

BScNets: Block Simplicial Complex Neural Networks

Yuzhou Chen,¹ Yulia R. Gel,² H. Vincent Poor¹

¹Department of Electrical and Computer Engineering, Princeton University

²Department of Mathematical Sciences, University of Texas at Dallas
{yc0774, poor}@princeton.edu, ygl@utdallas.edu

Abstract

Simplicial neural networks (SNN) have recently emerged as the newest direction in graph learning which expands the idea of convolutional architectures from node space to simplicial complexes on graphs. Instead of pre-dominantly assessing pairwise relations among nodes as in the current practice, simplicial complexes allow us to describe higher-order interactions and multi-node graph structures. By building upon connection between the convolution operation and the new block Hodge-Laplacian, we propose the first SNN for link prediction. Our new Block Simplicial Complex Neural Networks (BScNets) model generalizes the existing graph convolutional network (GCN) frameworks by systematically incorporating salient interactions among multiple higher-order graph structures of different dimensions. We discuss theoretical foundations behind BScNets and illustrate its utility for link prediction on eight real-world and synthetic datasets. Our experiments indicate that BScNets outperforms the state-of-the-art models by a significant margin while maintaining low computation costs. Finally, we show utility of BScNets as the new promising alternative for tracking spread of infectious diseases such as COVID-19 and measuring the effectiveness of the healthcare risk mitigation strategies.

Introduction

Graph Convolutional Networks (GCNs) is a powerful machinery for graph learning, allowing for efficient exploration of various pairwise interactions among graph nodes. However, most GCN-based approaches tend to be limited in their ability to efficiently exploit and propagate information across higher-order structures (Morris et al. 2019; Xiao and Deng 2020). In turn, many recent studies on cyber-physical, social, and financial networks suggest that relations among multi-node graph structures, as opposed to pairwise interaction among nodes, may be the key toward understanding hidden mechanisms behind structural organization of complex network systems. For example, disease transmission might be influenced not only by one-to-one interactions but also be largely impacted by various group relations and social reinforcement induced by each person’s social circle, e.g., COVID-19 anti-mask and anti-vaccine views. To enhance the obfuscation efforts in money laundering schemes, criminals involve complex interactions not only among multiple

individuals but also among multiple criminal groups. Such higher-order graph interactions beyond the node space and multi-node structures may be naturally modelled using simplicial complexes.

Simplicial neural networks (SNN) is the newly emerging direction in graph learning which extends convolutional operation to data that live on simplicial complexes. SNNs have recently been successfully applied to graph classification and trajectory forecasting (Ebli, Defferrard, and Spreemann 2020; Bunch et al. 2020; Roddenberry, Glaze, and Segarra 2021; Bodnar et al. 2021). The key engine behind SNNs is the Hodge-de Rham theory that generalizes the standard graph Laplacian which describes node-to-node interactions via edges to Hodge-Laplacian which allows us to model diffusion from edges to edges via nodes, edges to edges via triangles, triangles to triangles via edges, etc (Lim 2020; Schaub et al. 2020). Furthermore, Hodge-Laplacian establishes a natural connection between higher-order properties of discrete representations, e.g., graphs, and continuous representations, e.g., manifolds. As such, Hodge-Laplacian-based analytics opens multiple new perspectives for geometric representations at the interface of shape analysis, graph theory, and geometric deep learning (GDL) (Wang and Solomon 2019; Hajij, Istvan, and Zamzmi 2020).

In this paper we make the first step toward bridging SNNs with link prediction on graphs. We propose a new Block Simplicial Complex Neural Networks (BScNets) model, by building upon the connection between convolution operation and block Hodge-style representation. In contrast to other SNNs which tend to focus only on edge-to-edge information flows, or Hodge 1-Laplacian, BScNets allows us to simultaneously incorporate salient interactions among multiple simplicial complexes on graphs. Specifically, our BScNets scheme is composed of two components: the Adaptive Hodge Laplacian Based Block (AHLB) which describes multi-level structures and adaptively learns hidden dependencies among geometric representations of higher-order structures, and the Hodge-Style Adaptive Block Convolution (H-ABC) Module which automatically infers relations among multi-dimensional simplices. Our results indicate that this novel integration of information flows across not one but multiple higher-order structures via the Block Hodge Laplacian analytics yields the highest performance in link prediction tasks.

Significance and novelty of our contributions can be summarized as follows:

- Our BScNets is the first SNN for link prediction on graphs, bridging the recently emerging concepts of convolutional architectures on simplicial complexes with topological signal processing on graphs.
- We propose a new (random-walk based) adaptive Hodge Laplacian based block operator which simultaneously integrates knowledge on interactions among multiple higher-order graph structures and discuss its theoretical properties.
- We extensively validate BScNets on real-world and synthetic datasets, from such diverse domains as criminal, collaboration and transportation networks. Our results indicate that BScNets outperforms the state-of-the-art models by a significant margin while maintaining low computation costs.
- We discuss utility of BScNets and SNN tools as the new promising alternative for tracking spread of infectious diseases such as COVID-19 and evaluating healthcare risk mitigation strategies.

Related Work

Link Prediction GCN-based methods are known to be the powerful machinery for link prediction tasks. Specifically, the Graph Autoencoder (GAE) (Kipf and Welling 2016) and its variational version, i.e., Variational Graph Autoencoder (VGAE), are first employed to link prediction on citation networks. SEAL (Zhang and Chen 2018) extracts local enclosing subgraphs around the target links and learns a function mapping the subgraph patterns to link existence. In addition, the Hyperbolic Graph Convolutional Neural Networks (HGNN) (Chami et al. 2019) leverages both the hyperbolic geometry and GCN framework to learn node representations. Another interesting recent strategy is to use pairwise topological features to find latent representations of geometrical structure of graph using GCN (Yan et al. 2021). Our method differs from these approaches in that it explicitly models the higher-dimensional graph substructures and higher-order interactions via building an adaptive and interpretable Hodge block representation. Moreover, we propose a novel Hodge-style adaptive block convolution module to aggregate topological features encoded in the simplicial complexes by investigating relationships between simplices of different orders. This higher-order simplicial representation is substantially more general than the structural representation of node sets.

Simplicial Neural Networks Modeling higher-order interactions on graphs is an emerging direction in graph representation learning. While the role of higher-order graph structures for graph learning has been documented for a number of years (Agarwal, Branson, and Belongie 2006; Johnson and Goldring 2012) and involve such diverse applications as graph signal processing in electric, transportation and neuroscience systems, including link prediction (Benson et al. 2018), integration of higher-order graph substructures into deep learning on graphs has emerged only in 2020.

