

Algebraic Topological Networks via the Persistent Local Homology Sheaf

Gabriele Cesa

Arash Behboodi

Qualcomm AI Research, Amsterdam*

GCESA@QTI.QUALCOMM.COM

BEHBOODI@QTI.QUALCOMM.COM

Editors: Sophia Sanborn, Christian Shewmake, Simone Azeglio, Nina Miolane

Abstract

In this work, we introduce a novel approach based on algebraic topology to enhance graph convolution and attention modules by incorporating local topological properties of the data. To do so, we consider the framework of *sheaf neural networks*, which has been previously leveraged to incorporate additional structure into graph neural networks’ features and construct more expressive, non-isotropic messages. Specifically, given an input simplicial complex (e.g. generated by the cliques of a graph or the neighbors in a point cloud), we construct its *local homology sheaf*, which assigns to each node the vector space of its local homology. The intermediate features of our networks live in these vector spaces and we leverage the associated sheaf Laplacian to construct more complex linear messages between them. Moreover, we extend this approach by considering the *persistent* version of local homology associated with a weighted simplicial complex (e.g., built from pairwise distances of nodes embeddings). This *i)* solves the problem of the lack of a natural choice of basis for the local homology vector spaces and *ii)* makes the sheaf itself differentiable, which enables our models to directly optimize the topology of their intermediate features.

Keywords: Graph, Simplicial, Sheaf, Laplacian, Homology, Topology

1. Introduction

Many works in the literature extended standard *Graph Convolution Networks* (GCNs) [Kipf and Welling \(2016\)](#), which rely on isotropic message passing along a graph’s edges, to more expressive message passing operators. *Sheaf neural networks* [Hansen and Gebhart \(2020\)](#) provide a generic framework to encode more structure into the features attached to a graph’s nodes, which can be leveraged to define more expressive messages between the feature spaces of neighboring nodes via the sheaf’s restriction maps and the *sheaf Laplacian*. Briefly, a **sheaf** \mathcal{F} on a space X associates a (feature) vector space $\mathcal{F}(U)$ to each (open) set $U \subset X$ and a linear map $\mathcal{F}(U \subset V)$ to each pair $U \subset V$, i.e. the *restriction map*. Two restrictions $\mathcal{F}(W \subset U)^T \mathcal{F}(W \subset V)$ can be combined to send messages between U and V via their intersection $W = U \cap V$: this is the idea behind the *sheaf Laplacian*. While a sheaf should also satisfy *locality* and *gluing properties*, these are not necessary to construct the Laplacian and are usually ignored in neural networks; see [Apx. B](#) for more details. In practice, sheaf neural networks associate a feature vector space to each node in a graph and a linear map to each edge, relating the feature spaces of connected nodes. With respect to the graph Laplacian, this new Laplacian doesn’t enforce similarity between neighboring nodes’ features, thereby circumventing the homophily assumption [Bodnar et al. \(2022\)](#).

* Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

GCNs are the simplest example of sheaf neural networks: these architectures rely on a sheaf which associates the same vector space to each node and whose restriction maps are identities. This enables a simple weight sharing at the cost of less expressive message passing. Other works can be interpreted under this lens: [de Haan et al. \(2020\)](#) constructs a very expressive sheaf over graphs where each node has a feature dimension for each of its neighbors and restriction maps match dimensions corresponding to the same nodes¹. Alternatively, since datasets rarely come with a sheaf structure already defined, [Bodnar et al. \(2022\)](#) propose learning to predict restriction maps from input features during inference.

Contributions We use tools from *algebraic topology* [Hatcher \(2002\)](#) to construct a new sheaf for neural networks: the **Local Homology** sheaf in the *flag complex* of a graph [Robinson et al. \(2018\)](#). This sheaf catches local topological features of a space: it associates to each node a feature vector space with a component for each "relative cycle" in its neighborhood. Intuitively, an order k local relative cycle detects a subspace which locally looks like a k -dimensional manifold. For this reason, the local homology sheaf is typically used for **stratification detection** of triangulated spaces. Interestingly, sheaf diffusion along the edges is sufficient to detect higher order (local and global) homological properties of the space, with no need of higher-order simplicial message passing.

Unfortunately, the homology sheaf doesn't prescribe a natural choice of basis for the feature vector space, which makes constructing learnable linear and activation layers challenging. We tackle this limitation by considering *weighted graphs* and leveraging **persistent homology**, the standard tool in *Topological Data Analysis* [Carlsson \(2009\)](#). Finally, this new construction generates a sheaf whose Laplacian is *differentiable* with respect to the graph weights, which can be output of another learnable module (e.g. from learnable node embeddings): this enables our model to learn the sheaf structure or tune the weights in a topological informed way.

2. Simplicial Complexes, Homology and the Local Homology Sheaf

We first briefly review some essential concepts but see [Apx. C](#) for more details.

Simplicial Complexes Assume a *finite* set V of $|V| = N$ nodes. A **simplicial complex** is a collection $S \subset 2^V$ of subsets of V ; a subset $\sigma \in S$ with $k + 1$ elements is called a **k -simplex**. Simplicial complexes generalize the common notion of *graph* beyond pairwise relationships. For example, if $G = (V, E)$ is a graph, its **flag** (or **clique**) **complex** is a simplicial complex S with nodes V and containing a simplex for each *clique* in G , i.e. for each set of nodes in G which form a complete subgraph.

Chains and Boundaries The graph Laplacian can be constructed from the *incidence matrix* $\partial \in \mathbb{R}^{|V| \times |E|}$ as $\Delta_0 = \partial\partial^T$. This construction generalizes to simplicial complexes. A **k -chain** of S is a scalar signal over (oriented) k -simplicies; $C_k(S)$, or just C_k , is the vector space of all k -chains. The incidence matrix is generalized by the **boundary operator** $\partial_k: C_k \rightarrow C_{k-1}$, which models the relationship between each k -simplex and its *faces* (its k -dimensional subsets). The k -th *Hodge Laplacian* is defined as $\Delta_k := \partial_k^T \partial_k + \partial_{k+1} \partial_{k+1}^T: C_k \rightarrow C_k$ and has been used to construct a variety of simplicial neural networks [Papillon et al. \(2023\)](#).

1. Messages are actually constructed with something more similar to a cosheaf Laplacian by leveraging the union rather than the intersection of open sets. The work also supports more generic feature spaces.




Figure 1: Examples of homology and relative homology. The greyed out simplexes can be thought as being "collapsed" in a single point to compute relative homology: then, the blue area β turns into a 2-sphere while the red line γ turns into a 1d ring.

Cycles and Homology A classical result in topology is that *a boundary of a space has no boundary*: $\text{im } \partial_{k+1} \subset \ker \partial_k$. The k -th **homology group** is the *quotient vector space* $H_k(S) := \ker \partial_k / \text{im } \partial_{k+1}$. Its dimensionality $\dim H_k$ is an important invariant counting the k -dimensional holes in S and its basis can be thought as a set of independent k -dimensional cycles in S (0-cycles are connected components, 1-cycles are loops, 2-cycles are cavities).

