

Higher-Order Topological Directionality and Directed Simplicial Neural Networks

Manuel Lecha, Andrea Cavallo, Francesca Dominici, Elvin Isufi, Claudio Battiloro

Abstract—Topological Deep Learning (TDL) has emerged as a paradigm to process and learn from signals defined on higher-order combinatorial topological spaces, such as simplicial or cell complexes. Although many complex systems have an asymmetric relational structure, most TDL models forcibly symmetrize these relationships. In this paper, we first introduce a novel notion of higher-order directionality and we then design Directed Simplicial Neural Networks (Dir-SNNs) based on it. Dir-SNNs are message-passing networks operating on directed simplicial complexes able to leverage directed and possibly asymmetric interactions among the simplices. To our knowledge, this is the first TDL model using a notion of higher-order directionality. We theoretically and empirically prove that Dir-SNNs are more expressive than their directed graph counterpart in distinguishing isomorphic directed graphs. Experiments on a synthetic source localization task demonstrate that Dir-SNNs outperform undirected SNNs when the underlying complex is directed, and perform comparably when the underlying complex is undirected.

Index Terms—Topological Deep Learning, Directed Simplicial Complexes, Directed Simplicial Neural Networks

I. INTRODUCTION

A strong inductive bias for deep learning models is processing signals while respecting the relational structure of their underlying space. Topological Deep Learning (TDL) is an emerging paradigm to process and learn from signals defined on combinatorial topological spaces (CTS) like simplicial and cell complexes [1], [2]. Unlike traditional graphs and Graph Neural Networks (GNNs), which can capture only pairwise relationships, i.e., two nodes connected by an edge, combinatorial topological spaces and Topological Neural Networks (TNNs) can capture higher-order interactions [3]. Such higher-order interactions are essential in many interconnected systems, including biological networks, where multi-way links exist among genes, proteins, or metabolites [4]. In this work, we are interested in Simplicial Complexes, powerful CTS allowing for more sophisticated adjacency schemes among *simplices* (nodes or groups of nodes closed under inclusion) than graphs, thus leading to a richer topological characterization and inductive bias.

Related Works. *Topological Neural Networks* have been shown to be expressive [5] (using the WL criterion [6]), able to handle long-range interactions [7], and effective in heterophilic settings [8], [9]. Simplicial convolutional neural networks have been proposed in [10], [11]. Message-passing simplicial networks have been introduced in [5] along with a simplicial generalization of the WL test. Recurrent message-passing simplicial networks were explored in [12]. $E(n)$ equivariant message-passing simplicial networks have been introduced in [13]. Message-passing free simplicial neural networks have been proposed in [14], [15]. Finally, simplicial attention networks have been proposed in [16]–[19].

Current Gaps. Accounting for edge directionality *transformed* our understanding of networks modeled as graphs and led to better learning performance [20]–[23]. We anticipate that the emerging research

Part of this work was supported by the TU Delft AI Labs program, the NWO OTP GraSPA proposal #19497, and the NWO VENI proposal 222.032. M. Lecha is with the Istituto Italiano di Tecnologia, Genoa, Italy. Email: manuel.lecha@iit.it. A. Cavallo and E. Isufi are with the Delft University of Technology, Delft, Netherlands. Emails: {a.cavallo,e.isufi-1}@tudelft.nl F. Dominici and C. Battiloro are with Harvard University, Cambridge, MA, USA. Emails: {fdominic,cbattiloro}@hsph.harvard.edu

on directionality in higher-order networks modeled as CTS [24], [25] will have a similar transformative impact. However, extending the TDL machinery to account for higher-order directionality remains unexplored and challenging. The main reason for this is that the very concept of higher-order directionality is not yet well defined in the Topological Deep Learning literature.

Contribution. Here, we fill this gap with a three-fold contribution.

- C1. We motivate and introduce a novel notion of higher-order topological directionality, i.e., novel directed adjacencies among the simplices of a *directed simplicial complex*, hinging on the theory of directed simplicial paths [24]. Directed simplicial paths generalize directed paths among nodes to directed paths among higher-order simplices.
- C2. We introduce the first directed Topological Neural Networks: **Directed Simplicial Neural Networks** (Dir-SNNs), message-passing networks operating on directed simplicial complexes and leveraging the above directed adjacencies.
- C3. We theoretically and empirically prove that directed simplicial neural networks **can distinguish isomorphic directed graphs** better than directed graph neural networks.

