



Delft University of Technology

## Hodge-Aware Contrastive Learning

Möllers, Alexander; Immer, Alexander; Fortuin, Vincent; Isufi, Elvin

### DOI

[10.1109/ICASSP48485.2024.10447015](https://doi.org/10.1109/ICASSP48485.2024.10447015)

### Publication date

2024

### Document Version

Final published version

### Published in

Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2024

### Citation (APA)

Möllers, A., Immer, A., Fortuin, V., & Isufi, E. (2024). Hodge-Aware Contrastive Learning. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2024* (pp. 9746-9750). (ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings). IEEE. <https://doi.org/10.1109/ICASSP48485.2024.10447015>

### Important note

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

### Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

### Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***

***'You share, we take care!' - Taverne project***

**<https://www.openaccess.nl/en/you-share-we-take-care>**

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

# HODGE-AWARE CONTRASTIVE LEARNING

Alexander Möllers\*, Alexander Immer†, Vincent Fortuin‡§, Elvin Isufi\*

## ABSTRACT

Simplicial complexes prove effective in modeling data with multi-way dependencies, such as data defined along the edges of networks or within other higher-order structures. Their spectrum can be decomposed into three interpretable subspaces via the Hodge decomposition, resulting foundational in numerous applications. We leverage this decomposition to develop a contrastive self-supervised learning approach for processing simplicial data and generating embeddings that encapsulate specific spectral information. Specifically, we encode the pertinent data invariances through simplicial neural networks and devise augmentations that yield positive contrastive examples with suitable spectral properties for downstream tasks. Additionally, we reweight the significance of negative examples in the contrastive loss, considering the similarity of their Hodge components to the anchor. By encouraging a stronger separation among less similar instances, we obtain an embedding space that reflects the spectral properties of the data. The numerical results on two standard edge flow classification tasks show a superior performance even when compared to supervised learning techniques. Our findings underscore the importance of adopting a spectral perspective for contrastive learning with higher-order data.

**Index Terms**— Hodge Laplacian, simplicial filter, contrastive learning.

## 1. INTRODUCTION

The difficulty to represent many biological, social, and technological networks arises from the complexity of their inherent relational structures and the nuanced definitions of the associated data. Importantly, the data in such networks is often defined on higher-order components like edges and triangles, thus demanding an approach that incorporates interactions beyond the pairwise paradigm [1]. Simplicial complexes (SCs) have been proposed to model such dependencies in higher-order networks [2, 3]. Amongst others, SCs have shown particular efficacy in edge flow applications (e.g., mass, energy, information, or trajectories) for which traditional graph-based techniques do not induce a good inductive bias [4, 5]. Notably, they have been leveraged to alleviate the curse of dimensionality in autoregressive flow prediction for water networks [6] and to remove arbitrage opportunities in currency exchange markets [7].

An important property of SCs is that they enjoy an algebraic representation via the Hodge Laplacian matrices, ultimately allowing for spectral analysis. The latter is achieved through the Hodge decomposition of the spectrum and has been used to develop Hodge-aware signal processing and learning techniques for edge flows and other simplicial data [5, 8, 9]. Recent advances include developing a simplicial Fourier transform [4], convolutional filters [7], and neural networks that are trained in a (semi-)supervised manner. [10–13].

\* TU Delft, Delft, The Netherlands. † ETH Zürich, Zürich, Switzerland. ‡ TU Munich, Munich, Germany. § Helmholtz AI, Munich, Germany. EI is supported by the TU Delft AI Labs Programme. AI is supported by a Max Planck ETH Center for Learning Systems doctoral fellowship. VF was supported by a Branco Weiss Fellowship. AM conducted the work while at ETH Zürich. Code: [https://github.com/injurise/hodge\\_simpcl](https://github.com/injurise/hodge_simpcl). Correspondence to alexander.j.moellers@gmail.com.

While these supervised learning (SL) methods for simplicial data have their merits, they also come with limitations such as an over-reliance on labeled examples and suboptimal performance in data-scarce scenarios. SL methods are also mainly designed for a specific task and the resulting output is generally not reusable in different applications. Therefore, we propose a contrastive learning (CL) approach for simplicial data that addresses these issues. To this end, we employ a simplicial convolutional neural network (SCNN) to produce embeddings and optimise it in a self-supervised manner using the InfoNCE loss [14]. The resulting CL approach can use both labeled and unlabeled examples to create robust and reusable representations for simplicial data. These representations can subsequently be employed in various downstream tasks, offering improved accuracy especially in label-scarce scenarios.

To further enhance our method, we propose stochastic augmentations and introduce information about the Hodge Decomposition into the embeddings by i) optimizing the parameters of the augmentations so that they generate positive examples that respect the spectral properties of the data; and ii) weighting the negative samples in the InfoNCE loss by a Hodge-aware distance to the anchor (true datum). This approach results in a spectrally organized embedding space and facilitates downstream learning. We corroborate the latter in two edge flow classification settings and show superior performance compared with a fully-supervised model. Related to this, our contribution is threefold:

- C1) we propose simplicial contrastive learning (SCL), design related augmentations and experimentally validate all our approaches;
- C2) we show how augmentations in the simplicial domain can be optimized with respect to the Hodge decomposition;
- C3) we introduce a reweighting of the negative examples based on the similarity of their Hodge components to encourage a spectrally organized embedding space.

