^j, b^k db^l \rangle_{H_1} = \frac{1}{2} \langle b^j \Delta b^l + b^l \Delta b^j - \Delta(b^j b^l), b^i b^k \rangle_H \\ &= \frac{1}{2} \langle \lambda_l b^j b^l + \lambda_j b^l b^j - \Delta \left(\sum_s c_{ljs} b^s \right), b^i b^k \rangle_H \\ &= \frac{1}{2} (\lambda_l + \lambda_j) c_{ijkl} - \sum_s \lambda_s c_{ljs} c_{iks}, \\ (B.1) \quad &= \frac{1}{2} \sum_s (\lambda_l + \lambda_j - \lambda_s) c_{ljs} c_{iks}, \end{aligned}$$

where $c_{ijk} \equiv \langle b^i b^j, b^k \rangle_H$ and $c_{ijk}^0 \equiv \langle b^i b^j, b^k b^l \rangle_H = \sum_s c_{ijs} c_{kls}$ (note that these are invariant to permutations of indices). Notice that the Gramm matrix $\hat{G}_{ijkl} = \langle \hat{b}^{ij}, \hat{b}^{kl} \rangle_{H_1}$ of the antisymmetric elements is easy to compute from G_{ijkl} since

$$\hat{G}_{ijkl} = G_{ijkl} + G_{jilk} - G_{ijlk} - G_{jikl},$$

so next we will consider the computation of the values E_{ijkl} of the Dirichlet form.

Recalling the formula for the 1-Laplacian $\Delta_1 = d\delta + \delta d$, for 1-forms ω, ν we can write

$$(B.2) \quad \langle \omega, \Delta_1 \nu \rangle_{H_1} = \langle \delta \omega, \delta \nu \rangle_H + \langle d\omega, d\nu \rangle_{H_2}$$

where d is the exterior derivative, and δ its formal adjoint, the codifferential. For 1-forms $\omega = f dh$ and $\nu = \alpha d\beta$, we have

$$\langle d(f dh), d(\alpha d\beta) \rangle_{H_2} = \langle df \wedge dh, d\alpha \wedge d\beta \rangle_{H_2} = \langle 1, \eta(df \wedge dh, d\alpha \wedge d\beta) \rangle_{H_2}$$

$$\begin{aligned}
&= \left\langle 1, \det \begin{bmatrix} df \cdot d\alpha & df \cdot d\beta \\ dh \cdot d\alpha & dh \cdot d\beta \end{bmatrix} \right\rangle_H \\
(B.3) \quad &= \langle df \cdot d\alpha, dh \cdot d\beta \rangle_H - \langle df \cdot d\beta, dh \cdot d\alpha \rangle_H.
\end{aligned}$$

We recall that the codifferential acting on k -forms is $\delta = (-1)^{m(k-1)+1} \star d \star$, where d is the exterior derivative acting on $m-k$ forms. We use m instead of d for the dimension of the manifold in this section to avoid confusion with the exterior derivative. For a function f and 1-form ω , we have

$$\begin{aligned}
\delta(f\omega) &= -\star d \star (f\omega) = -\star d(f \star \omega) \\
&= -\star (df \wedge \star \omega + f d \star \omega) = -\star (df \wedge \star \omega) - f \star d \star \omega \\
(B.4) \quad &= -\eta(df, \omega) + f \delta \omega = -df \cdot \omega + f \delta \omega,
\end{aligned}$$

so in particular if $\omega = dh$ we find (using $\Delta = \delta d$)

$$\begin{aligned}
\delta(f dh) &= -df \cdot dh + f \Delta h \\
&= \frac{1}{2} (\Delta(fh) - f \Delta h - h \Delta f) + f \Delta h = \frac{1}{2} (\Delta(fh) + f \Delta h - h \Delta f).
\end{aligned}$$