Indeed, several most recent studies propose to leverage simplicial information to perform neural networks on graphs. For instance, Ebli, Defferrard, and Spremann (2020) develops a model called Simplicial Neural Networks, which integrates the simplicial Laplacian of a simplicial complex into neural network framework based on simplicial Laplacian. Message Passing Simplicial Networks (MPSNs) (Bodnar et al. 2021) is proposed by performing message passing on simplicial complexes for graph classification. Similarly, SCoNe (Roddenberry, Glaze, and Segarra 2021) uses the GCN architecture depended on simplicial complexes to extrapolate trajectories on trajectory data. Besides, the convolutional message passing scheme on cell complex (Hajij, Istvan, and Zamzmi 2020) is shown to facilitate representational learning on graphs. Moreover, to emphasize the unifying principles offered by the categorical and cochain notions, Hajij et al. (2021) brings deep learning on complexes and discrete exterior calculus via topological networks. Our approach further advances these recent results by explicitly exploiting local topological information encoded in simplicial complexes of multiple dimensions and extracting the key interaction relations among multiple higher-order graph structures, which leads to significant gains in link prediction accuracy.

Block Simplicial Complex Neural Networks

We consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, where \mathcal{V} is the set of nodes ($|\mathcal{V}| = n$) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges ($|\mathcal{E}| = m$). Furthermore, each node v_i is associated with a d -dimensional vector of attributes \mathbf{x}_i which is the i -th row of matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ (where d is the input feature dimension). Connectivity of \mathcal{G} can be encoded in a form of an adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with entries $[\mathbf{A}]_{ij} = 1$ if nodes i and j are connected and 0, otherwise. For undirected graph \mathcal{G} , \mathbf{A} is symmetric (i.e., $\mathbf{A} = \mathbf{A}^\top$).

Background on Hodge Theory One of the focal points of graph theory and, in virtue of it, GDL, is graph Laplacian. Laplacian allows us to establish a natural link between discrete representations, e.g., graphs, and continuous representations, e.g., manifolds (Chung and Graham 1997). The (unnormalized) graph Laplacian is defined as $\mathbf{L}_0 = \mathbf{D} - \mathbf{A}$, and \mathbf{L}_0 is a symmetric and positive-semidefinite matrix. The Laplacian \mathbf{L}_0 represents a discrete counterpart of the Laplacian operator in vector calculus. In particular, as divergence of the gradient of some twice-differentiable multivariate function in vector calculus, the Laplacian operator is the flux density of the gradient flow of this function. That is, Laplacian measures how much the value of the function at any given point differs from average values of the function evaluated at nearby points. In turn, in the discrete case similarly, graph Laplacian \mathbf{L}_0 measures diffusion from node to node through edges and, broadly speaking, assesses at what extent the graph \mathcal{G} differs at one node from the graph \mathcal{G} at surrounding nodes. While \mathbf{L}_0 contains some very important information on the topology of \mathcal{G} , the natural question arises *what if we are interested in diffusion dynamics on graph substructures beyond the node level?* For example, formation of money laundering activities within criminal networks by

default involves very complex multi-node interactions and, hence, cannot be well captured by methods that focus on the dyadic graph relationships. To assess such higher-order network properties, we can study graphs in terms of topological objects such as *simplicial complexes* and exploit the discrete Hodge theory, particularly, Hodge-Laplacian-based analytics as generalization of Laplacian dynamics to polyadic substructures of \mathcal{G} .

Definition 1 A family Δ of finite subsets of a set \mathcal{V} is an abstract simplicial complex if for every $\sigma \in \Delta$, $\tau \subseteq \sigma$ implies $\tau \in \Delta$. I.e., Δ is closed under the operation of taking subsets. If $|\sigma| = k + 1$, then σ is called a k -simplex. Every subset $\tau \subset \sigma$ such that $|\tau| = k$ is called a face of σ . All simplices in Δ that have σ as face are called co-faces. Dimension of Δ is the largest dimension of any of its faces, or ∞ if there is no upper bound on the dimension of the faces.

Hence, nodes of \mathcal{G} are 0-simplices, edges are 1-simplices, and triangles are 2-simplices. For a k -simplex of $k > 0$, we can also define its *orientation* by (arbitrary) selecting some order for its nodes, and two orderings are said to be equivalent if they differ by an even permutation. As a result, for a given k -simplex σ with orientation $[i_0, i_1, \dots, i_k]$, any face of σ is assigned its own orientation (or “identifier”) $[i_0, i_1, \dots, i_{j-1}, i_{j+1}, \dots, i_k]$ (i.e., we omit the j -th element). To study diffusion among higher-order substructures of \mathcal{G} , we now form a real-valued vector space C^k which is endowed with basis from the oriented k -simplices and whose elements are called k -chains. Diffusion through higher-order graph substructures can be then defined via linear maps among spaces C^k of k -chains on \mathcal{G} (Lim 2020).

Definition 2 A linear map $\partial_k : C^k \rightarrow C^{k-1}$ is called a boundary operator. The adjoint of the boundary map induces the co-boundary operator $\partial_k^T : C^k \rightarrow C^{k+1}$. Matrix representations of ∂_k and ∂_k^T are B_k and B_k^T , respectively.

An operator over oriented k -simplices $L_k : C^k \rightarrow C^k$ is called the k -Hodge Laplacian, and its matrix representation is given by

$$L_k = B_k^T B_k + B_{k+1} B_{k+1}^T, \quad (1)$$

where $B_k^T B_k$ and $B_{k+1} B_{k+1}^T$ are often referred to L_k^{down} and L_k^{up} , respectively.

As $B_0 = 0$, the standard graph Laplacian L_0 is a subcase of (1) which tracks diffusion process from nodes to nodes via edges. Indeed, $L_0 = B_1 B_1^T$, where B_1 is an $n \times m$ -incidence matrix of \mathcal{G} (i.e., $B_1[ij] = 1$ if node i and edge j are incident and 0, otherwise). In turn, $L_1^{\text{down}} = B_1^T B_1$ is often referred to as *edge Laplacian* and assesses diffusion from edges to edges via nodes. Finally, Hodge 1-Laplacian L_1 measures variation on functions defined on graph edges (i.e., 1-simplices) with respect to the incidence relations among edges and nodes (i.e., 0-simplices) and edges and triangles (i.e., 2-simplices). More generally, L_k measures variation on functions defined on k -simplices of \mathcal{G} with respect to incidence relations among k -simplices with $(k-1)$ - and $(k+1)$ -simplices. Figure 1 shows an example of the simplicial complex on a graph (more details are in Appendix A.1).

