

HODGELETS: LOCALIZED SPECTRAL REPRESENTATIONS OF FLOWS ON SIMPLICIAL COMPLEXES

T. Mitchell Roddenberry*, Florian Frantzen†, Michael T. Schaub†, Santiago Segarra*

*Rice University, Dept. of Electrical and Computer Engineering, Houston, TX, USA

†RWTH Aachen University, Dept. of Computer Science, Aachen, Germany

ABSTRACT

We develop wavelet representations for edge-flows on simplicial complexes, using ideas rooted in combinatorial Hodge theory and spectral graph wavelets. We first show that the Hodge Laplacian can be used in lieu of the graph Laplacian to construct a family of wavelets for higher-order signals on simplicial complexes. Then, we refine this idea to construct wavelets that respect the Hodge-Helmholtz decomposition. For these *Hodgelets*, familiar notions of curl-free and divergence-free flows from vector calculus are preserved. We characterize the representational quality of our Hodgelets for edge flows in terms of frame bounds and demonstrate the use of these spectral wavelets for sparse representation of edge flows on real and synthetic data.

Index Terms— Graph signal processing, Hodge Laplacian, Simplicial complex, Wavelet, Discrete calculus

1. INTRODUCTION

There has been substantial interest in graph-based techniques to understand data with a complex relational structure [1–3], with applications ranging from biology [4] to system robustness [5]. In this context, graph signal processing (GSP) has proven to be a useful way to understand the processing of signals defined on graphs, leveraging ideas from both signal processing and graph theory [6]. The primary focus of GSP has been on signals supported on the *nodes* of a graph. For such signals, the graph Laplacian and adjacency matrix are natural shift operators, from which we can define notions of filtering and Fourier transformations [6].

However, there has been a recent flurry of interest in studying *flows* on the edges of graphs and simplicial complexes [7–12], which can be used to model the motion of mass, energy, or information. Since flows carry a natural orientation that does not arise when studying signals on the nodes of a graph, recent works have leveraged tools from algebraic topology [13] and discrete exterior calculus [14] to form appropriate Laplace operators that respect the orientation of edge flows. This approach has allowed for the study of edge flows through the lens of the celebrated Hodge-Helmholtz decomposition [7–10]. This viewpoint has even been leveraged to define neural network architectures for edge flows [15–19].

In the literature thus far, the primary focus has been on understanding filtering and signal representation in the spatial and Fourier domains, where we take the Fourier modes to be the eigenvectors of a suitably defined Laplacian. However, just as in classical signal processing, the Fourier modes are highly *delocalized*. That is to

say, the support of a Fourier mode is typically not restricted to one small region of the graph. In GSP, this has motivated the development of *spectral graph wavelets* [20, 21], which proposes to use a dictionary of atoms for signal representation that is localized in both the spatial and frequency domains. Here, we introduce a family of *wavelets for edge flows*, seeking to balance localization in the spatial and frequency domains, while also respecting the Hodge decomposition [22].

Contributions and outline. We consider the design of *spectral wavelets* for edge flows on simplicial complexes. In particular, we discuss how the orthogonal decomposition of the space of edge flows in terms of the Hodge Laplacian can be leveraged to design interpretable wavelets that yield high-quality sparse and localized representations of edge flows.

We begin by discussing preliminaries in graph signal processing for edge flows in Section 2. Then, we propose a simple construction for spectral graph wavelets based on previous literature in Section 3, as well as a modification that respects the Hodge decomposition. Theoretical properties of both models are considered in Section 4. In particular, we state frame bounds for both models, in terms of the family of spectral kernels used in their definition. Finally, we demonstrate the utility of our constructions for sparse representation and flow clustering on real and synthetic data in Section 5.

2. NOTATION AND BACKGROUND

For a positive integer N , we denote the set of integers $\{1, 2, \dots, N\}$ by $[N]$. We use $\cong$ to denote isomorphism between vector spaces, and $\oplus$ to denote the orthogonal direct sum of vector spaces. For a linear operator $\mathbf{A}$ between two vector spaces, we denote the set of eigenvalues of $\mathbf{A}$ by $s(\mathbf{A})$.