## 2. DATA PROCESSING ON SIMPLICIAL COMPLEXES

In this section, we introduce the fundamental concepts behind simplicial data structures and related data processing techniques.

### 2.1. Simplicial Data and Structures

Let  $\mathcal{V} = \{1, \dots, N\}$  be a set of vertices. A  $k$ -simplex  $\mathcal{S}^k$  is a subset of  $\mathcal{V}$  that contains  $k + 1$  distinct elements. A simplicial complex  $\mathcal{X}^K$  of order  $K$  is a collection of simplices such that it contains at least one  $K$ -simplex and if  $\mathcal{S}^k \in \mathcal{X}^K$  we have that all subsets of  $\mathcal{S}^k$  are also elements of  $\mathcal{X}^K$ . From a geometric representation perspective, in a SC of order  $K = 2$ , nodes are 0-simplices, edges are 1-simplices and (filled) triangles are 2-simplices [2, 3].

Neighborhood relations in an SC can be represented via the incidence matrices and Hodge Laplacians. Specifically, the incidence matrix  $\mathbf{B}_k \in \mathbb{R}^{N_{k-1} \times N_k}$  captures the adjacencies between  $(k-1)$ - and  $k$ -simplices. Consequently,  $\mathbf{B}_1$  represents the node-to-edge incidence matrix and  $\mathbf{B}_2$  the edge-to-triangle incidence matrix. For an SC of order two,  $\mathcal{X}^2$ , the related Hodge Laplacians are defined as:

$$\begin{aligned} \mathbf{L}_0 &= \mathbf{B}_1 \mathbf{B}_1^\top, \\ \mathbf{L}_1 &= \mathbf{L}_{1,\ell} + \mathbf{L}_{1,u} := \mathbf{B}_1^\top \mathbf{B}_1 + \mathbf{B}_2 \mathbf{B}_2^\top, \\ \mathbf{L}_2 &= \mathbf{B}_2^\top \mathbf{B}_2 \end{aligned} \quad (2)$$

where  $\mathbf{L}_0, \mathbf{L}_1, \mathbf{L}_2$  represent the neighborhood relationships between nodes, edges, and triangles, respectively. Moreover,  $\mathbf{L}_0$  coincides with the standard graph Laplacian [3]. The lower-Laplacian  $\mathbf{L}_{1,\ell} = \mathbf{B}_1^\top \mathbf{B}_1$  and upper-Laplacian  $\mathbf{L}_{1,u} = \mathbf{B}_2 \mathbf{B}_2^\top$  split the edge adjacencies into relations that are due to common vertices and common triangles, respectively.

A  $k$ -simplicial signal  $x^k$  is a mapping from a  $k$ -simplex to the set of real numbers that formalizes the simplicial data. We collect all  $k$ -signals in vector  $\mathbf{x}^k = [x_1^k, \dots, x_{N_k}^k]^\top$ , where  $x_i^k$  is the signal on the  $i$ th simplex and  $N_k$  is the total number of  $k$ -simplices. For instance, we denote an edge flow as  $\mathbf{x} = [x_1^1, \dots, x_{N_1}^1]^\top$  with  $x_e^1$  being the flow on the edge  $e = (m, n)$  in  $S^1$ . In the sequel, we focus on edge flows to ease exposition and because of their wider applicability; thus we drop the superscript and denote them as  $\mathbf{x}$ .

## 2.2. The Hodge Decomposition

SCs allow for a spectral processing of simplicial signals via the Hodge decomposition, which decomposes the space  $\mathbb{R}^{N_k}$  as:

$$\mathbb{R}^{N_k} = \text{im}(\mathbf{B}_k^\top) \oplus \text{im}(\mathbf{B}_{k+1}) \oplus \ker(\mathbf{L}_k) \quad (1)$$

where  $\text{im}(\cdot)$  and  $\ker(\cdot)$  are the image and kernels spaces of a matrix and  $\oplus$  is the direct sum of vector spaces [4, 7]. Accordingly, we can decompose any edge flow  $\mathbf{x}$ , into three parts  $\mathbf{x} = \mathbf{x}_G + \mathbf{x}_C + \mathbf{x}_H$  each living in an orthogonal subspace known as the gradient space  $\mathbf{x}_G \in \text{im}(\mathbf{B}_1^\top)$ , the curl space  $\mathbf{x}_C \in \text{im}(\mathbf{B}_2)$ , and the harmonic space  $\mathbf{x}_H \in \ker(\mathbf{L}_1)$ . In turn, the 1-Hodge Laplacian can be eigen-decomposed as  $\mathbf{L}_1 = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$  with eigenvectors  $\mathbf{U}$  and eigenvalues  $\mathbf{\Lambda}$ . By grouping the eigenvectors as  $\mathbf{U} = [\mathbf{U}_G, \mathbf{U}_C, \mathbf{U}_H]$  and projecting the edge flow onto them gives rise to the embeddings  $\tilde{\mathbf{x}}_G = \mathbf{U}_G^\top \mathbf{x}$ ,  $\tilde{\mathbf{x}}_C = \mathbf{U}_C^\top \mathbf{x}$ ,  $\tilde{\mathbf{x}}_H = \mathbf{U}_H^\top \mathbf{x}$ . These embeddings encode different interpretable properties of the data, which we will exploit in Sec. 4 to design augmentations. Refer to [4] for more detail on the role of the Hodge decomposition on processing simplicial data.

