

SGAT: Simplicial Graph Attention Network

See Hian Lee, Feng Ji and Wee Peng Tay

Nanyang Technological University, Singapore

seehian001@e.ntu.edu.sg, {jifeng, wptay}@ntu.edu.sg

Abstract

Heterogeneous graphs have multiple node and edge types and are semantically richer than homogeneous graphs. To learn such complex semantics, many graph neural network approaches for heterogeneous graphs use metapaths to capture multi-hop interactions between nodes. Typically, features from non-target nodes are not incorporated into the learning procedure. However, there can be nonlinear, high-order interactions involving multiple nodes or edges. In this paper, we present Simplicial Graph Attention Network (SGAT), a simplicial complex approach to represent such high-order interactions by placing features from non-target nodes on the simplices. We then use attention mechanisms and upper adjacencies to generate representations. We empirically demonstrate the efficacy of our approach with node classification tasks on heterogeneous graph datasets and further show SGAT’s ability in extracting structural information by employing random node features. Numerical experiments indicate that SGAT performs better than other current state-of-the-art heterogeneous graph learning methods.

1 Introduction

Real world objects are often connected and interdependent, with graph-structured data collected in many applications. Recently, graph neural network (GNN), which learns tasks utilizing graph-structured data, has attracted much attention. Initial models such as GraphSage [Hamilton *et al.*, 2017], Graph Convolution Network (GCN) [Kipf and Welling, 2016] and Graph Attention Network (GAT) [Veličković *et al.*, 2017] focused on homogeneous graphs in which all nodes and edges are of the same type. Works to extend these models to heterogeneous graphs, which have multiple node or edge types, have then been proposed [Yun *et al.*, 2019] as the majority of real-world graphs are heterogeneous in nature. These models typically employ metapaths [Fu *et al.*, 2020].

A metapath is a predefined sequence of node types describing a multi-hop relationship between nodes [Dong *et al.*, 2017]. For instance, in a citation network with author, paper and venue node types, the metapath Author-

Paper-Author (APA) represents co-author relationship. Moreover, Author-Paper-Venue-Paper-Author (APVPA) connects two authors who published papers at the same venue. To generalize further, the concept of a metagraph is introduced [Fang *et al.*, 2019]. A metagraph is a nonlinear combination of metapaths or more specifically, a subgraph of node types defined on the network schema, describing a more elaborated pairwise relation between two nodes through auxiliary nodes. We can therefore view metapaths and metagraphs as multi-hop relations between two nodes.

Although metapath-based methods can achieve state of the art performances, they have at least one of the following limitations: (1) The model is sensitive to metapaths with sub-optimal paths leading to suboptimal results [Hussein *et al.*, 2018]. Metapath selection requires domain knowledge of the dataset and the task at hand. Likewise, metagraphs are predefined based on domain knowledge. Alternatively, heuristic algorithms can be employed to extract the most common structures to form metagraphs [Fang *et al.*, 2019]. Yet, the generated metagraphs may not necessarily be useful [Yang *et al.*, 2018]. (2) Metapath-based models do not incorporate features from non-target nodes along the metapaths, thus discarding potentially useful information [Fu *et al.*, 2020]. Examples include HERec [Shi *et al.*, 2019] and Heterogeneous graph Attention Network (HAN) [Wang *et al.*, 2019]. (3) The model relies on *structures between two nodes* to capture complex pairwise relations and semantics. This, however, is not always sufficient to capture more complex interactions in graphs which simultaneously involve multiple target nodes and cannot be reduced to pairwise relationships, especially for heterogeneous graphs given their richer semantics [Bunch *et al.*, 2020].

Several works have aimed to improve the expressive power of GNNs and generate more accurate representations of heterogeneous graph components without utilizing metapaths or metagraphs. One potential alternative is a simplicial complex. Simplicial complexes are natural extensions of graphs that can represent high-order interactions between nodes. A graph defines pairwise relationships between elements of a vertex set. Meanwhile, a simplicial complex defines higher order relations, e.g., a 3-tuple is a triangle, a 4-tuple is a tetrahedron and so on (for a proper definition, see Section 3). For instance, a triangle made up of three author vertices in a citation network indicates co-authorship of the same paper be-

tween the three authors, whereas in a graph, edges between pairs of these three authors only tell us that they are pairwise co-authors with each other on some papers. These structures are already employed in Topological Data analysis (TDA) to extract information from data and in other applications such as tumor progression analysis [Roman *et al.*, 2015] and brain network analysis [Giusti *et al.*, 2016] to represent complex interactions between elements. The GNN literature [Bunch *et al.*, 2020; Bodnar *et al.*, 2021] on simplicial complexes thus far have focused on homogeneous graphs. These works cannot be directly applied to heterogeneous graphs.

In this paper, we propose a general framework, SGAT, which extends GAT to heterogeneous graphs with simplicial complexes. SGAT learns on heterogeneous graphs using simplicial complexes to model higher order relations and passing messages between the higher-order simplices. We first describe a procedure to generate k -order homogeneous simplices from a heterogeneous graph since heterogeneous datasets do not always possess higher-order simplices, given their innate schemas. In order to avoid discarding potentially useful information when transforming the heterogeneous graph into homogeneous simplices, we populate the k -simplices, for $k \geq 1$, with non-target node features and learn the importance of each of the k -order simplicies through attention mechanisms with upper adjacencies. Overall, the contributions of this paper are as follows:

- We develop a procedure to construct a simplicial complex from a heterogeneous graph. Our proposed procedure converts the graph into a homogeneous simplicial complex without loss of feature information.
- We propose GAT-like attention mechanisms that operate on simplicial complexes. We utilize upper adjacencies to pass messages between higher-order simplices to learn effective embeddings that capture higher order interactions. We also introduce a variant model SGAT-EF that incorporates edge features.
- We apply SGAT to the node classification task on standard heterogeneous graph datasets, which demonstrate that our proposed approach outperforms current state-of-the-art models. We additionally assess the ability of our model in extracting structural information using random node features.