We numerically illustrate the potential of Dir-SNNs on a source localization task at the edge level, confirming that both directionality and the directed topological inductive bias play a role when compared with undirected SNNs and directed GNNs, respectively.

II. BACKGROUND

In this section, we define our domain of interest, *directed simplicial complexes*, our main tool to define higher-order directionality, *face maps*, and the signals we are interested into, *topological signals*.

Directed Simplicial Complexes. An (abstract) undirected simplicial complex is a pair $\tilde{\mathcal{K}} = (V, \Sigma)$, where V is a finite set of vertices, and Σ is a collection of subsets of vertices such that for every element $\sigma \in \Sigma$, if $\tau \subseteq \sigma$, then $\tau \in \Sigma$ (inclusivity property) [3]. An element σ of Σ is called a *simplex* of $\tilde{\mathcal{K}}$. A directed simplicial complex \mathcal{K} is a pair $\mathcal{K} = (V, \Sigma)$, where this time Σ is a collection of non-empty ordered tuples of vertices, i.e., the *directed simplices*. Again, Σ satisfies the inclusivity property, i.e., if a simplex σ is in Σ , then any non-empty ordered subtuple τ of σ is in Σ . The dimension $\dim(\sigma)$ of a directed simplex $\sigma = (v_0, \dots, v_k)$ is given by its arity minus one, i.e., $\dim(\sigma) = k$ when $|\sigma| = k + 1$. If $\dim(\sigma) = k$, σ is called a k -simplex. The dimension $\dim(\mathcal{K})$ of a directed simplicial complex \mathcal{K} is the maximal dimension of a directed simplex in \mathcal{K} . The k -skeleton \mathcal{K}_k of \mathcal{K} is the collection of directed simplices of dimension up to k . A directed graph is an example of a directed simplicial complex of dimension one, in which nodes and directed edges are directed 0- and 1-simplices, respectively. A directed simplicial complex of dimension two also comprises directed triangles.

Directed Flag Complexes. It is often useful to address (directed) graph-based problems by enriching the graph with higher-order relations, thereby mapping it to a higher-order simplicial complex while preserving its structure – nodes (0-simplices) and edges (1-simplices) remain those of the underlying graph. When such a




Fig. 1: From left to right: an undirected graph, its corresponding two-dimensional flag complex (nodes, edges, and triangles), a directed graph, and its corresponding directed flag complex (nodes, directed edges, and directed triangles). In the undirected flag complex, the triangles are $\sigma = \{0, 1, 2\}$ and $\tau = \{0, 2, 3\}$. In the directed flag complex, only the triangle $\sigma = (0, 1, 2)$ forms a directed triangle, as $(0, 2, 3)$ lacks the required edge $(0, 3)$.

structure-preserving transformation maps distinguishable graphs to distinguishable simplicial complexes (formally, it preserves isomorphisms) and is injective, meaning it maps indistinguishable graphs to indistinguishable simplicial complexes, it is called a *graph lifting* [5], [26]. A prominent example is the flag complex lifting, which maps an undirected graph \mathcal{G} to a flag complex $\tilde{\mathcal{K}}_{\mathcal{G}}$ [27]. A flag complex is a simplicial complex where the k -simplices correspond to the $(k+1)$ -cliques in \mathcal{G} – subsets of $k+1$ vertices where each pair of distinct vertices is connected by an edge. Directed flag complex liftings extend this transformation to accommodate directed graphs (digraphs). In this case, a directed flag complex $\mathcal{K}_{\mathcal{G}}$ is a directed simplicial complex where the k -simplices are ordered $(k+1)$ -cliques in \mathcal{G} . In this context, an ordered k -clique is a totally ordered tuple (v_1, \dots, v_k) such that (v_i, v_j) is a directed edge for $i < j$. A simple example showing a graph, its flag complex, a directed graph, and its directed flag complex is presented in Fig. 1.

Face Maps. In a directed simplicial complex \mathcal{K} , if $\tau \subseteq \sigma$ and $\sigma \in \mathcal{K}$, then τ is said to be a face of σ . Specifically, if $\dim(\tau) = \dim(\sigma) - 1$, τ is called a facet of σ . A directed k -simplex has $k+1$ facets. Face maps are a formal tool used to identify the faces of a simplex by systematically removing one of its vertices. Let $\mathcal{K} = (V, \Sigma)$ be a directed simplicial complex of dimension K . The face map d_i maps a k -simplex σ to the $(k-1)$ -simplex τ_i obtained by omitting the i -th vertex from σ , thus it is a mapping $d_i : \Sigma \rightarrow \Sigma$ defined as:

$$d_i(\sigma) = \tau_i = \begin{cases} (v_0, \dots, \hat{v}_i, \dots, v_k) & \text{if } i < k, \\ (v_0, \dots, v_{k-1}, \hat{v}_k) & \text{if } i \geq k. \end{cases} \quad (1)$$

where \hat{v}_i means that the vertex in the i -th position has been removed. The resulting $(k-1)$ -simplex τ_i preserves the original order of vertices, excluding the omitted vertex. For instance, applying the face maps to a directed triangle returns its directed edges. Face maps satisfy the simplicial identity $d_i \circ d_j = d_{j-1} \circ d_i$, where \circ is the composition operator, for $i < j$, ensuring that the order in which vertices are removed does not affect the resulting face. Face maps are an essential tool for understanding and structuring the relationships between directed simplices within a complex.

Topological Signals. Given a directed simplicial complex $\mathcal{K} = (V, \Sigma)$, a topological signal over \mathcal{K} is defined as a mapping $x : \Sigma \rightarrow \mathbb{R}$ from the set of simplices Σ to real numbers. Therefore, the *feature vectors* $\mathbf{x}_{\sigma} \in \mathbb{R}^F$ and $\mathbf{x}_{\tau} \in \mathbb{R}^F$ of simplices σ and τ are a collection of F topological signals, i.e.,

$$\mathbf{x}_{\sigma} = [x_1(\sigma), \dots, x_F(\sigma)]^{\top} \text{ and } \mathbf{x}_{\tau} = [x_1(\tau), \dots, x_F(\tau)]^{\top}. \quad (2)$$

For example, in a directed simplicial complex of dimension two, there are F signals (features) associated with nodes, edges, and triangles.




Fig. 2: (Left) A pair of non-isomorphic (distinguishable) digraphs, (Middle) their corresponding two-dimensional directed flag complexes, and (Right) their symmetrized undirected versions. Symmetrizing the directed flag complexes collapses the non-isomorphic digraphs into isomorphic flag complexes, while the directed flag complexes are non-isomorphic.

III. DIRECTED SIMPLICIAL NEURAL NETWORKS

In this section, we introduce Directed Simplicial Neural Networks (Dir-SNNs), message-passing networks operating on directed simplicial complexes. To do so, we first formally motivate and then introduce a consistent notion of higher-order topological directionality.

Motivation. A directed simplicial complex can be transformed into an undirected simplicial complex via symmetrization, wherein the order of the vertices in each directed simplex is disregarded. The symmetrization preserves isomorphisms, meaning that two isomorphic directed simplicial complexes will remain isomorphic as undirected complexes after symmetrization. Informally, this means that two indistinguishable directed simplicial complexes remain indistinguishable after the symmetrization. However, the process is not injective: non-isomorphic directed simplicial complexes may be mapped to isomorphic undirected simplicial complexes. Informally, this means that two distinguishable directed simplicial complexes can become indistinguishable after the symmetrization. For the same reason, composing symmetrization with lifting into (directed) flag complexes – whether by first symmetrizing a digraph and then lifting it into a flag complex, or by first lifting it into a directed flag complex and then symmetrizing it – can collapse distinct digraphs into the same undirected simplicial complex. This is just one of the possible formal hints showing that transitioning to higher-order undirected simplicial complexes is not always inherently beneficial, calling for a notion of higher-order directionality. Fig. 2 illustrates this fact.

Higher-order Topological Directionality. We define directed relations among simplices in a directed simplicial complex \mathcal{K} using face maps. Consider an ordered pair of simplices (σ, τ) with $\dim(\sigma) = \dim(\tau)$. Let (d_i, d_j) denote an ordered pair of the i -th and j -th face maps, as defined in (1). We define (σ, τ) as being *down* (k, i, j) -adjacent if there exists a simplex κ such that $\dim(\kappa) = \dim(\sigma) - k$ and $d_i(\sigma) \supseteq \kappa \subseteq d_j(\tau)$. To illustrate this, let $\sigma = (0, 1, 2)$ and $\tau = (1, 2, 3)$ be 2-simplices (triangles) as in Fig. 3. Consider the ordered pair of face maps (d_0, d_2) . Applying these maps, we obtain the directed edges $d_0((0, 1, 2)) = (1, 2) = \kappa$ and $d_2((1, 2, 3)) = (1, 2) = \kappa$, making the ordered pair of simplices (σ, τ) down $(1, 0, 2)$ -adjacent. Similarly, we define (σ, τ) as being *up* (k, i, j) -adjacent if there exists a simplex κ such that $\dim(\kappa) = \dim(\sigma) + k$ and $\sigma \subseteq d_i(\kappa)$ and $\tau \subseteq d_j(\kappa)$. Notably, for $i \neq j$, if an ordered pair (σ, τ) is up/down (k, i, j) -adjacent,




Fig. 3: Example of $(1, 0, 2)$ -adjacent directed 2-simplices $\sigma = (0, 1, 2)$ and $\tau = (1, 2, 3)$, sharing the edge $\kappa = (1, 2)$ (in red).