Our construction is similar to (Robinson et al., 2018), which first introduced the *Local Homology Sheaf* over simplicial complexes. Given a k -simplex $\sigma \in S$, define its **star** as $\text{star } \sigma = \{\tau \in S : \sigma \subset \tau\}$. An *open subset* $A \subseteq S$ is the union of sets of the form $\text{star } \sigma$; note that this is not necessarily a simplicial complex. Instead, a subset $A \subseteq S$ is *closed* if it is a *subcomplex* of S (the faces of every simplex in A are also in A). We also define the **closure** $\text{cl } A$ as the smallest subcomplex of S containing A , the **interior** $\text{int } A$ as the largest open set contained in A and the **frontier** as $\partial A = \text{cl } A \setminus A$.

Relative Homology Let $A \subseteq S$ be a *subcomplex* of S . The k -th **relative homology** $H_k(S, A)$ describes the k -th homology of the *quotient space* S/A obtained from S by identifying all its points within A , i.e. by "collapsing" all points in A in a single point. Fig. 1 shows a few examples. However, note that the relative homologies $H_k(S, S \setminus \text{star } v)$ doesn't depend on (most) gray simplexes in $S \setminus \text{star } v$, but only on those in $\text{star } v$ and its closest neighbors. This is the **Excision Principle**: if $A \subset B \subset S$ are subsets of S such that $\text{cl } A \subset \text{int } B$, then $H_k(S, B) \cong H_k(S \setminus A, B \setminus A)$. When $A \subset S$ is an open set, $H_k(S, S \setminus A) \cong H_k(\text{cl } A, \partial A)$.

Local Homology Sheaf As in Robinson et al. (2018), we consider the sheaf \mathcal{H}_* defined as $\mathcal{H}_*(A) = H_*(S, S \setminus A) \cong H_*(\text{cl } A, \partial A)$ for each open set $A \subset S$ ($S \setminus A$ is closed if A is open). The sheaf structure is naturally given by the following *long exact sequence*²:

$$\cdots \longrightarrow \mathcal{H}_k(A \cup B) \xrightarrow{i_*, j_*} \mathcal{H}_k(A) \oplus \mathcal{H}_k(B) \xrightarrow{k_* - l_*} \mathcal{H}_k(A \cap B) \longrightarrow \cdots \quad (1)$$

where k_*, l_*, i_* and j_* are the sheaf restriction maps. This is a special case of the well known *Mayer-Vietoris sequence*; see Apx. C.1. In particular, $\mathcal{H}_*(\text{star } v_i)$ is called the **local homology** of the vertex v_i . Intuitively, the local homology of a point in a topological space

2. An *exact sequence* is a sequence of maps s.t. the image of a map equals to the kernel of the consecutive one.

contains information about what the space looks like around that point. If the space is an n -manifold, the local neighborhood U of any point looks like a n -ball, whose boundary ∂U is isomorphic to a $n - 1$ -sphere \mathcal{S}^{n-1} . Then, like in Fig. 1, via excision the local homology is $H_*(U, \partial U) \cong \tilde{H}_*(\mathcal{S}^n)$, i.e. the (reduced) homology of an n -sphere, which only has one cycle of order n . Hence, local homology *detects the local dimensionality of a space*. Moreover, points at the boundary of the space have empty local homology. This idea was used in Robinson et al. (2018), among others, for *stratification detection*. Finally, note that the restriction maps constructed in Eq. 1 are identity maps on \mathcal{H}_n for points in the interior of an n -manifold³.

Finally, recall that **sheaf diffusion** minimizes the *sheaf Dirichlet energy* of a signal Bodnar et al. (2022). At zero energy, the signal is in the Laplacian’s kernel and, by the sheaf property, belongs to the **global sections** of $\mathcal{H}(S)$ Hansen and Ghrist (2021). Because $\mathcal{H}_k(S) = H_k(S, \emptyset) = H_k(S)$ (Corollary 20 Robinson et al. (2018)), diffusion converges towards the global homology classes of S of any order k while only relying on messages along edges.

Persistent Homology provides a richer structure than homology, by enriching homology classes with a (differentiable) notion of resolution; see Apx. D. Rather than building a single sheaf for a fixed complex S , we consider a *filtration*, i.e. a sequence of simplicial complexes $\{S_t\}_t$ related by inclusion, and build the local homology sheaf of the complex S_t at each time-step t . Cycles in the local homology at a step in the filtration can "*persist*" in the consecutive steps or disappear. This enriches the local homology with a notion of time or scale, i.e. each cycle is associated with a time-step where it emerges and a time-step where it disappears. In practice, we define the "filtered" neighborhood of a node i as $A_i^t = \text{star}^t v_i \subset S_t$ and *compute the persistent cycles* in the **persistent module** $\mathcal{H}_k^\bullet(A_i) = \bigoplus_t H_k(S^t, S^t \setminus A_i^t)$ as in Apx. E. Persistent cycles are shared among the time-steps between their births and deaths, see Eq. 9. This *feature sharing strategy* generates the **persistent relative homology** subspace $\mathcal{H}_k(A_i) \subset \mathcal{H}_k^\bullet(A_i)$. Columns in Fig. 2 are examples of persistent local homology.

3. Proposed Architecture

Given a graph $G = (V, E)$ with weighted edges (e.g. the distance matrix of a point cloud), we construct the *Vietoris-Rips filtration*⁴ $\{S_t\}_t$ of its flag complex S . Unfortunately, while the persistent module \mathcal{H}_k^\bullet forms a sheaf, persistent local homology $\mathcal{H}_k \subset \mathcal{H}_k^\bullet$ fails to be a sheaf Palser (2019). To preserve the sheaf diffusion properties described before, we prefer using the sheaf Laplacian of \mathcal{H}_k^\bullet . Hence, our **message passing** on \mathcal{H}_k first *embeds* persistent homology features in the sheaf \mathcal{H}_k^\bullet , then *applies* the sheaf Laplacian $\Delta_{\mathcal{H}_k^\bullet}$ and, finally, *projects* the output on \mathcal{H}_k by averaging the features of a cycle along its life span. Fig. 2 shows an example of Laplacian $\Delta_{\mathcal{H}_k^\bullet}$. See Apx. F for details on the implementation.

To complete our architecture, we need to include a **learnable layer** operating on each node’s feature space $\mathcal{H}_*(A_i)$. This involves two challenges: *i*) a persistent cycle is only *defined up to a sign* (the Laplacian constructed is equivariant to these sign changes) and *ii*) each node’s feature space looks different. *i*) is related to the spectral symmetries studied in Lim et al. (2023) and can be solved similarly: given $\mathbf{x} \in \mathcal{H}_*(A_i)$, we construct a sign equivariant layer of the form $\psi(\mathbf{x}) = \mathbf{x} \circ \rho(|\mathbf{x}|)$. The learnable operator ρ can be modeled by a simple MLP. To share ρ among different nodes and solve *ii*), we learn a separate MLP Ψ to output

3. The local homology sheaf \mathcal{H}_n is closely related to the *orientation sheaf* of an n -manifold.

4. A simplex appears in the filtration at a time step equal to the maximum weight of its edges.

the weights of ρ for each node individually. Note that each persistent cycle is uniquely identified by its order k and its birth and death times $s, t \in \mathbb{R}$. Then, we can parameterize a linear map on $\mathcal{H}_*(A_i)$ via Ψ as follows: for each pair (i, j) of input/output persistent cycles, the (i, j) -th entry of the weight matrix is parameterized by $\Psi(k_i, s_i, t_i, k_j, s_j, t_j) \in \mathbb{R}$. As usual, this approach can be integrated in a multi-channel network, where the features of the node include multiple copies of the vector space $\mathcal{H}_*(A)$.