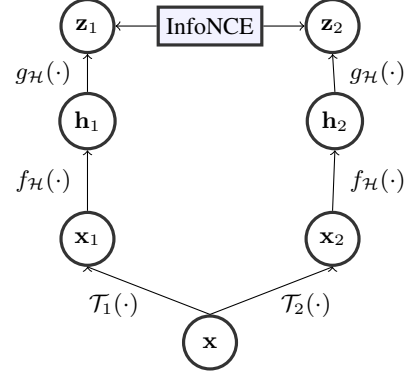
## 2.3. Simplicial Convolutional Neural Networks

Simplicial convolutional filters are linear parametric mappings that can process simplicial signals [7]. For an input edge flow  $\mathbf{x}$ , the filtered signal is  $\mathbf{y} = \mathbf{H}(\mathbf{L}_1) \mathbf{x}$  with simplicial filtering matrix:

$$\mathbf{H}(\mathbf{L}_1) := \left( \epsilon \mathbf{I} + \sum_{l_1=0}^{L_1} \alpha_{l_1} \mathbf{L}_{1,\ell}^{l_1} + \sum_{l_2=0}^{L_2} \beta_{l_2} \mathbf{L}_{1,u}^{l_2} \right). \quad (2)$$

Here,  $\{\epsilon, \alpha_{l_1}, \beta_{l_2}\}$  are filter coefficients and  $\mathbf{L}_{1,\ell}, \mathbf{L}_{1,u}$  propagate the edge signal via their respective neighbourhood relations. The filter is local in the sense that it moves information by at most  $L = \max\{L_1, L_2\}$  hops across the simplicial structure. Moreover, by exploiting the recursions  $\mathbf{L}_{1,\ell}^{l_1} \mathbf{x} = \mathbf{L}_{1,\ell}(\mathbf{L}_{1,\ell}^{l_1-1} \mathbf{x})$ ,  $\mathbf{L}_{1,u}^{l_2} \mathbf{x} = \mathbf{L}_{1,u}(\mathbf{L}_{1,u}^{l_2-1} \mathbf{x})$  and since the Laplacian matrices are sparse, we can obtain the output with a cost of order  $\mathcal{O}(LN_1)$ . The filter part related to the lower Laplacian acts on the signal gradient embedding, whereas that related to the upper Laplacian acts on the curl embedding. All the parameters contribute to processing the signal harmonic embedding. Refer to [7, Sec. IV] for the specific relation of filter (2) and decomposition (1).

The constant number of parameters and the linear computational complexity in the number of edges, make the filter (2) an appealing solution to learn representations from edge flows. To also learn non-linear mappings, simplicial convolutional neural networks (SCNNs)



**Fig. 1:** InfoNCE learning with augmentations. The data point (anchor)  $\mathbf{x}$  submits two transformations  $\mathcal{T}_{1/2}(\cdot)$  to generate positive augmented examples. The latter are first passed through the SCNN  $f_{\mathcal{H}}(\cdot)$  to generate the simplicial embeddings  $\mathbf{h}$  (cf. (3)) and then to a parametric map  $g_{\mathcal{H}}(\cdot)$  to obtain the final representation  $\mathbf{z}$ . These representations are contrasted in loss (4) to train both  $f_{\mathcal{H}}(\cdot)$  and  $g_{\mathcal{H}}(\cdot)$ .

have been developed as layered structures interleaving filters with pointwise non-linearities [13]. Specifically, given the SCNN input  $\mathbf{x}_0 := \mathbf{x}$ , the propagation rule at each layer  $t$  is:

$$\mathbf{x}_t = \sigma(\mathbf{H}_t(\mathbf{L}_1) \mathbf{x}_{t-1}) \quad (3)$$

where  $\sigma(\cdot)$  is a pointwise nonlinearity (e.g. rectified linear unit). The final layer constitutes the SCNN output which provides its embedding. Compactly, we will denote the SCNN input-output relation as  $\mathbf{h} := f_{\mathcal{H}}(\mathbf{x})$ , where  $\mathbf{h}$  is referred to as the SCNN embedding and set  $\mathcal{H} := \{\epsilon_t, \{\alpha_{\ell_1,t}\}, \{\beta_{\ell_2,t}\}\}_t$  collects the parameters of all layers. Furthermore, depending on the setting, a readout function is applied to  $\mathbf{h}$  to transform it for the task at hand (e.g., binary classification).