Fig. 4: Examples of edge down adjacencies. σ is the red edge in each subfigure. (a)/(b) $\mathcal{A}_{\downarrow,1}^{0,0}/\mathcal{A}_{\downarrow,1}^{1,1}$ connects σ with the edges with which it shares a target/source node; (c)/(d) $\mathcal{A}_{\downarrow,1}^{0,1}/\mathcal{A}_{\downarrow,1}^{1,0}$ connects σ with the edges whose source/target node is the target/source node of σ .

then (τ, σ) is up/down (k, j, i) -adjacent. Additionally, due to the symmetry in the face maps, if (σ, τ) is up/down (k, i, i) -adjacent, then (τ, σ) is also up/down (k, i, i) -adjacent, showing that not all the directed adjacencies are necessarily asymmetric. It is then natural to define notions of neighborhood among simplices of the same dimensions using the above adjacencies.

We define the *lower* (k, i, j) -adjacency $\mathcal{A}_{\downarrow,k}^{ij}$ of a simplex σ as

$$\mathcal{A}_{\downarrow,k}^{ij}(\sigma) = \left\{ \tau \in \Sigma \mid \begin{array}{l} \dim(\sigma) = \dim(\tau), \\ \exists \kappa : \dim(\kappa) = \dim(\sigma) - k, \\ d_i(\sigma) \supseteq \kappa \subseteq d_j(\tau) \end{array} \right\}. \quad (3)$$

We define the *upper* (k, i, j) -adjacency $\mathcal{A}_{\uparrow,k}^{ij}$ of a simplex σ as

$$\mathcal{A}_{\uparrow,k}^{ij}(\sigma) = \left\{ \tau \in \Sigma \mid \begin{array}{l} \dim(\sigma) = \dim(\tau), \\ \exists \kappa : \dim(\kappa) = \dim(\sigma) + k, \\ \sigma \subseteq d_i(\kappa) \text{ and } \tau \subseteq d_j(\kappa) \end{array} \right\}. \quad (4)$$

For the lower (k, i, j) -adjacency in (3) we assume $\dim(\kappa) = 0$ if $k > \dim(\sigma)$, while for the upper (k, i, j) -adjacency in (4) we assume $\dim(\kappa) = \dim(\mathcal{K})$ if $k > \dim(\mathcal{K}) - \dim(\sigma)$. We show some examples of edge lower and upper adjacencies in Fig. 4 and Fig. 5, respectively. The upper (k, i, j) -adjacency offers a complementary perspective to the lower (k, i, j) -adjacency, because the former captures the directed interactions where simplices σ and τ are both included in some higher-order simplices, while the latter when they both contain some lower-order simplices. Finally, face maps can also be used to define notions of neighborhood among simplices of different dimensions. We define the *boundary* \mathcal{B} and *coboundary* \mathcal{C} of $\sigma \in \mathcal{K}$ as

$$\mathcal{B}(\sigma) = \bigcup_{i=0}^{\dim(\sigma)} \{d_i(\sigma)\}, \quad \mathcal{C}(\sigma) = \bigcup_{i=0}^{\dim(\sigma)+1} d_i^{-1}(\sigma), \quad (5)$$

where d_i^{-1} is the preimage of d_i . The boundary and the coboundary of σ are then its facets and the simplices it is a facet of, respectively.