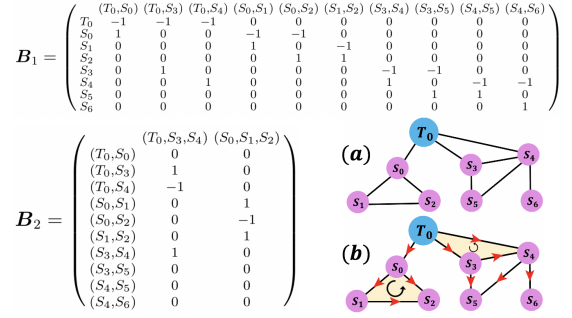


Figure 1: An example of generating simplicial complex from a classroom environment network, where T_0 is the teacher_0 and S_u is the student_ u ($u \in \{0, \dots, 6\}$). (a) Graph structure of the classroom environment network; (b) Simplicial complex of the classroom environment network; B_1 and B_2 are the corresponding node-to-edge and edge-to-face incidence matrices, respectively.

Random Walk Based Block Hodge-Laplacian Operator While the Hodge theory allows us to systematically describe diffusion across higher-order graph substructures, or k -simplices of any k , all current studies are restricted solely to Hodge 1-Laplacian L_1 (Schaub et al. 2020; Roddenberry, Glaze, and Segarra 2021; Bodnar et al. 2021). In this paper we propose a new random walk based block Hodge-Laplacian operator which enables us to simultaneously integrate knowledge on interactions among higher-order substructures of various orders into graph learning.

Definition 3 The K -block Hodge-Laplacian \mathfrak{L}_K is a divergence operator on the direct sum of vector spaces C^k , $k \in \mathbb{Z}_{\geq 0}$. That is, $\mathfrak{L}_K : \oplus_{k=0}^K C^k \rightarrow \oplus_{k=0}^K C^k$. Block Hodge-Laplacian \mathfrak{L}_K can be represented as a diagonal block matrix $\text{diag}\{L_0, L_1, \dots, L_K\}$. Furthermore, by selecting a subset of indices $\{k_1, \dots, k_J, \dots, k_J\} \in \mathbb{Z}_{\geq 0}$, we can consider a reduced block Hodge-Laplacian ${}^R\mathfrak{L}_J$, with a matrix representation $\text{diag}\{L_{k_1}, L_{k_2}, \dots, L_{k_J}\}$, which allows us to describe interaction relations among a particular subset of higher-order graph substructures. Here positive integers K and k_J are bounded by the dimension of the abstract simplicial complex Δ on \mathcal{G} and in practice, K and k_J are typically ≤ 2 .

The K -block Hodge-Laplacian \mathfrak{L}_K is related to the Dirac operator in differential geometry (i.e., Dirac operator is a square root of \mathfrak{L}_K) (Lloyd, Garnerone, and Zanardi 2016). As such, \mathfrak{L}_K has multiple implications for analysis of synchronization dynamics and coupling of various topological signals on graphs, with applications in physics and quantum information processing (Calmon et al. 2021).

Lemma 1 For $K, J \in \mathbb{Z}_{\geq 0}$, operators \mathfrak{L}_K and ${}^R\mathfrak{L}_J$ are symmetric and semipositive-definite, i.e. $\mathfrak{L}_K \geq 0$ and ${}^R\mathfrak{L}_J \geq 0$; $\mathfrak{L}_K = (\mathfrak{L}_K)^T$ and ${}^R\mathfrak{L}_J = ({}^R\mathfrak{L}_J)^T$.

Furthermore, we propose the random walk-based block Hodge-Laplacian, i.e., r -th power of block Hodge-Laplacian representation (where $r \in \mathbb{Z}_{\geq 0}$). Our random walk-based block Hodge-Laplacian is inspired by the recent success in

random walk based graph embeddings and simplicial complexes but is designed to conduct informative *joint* random walks on higher-order Hodge Laplacians instead of limited powering Hodge 1-Laplacian (Benson et al. 2018; Bodnar et al. 2021). Indeed, successfully travelling through higher-order topological features will provide us with additional feature information which is valuable for learning edge embeddings. The following Lemma 2 provides insight on the interplay among random walks on various k -simplices and their respective Hodge-Laplacian representation.

Lemma 2 *For any $r \in \mathbb{Z}_{>0}$ and $K, J \in \mathbb{Z}_{\geq 0}$, the r -power of $(\mathfrak{L}_K)^r$ and $({}^R\mathfrak{L}_J)^r$ are given by $\oplus_{k=0}^K ((\mathbf{L}_k^{\text{down}})^r + (\mathbf{L}_k^{\text{up}})^r)$ and $\oplus_{j=1}^J ((\mathbf{L}_{k_j}^{\text{down}})^r + (\mathbf{L}_{k_j}^{\text{up}})^r)$, respectively.*

Adaptive Hodge Laplacian Based Block Analytics Our new approach to learning higher-order graph topology is motivated by the following question: *Can we capture the interaction relations among k -simplices on the graph \mathcal{G} whose orders are farther than $k-1$ and $k+1$ apart?* While addressing this problem falls outside the Hodge-de Rham theory on the simplicial Hodge Laplacians, we provide an affirmative answer to this question through building an adaptive Hodge Laplacian Based Block Operator.