2 Related Work

In this section, models that are designed for heterogeneous graphs and related to SGAT are reviewed. Learning on heterogeneous graphs mainly utilizes predefined metapaths or automatically learning the optimal metapaths during the end-to-end training process [Hussein *et al.*, 2018]. Works along this line include HAN, HERec, metapath2vec [Dong *et al.*, 2017], Graph Transformer Network (GTN) [Yun *et al.*, 2019] and REGATHER [Lee *et al.*, 2021].

HAN utilizes predefined metapaths to transform heterogeneous graphs into metapath-based homogeneous graphs, discarding features of non-target nodes. After which, dual-level attention is applied to the metapath-based homogeneous graphs. On the other hand, HERec employs metapaths to generate metapath-based random walks, then filters the walks to

contain only the node type of the starting node. This avoids representing nodes of different node types in the same unified space. Similarly, metapath2vec utilizes metapaths to guide random walks to incorporate semantic relations and avoid bias towards more visible node types. Given the difficulty in selecting optimal metapaths, GTN proposes to learn metapath-based graph structures by softly selecting candidate adjacencies and multiplying them together. The resulting weighted adjacency matrices represent metapath-based graph structures. Meanwhile, REGATHER generates a set of multi-hop relation-type subgraphs, which indirectly includes metapath-based graph structures by multiplying adjacency matrices of first order relation-types during data preprocessing. Attention mechanism is then utilized to assign different weights to the subgraphs. All of these approaches do not learn high-order relations in the TDA context [Fang *et al.*, 2019].

Metagraphs have been proposed to express finer-grained, non-linear semantics. Approaches that utilize metagraphs are M-HIN [Fang *et al.*, 2019], mg2vec [Zhang *et al.*, 2020] and Meta-GNN [Sankar *et al.*, 2019]. M-HIN constructs triplets to illustrate the relationship between nodes and metagraphs and subsequently, applies the Hadamard function in complex space to encode the relationship between nodes and metagraphs. By utilizing metagraphs and a complex embedding scheme, M-HIN is able to capture more accurate node features. As for mg2vec, embeddings of metagraphs and nodes are simultaneously learned and mapped into a common low-dimensional space. Metagraphs are utilized to guide the learning of node representations and to capture the latent relationships between nodes. Lastly, Meta-GNN introduces a metagraph convolution layer, employing metagraphs to define the receptive field of nodes. Attention mechanism is also used to combine node features from different metagraphs.

Homogeneous GNN approaches involving simplicial complexes include Message Passing Simplicial Network (MPSN) [Bodnar *et al.*, 2021] and Simplicial Neural Network (SNN) [Ebli *et al.*, 2020]. MPSN introduces a general message passing framework on simplicial complexes, describing four different adjacencies that simplices can have. Besides that, they also introduced the Simplicial Weisfeiler-Lehman (SWL) test and showed that the SWL test is strictly more powerful than the WL test [Douglas, 2011] in distinguishing non-isomorphic graphs. Since standard GNNs with local neighborhood aggregation are shown to be equivalent to the WL test in their expressive powers, complementing GNNs with simplicial complexes can help increase their expressive powers. Nevertheless, MPSN is specific to homogeneous graphs and the application of simplicial complexes to heterogeneous graphs is unexplored by [Bodnar *et al.*, 2021].

SNN is a framework utilizing simplicial complexes on heterogeneous bipartite graphs. Heterogeneous bipartite graphs are transformed into a homogeneous complex with corresponding co-chains' data. The co-chain features are then refined by spectral simplicial convolution with Hodge Laplacians [Barbarossa and Sardellitti, 2020] of the simplicial complex. Although SNN also populates features on high-order structures, it is limited to bipartite graphs and its objective is to impute missing data. In addition, SNN operates on the spectral domain, which is different from our proposed ap-

proach that operates on the graph domain directly. To the best of our knowledge, our proposed framework is the first study to learn heterogeneous graphs (which are not limited to bipartite graphs) with simplicial complexes.

3 Simplicial Graph Attention Network

We first introduce the concept of a simplicial complex, its adjacencies, and related notations. The reader is referred to the Appendix E and [Hatcher, 2000] for further details.

Simplicial Complexes. A simplicial complex χ is a collection of finite sets (simplices) closed under subsets. A member of χ is called a k -simplex and denoted as σ^k if it has cardinality $k + 1$. A k -simplex has $k + 1$ faces of dimension $k - 1$. A face is obtained by omitting one element from the simplex. Specifically, by omitting the j -th element, a face of a k -simplex is a set containing elements of the form $(v_0, \dots, v_{j-1}, v_{j+1}, \dots, v_k)$. Subsequently, it is possible to think of 0-simplices σ^0 as vertices, 1-simplices σ^1 as edges, 2-simplices σ^2 as triangles, 3-simplices σ^3 as tetrahedron and so on. Lastly, the dimension of χ is the largest dimension of any simplex in it and is denoted by K . The set of all simplices in χ of dimension k is denoted by χ^k with cardinality $|\chi^k|$.

Simplicial Adjacencies. In this paper, we utilize the notion of upper-adjacency to denote the neighborhood of simplices. Two k -simplices in χ^k are upper-adjacent if they are faces of the same $(k + 1)$ -simplex. For instance, two edges are upper-adjacent if they are part of a same triangle. Hence, the neighborhood of a k -simplex, $N(\sigma_i^k)$ is the set of k -simplices upper-adjacent to it. We include self-loops to the upper adjacency matrix so that a k -simplex is upper-adjacent to itself. We have

$$N(\sigma_i^k) = \left\{ \sigma_j^k \in \left\{ \sigma_1^k \dots \sigma_{|\chi^k|}^k \right\} : A^k(i, j) = 1 \right\}, \quad (1)$$

where $A^k \in \mathbb{R}^{|\chi^k| \times |\chi^k|}$ denotes the upper adjacency matrix indicating whether pairs of k -simplices are upper-adjacent.