Simplicial complexes and the Hodge Laplacian. We consider data supported on (*abstract*) *simplicial complexes*, which generalize graphs to allow for higher-order connectivity. An (abstract) simplicial complex X is a finite collection of finite sets that is closed under restriction: that is to say, for any σ in X , all nonempty subsets of σ are also contained in X . We call the elements of X *simplices* and denote by X_k the set of all simplices in X with cardinality $k + 1$, also referred to as *k-simplices*. In particular, X_0 is the set of all singleton sets in X , X_1 is the set of all simplices with cardinality 2, and so on. Grounded in our intuition for graphs, we call X_0 the set of *nodes* in X , X_1 the set of *edges* in X , and X_2 the set of *triangles*.

We identify the set X_k with the integers $[N_k]$, and denote the cardinality of X_k by N_k . By convention, we label the nodes with $1, 2, \dots, N_0$. Further, we assign to each k -simplex an *orientation*¹, or a canonical ordering, following the ordering induced by the node

¹The choice of orientation is arbitrary and distinct from the notion of direction, e.g., in a directed graph. See [22, 23] for details.

This work was partially supported by USA NSF under award CCF-2008555. FF and MTS acknowledge partial support from the Excellence Strategy of the Federal Government and the Länder in Germany, and the NRW Rückkehrprogramm.

labeling, e.g., a triangle $\{i, j, k\}$ is given the orientation $[i, j, k]$, where $i < j < k$. Given this reference orientation, we define the space of k -cochains, denoted by $C^k(X)$, as the vector space of functions from the oriented simplices in X_k to $\mathbb{R}$. One can check that for each k , $C^k(X)$ is naturally isomorphic to $\mathbb{R}^{N_k}$. For a given $C^k(X)$, we take as an orthonormal basis for that space the set of functions $\{\mathbf{e}_\sigma\}_{\sigma \in X_k}$ taking unit value on each oriented k -simplex σ , and zero elsewhere. Of particular interest is the space $C^1(X)$, which models flows on the edges of a graph or simplicial complex [7–10].

The spaces $C^k(X)$ are related by the set of *incidence matrices*. The 2nd incidence matrix $\partial_2 : C^2(X) \rightarrow C^1(X)$ is a linear map defined over the standard orthonormal basis for $C^2(X)$ as follows: $\partial_2 \mathbf{e}_{[i,j,k]} = \mathbf{e}_{[i,j]} + \mathbf{e}_{[j,k]} - \mathbf{e}_{[i,k]}$. Similarly, the 1st incidence matrix $\partial_1 : C^1(X) \rightarrow C^0(X)$ is the familiar node-edge incidence matrix, defined according to $\partial_1 \mathbf{e}_{[i,j]} = \mathbf{e}_{[j]} - \mathbf{e}_{[i]}$. We also call the matrices ∂_2, ∂_1 the *boundary maps*, since they map each k -simplex to the $k-1$ -simplices on its boundary, as well as obeying the important property $\partial_1 \partial_2 = 0$. The following well-known result characterizes the space $C^1(X)$ in terms of these maps.

Theorem 1 (Hodge Decomposition). *Let X be a finite simplicial complex. Define the first Hodge Laplacian as $\Delta_1 = \partial_2 \partial_2^\top + \partial_1^\top \partial_1$. The vector space $C^1(X)$ can be written as the direct sum of orthogonal subspaces:*

$$C^1(X) \cong \text{Im}(\partial_2) \oplus \text{Im}(\partial_1^\top) \oplus \ker(\Delta_1). \quad (1)$$

See [23] for the proof. Viewing the standard graph Laplacian $\Delta_0 = \partial_1 \partial_1^\top$ as an operator that measures smoothness for node signals based on their edgewise incidence, the Hodge Laplacian Δ_1 similarly measures “smoothness” for 1-cochains based on nodewise and trianglewise incidence. For convenience, we define the *upper Hodge Laplacian* $\Delta_1^U = \partial_2 \partial_2^\top$ and the *lower Hodge Laplacian* $\Delta_1^L = \partial_1^\top \partial_1$, so that $\Delta_1 = \Delta_1^U + \Delta_1^L$.