## 3. PROBLEM STATEMENT

Learning representations via the SCNN is typically done in a supervised manner, but we often have just a few labeled examples or we are missing labels at all. To tackle this challenge, we resort to contrastive learning and propose a self-supervised learning approach for simplicial complexes. Our problem statement reads as:

*Given a set of unlabeled edge flows, we want to train a simplicial convolutional neural network in a self-supervised manner to generate embeddings that reflect the Hodge-properties of the data and that can be used in a downstream task.*

We approach this problem by training the network with the contrastive InfoNCE loss and by designing augmentations that preserve the desired spectral properties. We also reweight the negative samples in the loss to push apart spectrally-different embeddings.

## 4. SIMPLICIAL CONTRASTIVE LEARNING

To train an SCNN in a self-supervised manner, we resort to the contrastive learning framework [15, 16]. In the simplicial setting, this principle suggests that for each edge flow datum  $\mathbf{x}$  we create both positive and negative examples and train the SCNN (a.k.a. the encoder) to map the positive embeddings close to each other and the negative ones further apart. Specifically, we consider an edge flow datum  $\mathbf{x}$  with its final representation  $\mathbf{z} = g_{\mathcal{H}}(\mathbf{h}) = g_{\mathcal{H}}(f_{\mathcal{H}}(\mathbf{x}))$ , where  $f_{\mathcal{H}}(\cdot)$  is the SCNN [cf. (3)] and  $g_{\mathcal{H}}(\cdot)$  is a parametric map (typically a fully connected layer). Then, we create a positive pair for  $\mathbf{x}$  via two topological augmentations  $\mathbf{x}'_i = \mathcal{T}_1(\mathbf{x})$  and  $\mathbf{x}'_j = \mathcal{T}_2(\mathbf{x})$  with respective representations  $\mathbf{z}'_i = g_{\mathcal{H}}(f_{\mathcal{H}}(\mathbf{x}'_i))$  and  $\mathbf{z}'_j = g_{\mathcal{H}}(f_{\mathcal{H}}(\mathbf{x}'_j))$ . The negative examples w.r.t. to  $\mathbf{x}$  consists of

other edge flows from the dataset  $\mathbf{x}_m \neq \mathbf{x}$  and their augmentations. With reference to Fig. 1, the overall network is trained to minimize the so-called temperature-scaled InfoNCE objective:

$$\mathcal{L}_{\text{InfoNCE}} = - \sum_{(\mathbf{x}'_i, \mathbf{x}'_j) \in \mathbb{P}} \log \left( \frac{e^{\text{sim}(\mathbf{z}'_i, \mathbf{z}'_j)/\tau}}{\sum_{m=1}^M e^{\text{sim}(\mathbf{z}'_i, \mathbf{z}_m)/\tau}} \right) \quad (4)$$

where  $\mathbb{P}$  is the set of all positive pairs in the data,  $\text{sim}(\mathbf{u}, \mathbf{v}) = \mathbf{u}^\top \mathbf{v} / \|\mathbf{u}\|_2 \|\mathbf{v}\|_2$  is the cosine similarity,  $\tau$  is a temperature parameter, and  $M$  is the number of negative examples  $\mathbf{x}_m$  with representations  $\mathbf{z}_m$ . The numerator encourages the network to map the positive embeddings close to each other, while the denominator repulses the negative embedding  $\mathbf{z}_m$  from the positive one  $\mathbf{z}'_i$ .

The InfoNCE optimizes a lower bound on the mutual information between the representations of the positive pairs [14]. Hence, augmentations should only preserve the information necessary to perform well on a downstream task [17]. Consequently, the irrelevant information is destroyed and the mutual information encoded in the representations is optimized for the task at hand. Common augmentations in contrastive learning are stochastic and include masking part of the data (pixels, vertices, edges, node features, etc.), which expresses the belief that the most important parts of the data are preserved when a few connections or feature values are removed [16, 18]. While this preserves, in probability, crucial information and works well empirically, the most relevant parts of a data point may not always remain intact. To mitigate this, the dropout probabilities are often chosen such that the most important properties of the data are conserved more often (e.g., in graphs nodes with a high centrality) instead of randomly removing information [19–22]. With this in place, we propose the following augmentation method.

**Edge flow masking.** This method masks edge flows with probabilities  $\mathbf{p}$  to generate a positive example. That is,  $\mathbf{x}' = \mathcal{T}(\mathbf{x}) := \mathbf{x} \circ \mathbf{e}$  where  $\mathbf{e}$  is a random Bernoulli vector with entry  $\mathbf{e}_i \sim \text{Ber}(\mathbf{p}_i)$  and  $\circ$  is the elementwise product. The standard approach is to pick the same masking probability for all edges.

Such an augmentation is more effective in settings with binary flow types  $\{-1, 1\}$  or when a zero value is not indicative. This is because the masked flow attains a zero value and the augmentation effectively drops parts of the flow. Next, we show how to optimize the masking probabilities w.r.t. flow Hodge-related information.