3.1 Construction of Simplices

Many real-world heterogeneous graph datasets do not have predefined higher-order k -simplices for $k \geq 2$ due to their unique schemas. We propose a procedure to transform heterogeneous graphs into homogeneous simplices where the 0-simplices are the singleton subsets of the target node type’s vertices. The target node type is the node type whose labels we want to predict (node classification). Meanwhile, non-target nodes features reside on higher order simplices.

Consider a heterogeneous graph $G = (V, E)$ with a edge-types and b node types, where $a + b > 2$. The set of nodes is V and the set of edges is given by E . Let $\tilde{V} \subset V$ be the set of target nodes. We assume that each $v \in \tilde{V}$ has a feature vector h_v . To simplify the presentation, we assume that every non-target node $u \in V \setminus \tilde{V}$ is also associated with a feature vector h_u (if a non-target node does not come with a feature vector, we can associate with it a one-hot encoding of its node type).

For $1 \leq k \leq K$, where K is a hyperparameter, we choose the k -simplices to be the sets of $k + 1$ target vertices that share at least ϵ_η^k common non-target neighbors which are exactly η hops away, where $\eta > 0$ is a chosen parameter. This

is performed for each $\eta \leq \eta_{\max}$, where η_{\max} is a hyperparameter, to generate η number of K dimensional simplicial complexes. Consider a k -simplex σ^k , where $k > 1$, whose vertices share $m \geq \epsilon_\eta^k$ common η -hop non-target neighbors, s_1, \dots, s_m . The k -simplex is assigned the feature $h_{\sigma^k} = \phi(h_{s_1}, \dots, h_{s_m})$ where ϕ is the average operator. For 1-simplices, the features of the intermediate node(s) along a path between the 0-simplices are first summed to form the path feature, $\Theta_{\sigma^1} = \sum_{u \in \theta} h_u$ where θ is the set of intermediate nodes along the path connecting the faces of σ^1 . The path features between the same pair of 0-simplices are then averaged and assigned as the feature of the 1-simplex. In addition, we place the simplex’s own feature on the connecting simplex when a self-loop is present.

Taking the IMDB¹ dataset as an example in Fig. 1, we set $\eta = 1$ and $\epsilon_1^1 = 1$. The constructed simplicial complex is illustrated in Fig. 1. We note that the set of 1-simplices is different from the original edges in the graph.




Figure 1: Constructing simplicial complexes from the heterogeneous IMDB graph (a)-(c) with $\eta = 1$. (a) IMDB node types. (b) IMDB graph illustration. (c) Resulting movie one-hop sharing complex with $\epsilon_1^1 = 1$.

The technical details to construct all the k -simplices of a single simplicial complex within a network is given in Algorithm 1 in Appendix B

Incorporation of Edge Features

In the case where each edge $(i, j) \in E$ has a feature vector e_{ij} , we concatenate the feature generated as described above for a 1-simplex $h_{\sigma^1} = \phi(\Theta_{\sigma^1, 1}, \dots, \Theta_{\sigma^1, t})$ where t is the number of paths between the faces of σ^1 , with the average of the edge features along the original path in the graph that gives rise to σ^1 . For σ^1 with end nodes \tilde{v} and v , we have

$$h_{\sigma^1} = \phi(\Theta_{\sigma^1, 1}, \dots, \Theta_{\sigma^1, t}) \parallel \phi(\Phi_1, \dots, \Phi_m), \quad (2)$$

$$\Phi_i = \phi(e_{\tilde{v}s_{i,1}}, \dots, e_{s_{i,j}s_{i,j+1}}, \dots, e_{s_{i,\kappa}v}), \quad (3)$$

where κ is the number of intermediate nodes between the end nodes and Φ_i is the summarized edge feature along each of the original m paths. We note that many heterogeneous GNNs such as GTN and HAN are not designed to incorporate edge features during their learning process.

When $n > k$ target nodes share more than ϵ_η^k non-target nodes among them, a total of $\binom{n}{k+1}$ k -simplices are constructed, which may cause memory overloading issues for large values of n . To further control the the amount of constructed simplices, a hyperparameter $\lambda > K$ is introduced.

¹<https://www.imdb.com/interfaces/>




Figure 2: Combining edge features with features of 1-simplices to enrich features of 1-simplices for $\eta = 1$.

We do not construct the simplices (and their corresponding faces) if $n \geq \lambda$. This can also be interpreted as disregarding the k -simplices (and their associated faces) where $k \geq \lambda$.

3.2 Simplicial Attention Layer

After constructing the k -simplices, $0 \leq k \leq K$, from the original graph, the inputs to the SGAT model are $K + 1$ different sets of features, $\{H_{\chi^0}, H_{\chi^1}, \dots, H_{\chi^K}\}$, where $H_{\chi^k} = \{h_{\sigma_1^k}, \dots, h_{\sigma_{|\chi^k|}^k}\}$, $h_{\sigma_i^k}$ is the feature vector of the i -th k -simplex, $i = 1, \dots, |\chi^k|$. Moreover, each of the $K + 1$ feature sets is associated with an upper adjacency matrix, A^k .

We employ a simplicial graph attention layer to compute the importance of the neighboring simplices and update the various simplices' features. A simplicial graph attention layer is made up of K graph attention layers, one for each of the k -order simplices, $0 \leq k < K$, excluding the K -order. This captures the interactions between higher-order structures of the same order by passing messages between themselves through upper adjacencies. For each of the K graph attention layers, given a k -simplex pair σ_i^k, σ_j^k , the unnormalized attention score between the two simplicies is defined as

$$\zeta_{ij}^k = \text{LeakyReLU}\left(\left(\mathbf{a}^k\right)^\top \left[\mathbf{W}^k h_{\sigma_i^k} \parallel \mathbf{W}^k h_{\sigma_j^k} \parallel \mathbf{W}^{k+1} h_{\sigma_{ij}^{k+1}}\right]\right), \quad (4)$$

where \parallel represents concatenation, \top denotes transposition, \mathbf{a}^k is the learnable attention vector and \mathbf{W}^k is the weight matrix for a linear transformation specific to the k -order and are shared by all the k -simplices. The feature vector $h_{\sigma_{ij}^{k+1}}$ belongs to the $k + 1$ -simplex that both σ_i^k and σ_j^k are faces of.