Without loss of generality, let us consider a 2-block (reduced) Hodge-Laplacian

$$\mathfrak{L}_2 = \left[\begin{array}{c|c} \mathbf{L}_{k_1} & \mathbf{O} \\ \hline \mathbf{O} & \mathbf{L}_{k_2} \end{array} \right]. \quad (2)$$

Our goal is to construct a new linear operator such that we describe interaction relations among \mathbf{L}_{k_1} and \mathbf{L}_{k_2} . This problem can be addressed by defining and learning a similarity function $f(\mathbf{L}_{k_1}, \mathbf{L}_{k_2})$. The natural choice for such similarity function is the inner product among elements in C^{k_1} (i.e., k_1 -chains) and elements in C^{k_2} (i.e., k_2 -chains). However, in general, $\dim(C^{k_1}) \neq \dim(C^{k_2})$. Assume that $d_1 = \dim(C^{k_1}) > d_2 = \dim(C^{k_2})$. Hence, instead we can consider a similarity function based on the inner product between elements in C^{k_2} and elements in $\mathbb{P}_{d_2} C^{k_1}$, i.e., the projection of C^{k_1} on the lower dimensional space. Here \mathbb{P}_{d_2} is the corresponding orthogonal projector and, given symmetry of \mathbf{L}_{k_1} , can be formed by the eigenvectors of \mathbf{L}_{k_1} corresponding to the top d_2 largest eigenvalues. Our new Adaptive Hodge Laplacian Based (AHLB) Block operator then takes the form

$$\mathbb{L}_2^B = \left[\begin{array}{c|c} \mathbf{L}_{k_1} & f(\mathbb{P}_{d_2} \mathbf{L}_{k_1}, \mathbf{L}_{k_2}) \\ \hline f(\mathbb{P}_{d_2} \mathbf{L}_{k_1}, \mathbf{L}_{k_2})^\top & \mathbf{L}_{k_2} \end{array} \right]. \quad (3)$$

Note that in general, linear operator $\mathbb{L}_2^B \in \mathbb{R}^{(d_1+d_2) \times (d_1+d_2)}$ is no longer a Laplacian. For example, while \mathbb{L}_2^B is symmetric (by construction), it may not satisfy the condition of positive semidefiniteness (see Lemma 3 in Appendix A.2). However, as shown below, AHLB opens multiple opportunities to better describe higher-order interactions on graphs that are inaccessible not only with individual k -Hodge Laplacians but even with the K -block Hodge-Laplacian operator (see the ablation study for more details). In addition, followed by Lemma 2, we also consider a case when AHLB is constructed based on the r -th power of \mathfrak{L}_2 , and our studies in-

dicate that on average the best results are achieved for $r = 2$ (see Table 3 and Appendix A.3).

Armed with the AHLB operator (Eq. 3), we can utilize the non-local message operation $f(\cdot, \cdot)$ (Wang et al. 2018) to capture long-range relations and intrinsic higher-order connectivity among entities in the 2-block (reduced) Hodge-Laplacian \mathfrak{L}_2 , i.e., $f(\mathbb{P}_{d_2} \mathbf{L}_{k_1}, \mathbf{L}_{k_2})$ placed in the off-diagonal. We now describe the choices of non-local message passing functions which can be used for $f(\cdot, \cdot)$ in the following. Specifically, given two higher-order Hodge Laplacians \mathbf{L}_{k_1} and \mathbf{L}_{k_2} , we define two types of the non-local message passing functions to capture the relations between $[\mathbb{P}_{d_2} \mathbf{L}_{k_1}]_i \in \mathbb{R}^{d_2}$ and $[\mathbf{L}_{k_2}]_j \in \mathbb{R}^{d_2}$ (i.e., topological embedding of the i -th graph substructure in $\mathbb{P}_{d_2} \mathbf{L}_{k_1}$ and topological embedding of the j -th graph substructure in \mathbf{L}_{k_2} respectively) as

(1) *Inner-Product*

$$f([\mathbb{P}_{d_2} \mathbf{L}_{k_1}]_i, [\mathbf{L}_{k_2}]_j) = \langle [\mathbb{P}_{d_2} \mathbf{L}_{k_1}]_i, [\mathbf{L}_{k_2}]_j \rangle.$$

Since all considered Laplacians in our case are over the real field \mathbb{R} , we can consider a dot-product. However, in a more general case of multivariable functions on graph simplices, $f(\cdot, \cdot)$ can be an inner-product.

(2) *Embedded Inner-Product*

$$f([\mathbb{P}_{d_2} \mathbf{L}_{k_1}]_i, [\mathbf{L}_{k_2}]_j) = \langle \Theta_\xi [\mathbb{P}_{d_2} \mathbf{L}_{k_1}]_i, \Theta_\psi [\mathbf{L}_{k_2}]_j \rangle,$$

where $\Theta_\xi \in \mathbb{R}^{d_c \times d_2}$ and $\Theta_\psi \in \mathbb{R}^{d_c \times d_2}$ are weight matrices to be learned, and d_c is the embedding dimension. As such, we infer the relation value between two simplices.

Finally, the adaptive Hodge block can be formulated as

$$\tilde{\mathbb{L}}_2^B = \text{softmax} \left(\text{ReLU} \left(\mathbb{L}_2^B \right) \right), \quad (4)$$

where softmax function is applied to normalize the block operator \mathbb{L}_2^B , and the ReLU activation function (where $\text{ReLU}(\cdot) = \max(0, \cdot)$) eliminates both weak pairwise relation among higher-dimensional simplices and weak connections. As a result, the adaptive Hodge block $\tilde{\mathbb{L}}_2^B$ can dynamically learn interactions among simplices of different dimensions.

Hodge-Style Adaptive Block Convolution Module

Armed with the AHLB operator, we now discuss how to construct the Hodge-style adaptive block convolution and use it to learn the distance between two nodes in the embedding space. Instead of using graph convolutional operator, i.e., \mathbf{L}_0 , we adopt the AHLB operator $\tilde{\mathbb{L}}_2^B$ to demonstrate the effectiveness of the proposed novel higher-order Hodge convolution operator. The Hodge-style adaptive block convolution (H-ABC) module is given by

$$\mathbf{Z}^{(\ell+1)} = \left(\tilde{\mathbb{L}}_2^B \mathbf{Z}^{(\ell)} \Theta_1^{(\ell)} \right) \Theta_2^{(\ell)}, \quad (5)$$

where $\mathbf{Z}^{(0)} = \mathbf{X} \in \mathbb{R}^{n \times d}$ is node features matrix, $\Theta_1^{(\ell)} \in \mathbb{R}^{d \times (d_1+d_2)}$ and $\Theta_2^{(\ell)} \in \mathbb{R}^{(d_1+d_2) \times d_{\text{out}}}$ are two trainable weight matrices at layer ℓ , and d_{out} denotes the dimension of node embedding at the (ℓ) -th layer through the H-ABC operation. For the link prediction task, we use the Fermi-Dirac

decoder (Nickel and Kiela 2017) to compute the distance between the two nodes. Formally,

$$\text{dist}_{uv}^{\text{H-ABC}} = (\mathbf{Z}_u^{(\ell+1)} - \mathbf{Z}_v^{(\ell+1)})^2, \quad (6)$$

where $\text{dist}_{uv}^{\text{H-ABC}} \in \mathbb{R}^{1 \times d_{\text{out}}}$ is the distance between nodes u and v in a local spatial domain.