4. Limitations and Complexity

Persistent homology is computed by reducing the boundary matrices, with a worst case complexity cubic in the number of simplices. Assuming N nodes and by considering only homology up to order K (typically $K = 2$ or 3), there are at worst $O(N^{K+1})$ simplices so the complexity is $O(N^{3K+3})$. However, thanks to the excision principle, local homology can be computed by using only a limited number of neighboring nodes. Assuming each node has $O(n)$ neighbors, computing the local homology of each node costs only $O(Nn^{3K+3})$. Moreover, the computation of each local homology can be fully parallelized. For each pair of nodes, the sheaf Laplacian is also computed via a matrix reduction using the union of their local neighbors (with $O(2n)$ nodes): with a similar worst case complexity $O((2n)^{3K+3})$ for each pair of nodes $O(nN)$, the overall complexity is then $O(Nn^{3K+4}2^{3K+3})$. Still, we note that there exists optimized algorithms like **Ripser** [Bauer \(2021\)](#), which are much faster on average by leveraging a number of smart heuristics. Additionally, the number of neighbors n can be chosen sufficiently low to control the overall complexity. The main limitation we currently see is the fact that these computations can not be performed on a GPU in a straightforward way. As a result, computing the sheaf structure requires moving the edge weights to the CPU during inference and, then, move the sheaf Laplacian data back on GPU.

5. Conclusions

The proposed local homology sheaf Laplacian can be used to enhance existing deep learning architectures by making them aware of the local and global topology of the underlying data structure during inference. We expect this to be especially useful in tasks such as graph link prediction, mesh reconstruction or simply where the data presents a variety of topologies. We plan to experimentally evaluate this method on similar tasks in future works.

Acknowledgments

We thank Giovanni Luca Marchetti for the very insightful discussions about how to efficiently compute the sheaf laplacian, the Mayer-Vietoris sequences and other algebraic topology ideas.

References

- Ulrich Bauer. Ripser: efficient computation of vietoris–rips persistence barcodes. *Journal of Applied and Computational Topology*, 5(3):391–423, 2021.
- Nello Blaser and Morten Brun. Relative persistent homology. *Discrete & Computational Geometry*, pages 1–15, 2022.
- Cristian Bodnar, Francesco Di Giovanni, Benjamin Chamberlain, Pietro Liò, and Michael Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and over-smoothing in gnns. *Advances in Neural Information Processing Systems*, 35:18527–18541, 2022.
- Rickard Brüel-Gabrielsson, Bradley J Nelson, Anjan Dwaraknath, Primož Skraba, Leonidas J Guibas, and Gunnar Carlsson. A topology layer for machine learning. *arXiv preprint arXiv:1905.12200*, 2019.
- Gunnar Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2): 255–308, 2009.
- Pim de Haan, Taco S Cohen, and Max Welling. Natural graph networks. *Advances in neural information processing systems*, 33:3636–3646, 2020.
- Jakob Hansen and Thomas Gebhart. Sheaf neural networks. *arXiv preprint arXiv:2012.06333*, 2020.
- Jakob Hansen and Robert Ghrist. Opinion dynamics on discourse sheaves. *SIAM Journal on Applied Mathematics*, 81(5):2033–2060, 2021.
- Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2016.
- Derek Lim, Joshua Robinson, Stefanie Jegelka, Yaron Lipman, and Haggai Maron. Expressive sign equivariant networks for spectral geometric learning. In *ICLR 2023 Workshop on Physics for Machine Learning*, 2023.
- Megan Palser. An excision theorem for persistent homology. *arXiv preprint arXiv:1910.03348*, 2019.
- Mathilde Papillon, Sophia Sanborn, Mustafa Hajj, and Nina Miolane. Architectures of topological deep learning: A survey on topological neural networks. *arXiv preprint arXiv:2304.10031*, 2023.
- Michael Robinson, Chris Capraro, Cliff Joslyn, Emilie Purvine, Brenda Pragastis, Stephen Ranshous, and Arun Sathanur. Local homology of abstract simplicial complexes. *arXiv preprint arXiv:1805.11547*, 2018.
- Blaž Sovdat. Text mining via homology, 2016.

Christopher Tralie, Nathaniel Saul, and Rann Bar-On. Ripser.py: A lean persistent homology library for python. *The Journal of Open Source Software*, 3(29):925, Sep 2018. doi: 10.21105/joss.00925. URL <https://doi.org/10.21105/joss.00925>.

Appendix A. Example of persistent sheaf Laplacian

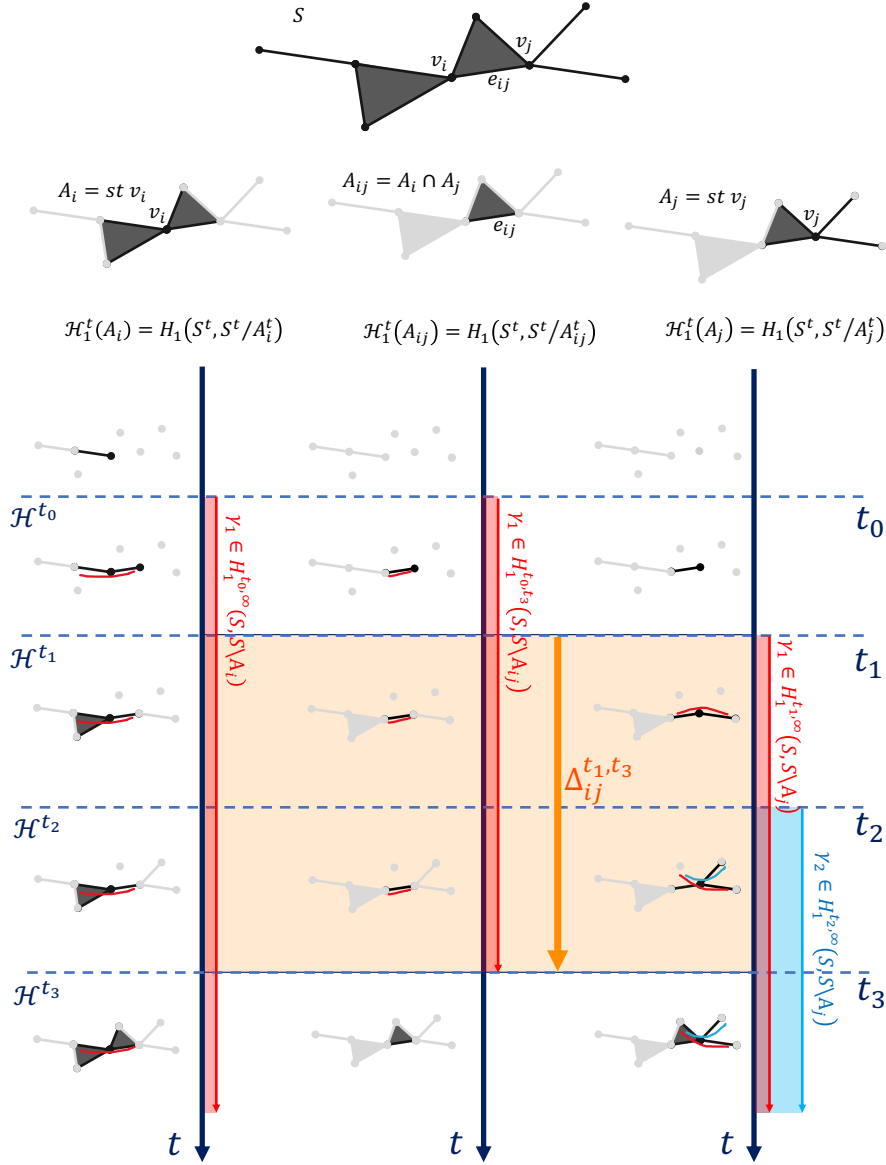


Figure 2: Example of persistent sheaf Laplacian. The three columns depict the time evolution of the filtrations of the local neighborhood of three simplices. At different time steps, some new relative cycles appear or disappear and each cycle "persists" for an interval of time. In our architecture, a single feature is stored for each persistent cycle; this feature can be thought as been shared over all time steps within the cycle's life span. Moreover, the three columns share a **relative 1-cycle** γ_1 . Note that this cycle exists at different times intervals in the three columns and, therefore, there exists a sheaf Laplacian only during the **intersection of these intervals** $[t_1, t_3]$.