The normalized attention score for each of the k -simplices can then be obtained by applying the softmax function as given by:

$$\alpha_{ij}^k = \frac{\exp(\zeta_{ij}^k)}{\sum_{r \in N(\sigma_i^k)} \exp(\zeta_{ir}^k)}, \quad (5)$$

where $N(\sigma_i^k)$ refers to the neighborhood of σ_i^k (cf. (1)). The weight coefficients learned can then be applied along with the corresponding neighbors' features to update the feature for each of the k -simplices as given by

$$z_{\sigma_i^k} = \mu\left(\sum_{j \in N(\sigma_i^k)} \alpha_{ij}^k \mathbf{W}^k h_{\sigma_j^k}\right), \quad (6)$$

where $\mu(\cdot)$ denotes an activation function. To stabilise the training, we employ multi-head attention where P independent attention mechanisms are performed. The results are then concatenated to generate the required learned simplex features:

$$z_{\sigma_i^k} = \parallel_{p=1}^P \mu\left(\sum_{j \in N(\sigma_i^k)} \alpha_{ij}^{k,p} \mathbf{W}^{k,p} h_{\sigma_j^k}\right), \quad (7)$$

where $\alpha_{ij}^{k,p}$ and $\mathbf{W}^{k,p}$ are the attention and transformation weights corresponding to α_{ij}^k and \mathbf{W}^k , respectively, in (6) for the p -th attention head.

The output for the k -order simplices χ^k is a set of simplicial complex specific embeddings, $Z_{\chi^k} = \{z_{\sigma_1^k}, \dots, z_{\sigma_{|\chi^k|}^k}\}$. Since we have K graph attention layers, there are K sets of simplicial complex specific embeddings $Z_{\chi^0}, \dots, Z_{\chi^{K-1}}$. Depending on the downstream task, not all of the embeddings are utilized. For instance, in the last softmax layer for node classification tasks, only the Z_{χ^0} embeddings are employed.

3.3 Attention to Fuse Simplicial Complexes

Recall that for $\eta_{\max} > 1$ and for each $\eta = 1, \dots, \eta_{\max}$, we construct k -simplices for $0 \leq k \leq K$. Applying the simplicial attention layer in Section 3.2 to each of the η number of K dimensional simplicial complexes, we obtain the embeddings $Z_{\chi^{0,\eta}}, \dots, Z_{\chi^{K-1,\eta}}$ for each η -specific simplicial complex. In our framework, a k -simplex σ^k is a unique k -tuple that is independent of η . We arbitrarily order all k -simplices as $\sigma_1^k, \sigma_2^k, \dots, \sigma_\tau^k$, where $\tau = \max_\eta |\chi^{k,\eta}|$. The collection of embeddings for $\chi^{k,\eta}$ can then be written as $Z_{\chi^{k,\eta}} = \{z_{\sigma_1^k}^\eta, \dots, z_{\sigma_\tau^k}^\eta\}$ where $z_{\sigma_i^k}^\eta$ is the feature vector of the i -th k -simplex for the η -specific simplicial complex, which is taken to be null if $\sigma_i^k \notin \chi^{k,\eta}$.

In this subsection, we propose an attention layer to account for the importance of different simplicial complexes in describing the respective k -dimensional simplices. An attention layer is required to fuse the η groups of k -simplex embeddings.

The unnormalized attention score to fuse the η -specific embeddings of k -simplices is defined as follows:

$$w^{k,\eta} = \frac{1}{|\chi^{k,\eta}|} \sum_{i=1}^{\tau} (\mathbf{q}^{k,\eta})^\top \tanh(\mathbf{F}^{k,\eta} z_{\sigma_i^k}^\eta + \mathbf{b}^{k,\eta}) \mathbf{1}\{\sigma_i^k \in \chi^{k,\eta}\}, \quad (8)$$

where $\mathbf{q}^{k,\eta}$ is the learnable attention vector, $\mathbf{F}^{k,\eta}$ is a learnable weight matrix and $\mathbf{b}^{k,\eta}$ is a learnable bias vector. Here, $\mathbf{1}\{A\}$ is the indicator function, which takes value 1 if A is true and 0 otherwise. The normalized weight coefficients can then be attained through the softmax function as

$$\beta^{k,\eta} = \frac{\exp(w^{k,\eta})}{\sum_{j=1}^{\eta_{\max}} \exp(w^{k,j})}. \quad (9)$$

We fuse the features $z_{\sigma_i^k}^\eta$ across $\eta = 1, \dots, \eta_{\max}$ by a linear combination using the learned weights to generate the final embedding of σ_i^k as

$$z_{\sigma_i^k} = \sum_{\eta=1}^{\eta_{\max}} \beta^{k,\eta} z_{\sigma_i^k}^\eta \mathbf{1}\{\sigma_i^k \in \chi^{k,\eta}\}. \quad (10)$$




Figure 3: The overall architecture of SGAT (self-loops are omitted for clarity). (1) Given the heterogeneous graph G and node features X as inputs, we first construct K simplices for each of the η simplicial complex(es) and the simplices' respective features. For each simplicial complex χ^i , (2) its associated simplices, features and upper adjacency matrices are fed into a simplicial attention layer which is made up of K graph attention layers. Then, (3) η sets of simplicial complex specific embeddings are generated. For node classification task, (4) we fuse the η sets of 0-simplices' embeddings using attention to obtain the final embedding which is then fed into a layer of MLP (the classifier).

Let $Z_{\chi^k} = \{z_{\sigma_1^k}, \dots, z_{\sigma_\tau^k}\}$ be the embedding of $\chi^k = \{\sigma_1^k, \dots, \sigma_\tau^k\}$.