Spectral graph wavelets. A key component of many modern signal processing and machine learning tasks is the choice of a proper representation for the data. In graph signal processing, this often amounts to using a spatial representation or a frequency domain representation. In time-domain signal processing, one can construct a dictionary of *wavelets* that interpolates between these two extremes [24]. Similarly, in [20, 21], methods to construct *spectral graph wavelets* are proposed, in which a family of kernel functions $\{g_m : \mathbb{R} \rightarrow \mathbb{R}\}_{m \in \Gamma}$ for some index set Γ is applied to the standard basis for $C^0(X)$ on a graph X to yield a dictionary for localized, bandlimited representations of graph signals as follows:

$$\mathcal{D} = \{\psi_{j,m} = g_m(\Delta_0) \mathbf{e}_j : m \in \Gamma, j \in [N_0]\}, \quad (2)$$

where Δ_0 indicates the graph Laplacian.

3. AN INTERPRETABLE SPECTRAL WAVELET MODEL

Much like in graph signal processing, we seek useful representations of edge flows that balance spatial localization and bandlimitedness in the frequency domain. As done by [11], we can define a *simplicial Fourier transform* by projecting an edge flow onto each eigenvector of Δ_1 . To illustrate this, we construct a simplicial complex by picking 40 points randomly distributed in the unit square, then taking their Delaunay triangulation [25]. We then create one hole in this simplicial complex, and construct the Hodge Laplacian Δ_1 for the remaining structure. We illustrate this complex as well as one of the eigenvectors in its nullspace in Fig. 1 (a). We see that the flow wraps around the hole of the complex, a global phenomenon, but




Fig. 1. Dictionary atoms for representing flows on simplicial complexes. All signals are scaled to have unit ℓ_∞ -norm. Edge orientations are indicated by the direction of the arrowheads. (a) An eigenvector in the nullspace of the Hodge Laplacian Δ_1 . (b) Jointly designed wavelet based on Δ_1 . (c1) Upper wavelet based on Δ_1^U . Triangles are colored according to the 2-cochain $\mathbf{c}$ such that the wavelet atom is equal to $\partial_2 \mathbf{c}$, with orientation of the flow induced by each triangle indicated by arrows (space permitting). (c2) Lower wavelet based on Δ_1^L . Nodes are colored according to the 0-cochain $\mathbf{x}$ such that the wavelet atom is equal to $\partial_1^\top \mathbf{x}$.

does not have *localized* support. In the ensuing discussion, we aim to construct dictionaries for edge flows that find a balance in this local-global tradeoff.

Jointly designed wavelets. As the Hodge Laplacian generalizes the graph Laplacian for higher-order signals, it is natural to apply the methods of [20, 21] for constructing graph wavelets by substituting the Hodge Laplacian for the graph Laplacian. In that direction, let $\{\mathbf{e}_j\}_{j=1}^{N_1}$ be the standard orthonormal basis for $C^1(X)$, and let $\{g_m\}_{m=1}^M$ be a set of continuous, non-negative functions on the real line. For each $j \in [N_1]$, $m \in [M]$, we define the atom $\psi_{j,m}$ by applying the polynomial $g_m(\Delta_1)$ to $\mathbf{e}_j$, where we have implicitly labeled the oriented 1-simplices with the integers $[N_1]$. That is, if the Hodge Laplacian Δ_1 admits an eigendecomposition $\{(\lambda_i, \mathbf{v}_i)\}_{i=1}^{N_1}$,

$$\psi_{j,m} = g_m(\Delta_1) \mathbf{e}_j = \left(\sum_{i=1}^{N_1} g_m(\lambda_i) \mathbf{v}_i \mathbf{v}_i^\top \right) \mathbf{e}_j. \quad (3)$$

Since this construction uses the sum of the upper and lower Hodge Laplacians and thus both components of the Laplacian, we refer to such atoms as *joint Hodgelets*.