Appendix B. Sheaves

Given a space X , a **pre-sheaf** \mathcal{F} associates to each open set $U \subset X$ a space $\mathcal{F}(U)$ and to each pair $U \subset V (\subset X)$ a map $\mathcal{F}(U \subset V) : \mathcal{F}(V) \rightarrow \mathcal{F}(U)$ (*restriction map*), such that $\mathcal{F}(U \subset U)$ is the identity and $\mathcal{F}(U \subset V)\mathcal{F}(V \subset W) = \mathcal{F}(U \subset W)$ if $U \subset V \subset W$. We are mostly interested in the cases where $\mathcal{F}(U)$ are vector spaces. An element of $\mathcal{F}(U)$ is called a "**local section**", while an element of $\mathcal{F}(X)$ is called a "**global section**".

Given an open cover $\{U_i \subset U\}_i$ of $U \subset X$, a **sheaf** is a pre-sheaf satisfying two additional axioms:

1. *locality*: if two local sections $s, t \in \mathcal{F}(U)$ agree when restricted on all $\{U_i\}_i$, then they are identical
2. *gluing*: if a set of local sections $\{s_i \in \mathcal{F}(U_i)\}_i$ agree on all their overlaps, then there exists section $s \in \mathcal{F}(U)$ which agrees with s_i when restricted on U_i , for all i

Given a sheaf \mathcal{F} , with $\mathcal{F}(U)$ vector spaces, we can construct the *sheaf Laplacian* Hansen and Gebhart (2020). To do so, consider an open cover $\{U_i\}_i$ of the space X . For any i, j s.t. $U_i \cap U_j \neq \emptyset$, define the linear map

$$\delta_{ij} : \mathcal{F}(U_i) \times \mathcal{F}(U_j) \rightarrow \mathcal{F}(U_i \cap U_j), \quad x_i, x_j \mapsto \mathcal{F}(U_i \cap U_j \subset U_i)x_i - \mathcal{F}(U_i \cap U_j \subset U_j)x_j$$

Then, the sheaf Laplacian is a block matrix defined as $L_{\mathcal{F}} = \delta^T \delta$. Note that, if $i \neq j$, the (i, j) -th block is defined as $[\Delta_{\mathcal{F}}]_{ij} = -\mathcal{F}(U_i \cap U_j \subset U_i)^T \mathcal{F}(U_i \cap U_j \subset U_j)$.

Given a sheaf defined over a graph, the sheaf Laplacian generalizes the classical graph Laplacian and provides a useful tool to build more expressive message passing operators for neural networks.

To build message passing, the restriction map of a pre-sheaf is sufficient and we do not actually need the additional two axioms of a sheaf. Still, since local homology in Sec. 2 forms a sheaf with some interesting properties and to keep the notation simpler, we use the word "sheaf" also in the message passing architectures which don't enforce these axioms.

Appendix C. Simplicial Complexes, Boundary Maps and Homology

Simplicial Complex Give a *finite* set of nodes V , with $|V| = N \in \mathbb{N}$, a **simplicial complex** S is a mathematical objects that can be thoughts as a collection of subsets of V , i.e. $S \subset 2^V$, such that $\forall \sigma \in S, \tau \subset \sigma \implies \tau \in S$. Each such subset $\sigma \in S$ is called a **simplex**. We usually refer to simplices with $k + 1$ elements as k -simplices. Simplicial complexes generalize the common notion of *graph*, by thinking of an edge as a set containing two nodes. A k -simplex $\sigma = \{v_0, \dots, v_k\} \subset V$, like edges, is typically associated with an **orientation**, i.e. a particular choice of ordering of its elements $\sigma = [v_0, \dots, v_k] \in S$. Two k -simplices σ, σ' containing the same subset of nodes share the same orientation if they differ by an *even permutation* but have opposite orientation if they differ by an *odd permutation*.

Chain Complexes and Boundary Operators Let S be a simplicial complex. A **k -chain** of S is a scalar function $f : S \rightarrow \mathbb{R}$ on the oriented k -simplices of S such that $f(\sigma) = f(\sigma')$ if σ and σ' have the same orientation (differ by an even permutation) and $f(\sigma) = -f(\sigma')$ if they have opposite orientation (differ by an odd permutation). A **chain**

complex $C_\bullet(S)$ is a sequence of vector spaces $C_0(S), C_1(S), \dots$, where $C_k(S)$ is the vector space of all k -chains. A chain complex is associated with a *linear boundary operator* (or *differential*) $\partial_k : C_k \rightarrow C_{k-1}$, defined on a k -simplex $\sigma = [v_0, v_1, \dots, v_k] \in S$ (intended as one of the basis elements of C_k) as⁵

$$\partial_k \sigma := \sum_{i=0}^k (-1)^i [v_0, \dots, \hat{v}_i, \dots, v_k] \quad (2)$$

where $[v_0, \dots, \hat{v}_i, \dots, v_k]$ is a $k - 1$ -simplex obtained from σ by removing the node v_i . We often use $\partial_\bullet : C_\bullet \rightarrow C_\bullet$ to denote the operator acting on each subspace C_k of C_\bullet with the corresponding operator ∂_k .

Example If S is just a graph $G = (V, E)$, C_0 are functions over the nodes V while C_1 are functions over the (oriented) edges E . Moreover, the operator $\delta_1 : C_1 \rightarrow C_0$ maps an edge $e = (v_0, v_1) \in E$ to $\delta_1(e) = [v_1] - [v_0]$ and, therefore, if $f \in C_1$, then

$$(\delta_1 f)(v_i) = \sum_{e=[v_j, v_i] \in E} f(e) - \sum_{e=[v_i, v_j] \in E} f(e) \quad (3)$$

This boundary operator δ_1 can be used to construct the *Graph Laplacian* as $\Delta_0 := \delta_1 \delta_1^T$, which is typically used to perform message passing in GCNs. The boundary operators can be used to generalize this construction to a *Hodge Laplacian* over a simplicial complex, defined as $\Delta_k := \partial_k^T \partial_k + \partial_{k+1} \partial_{k+1}^T : C_k \rightarrow C_k$, which can be used to construct a variety of higher-order simplicial neural networks [Papillon et al. \(2023\)](#).