Lastly, when L layers are utilized, let $Z_{\chi^0}^{(l)}$ be the embedding of χ^0 in the l -th layer, for $l = 1, \dots, L$. We concatenate the learned 0-simplex features of each layer at the last layer to obtain

$$Z_{\chi^0}^{(L)} = \left\| \right\|_{l=1}^L Z_{\chi^0}^{(l)}, \quad (11)$$

and feed it into a linear layer for our semi-supervised, node classification task along with cross-entropy loss. Nonetheless, we note that the final learned features, $Z_{\chi^0}^{(L)}$ can be used for other downstream tasks similarly optimising the model via backpropagation. The overall architecture of SGAT is as shown in Fig. 3.

4 Numerical Experiments

In this section, we verify the empirical performance of our proposed method against several state-of-the-art methods on node classification task for heterogenous graphs.

4.1 Datasets

The heterogeneous datasets utilized are two citation network datasets DBLP² and ACM, and a movie dataset IMDB. Since these heterogeneous datasets do not consist of edge features, when edge features are required, we form h_e for each $e \in E$ of the respective datasets by concatenating its starting node feature, ending node feature and a one-hot encoding of its edge-type. The dataset statistics are in Appendix A.

4.2 Baselines and Settings

We compare SGAT and SGAT-EF with six state-of-the-art GNN models. For all models, the hidden units are set to 64, the Adam optimizer was used and its hyperparameters such as learning rate and weight decay, are respectively chosen to yield best performance. When metapaths are required, the metapaths employed in [Wang *et al.*, 2019] are utilized. As for metagraphs, the relevant metapaths and 3-node metagraphs as in [Sankar *et al.*, 2019] are used.

²<https://dblp.uni-trier.de/>

For SGAT, we set K , the dimension of the simplicial complexes to be 2, the number of layers to be 2 for all the datasets. Moreover, when $\eta \geq 2$, the dimension of the attention vector $\mathbf{q}^{k,\eta}$ (cf. (8)) is set to 128. Besides the parameters mentioned above, ϵ_η^k , η and λ are tuned for each dataset. Specifically, for ACM, we choose $\eta = 1$, $\epsilon_1^1 = 1$, and $\lambda = 20$. For DBLP, $\eta = 2$, $\epsilon_1^1 = 3$, $\epsilon_2^1 = 4$, and $\lambda = 10$. For IMDB, $\eta = 1$, $\epsilon_1^1 = 1$, and $\lambda = 10$.

4.3 Node Classification Performance

Table 1 shows the performances of SGAT, SGAT-EF and other node classification baselines. Our empirical results demonstrate that SGAT and SGAT-EF outperform the baselines for node classification task. We observe that SGAT performs better than GTN, HAN and REGATHER. This demonstrates that nonlinear structures encoded by the simplicial complexes are useful for learning more effective node representations and more complicated relationships, which may not be possible using linear structures such as metapaths.

5 Model Analysis

5.1 Ablation Study

We conduct experiments on different SGAT variants to validate the effectiveness of the components of our model. SGAT considers message passing between higher-order simplices without incorporating edge features while SGAT-EF is the equivalent model utilizing edge features. Moreover, since SGAT generalizes to the GAT model when $K = 1$, the GAT model can also be taken as an ablation study of SGAT without message passing on higher order structures. As observed in Table 1, by employing message passing on higher-order simplices, SGAT obtained a significant improvement over GAT. We also observe that SGAT-EF frequently performs better than SGAT. This show that our method incorporating edge features may bring benefits when the task of interest is edge dependent or when the dataset has meaningful edge features.

A parameter sensitivity study (for λ and ϵ_η^k) is given in Appendix D

5.2 Meta-GNN, SGAT and SGAT-EF

To evaluate the advantage of utilising simplicial complexes over metagraphs, we conduct experiments to compare Meta-GNN, SGAT and SGAT-EF. The results are shown in Fig. 4.

Datasets	Metrics	GCN	GAT	Meta-GNN	HAN	REGATHER	GTN	SGAT	SGAT-EF
DBLP	Macro-F1	87.65 ± 0.29	91.69 ± 0.27	92.47 ± 0.92	91.93 ± 0.27	91.79 ± 0.69	93.59 ± 0.40	93.80 ± 0.20	93.73 ± 0.23
	Micro-F1	88.71 ± 2.74	92.65 ± 0.25	93.48 ± 0.75	92.51 ± 0.24	92.70 ± 0.67	94.17 ± 0.26	94.58 ± 0.20	94.51 ± 0.19
ACM	Macro-F1	91.46 ± 0.48	92.16 ± 0.32	88.74 ± 0.99	91.01 ± 0.76	92.38 ± 0.57	92.23 ± 0.60	92.41 ± 0.36	92.91 ± 1.79
	Micro-F1	91.33 ± 0.47	92.06 ± 0.33	88.71 ± 1.01	90.93 ± 0.73	92.30 ± 0.56	92.12 ± 0.62	92.35 ± 0.36	92.86 ± 1.75
IMDB	Macro-F1	56.72 ± 0.49	57.32 ± 0.88	56.19 ± 0.97	56.56 ± 0.77	56.34 ± 0.54	59.12 ± 1.58	59.97 ± 0.41	60.36 ± 0.53
	Micro-F1	58.31 ± 0.51	58.75 ± 0.98	58.68 ± 1.49	57.83 ± 0.93	57.59 ± 0.64	60.58 ± 2.10	62.51 ± 0.64	62.74 ± 0.77

Table 1: Node classification result on heterogeneous datasets (Standard split). Averaged over five runs, best performance boldfaced.

We can clearly observe that SGAT and SGAT-EF, which performs message passing between simplices, consistently outperforms Meta-GNN. Meta-GNN uses metagraphs, which despite being a subgraph pattern, solely represents relationships between *two* target nodes and not *multiple* target nodes as in the case of simplices. The result demonstrates that metagraphs cannot adequately represent high-order interactions between multiple target nodes.