We illustrate this approach in Fig. 1 (b), using the same simplicial complex as before. We construct wavelets using the log-scaled Hann kernels of [21], and illustrate a single wavelet atom. Note that this atom is spatially localized, due to the spectral kernel being well-approximated by a low-order polynomial of the Hodge Laplacian.

Separately designed wavelets. The approach of directly applying the wavelet construction of [20, 21] using the Hodge Laplacian presents some shortcomings. Importantly, it fails to differentiate between the spectral features of the Hodge Laplacian due to the upper and lower components being present in each atom. To address this issue, we treat each subspace of the Hodge decomposition [cf. Theorem 1] separately, by defining *upper* and *lower* wavelets for $C^1(X)$. A similar approach was taken in the design of filters for signals on

simplicial complexes by [11].

As before, let $\{e_j\}_{j=1}^{N_1}$ be the standard orthonormal basis for $C^1(X)$, and let $\{g_m^U\}_{m=1}^{M_U}$, $\{g_m^L\}_{m=1}^{M_L}$ be sequences of continuous functions on the real line. For each $j \in [N_k]$, $m \in [M_U]$, $m' \in [M_L]$, define

$$\begin{aligned}\psi_{j,m}^U &= g_m^U(\Delta_1^U) e_j \\ \psi_{j,m'}^L &= g_{m'}^L(\Delta_1^L) e_j,\end{aligned}\quad (4)$$

where $g_m^U(\Delta_1^U)$, $g_{m'}^L(\Delta_1^L)$ are defined via the functional calculus as before. If the set of M kernels are polynomials of degree K , the complexity of computing the wavelet atoms is $\mathcal{O}(MK(N_0 + N_2))$. The set $\{\psi_{j,m}^U\}$ forms what we call the *upper atoms*, and similarly $\{\psi_{j,m'}^L\}$ forms the set of *lower atoms* for $C^1(X)$. Since this construction separates the upper and lower components of the Hodge Laplacian, we refer to such atoms as *separate Hodgelets*.

By separately treating the upper and lower Hodge Laplacians, we can construct atoms with greater interpretability than those designed jointly. In particular, we can show the following result:

Proposition 1. Suppose g_m^U and $g_{m'}^L$ are kernels that take value 0 at 0. Then, for all $j, j' \in [N_1]$,

$$\psi_{j,m}^U \in \text{Im}(\partial_2) \quad \text{and} \quad \psi_{j',m'}^L \in \text{Im}(\partial_1^\top). \quad (5)$$

Thus, the upper wavelets are dictated by the boundaries of 2-simplices (triangles), and the lower wavelets are dictated by the coboundaries of 0-simplices (nodes). We omit the proof for space reasons.² Indeed, Proposition 1 reflects the properties of [26, Theorem 3], in which curl and divergence wavelets are constructed for differential forms in Euclidean space.

We illustrate this in Fig. 1 (c1,c2), by plotting wavelet atoms with the same kernels as the joint wavelet in Fig. 1 (b), except with a separated construction. One can see that there is a clear distinction between the upper wavelet (c1) which corresponds to a curl around triangles, and the lower wavelet (c2) which corresponds to the gradient of a node signal.

4. FRAME BOUNDS ON DICTIONARIES

In signal processing on graphs, the graph Fourier transform has the appealing property of being an orthogonal transform, thus preserving the norm of the signal it acts upon. Since wavelet dictionaries are typically overcomplete in their construction, we do not have orthogonality, but rather have *frame bounds* for the dictionary. For a Hilbert space V , a dictionary of vectors $\mathcal{D}$ with at most countably many elements is said to be an (A, B) -frame with $0 \leq A \leq B < \infty$ if for all $v \in V$, we have

$$A\|v\|^2 \leq \sum_{\psi \in \mathcal{D}} |\langle \psi, v \rangle|^2 \leq B\|v\|^2. \quad (6)$$

If $A = B$, we say that $\mathcal{D}$ forms a *tight frame*. We allow for the case where $A = 0$, in which case $\mathcal{D}$ is a degenerate frame. The frame bounds of a dictionary dictate its representational quality, as well as the performance of reconstruction algorithms [24]. Moreover, if $A = B = 1$, the coefficients $|\langle \psi, v \rangle|^2$ are analogous to the spectrogram representation of a signal [21]. Here, in the same vein as [20, 21], we characterize the frame bounds for both Hodgelet constructions in terms of the spectral properties of the kernels g_m .