Cycles, Boundaries and Homology A k -chain is said to be a **boundary** if it is the boundary of a $k + 1$ -chain; the subspace of k -boundaries is indicated by $B_k := \text{im } \partial_{k+1}$. A k -chain is said to be a **cycle** if its boundary is zero; the subspace of k -cycles is indicated by $Z_k := \ker \partial_k$. A classical result in topology is that *a boundary of a space has no boundary*, i.e. $\partial_\bullet \circ \partial_\bullet = \partial_\bullet^2 = 0$. It follows that $B_k = \text{im } \partial_{k+1} \subset Z_k = \ker \partial_k$. The **k -th homology group** is defined as the *quotient vector space* $H_k(S) := Z_k / B_k$. The dimensionality $\dim H_k$ is an important invariant and is equal to the k -th *Betti number* β_k of S , which counts the k -dimensional holes in S .

Topology, open sets and subcomplexes of a simplicial complex Given a finite simplicial complex S , a subset $A \subseteq S$ is said to be *closed* if it is also a simplicial complex (i.e. for each simplex in A , all its faces are also in A), i.e. it is a *subcomplex* of S . Instead, an *open subset*⁶ $A \subseteq S$ is the union of sets of the form $\text{star } \sigma = \{\tau \in S : \sigma \subset \tau\}$; note that this is not necessarily a simplicial complex. Finally, we define a few useful operations on a subset $A \subseteq S$:

- $S \setminus A$ indicates the standard set difference.
- the *closure* $\text{cl } A$ is the smallest subcomplex of S containing A .
- the *star* $\text{star } A$ is the set of all simplices in S which contain a simplex in A

5. This definition should be extended linearly to the full space C_k .

6. Formally, we consider the *Alexandrov topology* of the simplicial complex like in [Robinson et al. \(2018\)](#)

- the *boundary* (or frontier) $\partial A = \text{cl } A \cap \text{cl}(S \setminus A)$
- the *interior* $\text{int } A$ is the largest open set contained in A

Relative Homology Let $A \subseteq S$ be a *subcomplex* (i.e. a closed subset) of the simplicial complex S . The *relative k -chain space* $C_k(S, A) \cong C_k(S)/C_k(A)$ is the vector space of k -chains over S which are zeros over the simplices in A . Clearly, $C_k(S, A)$ is a subspace of $C_k(S)$ so the map $\partial_k : C_k \rightarrow C_{k-1}$ can be generalized to $\partial_k : C_k(S, A) \rightarrow C_{k-1}(S, A)$. Then, the sub-space of relative k -boundaries is indicated by $B_k(S, A) := \text{im } \partial_{k+1}$ and the subspace of relative k -cycles is indicated by $Z_k(S, A) := \ker \partial_k$. Finally, the k -th **relative homology** is defined as $H_k(S, A) = Z_k(S, A)/B_k(S, A)$. Intuitively, $H_k(S, A)$ describes the k -th homology of the **quotient space** S/A obtained from S by identifying all its points within A , i.e. by "collapsing" all points in A in a single point⁷.

C.1. Properties of Homology and Long Exact Sequences

Long Exact Sequence for the Relative Homology If $A \subset S$ is a subcomplex of S , the relative chains give rise to a chain complex of relative homology groups with the following *short exact sequence*:

$$\dots \rightarrow H_k(A) \xrightarrow{i_k} H_k(S) \xrightarrow{j_k} H_k(S, A) \xrightarrow{\partial} H_{k-1}(A) \rightarrow \dots \quad (4)$$

The map i_k comes from the inclusion of $C_k(A)$ into $C_k(S)$ and, intuitively, is relating the k -dimensional holes in A with their copy in S . The map j_k comes from the projection of $C_k(S)$ into $C_k(S, A)$ and, intuitively, relates the holes in S outside of A with their copies in S/A . Finally, the last map ∂ detects the k -dimensional holes in S/A , not present in S , which have appeared by collapsing A in a single point. These k -dimensional holes can be related with A 's $k - 1$ -dimensional boundary $\partial A \subset A$ and, therefore, included in $H_{k-1}(A)$.

Mayer-Vietoris Sequence Given two subcomplexes A, B and the union $S = \text{int } A \cup \text{int } B$, there is another important *long exact sequence*:

$$\dots \rightarrow H_{k+1}(S) \xrightarrow{\partial^*} H_k(A \cap B) \xrightarrow{i_*, j_*} H_k(A) \oplus H_k(B) \xrightarrow{k_* - l_*} H_k(S) \rightarrow \dots \quad (5)$$

Intuitively, if a $k + 1$ cycle in S is "broken" when S is split into A and B , the cycle splits into two $k + 1$ chains in A and B which overlap in $A \cap B$. The boundaries of the two $k + 1$ chains are homologous, i.e. they are a k -cycle in $A \cap B$, that is an element of $H_k(A \cap B)$.

This sequence holds also for relative homology, i.e. if $T = \text{int } C \cup \text{int } D \subset S$, with $C, D \subset S$, then

$$\dots \rightarrow H_{k+1}(S, T) \xrightarrow{\partial^*} H_k(S, C \cap D) \xrightarrow{i_*, j_*} H_k(S, C) \oplus H_k(S, D) \xrightarrow{k_* - l_*} H_k(S, T) \rightarrow \dots \quad (6)$$

Eq. 6 can also be used to construct the sequence in Eq. 1 by replacing $C = S \setminus A, D = S \setminus B$ and, therefore, $T = C \cup D = S \setminus (A \cap B), C \cap D = S \setminus (A \cup B)$:

$$\dots \rightarrow H_k(S, S \setminus (A \cup B)) \xrightarrow{i_*, j_*} H_k(S, S \setminus A) \oplus H_k(S, S \setminus B) \xrightarrow{k_* - l_*} H_k(S, S \setminus (A \cap B)) \rightarrow \dots \quad (7)$$

7. Note the difference between the set difference $S \setminus A$ and the quotient space S/A .

Note that the maps in this sequence are given by the restriction maps of the sheaf and the exactness of the sequence proves exactly the gluing property of a sheaf. See Proposition 19 [Robinson et al. \(2018\)](#) for a more precise proof.

Appendix D. Persistent Homology

Given a finite simplicial complex S and a function $f : S \rightarrow \mathbb{R}$ s.t. $f(\sigma) \leq f(\tau)$ if $\sigma < \tau$, define the simplicial complex $S_t = \{\sigma \in S : f(\sigma) \leq t\} \subset S$. Note that $S_{t_1} \subset S_{t_2}$ if $t_1 \leq t_2$ and there exists t^-, t^+ such that $S_t = \emptyset$ for any $t \leq t^-$ and $S_t = S$ for any $t \geq t^+$. Moreover, the sequence of simplicial complexes $\{S_t\}_{t \in \mathbb{R}}$ only contains a finite number of different complexes, so it can be replaced by a finite sequence $\{S_t\}_{t \in R}$ indexed by a subset $R \subset \mathbb{R}$. This sequence is called a **filtration** of simplicial complexes.

The inclusion $S_{t_1} \subset S_{t_2}$ induces an homomorphism $i_k^{t_1, t_2} : H_k(S_{t_1}) \rightarrow H_k(S_{t_2})$, whose image $\text{im } i_k^{t_1, t_2}$ is the **persistent homology group** $H_k^{t_1, t_2}(S)$ and detects k -cycles in S_{t_1} which are still present in S_{t_2} . In particular, any k -cycles is born at a certain "time" t_1 (is not in the image of i_k^{t, t_1} for any $t < t_1$). It can also disappear at a time t_2 (it is in the kernel of $i_k^{t_1, t}$ for any $t \geq t_2$) or persist forever (it is a cycle in $H_k(S)$).