Figure 4: Node classification F1 scores (%) for Meta-GNN, SGAT and SGAT-EF

5.3 Random Node Features

A survey [Borgwardt *et al.*, 2020] has shown that node features of benchmark graphs usually contain substantial information, allowing models to achieve close to optimal results. Hence, we prepare an additional dataset to classify nodes solely based upon graph structure. We replace node features of IMDB dataset by random ones sampled from standard normal distribution, naming it IMDB-RNF(N). The random node features are the same size as the replaced node features. This leaves only the graph structure for the classification task and thus, we can utilise the resulting performance to interpret the models’ capability in extracting structural information. This approach is also used in [Horn *et al.*, 2022]. Note that for IMDB-RNF(N) dataset, the models are observed to be highly sensitive, thus we re-tuned all the models to ensure a fair comparison. Table 2 depicts the results on such dataset.

We observe clear advantage of SGAT and SGAT-EF over their comparison partners. For instance, SGAT and SGAT-EF surpass the baselines by up to approximately 10% F1 scores on the IMDB-RNF(N) dataset. In addition, we find that the performance of all models decreased when the bag-of-words node features are replaced with uninformative ones. This shows that the original node features indeed contain useful information for the task, aiding all models to garner almost optimal results. Subsequently, it also implies that our models are able to effectively extract and utilise structural infor-

Methods	IMDB-RNF(N)	
	Macro-F1	Micro-F1
GAT	36.45 ± 2.03	37.94 ± 2.07
GCN	40.10 ± 1.09	41.84 ± 1.24
Meta-GNN	32.58 ± 1.62	43.75 ± 5.62
HAN	39.47 ± 1.03	40.97 ± 1.13
REGATHER	28.60 ± 3.45	42.62 ± 8.10
GTN	33.71 ± 0.51	36.52 ± 1.38
SGAT	49.29 ± 1.34	52.34 ± 1.81
SGAT-EF	48.90 ± 1.00	52.37 ± 1.09

Table 2: Node classification result on heterogeneous dataset with random node features, best performance boldfaced.

mation from graphs instead of relying on the original node features that may already contain important information.

6 Conclusion

In this paper, we introduced SGAT, a novel extension of GAT to heterogeneous graphs using simplicial complexes and upper adjacencies. SGAT does not utilize metapaths or metagraphs. Hence, SGAT is not subjected to the three characteristic limitations of metapath-based methods, namely (1) performance being sensitive to the choice of metapaths, (2) discarding non-target node features and (3) limited expressiveness when using structures between two nodes to capture complex interactions. Specifically, we utilize simplicial complex to capture higher order relations among multiple target nodes. We avoid discarding non-target node features when transforming the graph into homogeneous simplices by placing those features on simplices. We also introduced a variant incorporating edge features that can boost embedding performance. Empirically, we demonstrated that SGAT performs favorably on node classification task for heterogeneous datasets and even when random node features was employed.

A Dataset Statistics

We perform node classification task on three heterogeneous benchmark datasets. Characteristics of the datasets are summarised in Table 3. DBLP has three node types (author(A), paper(P) and conference(C)), and the research area of author serve as labels. ACM consists of three node types (paper(P), author(A) and subject(S)), and categories of papers are to be predicted. IMDB has three node types (movie(M), actor(A), and director(D)). The labels to be determined are the genres

of the movies. Each node in DBLP, ACM and IMDB has an associated, bag-of-words node feature.

Dataset	# Nodes	# Edges	# Node type	# Classes	# Features
DBLP	18405	67946	3	4	334
ACM	8994	25922	3	3	1902
IMDB	12772	37288	3	3	1256

Table 3: Summary of datasets

B Pseudocode

For simplicity, we assume that the ϵ_η^k is the same for all k -order simplices and the specified η . The Geometric Understanding in Higher Dimensions library [GUDHI Project, 2015] for computational topology is utilised to ensure the constructed simplices form simplicial complex(es) that adhere to the inclusion property.

C Relation to Hypergraphs

In this paper, the primary geometric concept is that of simplicial complexes, as a generalization of graphs. A simplicial complex is a special class of hypergraph. We choose to work exclusively with simplicial complexes instead hypergraphs for a few reasons. A simplicial complex has rich algebraic structure and plays a central role in algebraic topology. Many tools have been developed to study simplicial complexes which are unavailable for general hypergraphs, such as the Hodge-de Rham theory. Moreover, for any given hypergraph, we may construct a simplicial complex by treating each hyperedge as a simplex and including hyperedges associated with its faces. The resulting simplicial complex has many important geometric properties identical to the initial hypergraph. In this respect, using simplicial complexes is sufficient to extract useful geometric information from the datasets.

D Parameter Sensitivity

Given that SGAT involves a number of parameters to control the construction of simplices, we further examine how λ (the maximum simplex order to construct, including their faces. The simplex order refers to its dimension) and ϵ_η^k (the minimum number of shared non-target nodes that is k and η -specific) affect the performance of SGAT on the DBLP dataset. We first keep all the constructed edges by setting $\epsilon_1^1 = \epsilon_2^1 = 1$ and measure the ratio of constructed triangles to constructed edges, γ as a function of λ . The Micro-F1 and Macro-F1 scores are then obtained as a function of γ as seen in Fig. 5. We observe that an optimal γ where the information included from the amount of simplices considered is most beneficial and does not negatively affect the model’s performance. Increasing λ beyond the resulting optimal γ is likely to introduce noise by including unnecessary messages being passed between the simplices.

We also analyse how ϵ_η^k affect SGAT’s performance by setting λ to be fixed at 10 and the default values of $\epsilon_1^1 = 3$ and

Algorithm 1 Construction of k -simplices

Input: The adjacency list of heterogeneous graph Adj_list, Node features X ,

Number of shared non-target neighbours ϵ ,

Number of hops away η ,

Maximal k -order considered, K ,

The maximum simplex order to construct λ .

Output: Set of all k -simplices, AllKSimplices.