Using this formula we can simplify the inner product

$$\begin{aligned}
\langle \delta(f dh), \delta(\alpha d\beta) \rangle_H &= \frac{1}{4} \langle \Delta(fh) + f \Delta h - h \Delta f, \Delta(\alpha\beta) + \alpha \Delta \beta - \beta \Delta \alpha \rangle_H \\
&= \langle -df \cdot dh + f \Delta h, -d\alpha \cdot d\beta + \alpha \Delta \beta \rangle_H \\
&= \langle df \cdot dh, d\alpha \cdot d\beta \rangle - \langle df \cdot dh, \alpha \Delta \beta \rangle_H - \langle d\alpha \cdot d\beta, f \Delta h \rangle_H \\
(B.5) \quad &+ \langle f \Delta h, \alpha \Delta \beta \rangle_H.
\end{aligned}$$

Now that both summands in (B.2) have been written in terms of Δ , we can combine (B.3) and (B.5) and several cancellations yield,

$$\begin{aligned}
\langle f dh, \Delta_1(\alpha d\beta) \rangle_{H_1} &= \langle df \cdot d\alpha, dh \cdot d\beta \rangle_H - \langle df \cdot d\beta, dh \cdot d\alpha \rangle_H \\
&+ \langle df \cdot dh, d\alpha \cdot d\beta \rangle_H - \langle df \cdot dh, \alpha \Delta \beta \rangle_H \\
&- \langle d\alpha \cdot d\beta, f \Delta h \rangle_H + \langle f \Delta h, \alpha \Delta \beta \rangle_H.
\end{aligned}$$

When the functions are eigenfunctions of Δ , we find

$$\begin{aligned}
db^i \cdot db^j &= \frac{1}{2} (b^i \Delta b^j + b^j \Delta b^i - \Delta(b^i b^j)) = \frac{1}{2} ((\lambda_i + \lambda_j) b^i b^j - \Delta(b^i b^j)) \\
&= \frac{1}{2} \sum_s (\lambda_i + \lambda_j - \lambda_s) c_{ijs} b^s,
\end{aligned}$$

so that we can define

$$\begin{aligned}
F_{ijkl} &= \langle db^i \cdot db^j, db^k \cdot db^l \rangle_H = \frac{1}{4} \sum_s (\lambda_i + \lambda_j - \lambda_s) (\lambda_k + \lambda_l - \lambda_s) c_{ijs} c_{kls} \\
&= \frac{1}{4} \left[(\lambda_i + \lambda_j) (\lambda_k + \lambda_l) c_{ijkl} - (\lambda_i + \lambda_j + \lambda_k + \lambda_l) \sum_s \lambda_s c_{ijs} c_{kls} \right]
\end{aligned}$$

$$+ \sum_s \lambda_s^2 c_{ijs} c_{kls} \Big].$$

Moreover, recalling that

$$G_{ijkl} = \frac{1}{2} \sum_s (\lambda_l + \lambda_j - \lambda_s) c_{ljs} c_{iks} = \frac{1}{2} (\lambda_l + \lambda_j) c_{ijkl}^0 - \frac{1}{2} \sum_s \lambda_s c_{ljs} c_{iks},$$

we simplify the above formulas as

$$\langle d(b^i db^j), d(b^k db^l) \rangle_{H_2} = F_{ikjl} - F_{iljk}$$

and

$$\langle \delta(b^i db^j), \delta(b^k db^l) \rangle_H = F_{ijkl} - \lambda_l G_{kilj} - \lambda_j G_{ikjl} + \lambda_j \lambda_l c_{ijkl}^0,$$

so that

(B.6)