Note also that, if the complexes in the filtration only differ by a single simplex (i.e. the function f gives a total ordering of the simplexes), each time step a single k -simplex is added, which either creates a new k -cycle or destroys a $k - 1$ cycle. This is useful since the homology group $H_k(S)$ does not come with a natural choice of basis⁸; in this case, instead, cycles are uniquely identified by their birth and death times, which indirectly provides a choice of basis.

Relative persistent homology has also been studied in the literature, e.g. see [Robinson et al. \(2018\)](#); [Blaser and Brun \(2022\)](#). However, as far as we know, these works considered a slightly different formulation, assuming a filtration of pairs (S, A_t) , with $A_t \subset A_{t+1} \subset S$.

Instead, in this work, we consider a filtration of pairs in the following form. Let $S^\infty = S$ and $A^\infty = A \subset S$. Let $\mathbb{S} = (\dots, S_t, \dots, S^\infty = S)$ be a filtration of S and $\mathbb{A} = (\dots, A_t, \dots, A^\infty = A)$ be a filtration of A , with $A_t = S_t \cap A$ (and, clearly, $S_t \subset S_{t+1}$ and $A_t \subset A_{t+1}$). To simplify the notation, sometimes we just write S instead of \mathbb{S} to indicate a filtration.

As earlier, the inclusion $S_{t_1} \subset S_{t_2}$ induces an homomorphism $i_k^{t_1, t_2} : H_k(S_{t_1}, A_{t_1}) \rightarrow H_k(S_{t_2}, A_{t_2})$, whose image $\text{im } i_k^{t_1, t_2}$ is the **persistent relative homology** $H_k^{t_1, t_2}(S, A)$ and detects relative k -cycles in S_{t_1}/A_{t_1} which are still present in S_{t_2}/A_{t_2} . Given an open set $U \subset S_\infty$, define the **persistence module**

$$\mathcal{H}_k^\bullet(U) = \bigoplus_t H_k(S_t, S_t \setminus U) \quad (8)$$

Then, our persistent homology feature spaces can be formally defined as the quotient

$$\mathcal{H}_k(U) = \left(\bigoplus_t H_k(S_t, S_t \setminus U) \right) / \left(\bigoplus_{t_1 < t_2} \text{im } i_k^{t_1, t_2} \right) = \mathcal{H}_k^\bullet(U) / \left(\bigoplus_{t_1 < t_2} \text{im } i_k^{t_1, t_2} \right) \quad (9)$$

8. A basis for $H_k(S)$ can be computed as the 0-eigenvectors of the k -th Hodge Laplacian $\Delta_k = \partial_k^T \partial_k + \partial_{k+1} \partial_{k+1}^T$. However, this basis is not unique and numerical algorithms are not guaranteed to return the same solution consistently. The lack of a choice of basis is problematic to construct learnable neural operations like linear layers and non-linearities, which depend on a specific basis.

The quotient removes the copies of a persistent cycle through its life interval. Hence, the resulting space has a dimension for each unique persistent cycle.

[Sovdat \(2016\)](#) studied a similar sequence where $A_t \subset S_t$ but not necessarily $A_t = S_t \cap A_\infty$ (i.e. a simplex can enter S at a time step but also enter in A at a later time step) and proposed an algorithm to compute this relative persistent (co)homology.

Apx. [E](#) describes how persistent relative homology can be computed while Apx. [F](#) describes a method to construct the corresponding sheaf Laplacian.

Appendix E. Computing Relative Homology and Relative Persistent Homology

Computing Persistent Homology The `Ripser` library implements an efficient algorithm to compute persistent homology [Bauer \(2021\)](#); [Tralie et al. \(2018\)](#). This algorithm can be easily adapted to also return the indices of the simplices which created and destroyed each homology class / persistent cycle; indeed, these indices are needed to implement a differentiable version of persistent homology [Brüel-Gabrielsson et al. \(2019\)](#). Note that this software actually computes persistent co-homology and also returns representative cochains, which can be thought simply as the transpose of representative chains. In the rest of this section, we will work with **co-homology** groups $H^k(\cdot)$ rather than homology groups $H_k(\cdot)$ to better reflect the algorithm but we first emphasize that these groups are isomorphic.

Unfortunately, `Ripser` only compute absolute (co)homology. [Sovdat \(2016\)](#) previously described a very similar algorithm to compute the persistent *relative* homology of a sequence of pairs $\{(S_t, A_t)\}_t$. As discussed in Apx. [D](#), they consider more general filtrations than ours and, therefore, their algorithm is unnecessarily complicated for us.

Instead, we note that the `Ripser` algorithm from [Bauer \(2021\)](#) essentially performs an (optimized) *Gauss reduction* of the co-boundary matrix $\partial_S^\bullet : C^\bullet(S) \rightarrow C^\bullet(S)$, with rows and columns (corresponding to different simplices in the filtration) sorted by decreasing weight / birth time. This algorithm can be used to compute the relative (co)homology $H^\bullet(S, A)$ by simply removing those rows and columns of ∂_S^\bullet which belongs to A ; indeed, by definition one obtains precisely the relative co-boundary map $\partial_{S,A}^\bullet : C^\bullet(S, A) \rightarrow C^\bullet(S, A)$ which defines relative (co)homology.

Moreover, as most existing persistent homology tools, `Ripser` only supports finite fields $\mathbb{F} = \mathbb{Z}/p\mathbb{Z}$ (for p prime), while our sheaf requires features in the real field $\mathbb{F} = \mathbb{R}$. Fortunately, the algorithm described in [Bauer \(2021\)](#) works for any generic field \mathbb{F} , so `Ripser` can be easily adapted to compute (co)homology with $\mathbb{F} = \mathbb{R}$ coefficients.

Appendix F. Computing the sheaf Laplacian

Let $A', B' \subset S$ be two open sets and $C' = A' \cap B' \subset S$ their intersection. To construct the sheaf Laplacian between these two open sets $\Delta_{B',A'}^k = -[\mathcal{H}^k(C' \subset A')]^T \circ \mathcal{H}^k(C' \subset B')$ we need to construct the two restriction maps $\mathcal{H}^k(C' \subset A'), \mathcal{H}^k(C' \subset B')$ and then find equivalent cocycles in their images.

The following Mayer-Vietoris sequence for relative cohomology suggests a way to perform this computation. Let $D' = A' \cup B'$ the union of the two open sets and $D = S \setminus D'$ its

complementary; then the following sequence is exact:

$$\dots \rightarrow H^{k-1}(S, D) \xrightarrow{\partial^{k-1}} H^k(S, C) \xrightarrow{i_* \oplus -j_*} H^k(S, A) \oplus H^k(S, B) \xrightarrow{k_* + l_*} H^k(S, D) \rightarrow \dots \quad (10)$$

where the maps i_* and j_* are adjoint of the restriction maps $\mathcal{H}^k(C' \subset A')$, $\mathcal{H}^k(C' \subset B')$. The co-boundary map ∂^{k-1} detects the k relative cycles in C' , not present in neither A' nor B' , which have appeared when collapsing $C \setminus D$ in a single point (e.g. a line with its extremes in $D' \setminus C'$ is a connected component, i.e. a 0-cycle, in $H^0(S, D)$, but when $D' \setminus C'$ is collapsed, the two extremes merge and the 0-cycle becomes a 1-cycle in $H^1(S, C)$).