#### 4.1. Hodge-Aware Spectral Augmentations

Simplicial data often have particular properties in one of the three Hodge embeddings [3, 23], which may be wrongly affected by augmentations if ignored. Hence, following the InfoMin principle simplicial augmentations should destroy information on irrelevant Hodge embeddings and preserve it on the others. To account for the latter, we cast the dropout probabilities as a stochastic optimization problem, considering the expected value of the difference of the generated Hodge embeddings to the embeddings of the anchor. Specifically, consider the embeddings of an augmented example  $\tilde{\mathbf{x}}'_G = \mathbf{U}_G^\top \mathbf{x}'$ ,  $\tilde{\mathbf{x}}'_C = \mathbf{U}_C^\top \mathbf{x}'$ ,  $\tilde{\mathbf{x}}'_H = \mathbf{U}_H^\top \mathbf{x}'$ . Then the expressions  $\mathcal{L}_G(\mathbf{p}) = \mathbb{E}[\|\tilde{\mathbf{x}}_G - \tilde{\mathbf{x}}'_G\|_2^2]$ ,  $\mathcal{L}_C(\mathbf{p}) = \mathbb{E}[\|\tilde{\mathbf{x}}_C - \tilde{\mathbf{x}}'_C\|_2^2]$ ,  $\mathcal{L}_H(\mathbf{p}) = \mathbb{E}[\|\tilde{\mathbf{x}}_H - \tilde{\mathbf{x}}'_H\|_2^2]$  quantify the expected quadratic differences between the original and the augmented hodge embeddings.<sup>1</sup>

<sup>1</sup>To see that these indeed depend on  $\mathbf{p}$  use the equality  $\text{Tr}(\mathbf{X}\mathbf{X}^\top) = \|\mathbf{X}\|_2^2$  and recall  $\mathbb{E}[\mathbf{x} \circ \mathbf{e}] = \mathbf{x} \circ \mathbf{p}$ . Then  $\mathbb{E}[\|\tilde{\mathbf{x}}_G - \tilde{\mathbf{x}}'_G\|_2^2] = \|\tilde{\mathbf{x}}_G\|_2^2 - \text{Tr}(\mathbf{U}_G \mathbf{x} (\mathbf{x} \circ \mathbf{p})^\top \mathbf{U}_G^\top) - \text{Tr}(\mathbf{U}_G (\mathbf{x} \circ \mathbf{p}) \mathbf{x}^\top \mathbf{U}_G^\top) + \text{Tr}(\mathbf{U}_G \mathbf{P} \mathbf{U}_G^\top)$ , where  $\mathbf{P}$  is a matrix with entries  $\mathbf{P}_{i,j} = \mathbf{x}_i \mathbf{x}_j \mathbf{p}_i \mathbf{p}_j$  for  $i \neq j$  and  $\mathbf{P}_{i,i} = (\mathbf{x}_i)^2 \mathbf{p}_i$  for  $i = j$ .

We use these distances to optimize the probabilities based on prior knowledge such that one or more of the augmented embeddings are similar while the remaining ones differ. For example, when the curl and harmonic embeddings are important, such as in the trajectory prediction task that we will touch in the experiments (Sec. 5), we could design  $\mathbf{p}$  by solving:

$$\min_{\mathbf{p}} \quad -\mathcal{L}_G(\mathbf{p}) + \mathcal{L}_C(\mathbf{p}) + \mathcal{L}_H(\mathbf{p}) \quad (5a)$$

$$\text{subject to} \quad \mathbf{p} \in \mathcal{G}_{\mathbf{p}} := \left\{ \mathbf{p} \mid \mathbf{p} \in [0, 1]^{N_1}, \|\mathbf{p}\|_1 \leq \epsilon_{\mathbf{p}} \right\}, \quad (5b)$$

where set  $\mathcal{G}_{\mathbf{p}}$  puts a maximum budget  $\epsilon_{\mathbf{p}}$  on the allocated drop probabilities in a sparse manner (i.e., the  $\ell_1$ -norm  $\|\cdot\|_1$  imposes sparse probabilities on a few flows). The budget  $\epsilon_{\mathbf{p}}$  is tuned as a hyperparameter.<sup>2</sup> By solving (5), we find the dropout probabilities  $\mathbf{p}$  that generate examples with similar curl and harmonic components to the original data point but with a different gradient component. We solve this optimization problem with projected gradient descent, projecting  $\mathbf{p}$  onto the constraint set  $\mathcal{G}_{\mathbf{p}}$  after every step.

#### 4.2. Hodge-Aware Debiasing

Problem (5) influences the embedding space by acting on the augmentation functions  $\mathcal{T}_{1 \setminus 2}(\cdot)$  to generate better positive examples. To further improve the organization of the embedding space, we shall also act on the negative samples. This is known as a *debiasing* technique and consists of reweighting the negative samples in the InfoNCE loss [24–26]. I.e., by optimizing w.r.t. the loss:

$$\mathcal{L}_{\text{weighted}} = - \sum_{(\mathbf{x}'_i, \mathbf{x}'_j) \in \mathbb{P}} \log \left( \frac{e^{\text{sim}(\mathbf{z}'_i, \mathbf{z}'_j)/\tau}}{\sum_{m=1}^M w(\mathbf{x}_i, \mathbf{x}_m) e^{\text{sim}(\mathbf{z}'_i, \mathbf{z}_m)/\tau}} \right) \quad (6)$$

where  $w(\mathbf{x}_i, \mathbf{x}_m)$  is the weighting term between the anchor  $\mathbf{x}_i$  and the negative example  $\mathbf{x}_m$ . For Hodge-aware learning, this weight should reflect the spectral properties of the data; thus, we would like to push spectrally different samples further away from the anchor.