$$\begin{aligned} E_{ijkl} &= \langle b^i db^j, \Delta_1(b^k db^l) \rangle_{H_2} \\ &= F_{ikjl} - F_{iljk} + F_{ijkl} - \lambda_l G_{kilj} - \lambda_j G_{ikjl} + \lambda_j \lambda_l c_{ijkl} \\ &= c_{ijkl}^0 \left(\frac{1}{4} (\lambda_i + \lambda_k) (\lambda_j + \lambda_l) - \frac{1}{4} (\lambda_i + \lambda_l) (\lambda_j + \lambda_k) + \frac{1}{4} (\lambda_i + \lambda_j) (\lambda_k + \lambda_l) \right) \\ &\quad + c_{ijkl}^0 \left(\lambda_j \lambda_l - \frac{1}{2} \lambda_l (\lambda_i + \lambda_j) - \frac{1}{2} \lambda_j (\lambda_k + \lambda_l) \right) \\ &\quad + \frac{1}{4} (\lambda_i + \lambda_j + \lambda_k + \lambda_l) (c_{iljk}^1 - c_{ikjl}^1 - c_{ijkl}^1) + \frac{1}{2} (\lambda_j + \lambda_l) c_{ijkl}^1 \\ &\quad + \frac{1}{4} (c_{ikjl}^2 + c_{ijkl}^2 - c_{iljk}^2) \\ &= \frac{1}{4} (\lambda_i + \lambda_j + \lambda_k + \lambda_l) (c_{iljk}^1 - c_{ikjl}^1 - c_{ijkl}^1) + \frac{1}{2} (\lambda_j + \lambda_l) c_{ijkl}^1 \\ &\quad + \frac{1}{4} (c_{ikjl}^2 + c_{ijkl}^2 - c_{iljk}^2) \\ &= \frac{1}{4} [(\lambda_i + \lambda_j + \lambda_k + \lambda_l) (c_{iljk}^1 - c_{ikjl}^1) + (\lambda_j + \lambda_l - \lambda_i - \lambda_k) c_{ijkl}^1 \\ &\quad + (c_{ikjl}^2 + c_{ijkl}^2 - c_{iljk}^2)] \end{aligned}$$

A straightforward computation then shows that the c_{ijkl}^0 coefficients exactly cancel, and we define

$$c_{ijkl}^1 \equiv \sum_s \lambda_s c_{ijs} c_{skl} \quad \text{and} \quad c_{ijkl}^2 \equiv \sum_s \lambda_s^2 c_{ijs} c_{skl},$$

which also allows us to write $G_{ijkl} = \frac{1}{2} ((\lambda_l + \lambda_j) c_{ijkl}^1 - c_{ikjl}^1)$. Finally, to compute the Dirichlet form for the antisymmetric elements, we note that

$$E_{ijkl} - E_{jikl} = \frac{1}{2} [(\lambda_i + \lambda_j + \lambda_k + \lambda_l) (c_{iljk}^1 - c_{ikjl}^1) + (\lambda_j - \lambda_i) c_{ijkl}^1 + (c_{ikjl}^2 - c_{iljk}^2)]$$

and

$$E_{ijkl} - E_{jilk} = \frac{1}{2} \left[-(\lambda_i + \lambda_j + \lambda_k + \lambda_l)(c_{iljk}^1 - c_{ikjl}^1) + (\lambda_j - \lambda_i)c_{ijkl}^1 \right. \\ \left. - (c_{ikjl}^2 - c_{iljk}^2) \right],$$

leading to

$$\hat{E}_{ijkl} \equiv \langle b^i db^j - b^j db^i, \Delta_1(b^k db^l - b^l db^k) \rangle = E_{ijkl} - E_{jikl} - (E_{ijlk} - E_{jilk}) \\ = (\lambda_i + \lambda_j + \lambda_k + \lambda_l)(c_{iljk}^1 - c_{ikjl}^1) + (c_{ikjl}^2 - c_{iljk}^2).$$

These formulas are summarized in Table 2.2 in Section 2.5.