Consistency of Higher-order Topological Directionality. The way we define the directed adjacencies among simplices is grounded in the notion of simplicial directed paths [24]. In a digraph $G = (V, E)$, a directed path is defined as a sequence of vertices (v_0, v_1, \dots, v_n) where each consecutive pair $(v_i, v_{i+1}) \in E$ forms a directed edge. In a directed simplicial complex \mathcal{K} , a (k, i, j) -simplicial path between an ordered pair of simplices (σ, τ) in \mathcal{K} is a sequence of simplices $\sigma = \alpha_0, \alpha_1, \alpha_2, \dots, \alpha_n, \alpha_{n+1} = \tau$ such that each consecutive pair (α_k, α_{k+1}) is (k, i, j) -adjacent along the face maps (d_i, d_j) . We show some examples of simplicial paths of triangles in Fig. 6. Higher-order directionality reveals novel, discriminative structural properties,




Fig. 5: Examples of edges upper adjacencies. σ is the red edge in each subfigure. (a) $\mathcal{A}_{\uparrow,1}^{2,0}$ connects σ with the edge on its right; (a) $\mathcal{A}_{\uparrow,1}^{2,1}$ connects σ with the edge on its left.




Fig. 6: Examples of simplicial paths (in red) of 2-simplices (triangles). (a) The $(1, 0, 2)$ path, showing the simplices are equidirected; (b) the $(1, 1, 2)$ path, revealing a circular flow around a source node.

as demonstrated again in Fig. 2, where circular flows (in red) emerge considering of $(1, 0, 2)$ -simplicial paths. Finally, simplicial paths can also traverse simplices of different orders. We decided to add the constraint $\dim(\sigma) = \dim(\tau)$ in (3)-(4) to keep a distinction between same-dimension and different-dimension neighbors.

Directed Simplicial Neural Networks Directed Simplicial Neural Networks (Dir-SNNs) are message-passing networks [28] leveraging the adjacencies from (3)-(4). Given (i) a directed simplicial complex \mathcal{K} , and (ii) sets $\{\mathcal{A}_{\downarrow,k}^{ij}\}_{i,j,k}$ and $\{\mathcal{A}_{\uparrow,k}^{ij}\}_{i,j,k}$ collecting *some* of the lower and upper (k, i, j) -adjacencies as in (3)-(4), respectively, the l -th layer of a Dir-SNN updates the feature vector \mathbf{x}_σ^l of a $\sigma \in \mathcal{K}$ as

$$\mathbf{m}_{\sigma, \downarrow^{ijk}}^{l+1} = \bigoplus_{\tau \in \mathcal{A}_{\downarrow,k}^{ij}(\sigma)} \psi_{\mathcal{A}_{\downarrow,k}^{ij}}(\mathbf{x}_\sigma^l, \mathbf{x}_\tau^l, \mathbf{x}_\kappa^l), \quad (6)$$

$$\mathbf{m}_{\sigma, \uparrow^{ijk}}^{l+1} = \bigoplus_{\tau \in \mathcal{A}_{\uparrow,k}^{ij}(\sigma)} \psi_{\mathcal{A}_{\uparrow,k}^{ij}}(\mathbf{x}_\sigma^l, \mathbf{x}_\tau^l, \mathbf{x}_\kappa^l), \quad (7)$$

$$\mathbf{m}_{\sigma, \mathcal{B}}^{l+1} = \bigoplus_{\tau \in \mathcal{B}(\sigma)} \psi_{\mathcal{B}}(\mathbf{x}_\sigma^l, \mathbf{x}_\tau^l), \quad (8)$$

$$\mathbf{m}_{\sigma, \mathcal{C}}^{l+1} = \bigoplus_{\tau \in \mathcal{C}(\sigma)} \psi_{\mathcal{C}}(\mathbf{x}_\sigma^l, \mathbf{x}_\tau^l), \quad (9)$$

$$\mathbf{x}_\sigma^{l+1} = \phi(\mathbf{x}_\sigma^l, \{\mathbf{m}_{\sigma, \downarrow^{ijk}}^{l+1}\}_{ijk}, \{\mathbf{m}_{\sigma, \uparrow^{ijk}}^{l+1}\}_{ijk}, \mathbf{m}_{\sigma, \mathcal{B}}^{l+1}, \mathbf{m}_{\sigma, \mathcal{C}}^{l+1}). \quad (10)$$

with κ in (6) as in (3) and in (7) as in (4), \bigoplus being an intra-neighborhood aggregator. The neighborhood-dependent message functions $\psi_{\mathcal{A}_{\downarrow,k}^{ij}}, \psi_{\mathcal{A}_{\uparrow,k}^{ij}}, \psi_{\mathcal{B}}$ and $\psi_{\mathcal{C}}$, and the update function ϕ are learnable functions. In other words, the feature vector of a simplex is updated in a learnable fashion through aggregated messages with its neighboring simplices. At the l -th layer, a simplex has collected information from simplices that are up to l steps away from it along the simplicial directed paths induced by the chosen adjacencies.