Graph Convolution Layer Similar to the process of computing distances between learnable node embeddings with the Hodge-style adaptive block convolution, we also use graph convolution operation (Eq. 7) to evaluate the distance (Eq. 8) between the node embeddings of nodes u and v as

$$\mathbf{H}^{(\ell+1)} = \tilde{\mathbf{L}}\mathbf{H}^{(\ell)}\boldsymbol{\Theta}_3^{(\ell)}, \quad (7)$$

$$\text{dist}_{uv}^{\text{GC}} = (\mathbf{H}_u^{(\ell+1)} - \mathbf{H}_v^{(\ell+1)})^2, \quad (8)$$

where $\tilde{\mathbf{L}} = \mathbf{D}_v^{-1/2}(\mathbf{A} + \mathbf{I})\mathbf{D}_v^{-1/2}$ (where \mathbf{D}_v is the degree matrix of $\mathbf{A} + \mathbf{I}$, i.e., $[\mathbf{D}_v]_{ii} = \sum_j [\mathbf{A} + \mathbf{I}]_{ij}$), $\mathbf{H}^{(0)} = \mathbf{X} \in \mathbb{R}^{n \times d}$ is node features matrix, $\boldsymbol{\Theta}_3^{(\ell)} \in \mathbb{R}^{d \times d_{\text{out}}}$ is the trainable weight matrix, d_{out} denotes the dimension of node embedding at the (ℓ) -th layer through the graph convolution operation, $\text{dist}_{uv}^{\text{GC}} \in \mathbb{R}^{1 \times d_{\text{out}}}$ is the distance between nodes u and v in global spatial domain.

Distance between Two Nodes We wrap the concatenation of Hodge-style adaptive convolution and graph convolution operation outputs (i.e., in Eqs. 6 and 8) into Multi-layer Perceptron (MLP) block

$$\text{dist} = \text{ReLU}(f_{\text{MLP}}([\pi_\alpha \times \text{dist}_{uv}^{\text{GC}}, \pi_\beta \times \text{dist}_{uv}^{\text{H-ABC}}])),$$

where $[\cdot, \cdot]$ denotes the concatenation of the outputs from H-ABC module and graph convolution operation, f_{MLP} is an MLP layer that maps the concatenated embedding to a d_o -dimensional space, and π_α and π_β are the hyperparameters representing the weight of each distance (i.e., dist) in the $(\ell + 1)$ -th layer. Based on the propagation rule above, the edges connection probability can be computed as $\text{prob}(u, v) = [\exp((\text{dist}_{uv} - \delta)/\eta) + 1]^{-1}$, where δ and η are hyperparameters. Then, training via standard back-propagation is performed via binary cross-entropy loss function using negative sampling. Figure 2 illustrates our proposed BScNets framework, which consists of the Hodge-style adaptive block convolution module and graph convolution operation (see the Appendix A.4 for details).

Experimental Study

Datasets We experiment on three types of networks (i) citation networks: Cora and PubMed (Sen et al. 2008); (ii) social networks: (1) flight network: Airport (Chami et al. 2019), (2) criminal networks: Meetings and Phone Calls (Cavallaro et al. 2020), and (3) contact networks: High School network and Staff Community (Salathé et al. 2010; Eleteby et al. 2020); (iii) disease propagation tree: Disease (Chami et al. 2019). The statistics and more details of the datasets are provided in Appendix B.1.1.

Baselines We compare against ten state-of-the-art (SOA) baselines, including (i) MLP, (ii) GCN (Kipf

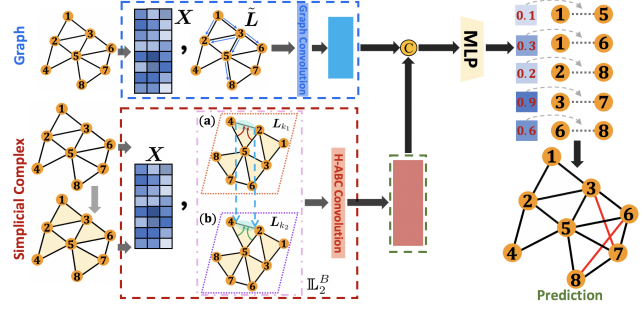


Figure 2: The architecture of BScNets (for more details see Appendix A.4).

and Welling 2017), (iii) Simplified Graph Convolution (SGC) (Wu et al. 2019), (iv) Graph Attention Networks (GAT) (Veličković et al. 2018), (v) GraphSAGE (SAGE), (vi) SEAL (Zhang and Chen 2018), (vii) Hyperbolic neural networks (HNN) (Ganea, Bécigneul, and Hofmann 2018), (viii) Hyperbolic Graph Convolutional Neural Networks (HGNN) (Chami et al. 2019), (ix) Persistence Enhanced Graph Neural Network (PEGN) (Zhao et al. 2020), and (x) Topological Loop-Counting Graph Neural Network (TLC-GNN) (Yan et al. 2021). Further details of baselines are contained in Appendix B.1.2.

Experiment Settings We implement our proposed BScNets with Pytorch framework on two NVIDIA RTX 3090 GPUs with 24 GiB RAM. Following (Chami et al. 2019), for all datasets, we randomly split edges into 85%/5%/10% for training, validation, and testing. Further, for all datasets, BScNets is trained by the Adam optimizer with the Cross Entropy Loss function. For all methods, we run 20 times with the same partition and report the average and standard deviation of ROC AUC values. More details about the experimental setup and hyperparameters are in Appendix B.2. Our datasets and codes are available on <https://github.com/BScNets/BScNets.git>.