We first consider a weighted embedding similarity between two data points:

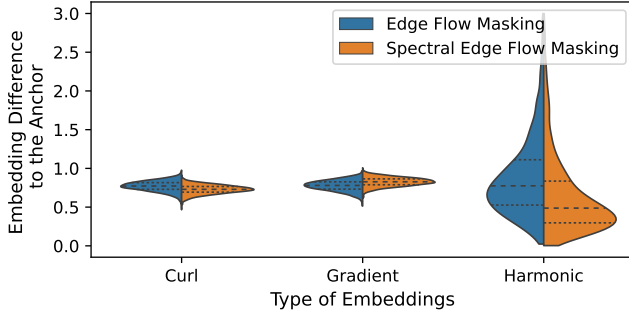
$$\mathcal{S}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_m) = \gamma_H \text{CD}(\tilde{\mathbf{x}}_{H,i}, \tilde{\mathbf{x}}_{H,m}) + \gamma_G \text{CD}(\tilde{\mathbf{x}}_{G,i}, \tilde{\mathbf{x}}_{G,m}) + \gamma_C \text{CD}(\tilde{\mathbf{x}}_{C,i}, \tilde{\mathbf{x}}_{C,m}) \quad (7)$$

$\text{CD}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_m) = 1 - \frac{\tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}}_m}{\|\tilde{\mathbf{x}}_i\|_2 \|\tilde{\mathbf{x}}_m\|_2}$  is the cosine distance and weights  $\gamma_H, \gamma_G, \gamma_C \geq 0$  are picked based on prior knowledge about the task or tuned as hyperparameters. Choosing the cosine distance leads to higher negative values for spectrally more dissimilar examples. Then, we compute the weight as the normalized similarity over the  $M$  negative samples:

$$w(\mathbf{x}_i, \mathbf{x}_m) = \frac{\mathcal{S}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_m)}{\sum_{m=1}^M \mathcal{S}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_m)} \quad (8)$$

which means that the weights for each datum sum to one and ensures that the loss terms for different data points are comparable. Summarizing the above, substituting (7) into (8) and the latter into (6) leads to an overall contrastive loss that pushes spectrally different samples away from the anchor based on this dissimilarity score. This encourages an embedding space that is organized with respect to the Hodge decomposition.

<sup>2</sup>In (5a), we could also weight the contributions of the different components (e.g.,  $\alpha_C \mathcal{L}_C(\mathbf{p}) + \alpha_H \mathcal{L}_H(\mathbf{p})$ ) with scalars  $\alpha_C, \alpha_H > 0$  when the signals have some contribution in each of them.



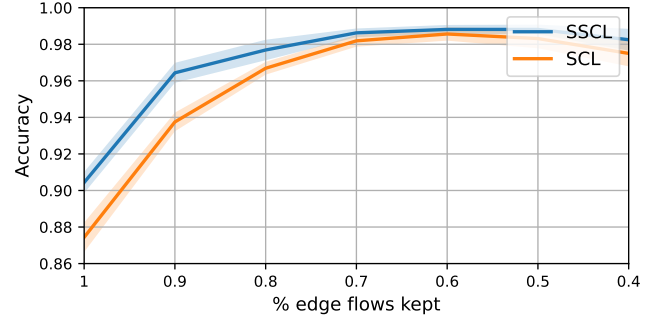
**Fig. 2:** Embedding distance for augmentations sampled with uniform probabilities (blue) and with the proposed spectrally optimized ones (orange). In the harmonic embedding, for the distribution associated with the spectral edge drop more probability mass lies over smaller differences and we are thus more likely to generate samples with more similar harmonic components than with the standard method. The dashed lines represent the three quartiles of the data, with the middle one being the median.

## 5. NUMERICAL RESULTS

We corroborate the proposed approach and compare it with supervised alternatives on two edge flow classification tasks that are common in the literature: i) a synthetic task considering trajectories on a simulated Simplicial Complex ; and ii) a real-data case that contains ocean drifters moving around the Madagascar island [27]. In i), the downstream learner has to differentiate between trajectories that pass through the bottom-right or top-left corner of the SC and in ii) it has to distinguish between clockwise and anti-clockwise moving drifters. Because of holes in the SC representations, the harmonic embedding captures important information for solving the tasks. Due to the limited space, we refer the reader to [28, Section 5] for more details about both datasets.