B.2 Extension to k -forms

Next, we briefly summarize how we can extend the SEC formulas derived in B.1 to higher-order forms. Let $b^I = b^{i_0} db^{\tilde{I}} \equiv b^{i_0} db^{i_1} \wedge \cdots \wedge db^{i_k}$ be a k -form frame element ($\tilde{I} = (i_1, \dots, i_k)$ and $I = (i_0, \dots, i_k)$). As we will see below, the Hodge inner product of the exterior derivatives is easily computed as a determinant of inner products of 1-forms. The more complex term is the Hodge inner product of the codifferential terms. To understand this, we need to generalize the product rule for the codifferential as

$$\delta(b^I) = (-1)^{m(k-1)+1} \star d \star (b^{i_0} db^{\tilde{I}}) = (-1)^{m(k-1)+1} \star d(b^{i_0} \star db^{\tilde{I}}) \\ = (-1)^{m(k-1)+1} \star (db^{i_0} \wedge \star db^{\tilde{I}} + b^{i_0} d \star db^{\tilde{I}}) \\ = (-1)^{m(k-1)+1} \star (db^{i_0} \wedge \star db^{\tilde{I}}) + b^{i_0} \delta(db^{\tilde{I}}).$$

We can now reduce $\delta(db^{\tilde{I}})$ by rewriting $db^{\tilde{I}} = d(b^{i_1} db^{\hat{I}})$, where $\hat{I} = (i_2, \dots, i_k)$, so that

$$\delta_k(db^{\tilde{I}}) = \delta_k d_{k-1}(b^{i_1} db^{\hat{I}}) = \Delta_{k-1}(b^{i_1} db^{\hat{I}}) - d_k \delta_{k-1}(b^{i_1} db^{\hat{I}}).$$

Notice that in the above formula we have reduced the problem of computing the k -codifferential to computing the $(k-1)$ -Laplacian and the $(k-1)$ -codifferential. Since we have shown how to compute the codifferential on 1-forms, this strategy can be used to lift the SEC formulas to higher-order forms, although the formulas become quite complicated. We carried this out for 2-forms, but the derivations are quite long and a closed formula for the general case remains elusive.

Thus, while in principle the iterative formula above allows us to lift our formulation to k -forms, we do not yet have a closed formula for the inner product

$$\langle b^I, \Delta_k b^J \rangle_{H_k} = \langle d(b^I), d(b^J) \rangle_{H_{k+1}} + \langle \delta(b^I), \delta(b^J) \rangle_{H_{k-1}}.$$

However, the first term above is simply the integral of the determinant of the matrix of pairwise inner products $db^{i_s} \cdot db^{j_r}$ for $s, r \in \{0, \dots, k\}$. The key to computing this term is the formula

$$(db^i \cdot db^j)(db^k \cdot db^l) = \left(\sum_q g_{ijq} b^q \right) \left(\sum_n g_{kl n} b^n \right) = \sum_{q,n} g_{ijq} g_{kl n} b^q b^n,$$

TABLE 2.1. Table of formulas for the SEC on k -forms. We abbreviate the multi-indices $I = (i_0, \dots, i_k)$ and $\tilde{I} = (i_1, \dots, i_k)$ and similarly $J = (j_0, \dots, j_k)$ and $\tilde{J} = (j_1, \dots, j_k)$. We use $[db^{i_q} \cdot db^{j_n}]$ to abbreviate the $k \times k$ matrix with (q, n) -th entry given by $db^{i_q} \cdot db^{j_n}$. Finally, note that $db^I = db^{i_0} \wedge \dots \wedge db^{i_k}$.