This sequence implies that $H^k(S, C)$ splits as the co-image $\text{coim}(i_* \oplus -j_*)$ (i.e. the image of the restriction maps) and the image $\text{im } \partial^{k-1}$. In other words, the restrictions of two cocycles in $H^k(S, A)$ and $H^k(S, B)$ are equivalent if their difference is zero modulo $\text{im } \partial^{k-1}$.

Hence, we set up an *extended coboundary matrix* \mathcal{B}^k whose reduction computes the sheaf Laplacian.

Columns The matrix columns are divided in two sets. First, it contains all columns of $\partial^\bullet(S, D)$ as used in Apx. E to compute the persistent relative cohomology $H^k(S, D)$. Second, it contains a column for each persistent cocycle found previously in $H^k(S, A)$ and $H^k(S, B)$. Like in Apx. E, the columns in the first set are sorted inversely by the weight of each $k-1$ simplex in D' . Instead, the columns in the second set are sorted inversely by their corresponding cocycle's birth time (cocycles of A' and B' are mixed by sorting). These two sets split the matrix in two sub-matrices $\mathcal{B}^k = [\mathcal{B}_D^k, \mathcal{B}_{AB}^k]$.

Rows in \mathcal{B}_D^k Columns in \mathcal{B}_D^k simply contain the coboundaries in D' of each simplex, sorted by decreasing weight, as in Apx. E.

Before defining the rows in \mathcal{B}_{AB}^k , let's first recall some details about the algorithm in Bauer (2021). A cocycle in $H^k(S, A)$ (or B) can be represented by the column of the reduction matrix used to reduce ∂_A^k . This vector expresses a k -cocycle as a linear combination of k -simplices in A' . The non-zero simplex with lowest weight defines the birth time of the cocycle. The corresponding reduced column contains the coboundary and the first non-zero $k+1$ simplex (the *pivot*) defines the death time of the cocycle (since, after that time, the cocycle doesn't belong to the kernel of the coboundary map anymore).

Rows in \mathcal{B}_{AB}^k The columns in \mathcal{B}_{AB}^k contain three sets of row. Each column, corresponding to a certain cocycle to restrict, has

1. one row for each k -simplex in D' : these rows contain a copy of the reduction vector representing the cocycle as above (note that $A', B' \subset D'$). These are also the same rows in \mathcal{B}_D^k
2. one row for each $k+1$ simplex in A' : these rows contain a copy of the coboundary of the cocycles in A'
3. another row for each $k+1$ simplex in B' : these rows contain a copy of the coboundary of the cocycles in B'

Note that each $k+1$ simplex in D' appears twice in the rows.

A linear combination of the columns of this extended reduction matrix is a linear combination of cocycles in $H^k(S, A)$, $H^k(S, B)$ and $H^{k-1}(S, D)$. This represents a pair of

cocycles $\gamma_A \in H^k(S, A)$ and $\gamma_B \in H^k(S, B)$ and the rows in the resulting column model the three constraints we are trying to enforce. Indeed, a non-zero value in a row implies

- if the row is a $k+1$ -simplex in A' (or B'), the cocycle $\gamma_A \in H^k(S, A)$ (or $\gamma_B \in H^k(S, B)$) is dead at this time step (and so must be also its restriction to $H^k(S, C)$ as proved in Theorem 1).
- if the row is a k -simplex in $C' \subset D'$, it means that the sum of γ_A and γ_B is not zero at this time step i.e. their restrictions are not equivalent cocycles.
- if the row is a k -simplex in $D' \setminus C' = (A' \setminus C') \cup (B' \setminus C')$, either γ_A or γ_B can not be restricted to $H^k(S, C)$ at this time step.

Then, the matrix reduction algorithm trying to find pairs of cocycles which satisfy these constraints for the longest time. Once this matrix is reduced, a column in \mathcal{B}_{AB}^k represents a pair of cocycles $\gamma_A \in H^k(S, A)$ and $\gamma_B \in H^k(S, B)$ whose sum is 0 when restricted to $H^k(S, C)$, modulo the coboundary of some cocycles in $H^{k-1}(S, D)$, until the time step the *pivot* of this column appears in the filtration. Then, the pivot corresponds to the time step one of the three constraints above is violated.

Hence, the reduced columns in \mathcal{B}_{AB}^k can be used to construct the sheaf Laplacian as follows. Let the i -th reduced column correspond to a pair of cocycles (γ_A^i, γ_B^i) which are obtained by linearly combining the persistent bases of $H^k(S, A)$ and $H^k(S, B)$ via the reduction vectors \mathbf{v}_A^i and \mathbf{v}_B^i , respectively. Note that these reduction vectors essentially construct the two restriction maps. Let t_i be the time the pivot of this column appear and let s_A^i be the birth time of γ_A^i (i.e. the lowest weight of its simplexes) and t_A^i its death time, and s_B^i and t_B^i those of γ_B^i . This pair restricts to the same cocycle in $H^k(S, C)$ only in the time interval $[s^i, t^i)$, with $s_i = \max(s_A^i, s_B^i)$ and $t_i \leq s_B^i, s_A^i$ due to Theorem 1. The pair (s^i, t^i) defines the time interval during which an i -th sheaf Laplacian persists:

$$[\Delta_{\mathcal{H}^k}^i]_{A', B'} = \mathbf{v}_A^i (\mathbf{v}_B^i)^T$$

We do not include the -1 sign since our constraint enforced $\gamma_A + \gamma_B \cong 0$, i.e. $\gamma_A \cong -\gamma_B$. This Laplacian is visualized also in Fig. 2.

Note that the non-zero coefficients in the vector \mathbf{v}_A^i or \mathbf{v}_B^i are associated with persistent cocycles of $H^k(S, A)$ or $H^k(S, B)$ which might appear and die at different time steps. It follows that each entry of $[\Delta_{\mathcal{H}^k}^i]_{A', B'}$ has an independent persistence interval given by the intersection of $[s^i, t^i)$ with the intervals of the two cocycles of A' and B' involved.

If we define $\mathbf{v}|_t$ as the components of \mathbf{v} which are "active" at time t , the sheaf Laplacian at a time step t can be constructed as

$$[\Delta_{\mathcal{H}^k}^t]_{A', B'} = \sum_{i: t \in [s^i, t^i)} \mathbf{v}_A^i |_t (\mathbf{v}_B^i |_t)^T$$

Finally, the embedding and projection operations mention in Sec. 3 can be easily implemented by weighting the entry (a, b) of the matrix $[\Delta_{\mathcal{H}^k}^i]_{A', B'}$ by its own life span $\min(t_i, t_a, t_b) - \max(s_i, s_a, s_b)$ divided by the output cocycle life span $t_a - s_a$.