- 1: \triangleright DFSFindPaths is a modified depth first search that breaks at 2η depth. To get paths that starts and ends with target node type. Each path should contain at least one non-target node.
- 2: \triangleright Each path is of form [target_src, node_k(s), target_dst], where node_k is in θ and θ is the set of intermediate nodes along the path.
- 3: Initialise PathList as empty list.
- 4: **for** $v \in V$ **do**
- 5: PathList \leftarrow DFSFindPaths(Adj_list, v , 2η)
- 6: \triangleright The middle node refers to the non-target node η hops away from the target_src and target_dst.
- 7: Initialise MidNodeNeighborDict, D_{mid}
- 8: \triangleright where key is MidNodeID and value is its set of unique Neighbors.
- 9: **for** path \in PathList **do**
- 10: D_{mid} [path.MidNodeID].insert(path.src)
- 11: D_{mid} [path.MidNodeID].insert(path.dst)
- 12: Initialise KSimplicesDict
- 13: \triangleright where key is the Neighbors and value is the list of MidNodeID
- 14: **for** MidNodeID, Neighbors $\in D_{mid}$ **do**
- 15: **if** $2 \leq \text{size}(\text{Neighbors}) \leq \lambda$ **then**
- 16: KSimplicesDict[Neighbors].insert(MidNodeID)
- 17: **for** Neighbors, MidNodeIDList in KSimplicesDict **do**
- 18: **if** $\text{size}(\text{MidNodeIDList}) \geq \epsilon$ **then**
- 19: \triangleright Simplex tree is a data structure in Gudhi library
- 20: SimplexTree.insert(Neighbors)
- 21: \triangleright Get up to K simplices from SimplexTree
- 22: AllKSimplices \leftarrow SimplexTree.get_simplices()
- 23: **return** AllKSimplices




Figure 5: λ against γ (left) and γ against performance score (right)

$\epsilon_2^1 = 4$. Besides the parameter being tested, the other parameters assume their default values. When ϵ_1^1 and ϵ_2^1 are increased, the number of constructed edges decreases. We observe that increasing ϵ_1^1 improves the performance until the

default value of $\epsilon_1^1 = 3$ is reached. After which, the performance deteriorates. This phenomenon is similar for ϵ_2^1 . Increasing ϵ_η^k reduces noise in the model as it removes weaker edges (those that shared fewer nodes).




Figure 6: ϵ_1^1 against γ (left) and ϵ_2^1 against γ (right)




Figure 7: γ against performance score

Lastly, we examine how different combinations of ϵ_1^1 and ϵ_2^1 , each chosen in the range of $[1, 5]$, affect SGAT’s performance. Interestingly, we found that the performance of SGAT fluctuates minimally across different values of these parameters when the ratio of triangles to edges, γ is similar.

E Preliminaries

In this section, we give a brief overview of simplicial complexes. We refer interested readers to [Hatcher, 2000] for more details. Here, we discuss some properties of simplicial complexes and the notion of upper adjacency, which we use in our proposed SGAT model.

E.1 Simplicial Complex

A simplicial complex is a set consisting of vertices, edges and other higher order counterparts, all of which are known as k -simplices, $k \geq 0$, defined as follows.

Definition 1 (k -simplex). A standard (unit) k -simplex is defined as

$$\sigma^k = \left\{ (\tilde{v}_0, \dots, \tilde{v}_k) \in \mathbb{R}_{\geq 0}^{k+1} : \sum_{i=0}^k \tilde{v}_i = 1 \right\}. \quad (12)$$

Any topological space that is homeomorphic to the standard k -simplex and thus, share the same topological characteristics is called a k -simplex.

As an example, let v_0, \dots, v_k be vertices of a graph embedded in a Euclidean space so that they are affinely independent (i.e., $v_1 - v_0, \dots, v_k - v_0$ are linearly independent). Then the set

$$\left\{ \sum_{i=0}^k \alpha_i v_i : \alpha_i \geq 0, \sum_{i=0}^k \alpha_i = 1 \right\}$$

is a k -simplex. In particular, the vertices $\{v_i\}$ are instances of σ^0 while the edges (v_i, v_j) , $i \neq j$, are instances of σ^1 . Subsequently, it is possible to geometrically interpret 0-simplices σ^0 as vertices, 1-simplices σ^1 as edges, 2-simplices σ^2 as triangles, 3-simplices σ^3 as tetrahedron and so on.

A face of a k -simplex is a $(k-1)$ -simplex that we obtain from (12) by restricting a fixed coordinate in $(\tilde{v}_0, \dots, \tilde{v}_k)$ to be zero. Specifically, a face of a k -simplex is a set containing elements of the form $(\tilde{v}_0, \dots, \tilde{v}_{j-1}, \tilde{v}_{j+1}, \dots, \tilde{v}_k)$. A k -simplex has $k+1$ faces.

A simplicial complex is a class of topological spaces that encodes higher-order relationships between vertices. The formal definition is as below.

Definition 2 (Simplicial complex). A simplicial complex χ is a finite set of simplices such that the following holds.

- Every face of a simplex in χ is also in χ .
- The non-empty intersection of any two simplices σ_1, σ_2 in χ is a face of both σ_1 and σ_2 .

It is also possible to produce a concrete geometric object for each simplicial complex χ . The geometric realisation of a simplicial complex is a topological space formed by glueing simplices together along their common faces [Ji *et al.*, 2022]. For instance, a simplicial complex of dimension 1, consists of two kinds of simplices σ^0 and σ^1 . By glueing σ^1 with common σ^0 , we obtain a graph in the usual sense. This means that χ^0 is the set of vertices V in the graph and χ^1 is the set of edges E in the graph.

E.2 Simplicial Adjacencies

Four kinds of adjacencies can be identified between simplices. They are boundary adjacencies, co-boundary adjacencies, lower-adjacencies and upper-adjacencies [Barbarossa and Sardellitti, 2020]. In this paper, we utilize only upper-adjacency so that our SGAT model is equivalent to GAT when $K=1$. We say that two k -simplices in χ^k are upper-adjacent if they are faces of the same $(k+1)$ -simplex. The upper adjacency matrix $A^k \in \mathbb{R}^{|\chi^k| \times |\chi^k|}$ of $\chi^k = \{\sigma_1^k \dots \sigma_{|\chi^k|}^k\}$ indicates whether pairs of k -simplices are upper-adjacent. The upper adjacency matrix of 0-simplices (nodes), A^0 , is the usual adjacency matrix of a graph.