Object	Symbolic	Spectral
Multiple Product	$c_I^0 = \langle b^{i_0} \dots b^{i_k}, 1 \rangle_H$	$c_I^0 = \sum_s c_{i_0 i_1 s} c_{s i_2 \dots i_k}^0$
Tensor Evaluation	$H^{\tilde{I}\tilde{J}} = (db^{i_1} \cdot db^{j_1}) \dots (db^{i_k} \cdot db^{j_k})$ $= \text{grad } b^{i_1} \otimes \dots \otimes \text{grad } b^{i_k} (b^{j_1}, \dots, b^{j_k})$	$\hat{H}_I^{\tilde{I}\tilde{J}} \equiv \langle H^{\tilde{I}\tilde{J}}, b^I \rangle_H$ $= \sum_{n=1}^{k^2} \prod_{s,r=1}^k g_{i_s j_r m_n} c_{l m_1 \dots m_{k/2}}$
Tensor Product	$b_J = b^{j_0} \text{grad } b^{j_1} \otimes \dots \otimes \text{grad } b^{j_k}$	$\langle b_J (b^{i_1}, \dots, b^{i_k}), b^I \rangle = \sum_s \hat{H}_s^{\tilde{I}\tilde{J}} c_{s j_0 l}$
Frame Elements	$b^I = b^{i_0} db^{i_1} \wedge \dots \wedge db^{i_k}$	$\langle b^I (b_J), b^I \rangle_H = \langle b^I \cdot b^J, b^I \rangle_H$
Riemannian Metric	$b^I \cdot b^J = b^{i_0} b^{j_0} \det([db^{i_q} \cdot db^{j_n}])$	$\langle b^I \cdot b^J, b^I \rangle_H =$ $\sum_s \sum_{\sigma \in S_k} \text{sgn}(\sigma) c_{s i_0 j_0 l} \hat{H}_s^{\tilde{I}\sigma(\tilde{J})}$
Hodge Grammian	$G_{IJ} = \langle b^I, b^J \rangle_{H_k}$	$\langle b^I \cdot b^J, 1 \rangle_H =$ $\sum_s \sum_{\sigma \in S_k} \text{sgn}(\sigma) c_{s i_0 j_0} \hat{H}_s^{\tilde{I}\sigma(\tilde{J})}$
d -Energy	$E_{IJ}^d = \langle db^I, db^J \rangle_{H_{k+1}}$	$\langle db^I \cdot db^J, 1 \rangle_{H_{k+1}} = \hat{H}_0^{IJ}$

and iterating the above we can expand the product as

$$\prod_{s,r=0}^k db^{i_s} \cdot db^{j_r} = \sum_{n=0}^{k^2} \prod_{s,r=0}^k g_{i_s j_r q_n} b^{q_n}.$$

The integral of these products can then be represented in terms of the integrals of products of eigenfunctions, which can be computed from the c tensor. We briefly summarize these formulas for the SEC on general k -forms in Table 2.1.

Acknowledgment.

Tyrus Berry acknowledges support by NSF grant DMS-1723175. Dimitrios Giannakis acknowledges support by ONR grant N00014-14-0150, YIP grant N00014-16-1-2649, NSF grant DMS-1521775, and DARPA grant HR0011-16-C-0116. We would like to thank Jeff Cheeger for stimulating conversations.

Bibliography

- [1] Abraham, R.; Marsden, J. E.; Ratiu, T. *Manifolds, Tensor Analysis, and Applications, Applied Mathematical Sciences*, vol. 75, Springer, New York, 1988.
- [2] Arnold, D.; Falk, R.; Winther, R. Finite element exterior calculus: from Hodge theory to numerical stability. *Bull. Amer. Math. Soc.* **47** (2010), no. 2, 281–354.
- [3] Arnold, D. N.; Falk, R. S.; Winther, R. Finite element exterior calculus, homological techniques, and applications. *Acta Numer.* **15** (2006), 1–155.