Remark. In (8) and (10), single message functions $\psi_{\mathcal{B}}$ and $\psi_{\mathcal{C}}$ are used for computational efficiency. However, face map- and preimage-dependent message functions, i.e. ψ_{d_i} and $\psi_{d_i^{-1}}$, could leverage more fine-grained directed information, e.g. edges communicating with their source or target nodes using different sets of weights.

Expressiveness of Dir-SNNs. The expressive power of topological neural networks (including GNNs) is usually measured by their capacity to distinguish non-isomorphic objects within their underlying domain [6]. Here, we introduce the following result:

Theorem 1. There exist Dir-SNNs that are more powerful than Dir-GNNs [20] at distinguishing non-isomorphic digraphs using a directed flag complex lifting.





Fig. 7: SNR vs accuracy of directed and undirected TNNs and GNNs on directed (left) and undirected (right) synthetic flag complexes.




Fig. 8: A pair of non-isomorphic directed graphs (a) and (b), along with their corresponding flag complexes (c) and (d). These digraphs can be distinguished by Dir-SNNs but not by Dir-GNNs [20].

Proof. See Appendix.

An example is shown in Fig. 8. Intuitively, our proof follows the same strategy of [5], thus it relies on the definition of a directed simplicial isomorphism test that is proven to be (i) an upper bound on the expressiveness of Dir-SNNs, and (ii) more powerful than the directed graph isomorphism test [29], being an upper bound on the expressiveness of Dir-GNNs [20].

Remark. Dir-SNNs generalize Dir-GNNs [20]. A Dir-GNN is a Dir-SNNs operating on a digraph, i.e., a directed simplicial complex of dimension one, which updates the node feature vectors using only the $(1, 0, 1)$ - and the $(1, 1, 0)$ -upper adjacencies in (4).

IV. NUMERICAL RESULTS

We first provide some preliminary results to validate the effectiveness of Dir-SNN on a synthetic source localization task at the edge level [30]. Then, we perform a simple experiment to validate Theorem 1.¹

A. Source Localization

Dataset. We generate directed and undirected graphs following a Stochastic Block Model [31]. Each graph has 70 nodes uniformly divided into 10 communities, with intra- and inter-community edge probabilities of 0.9 and 0.01, respectively. Intra-community edges are grouped into 10 edge communities, and the remaining inter-community edges form an 11th partition. We generate 1000 edge signals from a zero-mean Gaussian distribution with variance $1/N_{\text{edges}}$, and then, for each signal, we introduce spikes to randomly selected source edges belonging to a single community with intensity $\alpha \sim \mathcal{N}(0, 1)$. In the directed case, the graphs are lifted in their corresponding directed flag complexes, and the spikes are diffused over the graph following $\mathbf{x}' = \mathbf{S}^t \mathbf{x} + \mathbf{n}$, where $\mathbf{S} \in \mathbb{R}^{N_{\text{edges}} \times N_{\text{edges}}}$ is a non-symmetric binary matrix encoding $\mathcal{A}_{\downarrow, 1}^{0, 1}$ from (3) for each edge, t is the order of diffusion sampled from a Student-T distribution with 10 degrees of freedom and capped at 100, \mathbf{x} is the original signal with the added spikes and \mathbf{n} is additive white Gaussian noise inducing a specific SNR. In the undirected case, the graphs are lifted in their corresponding flag complexes, and \mathbf{S} is the symmetric lower edge adjacency, analogously to [30]. The task is to identify the community originating the spikes, thus a classification problem with 11 classes.

Experimental setup. Since the designed task considers only edge signals, we decided to employ a specific instance of Dir-SNN that

Model	Accuracy
Dir-GNN	50%
Dir-SNN	100%

TABLE I: Discrimination accuracies of Dir-GNN and Dir-SNN.

operates on the four directed adjacency relations described by $\mathcal{A}_{\downarrow, 1}^{0, 0}$, $\mathcal{A}_{\downarrow, 1}^{0, 1}$, $\mathcal{A}_{\downarrow, 1}^{1, 0}$, $\mathcal{A}_{\downarrow, 1}^{1, 1}$, without considering boundary and coboundary. The aggregators, update, and message functions are chosen to have a convolutional architecture [32]. Therefore, we compare with an undirected convolutional SNN using the undirected lower adjacency of the edges [11], and with two GNNs that operate on the projection of edge features on their endpoint nodes, one directed (Dir-GNN [20]) and the other undirected (GCN [32]). We perform a grid search for all the models among the following values: $\{1, 2, 3\}$ layers of size $\{16, 32, 64\}$. The embeddings are max-pooled and fed into an MLP to perform classification. We average over 5 seeds.