Experiment Results Tables 1 and 2 show the comparison of our proposed BScNets with SOAs for link prediction tasks on six network datasets (i.e., Cora, PubMed, Meetings, Phone Calls, Airport, and Disease) and two contact network datasets (i.e., High School and Staff Community), respectively. The results indicate that our BScNets consistently achieves the best performance on all network datasets compared to all SOAs. Particularly, from Table 1, we find that: (i) Compared to the spectral-based ConvGNNs (i.e., GCN and SGC), BScNets yields more than 3.32% relative improvements to the existing best results for all six datasets; (ii) BScNets outperforms spatial-based ConvGNNs (i.e., GAT, SAGE, and SEAL) with a significant margin; (iii) Compared to the hyperbolic-based NNs (i.e., HNN and HGNN), BScNets improves upon HGNN (which performs best among the hyperbolic-based models) by a margin of 2.00%, 1.29%, 5.51%, 11.62%, 1.20% and 7.91% on datasets Cora, PubMed, Meetings, Phone Calls, Airport, and Disease, respectively; (iv) BScNets further improves PH-based ConvGNNs (i.e., PEGN and TLC-GNN)

with a significant margin on all six datasets. Additionally, Table 2 shows performance of BScNets and baseline methods on High School network and Staff Community. BScNets obtains the superior results on all datasets, outperforming the representative spectral-, spatial-, PH-based models including GCN, SEAL, and TLC-GNN by a large margin. Based on the one-sided two-sample t -test, our BScNets also demonstrates statistically significant improvement in performance (marked by *) compared to the existing best results in ROC AUC for all eight datasets (see Appendix B.5 for additional experimental results of BScNets).

Overall, the results show that BScNets can accurately capture the key structural information on the graph, both at the dyadic and polyadic levels of interactions, and achieves a highly promising performance in link prediction.

Model	High School	Staff Community
GCN (Kipf and Welling 2017)	63.55 \pm 1.72	64.97 \pm 6.88
SEAL (Zhang and Chen 2018)	68.13 \pm 1.50	65.60 \pm 3.46
HGCN (Chami et al. 2019)	67.30 \pm 1.28	66.75 \pm 1.38
TLC-GNN (Yan et al. 2021)	69.15 \pm 1.49	67.35 \pm 5.29
BScNets (ours)	*71.68 \pm 1.72	*79.13 \pm 2.95

Table 2: ROC AUC for link prediction on contact networks.

Ablation Study To further investigate the importance of the different components in BScNets, we have conducted an ablation study of our proposed model on Cora and Disease and results are presented in Table 3 (* means statistically significant result). We compare our BScNets with three ablated variants, i.e., (i) BScNets without random walk on block Hodge-Laplacian (W/o Random walks), (ii) BScNets without relation modeling (off-diagonal terms) (W/o Relation), and (iii) BScNets without the AHLB operator but with Hodge 1-Laplacian (With only L_1 ; i.e., instead of using the block structure, we directly incorporate Hodge 1-Laplacian L_1 into the Hodge-style convolution module by replacing $\tilde{\mathbb{L}}_2^B$ with L_1 in Eq. 5). The results indicate that, when ablating the above components, the ROC AUC score of BScNets drops significantly. For both datasets, random walk mechanism on the block Hodge-style representation significantly improves the results as it utilizes higher-order relationships expressed in the graph data and learns embeddings beyond the node-space. In addition, we show that BScNets outperforms BScNets without relation modeling due to the fact that relation modeling integrates the learnt relationship between information in multi-dimensional simplices into the Hodge-style adaptive block convolutional encoder. Moreover, as expected, we find that replacing the AHLB operator by only Hodge 1-Laplacian results in a significant decrease in performance that indicates that block structure and off-diagonal higher-order relationships consistently boost the performance of link prediction.

Computational Complexity Incidence matrices B_1 and B_2 can be calculated efficiently with the computational complexity $\mathcal{O}(n + m)$ and $\mathcal{O}(m + q)$ respectively, where n is the number of 0-simplices (i.e., nodes), m is the number of 1-simplices (i.e., edges), and q is the number of 2-

simplices (filled triangles). For large-scale graphs, we sample $\epsilon \times m$ edges (where $\epsilon \in (0, 1]$) from edge set \mathcal{E} and then pass these $\epsilon \times m$ edges to construct higher-order Hodge Laplacians. Thus this sparse sampling reduces the computation complexity of constructing B_1 and B_2 to $\mathcal{O}(n + \epsilon \times m)$ and $\mathcal{O}(\epsilon \times m + q')$ respectively (where q' denotes the number of 2-simplices after sparse sampling). This sampling method is motivated by the DropEdge technique (Rong et al. 2019) which has been used to prevent over-fitting and over-smoothing in GCNs via randomly removing edges from the graph. Equipping BScNets with edge sampling results in substantial reductions in the size of the input data and, as such, computational complexity. Table in Appendix B.4 reports the average running time of incidence matrices generation for target network, training time per epoch of our BScNets model on all datasets, and running time comparison between BScNets and baselines.

Infectious Disease Transmission As recently noted by Iacopini et al. (2019); Schaub et al. (2020); Li et al. (2021), higher-order relationships such as multi-member group interactions, may be the key driving factors behind contagion dynamics, and simplicial complexes offer a natural way to describe such polyadic group relations. Here we are motivated by the two interlinked questions: *Can we use link prediction with and without simplicial complexes to reconstruct unobserved social interactions and then to approximate the underlying contagion dynamics in the whole population? How accurate are the link prediction methods, with and without simplicial complexes, in reflecting the impact of disease mitigation strategies?*

Indeed, in practice public healthcare professionals might have records on most registered residents (i.e., nodes) via, e.g., social security offices. However, there is only partial information on social interactions (i.e., edges) among these individuals. Our goals are to (i) predict social links among individuals where the interaction information is unknown, (ii) assess how mitigation strategies will impact the epidemic spread predicted from the reconstructed network, and (iii) evaluate how close the infection curves on the reconstructed data are to the infection curve on the whole true population. More details about the procedure of infectious disease transmission evaluation are provided in Appendix B.3.