Joint wavelets. Given that the jointly designed wavelets are a direct adaptation of those proposed in [20, 21], we can show a similar result in our context:

²Proofs can be found at arxiv.org/abs/2109.08728.

Theorem 2. (based on [20, 21]) Let $\{e_j\}_{j=1}^{N_1}$ be the standard orthonormal basis of $C^1(X)$, and let $\{g_m\}_{m=1}^M$ be continuous non-negative functions on the real line. Let $\mathcal{D}$ be the dictionary of atoms defined by (3), and define

$$G(\lambda) = \sum_{m=1}^M |g_m(\lambda)|^2. \quad (7)$$

Then, $\mathcal{D}$ forms an (A, B) -frame for $C^1(X)$, where

$$A = \min_{\lambda \in s(\Delta_1)} G(\lambda) \quad B = \max_{\lambda \in s(\Delta_1)} G(\lambda). \quad (8)$$

We omit the proof, as it directly mirrors that of [21]. In particular, if G is constant on $s(\Delta_0)$, then $\mathcal{D}$ is a tight frame.

Separate wavelets. We now state frame bounds for the separate Hodgelet construction, keeping the mutual orthogonality of the upper and lower Hodge Laplacians in mind.

Theorem 3. Let $\{e_j\}_{j=1}^{N_1}$ be the standard orthonormal basis for $C^1(X)$, and let $\{g_m^U\}_{m=1}^{M_U}$, $\{g_{m'}^L\}_{m'=1}^{M_L}$ be collections of continuous non-negative functions on the real line. Let $\mathcal{D}$ be the dictionary of separate Hodgelets defined by (4), and define

$$G(\mu, \nu) = \sum_{m=1}^{M_U} |g_m^U(\mu)|^2 + \sum_{m'=1}^{M_L} |g_{m'}^L(\nu)|^2 \quad (9)$$

Then, $\mathcal{D}$ forms an (A, B) -frame for $C^1(X)$, where

$$\begin{aligned}A &= \min \left\{ \min_{\mu \in s(\Delta_1^U)} G(\mu, 0), \min_{\nu \in s(\Delta_1^L)} G(0, \nu) \right\} \\ B &= \max \left\{ \max_{\mu \in s(\Delta_1^U)} G(\mu, 0), \max_{\nu \in s(\Delta_1^L)} G(0, \nu) \right\}.\end{aligned} \quad (10)$$

By Theorem 3, we see that the frame bounds are determined by the quality of the kernels for the upper and lower parts of the spectrum independently. In particular, if $G(\cdot, 0)$ and $G(0, \cdot)$ are constant on $s(\Delta_1^U)$ and $s(\Delta_1^L)$, respectively, then the dictionary forms a tight frame.

In the context of Proposition 1, Theorem 3 indicates that we can construct norm-preserving representations of edge flows that are also interpretable in terms of harmonic flows in $\ker(\Delta_1)$, curl flows in $\text{Im}(\partial_2)$, and divergence flows in $\text{Im}(\partial_1^\top)$. This is aligned with the development in [11], where edge flow filters were designed using Δ_1^U and Δ_1^L separately, rather than the total Hodge Laplacian Δ_1 .

5. EXPERIMENTS

We demonstrate the utility of applying spectral wavelets based on the Hodge Laplacian for sparse, localized representations of flow data. For all experiments,³ we take the spectral kernels $\{g_m\}_{m=1}^M$ to be the log-scaled Hann kernels proposed by [21] with $R = 3$, where R dictates the degree of overlap between each kernel, and M is chosen based on the particular task.