Discussion. In Fig. 7, Dir-SNN consistently and largely outperforms all the baselines for all levels of SNR on directed flag complexes (Left). Moreover, Dir-SNN performs comparably to SNN on flag complexes (Right), showing their robustness to model mismatching.

B. Expressivity Validation

Dataset. We employ a toy dataset containing only the two graphs in Fig. 8 (a)-(b), and we assign them two different classes. The task is then to correctly classify them.

Experimental Setup. We compare Dir-SNN and Dir-GNN. The parameters are the same of IV-A, but for Dir-SNN we also use $\mathcal{A}_{\uparrow, 1}^{2, 0}$ (to leverage the directed triangles). The two graphs in Fig. 8 (a)-(b) are fed to the Dir-GNN with constant features on each node, while the two corresponding directed flag complexes in Fig. 8 (c)-(d) are fed to the Dir-SNN with constant features on each simplex.

Discussion. The results in Table I numerically validate the improved expressivity of Dir-SNN stated in Theorem 1. Indeed, Dir-GNN cannot discriminate the two graphs, while Dir-SNN can discriminate them thanks to the lifting into directed flag complexes.

V. CONCLUSIONS

We introduced *Directed Simplicial Networks* (Dir-SNNs), the first family of message-passing networks operating on directed simplicial complexes and leveraging *novel notions of higher-order topological directionality*. Dir-SNNs effectively model complex directed and possibly asymmetric relationships that are inaccessible to traditional directed graph-based or undirected topological models, showing improved expressivity. Numerical experiments validated the effectiveness and the expressivity of Dir-SNNs. In the journal version of this paper, we will characterize the expressivity of Dir-SNNs more comprehensively. Moreover, we will develop a spectral theory to characterize better the (implicit, at the moment) filtering operations in Dir-SNNs. Finally, we will exhaustively evaluate Dir-SNNs on a wide range of real data, both on graph and higher-order tasks.