**Setup.** For the trajectory dataset, we generate 200 training, 100 validation, and 100 test data points while there are 160 training and 40 test data points for the ocean drifter. We train the unsupervised simplicial contrastive learner (SCL) on all available unlabeled data points and fit a linear support vector machine (SVM) on the obtained embeddings. For the ocean drifters, we use a 10-fold cross-validation on the training set to estimate the penalty parameter for the SVM. We report the average accuracy over 16 data splits. We optimize the network with stochastic gradient descent and grid search the learning rate and weight decay in the interval  $[10^{-5}, 1]$  in decimal steps. Furthermore, we select the edge flow drop probabilities  $\mathbf{p}$  and perturbation budget  $\epsilon_p$  from  $[0.1, 0.7]$ . All models are trained for 200 epochs with a batch size of 100. For the encoder we follow the setting in [12] (which is supervised in there) and use a Tanh-activation and a hidden-layer of size 64. We tune the number of layers and the convolutional orders  $L_1 = L_2$  in  $[1, 2, 3]$ . We compare the proposed approach with a fully-supervised SCNN.

**Results.** Table 1 depicts the overall performance on the downstream tasks. The spectral simplicial contrastive learner (SSCL<sub>Spec</sub>) trained with reweighted negative samples and spectrally optimized probabilities achieves the best downstream accuracy on both datasets. This shows the ability of the proposed approach to effectively encode more relevant Hodge-related information into the embeddings, facilitating the subsequent linear learner. Fig. 2 further reinforces this aspect by showing the embedding distance between the anchor and two different augmentation techniques (random edge drop and proposed). The proposed approach generates more similar harmonic embeddings, which is key to the obtained results for the task. In



**Fig. 3:** Performance comparison between a model trained with the standard loss (SCL) (cf. (4)) vs. a model trained with a spectrally reweighted loss (SSCL) (cf. (6)) for an edge drop augmentation. The SSCL outperforms the SCL irrespective of the augmentation quality.

**Table 1:** Test Accuracies for the Trajectory and Ocean Drifter datasets. SCL denotes models trained with the standard InfoNCE loss (cf. (4)), while SSCL models are trained with spectrally reweighted negatives (cf. (6)). The subscript Spec denotes that the augmentation probabilities are spectrally optimized.

Model	Trajectory Task	Ocean Drifters
SSCL <sub>Spec</sub> (ours)	97.9 $\pm$ 0.3	90.3 $\pm$ 1.4
SCNN (supervised)	95.2 $\pm$ 0.5	78.5 $\pm$ 1.1
SSCL	96.8 $\pm$ 0.4	89.1 $\pm$ 1.0
SCL <sub>Spec</sub>	98.2 $\pm$ 0.4	83.1 $\pm$ 1.1
SCL	96.1 $\pm$ 0.6	81.6 $\pm$ 1.6
SCL <sub>low</sub>	91.0 $\pm$ 0.2	77.1 $\pm$ 1.2

Fig. 3, we show the proposed approach consistently achieves a superior performance independent of the augmentation quality.

Notably, even for models trained without a reweighted loss, the incorporation of spectrally optimized augmentations (SCL<sub>Spec</sub>) improves the accuracy over uniform probabilities (SCL). This substantiates the importance of the spectral augmentations as a standalone feature. Furthermore, to evaluate the impact of the encoder, we tested a learner that uses only lower Laplacian encoding (SCL<sub>low</sub>), omitting triangle relationships. Compared to its simplicial counterpart SCL under identical conditions, SCL<sub>low</sub>, manifests a noticeable decrease in performance. This demonstrates that the structural advantages of simplicial networks to process flow data transfer to the contrastive learning setting.

## 6. CONCLUSION

We showed that a simplicial contrastive framework is effective for generating representations for edge flow data that contain hodge-related information. Related to this, we demonstrated that positive examples with specific spectral properties can be generated by casting the task as an optimization problem on the underlying probabilities of a dropout augmentation. Once these probabilities are optimized, they can be used to generate examples with desired spectral characteristics. We also introduce Hodge-related information into the problem by reweighting the negative examples in the loss based on their spectral difference to the anchor. This pushes spectrally different examples further apart and results in an embedding space that takes the relevant hodge information into account. Empirical results demonstrate that these optimized embeddings can be used to significantly outperform a fully-supervised model on two edge flow classification tasks. For future work, exploring other types of data augmentation methods and conducting experiments with simplicial complexes of varying dimensions remain as key areas.