Sparse representations. To illustrate the advantage of using wavelets based on the upper and lower Hodge Laplacians separately, we consider a vector field on $[-2, 2]^2$, given by

$$F(x, y) = \begin{cases} [\cos(x + y), \sin(x - y)] & (x, y) \in B_1 \cup B_2, \\ [0, 0] & \text{otherwise,} \end{cases} \quad (11)$$

³Code is available at git.rodnenberry.xyz/hodgelets.




Fig. 2. Representation of flows on simplicial complexes with spectral wavelets. (a) Vector field $F(x, y)$. (b) Sparsity of representing a discretization of the vector field using different dictionaries as error tolerance increases. ‘LG’ indicates the use of the linegraph Laplacian. (c) Sum of all buoy trajectories in the Global Drifter dataset around Madagascar. Signal is scaled to have unit ℓ_∞ -norm. (d) Average maximum inner product L between test set and centroids obtained via sparse k -means clustering using standard basis, Fourier basis, joint Hodgelets, and separate Hodgelets.

where we define B_1, B_2 to be closed balls of radius 0.7 centered at $(\pm\pi/4, \pm\pi/4)$, respectively, as illustrated in Fig. 2 (a). We discretize $[-2, 2]^2$ with a hexagonal grid, then construct a simplicial complex X ($N_0 = 225, N_1 = 629, N_2 = 405$) by treating each hexagon as a node, with edges for each pair of hexagons that share a side, and triangles for each set of three hexagons that share a corner. This vector field is converted to a vector $\mathbf{f} \in C^1(X)$ by taking the flow on each edge to be the total flow perpendicular to the corresponding side between hexagons. Then, dictionaries of joint and separated Hodgelets are constructed with $M = M_U = M_L = 4$. The joint dictionary has 2516 atoms and forms a tight frame, while the separated dictionary has 5032 atoms and also forms a tight frame.

We now consider the sparsity of the representation of $\mathbf{f}$ in each dictionary. For each dictionary $\mathcal{D}$ of atoms in $C^1(X)$, we construct the sparsest linear combination $\hat{\mathbf{f}}$ of atoms in $\mathcal{D}$ via orthogonal matching pursuit [27, 28] such that $\|\hat{\mathbf{f}} - \mathbf{f}\| \leq \epsilon$, for $\epsilon \in (0, \|\mathbf{f}\|]$ sampled on a logarithmic scale. As a baseline, we repeat this task using the Fourier basis vectors, *i.e.*, the eigenvectors of the Hodge Laplacian, as well as with Fourier basis vectors using the linegraph Laplacian [7], and wavelets constructed from the linegraph Laplacian. The results of this are plotted in Fig. 2 (b), where it is apparent that the separately designed dictionary outperforms the jointly designed one, with both performing better than the Fourier bases and the linegraph dictionary.

Clustering buoy trajectories. One setting in which 1-cochains find particular utility is in modeling flows and trajectories [7, 18, 22]. We consider a dataset from the Global Drifter Program dataset, restricted to the region around Madagascar.⁴ In this dataset, a set of buoys floating in the ocean have their location logged every 12 hours,

which we use to construct trajectories along the edges of a triangulation ($N_0 = 133, N_1 = 320, N_2 = 186$) of the region. Since these trajectories consist of a combination of oriented edges, they are naturally modeled as vectors in $C^1(X)$, following the approach taken by [18, 22]. We picture the sum of all $P = 334$ such trajectories in Fig. 2 (c), where the hole in the simplicial complex corresponds to the landmass of Madagascar.

We aim to find a good set of representative trajectories that captures both the local and global structure of the simplicial complex. For a given dictionary, we first transform each trajectory by taking the inner product with each element of the dictionary. Then, we perform sparse k -means [29, 30] with $K = 2$ clusters on $P_{tr} = 0.75P$ trajectories in the dataset.