We adopt the Susceptible-Exposed-Infected-Recovered (SEIR) (Kermack and McKendrick 1927) epidemic model. We set the infection curve delivered by SEIR on the whole population as the ground truth. We perturb certain fraction of links (i.e., remove the real links and add fake links) from the whole population to address the real world scenario when we do not have access to information on all social interactions. We then apply two link prediction methods to the network with perturbation, with and without simplicial complexes, BScNets and TLC-GNN, respectively. We now reconstruct the whole population by using BScNets and TLC-GNN and consider the two infection curves yielded by SEIR on the BScNets-based and TLC-GNN-based reconstructed networks. In addition, we assess sensitivity of the BScNets and TLC-GNN curves to mitigation strategies where more central nodes (individuals) are targeted to receive a vaccine, to

Model	Cora	PubMed	Meetings	Phone Calls	Airport	Disease
MLP	83.15 \pm 0.51	84.10 \pm 0.97	63.20 \pm 6.22	60.10 \pm 6.72	89.51 \pm 0.52	72.62 \pm 0.61
HNN (Ganea, Bécigneul, and Hofmann 2018)	89.00 \pm 0.10	94.87 \pm 0.11	71.00 \pm 3.28	60.90 \pm 4.25	90.78 \pm 0.22	75.10 \pm 0.35
GCN (Kipf and Welling 2017)	90.42 \pm 0.28	91.11 \pm 0.55	72.08 \pm 4.19	61.50 \pm 5.80	89.27 \pm 0.42	64.70 \pm 0.56
GAT (Veličković et al. 2018)	93.89 \pm 0.13	91.22 \pm 0.12	74.00 \pm 4.68	63.40 \pm 5.20	90.55 \pm 0.37	69.99 \pm 0.32
SAGE (Hamilton, Ying, and Leskovec 2017)	86.24 \pm 0.65	85.96 \pm 1.16	72.30 \pm 5.25	62.07 \pm 5.49	90.47 \pm 0.59	65.91 \pm 0.33
SGC (Wu et al. 2019)	91.67 \pm 0.20	94.10 \pm 0.20	73.38 \pm 3.49	63.80 \pm 5.71	90.01 \pm 0.32	65.21 \pm 0.23
SEAL (Zhang and Chen 2018)	92.55 \pm 0.50	92.42 \pm 0.12	71.09 \pm 7.50	62.96 \pm 4.17	95.16 \pm 0.39	85.23 \pm 0.79
HGCN (Chami et al. 2019)	93.00 \pm 0.45	96.29 \pm 0.18	83.20 \pm 4.15	70.20 \pm 3.77	96.40 \pm 0.19	90.80 \pm 0.30
PEGN (Zhao et al. 2020)	93.13 \pm 0.50	95.82 \pm 0.20	74.17 \pm 5.00	65.23 \pm 4.15	95.46 \pm 0.71	83.61 \pm 1.26
TLC-GNN (Yan et al. 2021)	94.22 \pm 0.78	97.03 \pm 0.10	73.20 \pm 5.32	66.17 \pm 3.90	96.60 \pm 0.69	86.19 \pm 1.23
BScNets (ours)	*94.90 \pm 0.70	*97.55 \pm 0.12	*88.05 \pm 5.51	*79.43 \pm 6.04	*97.57 \pm 0.67	*98.60 \pm 0.58

Table 1: ROC AUC and standard deviations for link prediction. Bold numbers denote the best results. * means statistically significant result.

Architecture	Cora	Disease
BScNets	*94.90 \pm 0.70	*98.60 \pm 0.58
W/o Random walk	93.79 \pm 0.85	93.38 \pm 5.85
W/o Relation	94.00 \pm 0.60	95.90 \pm 1.13
With L_1	93.73 \pm 0.44	93.43 \pm 1.89

Table 3: Ablation study of the network architecture (%).

quarantine or are persuaded to wear masks (Chen et al. 2021; Curiel and Ramírez 2021). Under the targeted mitigation strategies, a fraction of the most central nodes cannot transmit the disease. The ‘base’ infection curve is obtained from the whole population where such central nodes are removed from the disease spread. In turn, BScNets and TLC-GNN operate on the partially observed data where some edges are perturbed (as discussed above) and, upon reconstructing the unknown social interactions, BScNets and TLC-GNN are asked to re-determine their own individuals to be targeted through the disease mitigation strategy. We perform these experiments on the High School network (Salathé et al. 2010; Elestreby et al. 2020). We select TLC-GNN as the competing link prediction approach as it is a runner-up (see Table 2). We set the edge perturbation rate to 20%. We consider betweenness and degree centralities. We simulate each scenario 50 times and average the results. (For more experiments see Appendix B.3).

Figure 3 shows the infection curves during the time period of 180 days fitted to (i) the original network, (ii) the BScNets reconstructed network, and (iii) the TLC-GNN reconstructed network, under targeted mitigation strategy based on the betweenness centrality. We observe that the infection curve of the BScNets-reconstructed network (curve_{BScNets}) is significantly closer to the ‘base’ infection curve (i.e., curve_{base}) than the curve based on TLC-GNN (curve_{runner-up}). Most importantly, we find that while the peaks of the curve_{BScNets} and the curve_{base} are very close, i.e. 59.1% and 59.8%, respectively, the TLC-GNN curve (which does not account for multi-node group interactions) tends to substantially underestimate the number of infected individuals. For instance, at day $t = 30$, $t = 60$, $t = 90$, $t = 120$, and at day

$t = 180$, TLC-GNN suggests that only 0.5%, 2.2%, 10.1%, 29.5%, and 54.2% of population are infected, while the ground truth ‘base’ curve projects that 0.6%, 2.9%, 12.2%, 32.2%, and 59.8% are infected. That is, the difference between TLC-GNN and the ‘base’ curve, even for such small network, reaches 5% of the population, which eventually translates into significant societal implications such as shortages of health care facilities and overall system unpreparedness to pandemics like COVID-19. This phenomenon also confirms the most recent premise of epidemiological studies that higher-order interactions among multi-node groups (i.e., simplicial complexes) are the hidden driving factors behind the disease transmission mechanism. In turn, BScNets and other SNNs may be the most promising direction not only to capture such higher-order group interactions but also to reveal hidden polyadic relations in social networks.




Figure 3: SEIR infection curves based on the original High School (HS) network (purple), BScNets-reconstructed HS network (orange), and TLC-GNN-reconstructed HS network (green), under betweenness-based mitigation strategy.

Conclusion

We have proposed the first SNN for link prediction. We have shown that relations among multiple multi-node structures play a significant role in graph learning. In the future, we will extend the ideas of simplicial DL to dynamic networks.