## 7. REFERENCES

- [1] F.Battiston, G.Cencetti, I.Iacopini, V.Latora, M.Lucas, A.Patania, J.-G.Young, and G.Petri, "Networks beyond pairwise interactions: Structure and dynamics," *Physics Reports*, vol. 874, pp. 1–92, 2020.
- [2] M.Robinson, *Topological signal processing*, vol. 81, Springer, 2014.
- [3] M. T.Schaub, Y.Zhu, J.-B.Seby, T. M.Roddenberry, and S.Segarra, "Signal processing on higher-order networks: Livin' on the edge... and beyond," *Signal Processing*, vol. 187, pp. 108149, 2021.
- [4] S.Barbarossa and S.Sardellitti, "Topological signal processing over simplicial complexes," *IEEE Transactions on Signal Processing*, vol. 68, pp. 2992–3007, 2020.
- [5] T. M.Roddenberry, F.Frantzen, M. T.Schaub, and S.Segarra, "Hodgelets: Localized spectral representations of flows on simplicial complexes," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 5922–5926.
- [6] J.Krishnan, R.Money, B.Beferull-Lozano, and E.Isufi, "Simplicial vector autoregressive model for streaming edge flows," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [7] E.Isufi and M.Yang, "Convolutional filtering in simplicial complexes," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 5578–5582.
- [8] M.Yang, E.Isufi, M. T.Schaub, and G.Leus, "Finite impulse response filters for simplicial complexes," in *2021 29th European Signal Processing Conference (EUSIPCO)*.
- [9] M.Yang and E.Isufi, "Simplicial trend filtering," in *2022 56th Asilomar Conference on Signals, Systems, and Computers*, 2022.
- [10] S.Ebli, M.Defferard, and G.Spremann, "Simplicial neural networks," in *TDA & Beyond*, 2020.
- [11] T. M.Roddenberry, N.Glaze, and S.Segarra, "Principled simplicial neural networks for trajectory prediction," in *International Conference on Machine Learning*. PMLR, 2021, pp. 9020–9029.
- [12] C.Bodnar, F.Frasca, Y. G.Wang, N.Otter, G.Montufar, P.Liò, and M. M.Bronstein, "Weisfeiler and lehman go topological: Message passing simplicial networks," in *ICLR 2021 Workshop on Geometrical and Topological Representation Learning*, 2021.
- [13] M.Yang, E.Isufi, and G.Leus, "Simplicial convolutional neural networks," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 8847–8851.
- [14] A.van den Oord, Y.Li, and O.Vinyals, "Representation learning with contrastive predictive coding," *CoRR*, vol. abs/1807.03748, 2018.
- [15] T.Chen, S.Kornblith, M.Norouzi, and G. E.Hinton, "A simple framework for contrastive learning of visual representations," *CoRR*, vol. abs/2002.05709, 2020.
- [16] Y.You, T.Chen, Y.Sui, T.Chen, Z.Wang, and Y.Shen, "Graph contrastive learning with augmentations," in *Advances in Neural Information Processing Systems*, H.Larochelle, M.Ranzato, R.Hadsell, M.Balcan, and H.Lin, Eds. 2020, vol. 33, pp. 5812–5823, Curran Associates, Inc.
- [17] Y.Tian, C.Sun, B.Poole, D.Krishnan, C.Schmid, and P.Isola, "What makes for good views for contrastive learning?," in *Advances in Neural Information Processing Systems*, H.Larochelle, M.Ranzato, R.Hadsell, M.Balcan, and H.Lin, Eds. 2020, vol. 33, pp. 6827–6839, Curran Associates, Inc.
- [18] Y.Zhu, Y.Xu, F.Yu, Q.Liu, S.Wu, and L.Wang, "Deep Graph Contrastive Representation Learning," in *ICML Workshop on Graph Representation Learning and Beyond*, 2020.
- [19] Y.Zhu, Y.Xu, F.Yu, Q.Liu, S.Wu, and L.Wang, "Graph contrastive learning with adaptive augmentation," in *Proceedings of the Web Conference 2021*. apr 2021, ACM.
- [20] N.Liu, X.Wang, D.Bo, C.Shi, and J.Pei, "Revisiting graph contrastive learning from the perspective of graph spectrum," in *Advances in Neural Information Processing Systems*, A. H.Oh, A.Agarwal, D.Belgrave, and K.Cho, Eds., 2022.
- [21] L.Lin, J.Chen, and H.Wang, "Spectral augmentation for self-supervised learning on graphs," in *The Eleventh International Conference on Learning Representations*, 2023.
- [22] K.Ding, Z.Xu, H.Tong, and H.Liu, "Data augmentation for deep graph learning: A survey," *SIGKDD Explor. Newsl.*, vol. 24, no. 2, pp. 61–77, dec 2022.
- [23] H.Bhatia, G.Norgard, V.Pascucci, and P.-T.Bremer, "The helmholtz-hodge decomposition—a survey," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 8, pp. 1386–1404, 2013.
- [24] C.-Y.Chuang, J.Robinson, L.Yen-Chen, A.Torralba, and S.Jegelka, "Debiased contrastive learning," 2020.
- [25] Q.Sun, W.Zhang, and X.Lin, "Progressive hard negative masking: From global uniformity to local tolerance," 2022.
- [26] Y.Liu, X.Yang, S.Zhou, X.Liu, Z.Wang, K.Liang, W.Tu, L.Li, J.Duan, and C.Chen, "Hard sample aware network for contrastive deep graph clustering," 2022.
- [27] R.Lumpkin and L.Centurioni, "Global drifter program quality-controlled 6-hour interpolated data from ocean surface drifting buoys," *NOAA National Centers for Environmental Information*, 2019.
- [28] M. T.Schaub, A. R.Benson, P.Horn, G.Lippner, and A.Jadbabaie, "Random walks on simplicial complexes and the normalized hodge 1-laplacian," *SIAM Review*, vol. 62, no. 2, pp. 353–391, 2020.