- [4] Babuška, I.; Osborn, J. Eigenvalue problems. in *Finite Element Methods (Part 1), Handbook of Numerical Analysis*, vol. II, edited by P. G. Ciarlet; J. L. Lions, pp. 641–787, North-Holland, Amsterdam, 1991.
- [5] Baker, G. Combinatorial Laplacians and Sullivan-Whitney forms. in *Differential Geometry, Progr. Math.*, vol. 32, edited by R. Brooks; A. Gray; B. L. Reinhart, pp. 1–33, Birkhäuser, Boston, 1983.
- [6] Bakry, D.; Gentil, I.; Ledoux, M. *Analysis and Geometry of Markov Diffusion Operators*, Springer, Cham, 2014.
- [7] Belkin, M.; Niyogi, P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* **15** (2003), 1373–1396.
- [8] Belkin, M.; Niyogi, P.: Convergence of Laplacian eigenmaps, in *Advances in Neural Information Processing Systems*, 2007 pp. 129–136.
- [9] Belkin, M.; Niyogi, P.: Convergence of Laplacian eigenmaps, in *Advances in Neural Information Processing Systems*, MIT Press, 2007 pp. 129–136.
- [10] Berry, T.; Giannakis, D.; Harlim, J. Nonparametric forecasting of low-dimensional dynamical systems. *Phys. Rev. E.* **91** (2015), 032 915.
- [11] Berry, T.; Harlim, J. Variable bandwidth diffusion kernels. *Appl. Comput. Harmon. Anal.* **40** (2016), no. 1, 68–96.
- [12] Berry, T.; Sauer, T. Local kernels and the geometric structure of data. *Appl. Comput. Harmon. Anal.* **40** (2016), no. 3, 439–469.
- [13] Berry, T.; Sauer, T. Density estimation on manifolds with boundary. *Comput. Stat. Data An.* **107** (2017), 1–17.
- [14] Berry, T.; Sauer, T. Consistent manifold representation for topological data analysis. *Foundations of Data Science* **1** (2019), no. 1, 1–38.
- [15] Chicone, C.; Swanson, R. C. Spectral theory for linearizations of dynamical systems. *J. Differ. Equations* **40** (1981), 155–167.
- [16] Christensen, O. Frames and pseudo-inverses. *J. Math. Anal. Appl.* **195** (1995), 401–414.
- [17] Christensen, O. *Frames and bases: An introductory course*, Springer Science & Business Media, 2008.
- [18] Coifman, R. R.; Lafon, S. Diffusion maps. *Appl. Comput. Harmon. Anal.* **21** (2006), 5–30.
- [19] Das, S.; Giannakis, D. Delay-coordinate maps and the spectra of Koopman operators. *J. Stat. Phys.* **175** (2019), no. 6, 1107–1145.
- [20] de Oliveira, C. R. *Intermediate spectral theory and quantum dynamics*, *Progress in Mathematical Physics*, vol. 54, Birkhäuser, Basel, 2008.
- [21] Desbrun, M.; Kanso, E.; Tong, Y. Discrete differential forms for computational modeling. in *Discrete differential geometry*, pp. 287–324, Springer, 2008.
- [22] Dodziuk, J. Finite-difference approach to the Hodge theory of harmonic forms. *Amer. J. Math* **98** (1976), no. 1, 79–104.
- [23] Dodziuk, J.; Patodi, V. K. Riemannian structures and triangulations of manifolds. *J. Indian Math. Soc.* **40** (1976), no. 1-4, 1–52.
- [24] Duffin, R. J.; Schaeffer, A. C. A class of nonharmonic Fourier series. *Trans. Amer. Math. Soc.* **72** (1952), no. 2, 341–366.
- [25] Dunford, N.; Schwartz, J. T. *Linear operators*, vol. 2, Wiley, 1988.
- [26] El Karoui, N. On information plus noise kernel random matrices. *Ann. Statist.* **38** (2010), no. 5, 3191–3216.
- [27] El Karoui, N.; Wu, H.-T. Graph connection Laplacian methods can be made robust to noise. *Ann. Statist.* **44** (2016), no. 1, 346–372.
- [28] Ferreira, J. C.; Menegatto, V. A. Positive definiteness, reproducing kernel Hilbert spaces, and beyond. *Ann. Funct. Anal.* **4** (2013), 64–88.