Acknowledgements

This material is based upon work sponsored by the National Science Foundation under award numbers ECCS 2039701, ECCS 2039716, INTERN supplement for ECCS 1824716, DMS 1925346 and the Department of the Navy, Office of Naval Research under ONR award number N000142112226. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

References

- Agarwal, S.; Branson, K.; and Belongie, S. 2006. Higher order learning with graphs. In *ICLR*, 17–24.
- Benson, A. R.; Abebe, R.; Schaub, M. T.; Jadbabaie, A.; and Kleinberg, J. 2018. Simplicial closure and higher-order link prediction. *PNAS*, 115(48): E11221–E11230.
- Bodnar, C.; Frasca, F.; Wang, Y. G.; Otter, N.; Montúfar, G.; Lio, P.; and Bronstein, M. 2021. Weisfeiler and leman go topological: Message passing simplicial networks. In *ICML*.
- Bunch, E.; You, Q.; Fung, G.; and Singh, V. 2020. Simplicial 2-Complex Convolutional Neural Networks. In *NeurIPS 2020 Workshop on TDA and Beyond*.
- Calmon, L.; Restrepo, J. G.; Torres, J. J.; and Bianconi, G. 2021. Topological synchronization: explosive transition and rhythmic phase. *arXiv:2107.05107*.
- Cavallaro, L.; Ficara, A.; De Meo, P.; Fiumara, G.; Catanese, S.; Bagdasar, O.; Song, W.; and Liotta, A. 2020. Disrupting resilient criminal networks through data analysis: The case of Sicilian Mafia. *Plos One*, 15(8): e0236476.
- Chami, I.; Ying, Z.; Ré, C.; and Leskovec, J. 2019. Hyperbolic graph convolutional neural networks. In *NeurIPS*, volume 32, 4868–4879.
- Chen, J.; Hoops, S.; Marathe, A.; Mortveit, H.; Lewis, B.; Venkatramanan, S.; Haddadan, A.; Bhattacharya, P.; Adiga, A.; Vullikanti, A.; et al. 2021. Prioritizing allocation of COVID-19 vaccines based on social contacts increases vaccination effectiveness. *medRxiv PMID: 33564778*.
- Chung, F. R.; and Graham, F. C. 1997. *Spectral graph theory*. 92. AMS.
- Curiel, R. P.; and Ramírez, H. G. 2021. Vaccination strategies against COVID-19 and the diffusion of anti-vaccination views. *Scientific Reports*, 11(1): 1–13.
- Ebli, S.; Defferrard, M.; and Spreemann, G. 2020. Simplicial Neural Networks. In *NeurIPS 2020 Workshop on TDA and Beyond*.
- Eletreby, R.; Zhuang, Y.; Carley, K. M.; Yağan, O.; and Poor, H. V. 2020. The effects of evolutionary adaptations on spreading processes in complex networks. *PNAS*, 117(11): 5664–5670.
- Ganea, O.-E.; Bécigneul, G.; and Hofmann, T. 2018. Hyperbolic neural networks. In *NeurIPS*, 5350–5360.
- Hajij, M.; Istvan, K.; and Zamzmi, G. 2020. Cell complex neural networks. In *NeurIPS 2020 Workshop on TDA and Beyond*.
- Hajij, M.; Zamzmi, G.; Ramamurthy, K. N.; and Saenz, A. G. 2021. Data-Centric AI Requires Rethinking Data Notion. *arXiv preprint arXiv:2110.02491*.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *NeurIPS*, 1025–1035.
- Iacopini, I.; Petri, G.; Barrat, A.; and Latora, V. 2019. Simplicial models of social contagion. *Nature Communications*, 10(1): 1–9.
- Johnson, J. L.; and Goldring, T. 2012. Discrete hodge theory on graphs: a tutorial. *IEEE Comput Sci Eng*, 15(5): 42–55.
- Kermack, W. O.; and McKendrick, A. G. 1927. A contribution to the mathematical theory of epidemics. *Proc. of the Royal Soc. of London. Ser. A*, 115(772): 700–721.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *arXiv:1611.07308*.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Li, Z.; Deng, Z.; Han, Z.; Alfaro-Bittner, K.; Barzel, B.; and Boccaletti, S. 2021. Contagion in simplicial complexes. *Chaos, Solitons & Fractals*, 152: 111307.
- Lim, L.-H. 2020. Hodge Laplacians on graphs. *SIAM Review*.
- Lloyd, S.; Garnerone, S.; and Zanardi, P. 2016. Quantum algorithms for topological and geometric analysis of data. *Nature Communications*, 7(1): 1–7.
- Morris, C.; Ritzert, M.; Fey, M.; Hamilton, W. L.; Lenssen, J. E.; Rattan, G.; and Grohe, M. 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In *AAAI*, volume 33, 4602–4609.
- Nickel, M.; and Kiela, D. 2017. Poincaré embeddings for learning hierarchical representations. In *NeurIPS*, volume 30, 6338–6347.
- Roddenberry, T. M.; Glaze, N.; and Segarra, S. 2021. Principled simplicial neural networks for trajectory prediction. In *ICML*, 9020–9029.
- Rong, Y.; Huang, W.; Xu, T.; and Huang, J. 2019. DropEdge: Towards Deep Graph Convolutional Networks on Node Classification. In *ICLR*.
- Salathé, M.; Kazandjieva, M.; Lee, J. W.; Levis, P.; Feldman, M. W.; and Jones, J. H. 2010. A high-resolution human contact network for infectious disease transmission. *PNAS*, 107: 22020–22025.
- Schaub, M. T.; Benson, A. R.; Horn, P.; Lippner, G.; and Jadbabaie, A. 2020. Random walks on simplicial complexes and the normalized hodge 1-laplacian. *SIAM Review*, 62(2): 353–391.
- Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI Magazine*, 29(3): 93–93.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *ICLR*.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *CVPR*, 7794–7803.

- Wang, Y.; and Solomon, J. 2019. Intrinsic and extrinsic operators for shape analysis. In *Handbook of Numerical Analysis*, volume 20, 41–115. Elsevier.
- Wu, F.; Souza, A.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. 2019. Simplifying graph convolutional networks. In *ICML*, 6861–6871.
- Xiao, Z.; and Deng, Y. 2020. Graph embedding-based novel protein interaction prediction via higher-order graph convolutional network. *PloS One*, 15(9): e0238915.
- Yan, Z.; Ma, T.; Gao, L.; Tang, Z.; and Chen, C. 2021. Link Prediction with Persistent Homology: An Interactive View. In *ICML*, 11659–11669.
- Zhang, M.; and Chen, Y. 2018. Link prediction based on graph neural networks. In *NeurIPS*, volume 31, 5165–5175.
- Zhao, Q.; Ye, Z.; Chen, C.; and Wang, Y. 2020. Persistence enhanced graph neural network. In *AISTATS*, 2896–2906.