To evaluate the quality of the found centroids, we use the remaining $P_{ts} = 0.25P$ trajectories as a test set, and compute the average normalized inner product between each test trajectory and the nearest centroid flow. That is, for a set of centroids $\{\mathbf{c}_k\}_{k=1}^K$ in $C^1(X)$, we compute for the test set $\{\mathbf{f}_j\}_{j=1}^{P_{ts}}$ the value

$$L = \frac{1}{P_{ts}} \sum_{j=1}^{P_{ts}} \max_{1 \leq k \leq K} \frac{\langle \mathbf{f}_j, \mathbf{c}_k \rangle}{\|\mathbf{f}_j\|_2 \|\mathbf{c}_k\|_2}, \quad (12)$$

where the norms and inner products are taken with respect to the feature weighting kernel obtained from the sparse k -means procedure. If a set of centroids yields a large value of L , indicating that the set of nearest flows aligns to each centroid well, we interpret it to be a realistic model for the trajectories. We gather the results of this experiment when using the standard basis, the Fourier basis, the joint Hodgelet dictionary, and the separate Hodgelet dictionary in Fig. 2 (d). The wavelet dictionaries were constructed with $M = M^U = M^L = 16$ filter banks. We observe that the standard basis is not suitable for this task, since two trajectories that are close in space could have disjoint support, so the highly localized basis of edges does not perform well. The Fourier basis and joint dictionary perform similarly, while the separate dictionary significantly outperforms the other representations. This is due to the ability of the sparse clustering algorithm to both pick the proper scale for representation and discern between different qualitative features of the signal in this representation, as revealed by Proposition 1.

6. CONCLUSION

Signals supported on the edges of simplicial complexes have been of great interest lately, with applications in computer graphics, mobility analysis, and modeling of physical flow phenomena. We have considered the extension of spectral graph wavelets to such signals by replacing the graph Laplacian with the analogous Hodge Laplacian. In doing so, we open up the possibility of considering the upper and lower components of the Hodge Laplacian separately, in order to yield a dictionary of wavelet atoms that respects the Hodge decomposition. Based on these constructions, we state frame bounds for each type of wavelet dictionary, and then illustrate their utility for sparse representation on synthetic flow data. Leveraging the ability of these wavelet dictionaries to sparsely represent flow signals, we demonstrate how they can be used to find high-quality representative cluster centroids on real-world buoy trajectory data via a sparse k -means procedure.