- [29] Gannakis, D.; Slawinska, J.; Zhao, Z. Spatiotemporal feature extraction with data-driven Koopman operators. *J. Mach. Learn. Res. Proceedings* **44** (2015), 103–115.
- [30] Giannakis, D. Data-driven spectral decomposition and forecasting of ergodic dynamical systems. *Appl. Comput. Harmon. Anal.* **62** (2019), no. 2, 338–396.
- [31] Ginelli, F.; Poggi, P.; Turchi, A.; Chaté, H.; Livi, R.; Politi, A. Characterizing dynamics with covariant Lyapunov vectors. *Phys. Rev. Lett.* **99** (2007), 130 601.
- [32] Hein, M.; Audibert, J.-Y.; von Luxburg, U. From graphs to manifolds—weak and strong pointwise consistency of graph Laplacians. in *Learning Theory, Lecture Notes in Computer Science*, vol. 3559, pp. 470–485, Springer, Berlin, 2005.
- [33] Hirani, A.: Discrete exterior calculus, Ph.D. thesis, California Institute of Technology, 2003.
- [34] Hörmander, L. The spectral function of an elliptic operator. *Acta Math.* **121** (1968), 193–218.
- [35] Lorenz, E. N. Deterministic nonperiodic flow. *J. Atmos. Sci.* **20** (1963), 130–141.
- [36] Mallat, S. *A Wavelet Tour of Signal Processing*, Academic Press, Burlington, 2009, 3rd ed.
- [37] Meyer, F. G.; Shen, X. Perturbation of the eigenvectors of the graph Laplacian: Application to image denoising. *Appl. Comput. Harmon. Anal.* **36** (2014), 326–334.
- [38] Portegies, J.: Embeddings of Riemannian manifolds with heat kernels and eigenfunctions. 2014. [arXiv:1311.7568](https://arxiv.org/abs/1311.7568).
- [39] Rosenberg, S. *The Laplacian on a Riemannian Manifold, London Mathematical Society Student Texts*, vol. 31, Cambridge University Press, Cambridge, 1997.
- [40] Rufat, D. S.: Spectral exterior calculus and its implementation, Ph.D. thesis, California Institute of Technology, 2017.
- [41] Schulz, E.; Tsogtgerel, G.: Convergence of discrete exterior calculus approximations for Poisson problems. 2016. [arXiv:1611.03955](https://arxiv.org/abs/1611.03955).
- [42] Shi, Y.; Xu, B. Gradient estimate of an eigenfunction on a compact Riemannian manifold without boundary. *Ann. Glob. Anal. Geom.* **38** (2020), 21–26.
- [43] Shi, Z.: Convergence of Laplacian spectra from random samples. 2015. [1507.00151](https://arxiv.org/abs/1507.00151).
- [44] Singer, A. From graph to manifold Laplacian: The convergence rate. *J. Appl. Comput. Harmon. Anal.* **21** (2006), 128–134.
- [45] Singer, A.; Wu, H.-T. Vector diffusion maps and the connection Laplacian. *Commun. Pure Appl. Math.* **65** (2012), no. 8, 1067–1144.
- [46] Singer, A.; Wu, H.-T. Spectral convergence of the connection Laplacian from random samples. *Information and Inference: A Journal of the IMA* **6** (2016), no. 1, 58–123.
- [47] Strichartz, R. S. Analysis of the Laplacian on the complete Riemannian manifold. *J. Funct. Anal.* **52** (1983), 48–79.
- [48] Sullivan, D.: Differential forms and the topology of manifolds, in *Manifolds-Tokyo: Proceedings of the International Conference on Manifolds and Related Topics in Topology, Tokyo, 1973*, University of Tokyo Press, Tokyo, 1975 pp. 37–49.
- [49] Trillos, N. G.; Gerlach, M.; Hein, M.; Slepčev, D.: Error estimates for spectral convergence of the graph Laplacian on random geometric graphs towards the Laplace–Beltrami operator. 2018. [1801.10108](https://arxiv.org/abs/1801.10108).
- [50] Trillos, N. G.; Slepčev, D. A variational approach to the consistency of spectral clustering. *Appl. Comput. Harmon. Anal.* **45** (2018), no. 2, 239–281.
- [51] Tucker, W. The Lorenz attractor exists. *C. R. Acad. Sci. Paris, Ser. I* **328** (1999), 1197–1202.
- [52] von Luxburg, U.; Belkin, M.; Bousquet, O. Consistency of spectral clustering. *Ann. Statist.* **26** (2008), no. 2, 555–586.
- [53] Whitney, H. *Geometric Integration Theory*, Princeton University Press, Princeton, 1957.
- [54] Yan, B.; Sarkar, P.: On robustness of kernel clustering, in *30th Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain, 2016.