Acknowledgements

The first author is supported by Shopee Singapore Private Limited under the Economic Development Board Industrial Postgraduate Programme (EDB IPP). The programme is a collaboration between Shopee and Nanyang Technological University, Singapore. The last two authors are supported in part by the Singapore Ministry of Education Academic Research Fund Tier 2 grant MOE-T2EP20220-0002.

References

- [Barbarossa and Sardellitti, 2020] Sergio Barbarossa and Stefania Sardellitti. Topological signal processing over simplicial complexes. *IEEE Trans. on Signal Processing*, PP:1–1, 03 2020.
- [Bodnar et al., 2021] Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lió, and Michael Bronstein. Weisfeiler and lehman go topological: Message passing simplicial networks. In *Proc. of the 38th International Conference on Machine Learning (ICML)*, pages 1026–1037, 2021.
- [Borgwardt et al., 2020] Karsten Borgwardt, Elisabetta Ghisu, Felipe Llinares-López, Leslie O’Bray, and Bastian Rieck. *Graph Kernels: State-of-the-Art and Future Challenges*, volume 13. Foundations and Trends® in Machine Learning, 2020.
- [Bunch et al., 2020] Eric Bunch, Qian You, Glenn Fung, and Vikas Singh. Simplicial 2-complex convolutional neural networks. *NeurIPS 2020 Workshop on Topological Data Analysis and Beyond*, 2020.
- [Dong et al., 2017] Yuxiao Dong, Nitesh V. Chawla, and Ananthram Swami. Metapath2vec: Scalable representation learning for heterogeneous networks. In *Proc. of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 135–144, 2017.
- [Douglas, 2011] Brendan L. Douglas. The weisfeiler-lehman method and graph isomorphism testing. *arXiv preprint*, 2011. arXiv:1101.5211.
- [Ebli et al., 2020] Stefania Ebli, Michaël Defferrard, and Gard Spreemann. Simplicial neural networks. *NeurIPS Workshop in Topological Data Analysis and Beyond*, 2020.
- [Fang et al., 2019] Yang Fang, Xiang Zhao, Peixin Huang, Weidong Xiao, and Maarten de Rijke. M-hin: Complex embeddings for heterogeneous information networks via metagraphs. In *Proc. of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR’19, page 913–916, 2019.
- [Fu et al., 2020] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. MAGNN: Metapath aggregated graph neural network for heterogeneous graph embedding. In *Proc. of The Web Conference 2020*, page 2331–2341, 2020.
- [Giusti et al., 2016] Chad Giusti, Robert Ghrist, and Danielle S. Bassett. Two’s company, three (or more) is a simplex. *Journal of Computational Neuroscience*, 41(1):1–14, 2016.
- [GUDHI Project, 2015] GUDHI Project. *GUDHI User and Reference Manual*. GUDHI Editorial Board, 2015.
- [Hamilton et al., 2017] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proc. of the 31st International Conference on Neural Information Processing Systems*, page 1025–1035, 2017.
- [Hatcher, 2000] Allen Hatcher. *Algebraic topology*. Cambridge Univ. Press, Cambridge, 2000.
- [Horn et al., 2022] Max Horn, Edward De Brouwer, Michael Moor, Yves Moreau, Bastian Rieck, and Karsten Borgwardt. Topological graph neural networks. *International Conference on Learning Representations*, 2022.
- [Hussein et al., 2018] Rana Hussein, Dingqi Yang, and Philippe Cudré-Mauroux. Are meta-paths necessary? revisiting heterogeneous graph embeddings. In *Proc. of the 27th ACM International Conference on Information and Knowledge Management*, page 437–446, 2018.
- [Ji et al., 2022] Feng Ji, Giacomo Kahn, and Wee Peng Tay. Signal processing on simplicial complexes with vertex signals. *IEEE Access*, pages 1–1, 2022.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2016.
- [Lee et al., 2021] See Hian Lee, Feng Ji, and Wee Peng Tay. Learning on heterogeneous graphs using high-order relations. *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3175–3179, 2021.
- [Roman et al., 2015] Theodore Roman, Amir Nayyeri, Brittany Fasy, and Russell Schwartz. A simplicial complex-based approach to unmixing tumor progression data. *BMC bioinformatics*, 16:254, 12 2015.
- [Sankar et al., 2019] Aravind Sankar, Xinyang Zhang, and Kevin Chen-Chuan Chang. Meta-gnn: Metagraph neural network for semi-supervised learning in attributed heterogeneous information networks. In *Proc. of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2019.
- [Shi et al., 2019] Chuan Shi, Binbin Hu, Wayne Xin Zhao, and Philip S. Yu. Heterogeneous information network embedding for recommendation. *IEEE Trans. on Knowledge and Data Engineering*, 31(2):357–370, February 2019.
- [Veličković et al., 2017] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations*, 2017.
- [Wang et al., 2019] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu. Heterogeneous graph attention network. In *The World Wide Web Conference*, page 2022–2032, 2019.
- [Yang et al., 2018] Carl Yang, Yichen Feng, Pan Li, Yu Shi, and Jiawei Han. Meta-graph based HIN spectral embedding: methods, analyses and insights. In *ICDM*, 2018.
- [Yun et al., 2019] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J. Kim. Graph transformer networks. In *Proc. of the 33rd International Conference on Neural Information Processing Systems*, 2019.
- [Zhang et al., 2020] Wentao Zhang, Yuan Fang, Zemin Liu, Min Wu, and Xinming Zhang. mg2vec: Learning relationship-preserving heterogeneous graph representations via metagraph embedding. *IEEE Trans. on Knowledge and Data Engineering*, 2020.