⁴Data available at <http://www.aoml.noaa.gov/envids/gld/>

7. REFERENCES

- [1] Steven H Strogatz, "Exploring complex networks," *nature*, vol. 410, no. 6825, pp. 268–276, 2001.
- [2] Mark Newman, *Networks: An Introduction*, Oxford university press, March 2010.
- [3] Matthew O Jackson, *Social and economic networks*, Princeton university press, 2010.
- [4] Colin J Garroway, Jeff Bowman, Denis Carr, and Paul J Wilson, "Applications of graph theory to landscape genetics," *Evolutionary Applications*, vol. 1, no. 4, pp. 620–630, 2008.
- [5] Petter Holme, Beom Jun Kim, Chang No Yoon, and Seung Kee Han, "Attack vulnerability of complex networks," *Physical review E*, vol. 65, no. 5, pp. 056109, 2002.
- [6] David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 83–98, Apr. 2013.
- [7] Michael T. Schaub and Santiago Segarra, "Flow smoothing and denoising: Graph signal processing in the edge-space," in *IEEE Global Conference on Signal and Information Processing*, 2018, pp. 735–739.
- [8] Junteng Jia, Santiago Segarra, Michael T. Schaub, and Austin R. Benson, "Graph-based semi-supervised & active learning for edge flows," in *SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2019, pp. 761–771, ACM.
- [9] Sergio Barbarossa and Stefania Sardellitti, "Topological signal processing over simplicial complexes," *IEEE Transactions on Signal Processing*, vol. 68, pp. 2992–3007, 2020.
- [10] Sergio Barbarossa and Stefania Sardellitti, "Topological signal processing: Making sense of data building on multiway relations," *IEEE Signal Processing Magazine*, vol. 37, no. 6, pp. 174–183, 2020.
- [11] Maosheng Yang, Elvin Isufi, Michael T Schaub, and Geert Leus, "Finite impulse response filters for simplicial complexes," *arXiv preprint 2103.12587*, 2021.
- [12] Michael T. Schaub, Yu Zhu, Jean-Baptiste Seby, T. Mitchell Roddenberry, and Santiago Segarra, "Signal processing on higher-order networks: Livin' on the edge... and beyond," *Signal Processing*, vol. 187, pp. 108149, October 2021.
- [13] Allen Hatcher, *Algebraic Topology*, 2005.
- [14] Anil Nirmal Hirani, *Discrete Exterior Calculus*, Ph.D. thesis, California Institute of Technology, 2003.
- [15] T. Mitchell Roddenberry and Santiago Segarra, "HodgeNet: Graph neural networks for edge data," in *Asilomar Conference on Signals, Systems, and Computers*, November 2019, pp. 220–224.
- [16] Stefania Ebli, Michaël Defferrard, and Gard Spreemann, "Simplicial neural networks," in *NeurIPS Workshop on Topological Data Analysis and Beyond*, 2020.
- [17] Eric Bunch, Qian You, Glenn Fung, and Vikas Singh, "Simplicial 2-complex convolutional neural nets," in *NeurIPS Workshop on Topological Data Analysis and Beyond*, 2020.
- [18] T. Mitchell Roddenberry, Nicholas Glaze, and Santiago Segarra, "Principled simplicial neural networks for trajectory prediction," in *International Conference on Machine Learning*, Marina Meila and Tong Zhang, Eds. July 2021, vol. 139 of *Proceedings of Machine Learning Research*, pp. 9020–9029, PMLR.
- [19] Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lió, and Michael Bronstein, "Weisfeiler and Lehman go topological: Message passing simplicial networks," in *International Conference on Machine Learning*, Marina Meila and Tong Zhang, Eds. 18–24 Jul 2021, vol. 139 of *Proceedings of Machine Learning Research*, pp. 1026–1037, PMLR.
- [20] David K Hammond, Pierre Vandergheynst, and Rémi Gribonval, "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, 2011.
- [21] David I Shuman, Christoph Wismeyer, Nicki Holighaus, and Pierre Vandergheynst, "Spectrum-adapted tight graph wavelet and vertex-frequency frames," *IEEE Transactions on Signal Processing*, vol. 63, no. 16, pp. 4223–4235, 2015.
- [22] Michael T. Schaub, Austin R. Benson, Paul Horn, Gabor Lippner, and Ali Jadbabaie, "Random walks on simplicial complexes and the normalized Hodge 1-Laplacian," *SIAM Review*, vol. 62, no. 2, pp. 353–391, 2020.
- [23] Lek-Heng Lim, "Hodge Laplacians on graphs," *SIAM Review*, vol. 62, no. 3, pp. 685–715, 2020.
- [24] Stéphane Mallat, *A Wavelet Tour of Signal Processing*, Elsevier, 1999.
- [25] Boris Delaunay et al., "Sur la sphère vide," *Bulletin de l'Académie des Sciences de l'URSS, Classe des Sciences Mathématiques et Naturelles*, vol. 7, no. 793-800, pp. 1–2, 1934.
- [26] Christian Lessig, " Ψ ec: A local spectral exterior calculus," *Applied and Computational Harmonic Analysis*, vol. 51, pp. 56–103, 2021.
- [27] Yagyensh Chandra Pati, Ramin Rezaiifar, and Perinkulam Sambamurthy Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Asilomar Conference on Signals, Systems, and Computers*. IEEE, 1993, pp. 40–44.
- [28] Joel A Tropp and Anna C Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [29] Daniela M. Witten and Robert Tibshirani, "A framework for feature selection in clustering," *Journal of the American Statistical Association*, vol. 105, no. 490, pp. 713–726, 2010.
- [30] Daniela M Witten and Robert Tibshirani, "sparcl: Perform sparse hierarchical clustering and sparse k-means clustering," *R package*, vol. 1, no. 3, 2013.