F.1. Other properties of the Local (Co)Homology Sheaf

Theorem 1 (The restriction of a cocycle dies earlier) Consider the following commutative diagram for relative persistent cohomology and assume a single simplex is added to \mathbb{S} at each time step t :

$$\begin{array}{ccccccccccc}
 \cdots & \longrightarrow & H^{k-1}(S_t) & \xrightarrow{j_t^*} & H^{k-1}(A_t) & \xrightarrow{\partial_t^{k-1}} & H^k(S_t, A_t) & \xrightarrow{i_t^*} & H^k(S_t) & \longrightarrow & \cdots \\
 & & \downarrow f_S^{t,t+1} & & \downarrow f_A^{t,t+1} & & \downarrow f_{S,A}^{t,t+1} & & \downarrow f_S^{t,t+1} & & \\
 \cdots & \longrightarrow & H^{k-1}(S_{t+1}) & \xrightarrow{j_{t+1}^*} & H^{k-1}(A_{t+1}) & \xrightarrow{\partial_{t+1}^{k-1}} & H^k(S_{t+1}, A_{t+1}) & \xrightarrow{i_{t+1}^*} & H^k(S_{t+1}) & \longrightarrow & \cdots
 \end{array} \tag{11}$$

Let $\gamma \in \text{im } i_t^* \subset H^k(S_t)$ be a cocycle of S_t at time t which corresponds to a relative cocycle $\bar{\gamma} \in H^k(S_t, A_t)$, i.e. $\gamma = i_t^*(\bar{\gamma})$. Assume that at time $t+1$ a $k+1$ -simplex σ is added to S_t such that the cocycle γ dies in $H^k(S_{t+1})$, i.e. $\gamma \notin \text{im } f_S^{t,t+1}$. Then, $\bar{\gamma} \notin \text{im } f_{S,A}^{t,t+1}$ either and, therefore, the relative cocycle $\bar{\gamma}$ dies at time $t+1$, too.

Proof Let $\gamma \in \ker f_S^{t,t+1}$ and let σ be the $k+1$ simplex added in S_{t+1} which killed γ (i.e. $S_{t+1} = S_t \cup \{\sigma\}$). Assume $\exists \bar{\gamma} \in H^k(S_t, A_t)$ such that $\gamma = i_t^*(\bar{\gamma})$.

Since σ is a $k+1$ -simplex, $H^{k-1}(A_{t+1}) \cong H^{k-1}(A_t)$ and $H^{k-1}(S_{t+1}) \cong H^{k-1}(S_t)$. Because these cohomology groups did not change, $\text{im } j_t^* \cong \text{im } j_{t+1}^*$ and, therefore, $\ker \partial_t^{k-1} \cong \ker \partial_{t+1}^{k-1}$. It also follows that $\text{coim } \partial_t^{k-1} \cong \text{coim } \partial_{t+1}^{k-1}$, i.e. $\ker i_t^* = \ker i_{t+1}^*$.

Finally, because $\bar{\gamma} \notin \ker i_t^* \cong \ker i_{t+1}^*$, $\bar{\gamma} \in H^{k+1}(S_{t+1}, A_{t+1}) \cong \ker i_{t+1}^* \oplus \text{coim } i_{t+1}^*$ if and only if $\bar{\gamma} \in \text{coim } i_{t+1}^*$. This requires that $\exists \gamma' = i_{t+1}^*(\bar{\gamma}) \in H^k(S_{t+1})$. However, the commutativity of the diagram guarantees that $i_{t+1}^*(f_{S,A}^{t,t+1}(\bar{\gamma})) = f_S^{t,t+1}(i_{t+1}^*(\bar{\gamma})) = 0$, which is a contradiction. Hence, $\bar{\gamma} \in \ker f_{S,A}^{t,t+1}$, i.e. the relative cocycle $\bar{\gamma}$ must also die at time $t+1$. \blacksquare

A similar argument should work also for triples, i.e. projections $H^k(S, B) \rightarrow H^k(S, A)$ with $B \subset A \subset S$ by replacing $H^k(S)$ with $H^k(S, B)$ and $H^k(A)$ with $H^k(A, B)$.

Theorem 2 (The restriction of a cocycle appears earlier) Consider again the commutative diagram for relative persistent cohomology in Eq. 11 (here, shifted right by two steps):

$$\begin{array}{ccccccccccc}
 \cdots & \longrightarrow & H^k(S_t, A_t) & \xrightarrow{i_t^*} & H^k(S_t) & \xrightarrow{j_t^*} & H^k(A_t) & \xrightarrow{\partial_t^k} & H^{k+1}(S_t, A_t) & \longrightarrow & \cdots \\
 & & \downarrow f_{S,A}^{t,t+1} & & \downarrow f_S^{t,t+1} & & \downarrow f_A^{t,t+1} & & \downarrow f_{S,A}^{t,t+1} & & \\
 \cdots & \longrightarrow & H^k(S_{t+1}, A_{t+1}) & \xrightarrow{i_{t+1}^*} & H^k(S_{t+1}) & \xrightarrow{j_{t+1}^*} & H^k(A_{t+1}) & \xrightarrow{\partial_{t+1}^k} & H^{k+1}(S_{t+1}, A_{t+1}) & \longrightarrow & \cdots
 \end{array} \tag{12}$$

Again, assume a single simplex is added to \mathbb{S} at each time step t . Let $\gamma \in H^k(S_t)$ be a cocycle of S_t which persists to S_{t+1} , i.e. $\gamma \in \text{im } f_S^{t,t+1}$.

Assume that there exists a relative cocycle $\bar{\gamma} \in H^k(S_{t+1}, A_{t+1})$ such that $\gamma = i_{t+1}^*(\bar{\gamma})$.

Then, $\bar{\gamma} \in \text{im } f_{S,A}^{t,t+1}$, too. This implies that the projection $\bar{\gamma}$ must always appear in the filtration at the same time or earlier than the corresponding cocycle $\gamma = i^*(\bar{\gamma})$.

Proof

Assume $\gamma \notin \text{im } i_t^*$. Then, there exists a new relative cocycle $\bar{\gamma}'$ in $H^k(S_{t+1}, A_{t+1})$ appearing at time $t + 1$, with $\gamma = i_{t+1}^*(\bar{\gamma}')$. Let σ be the k -simplex added to $S_t \setminus A_t$ which gave birth to it (i.e. $S_{t+1} = S_t \cup \{\sigma\}$ and $A_t = A_{t+1}$). Since $\sigma \notin A_{t+1}$, $H^*(A_{t+1}) \cong H^*(A_t)$. Moreover, since σ is a k -simplex, $H^{k+1}(S_t, A_t) \cong H^{k+1}(S_{t+1}, A_{t+1})$, too.

It follows that $\ker \partial_t^k \cong \ker \partial_{t+1}^k$ and, therefore, $\text{coim } j_{t+1}^* \cong \text{coim } j_t^*$. Since $\gamma \in \text{im } i_{t+1}^*$, $\gamma \notin \text{coim } j_{t+1}^* \cong \text{coim } j_t^*$. Hence, $\gamma \in \ker j_t^* \cong \text{im } i_t^*$.

This a contradiction, so it must be the case that $\gamma \in \text{im } i_t^*$, too.

Now, let $\bar{\gamma} \in H^k(S_t, A_t)$ s.t. $\gamma = i_t^*(\bar{\gamma})$. Then, by commutativity of the diagram, $\gamma = f_S^{t,t+1}(i_t^*(\bar{\gamma})) = i_{t+1}^*(f_{S,A}^{t,t+1}(\bar{\gamma}))$, which implies $\bar{\gamma} \in \text{coim } f_{S,A}^{t,t+1}$. In other words, $\bar{\gamma}$ is also a persistent cocycle in $H^k(S, A)$. ■

As earlier, a similar argument should work also for triples $B \subset A \subset S$.