¹Code at <https://github.com/andrea-cavallo-98/DirSNN>

REFERENCES

- [1] M. Hajij, G. Zamzmi, T. Papamarkou, N. Miolane, A. Guzmán-Sáenz, K. N. Ramamurthy, T. Birdal, T. K. Dey, S. Mukherjee, S. N. Samaga, *et al.*, “Topological deep learning: Going beyond graph data,” *arXiv preprint arXiv:2206.00606*, 2022.
- [2] T. Papamarkou, T. Birdal, M. M. Bronstein, G. E. Carlsson, J. Curry, Y. Gao, M. Hajij, R. Kwitt, P. Lio, P. Di Lorenzo, *et al.*, “Position: Topological deep learning is the new frontier for relational learning,” in *Forty-first International Conference on Machine Learning*, 2024.
- [3] S. Barbarossa and S. Sardellitti, “Topological signal processing over simplicial complexes,” *IEEE Transactions on Signal Processing*, 2020.
- [4] R. Lambiotte, M. Rosvall, and I. Scholtes, “From networks to optimal higher-order models of complex systems,” *Nature physics*, 2019.
- [5] C. Bodnar, F. Frasca, Y. Wang, N. Otter, G. Montufar, P. Liò, and M. Bronstein, “Weisfeiler and lehman go topological: Message passing simplicial networks,” in *International Conference on Machine Learning*, 2021.
- [6] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How powerful are graph neural networks?,” in *International Conference on Learning Representations*, 2019.
- [7] L. Giusti, C. Battiloro, L. Testa, P. Di Lorenzo, S. Sardellitti, and S. Barbarossa, “Cell attention networks,” in *2023 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, IEEE, 2023.
- [8] C. Bodnar, F. D. Giovanni, B. P. Chamberlain, P. Lio, and M. M. Bronstein, “Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs,” in *Advances in Neural Information Processing Systems*, 2022.
- [9] C. Battiloro, Z. Wang, H. Riess, P. Di Lorenzo, and A. Ribeiro, “Tangent bundle convolutional learning: from manifolds to cellular sheaves and back,” *IEEE Transactions on Signal Processing*, 2024.
- [10] S. Ebli, M. Defferrard, and G. Spreemann, “Simplicial neural networks,” in *Advances in Neural Information Processing Systems Workshop on Topological Data Analysis and Beyond*, 2020.
- [11] M. Yang and E. Iusfi, “Convolutional learning on simplicial complexes,” 2023.
- [12] T. M. Roddenberry and S. Segarra, “Hodgenet: Graph neural networks for edge data,” *Computing Research Repository (CoRR)*, vol. abs/1912.02354, 2019.
- [13] F. Eijkelboom, R. Hesselink, and E. Bekkers, “ $E(n)$ equivariant message passing simplicial networks,” *arXiv 2305.07100*, 2023.
- [14] S. Gurugubelli and S. P. Chepuri, “Sann: Simple yet powerful simplicial-aware neural networks,” in *The Twelfth International Conference on Learning Representations*, 2023.
- [15] K. Maggs, C. Hacker, and B. Rieck, “Simplicial representation learning with neural $\mathbb{K}\mathbb{S}$ -forms,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [16] C. Battiloro, L. Testa, L. Giusti, S. Sardellitti, P. Di Lorenzo, and S. Barbarossa, “Generalized simplicial attention neural networks,” *arXiv preprint arXiv:2309.02138*, 2023.
- [17] L. Giusti, C. Battiloro, P. D. Lorenzo, S. Sardellitti, and S. Barbarossa, “Simplicial attention neural networks,” *arXiv:2203.07485v2*, 2022.
- [18] C. W. J. Goh, C. Bodnar, and P. Lio, “Simplicial attention networks,” in *International Conference on Learning Representations Workshop on Geometrical and Topological Representation Learning*, 2022.
- [19] S. H. Lee, F. Ji, and W. P. Tay, “SGAT: Simplicial graph attention network,” *arXiv 2207.11761*, 2022.
- [20] E. Rossi, B. Charpentier, F. D. Giovanni, F. Frasca, S. Günnemann, and M. Bronstein, “Edge directionality improves learning on heterophilic graphs,” 2023.
- [21] X. Zhang, Y. He, N. Brugnone, M. Perlmutter, and M. Hirn, “Magnet: A neural network for directed graphs,” *Advances in neural information processing systems*, vol. 34, pp. 27003–27015, 2021.
- [22] V. Thost and J. Chen, “Directed acyclic graph neural networks,” *arXiv preprint arXiv:2101.07965*, 2021.
- [23] Z. Tong, Y. Liang, C. Sun, D. S. Rosenblum, and A. Lim, “Directed graph convolutional network,” *arXiv preprint arXiv:2004.13970*, 2020.
- [24] H. Riihimäki, “Simplicial q -connectivity of directed graphs with applications to network analysis,” *SIAM Journal on Mathematics of Data Science*, vol. 5, no. 3, pp. 800–828, 2023.
- [25] X. Gong, D. J. Higham, K. Zygalakis, and G. Bianconi, “Higher-order connection laplacians for directed simplicial complexes,” *Journal of Physics: Complexity*, vol. 5, no. 1, p. 015022, 2024.
- [26] C. Bodnar, F. Frasca, N. Otter, Y. Wang, P. Liò, G. F. Montufar, and M. Bronstein, “Weisfeiler and lehman go cellular: Cw networks,” in *Advances in Neural Information Processing Systems*, 2021.
- [27] K. F. E. Chong and E. Nevo, “Flag complexes and homology,” *Journal of Combinatorial Theory, Series A*, vol. 182, p. 105466, 2021.
- [28] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, “Neural message passing for quantum chemistry,” in *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, p. 1263–1272, JMLR.org, 2017.
- [29] M. Grohe, K. Kersting, M. Mladenov, and P. Schweitzer, “Color refinement and its applications,” 2021.
- [30] S. Fiorellino, C. Battiloro, and P. Di Lorenzo, “Topological neural networks over the air,” in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 12986–12990, 2024.
- [31] P. W. Holland, K. B. Laskey, and S. Leinhardt, “Stochastic blockmodels: First steps,” *Social Networks*, vol. 5, no. 2, pp. 109–137, 1983.
- [32] T. N. Kipf and M. Welling, “Semi-Supervised Classification with Graph Convolutional Networks,” in *International Conference on Learning Representations*, 2017.
- [33] B. Zhang, J. Gai, Y. Du, Q. Ye, D. He, and L. Wang, “Beyond weisfeiler-lehman: A quantitative framework for gnn expressiveness,” *arXiv preprint arXiv:2401.08514*, 2024.
- [34] B. Weisfeiler and A. Leman, “The reduction of a graph to canonical form and the algebra which appears therein,” *NTI, Series*, 1968.
- [35] Q. Truong and P. Chin, “Weisfeiler and lehman go paths: Learning topological features via path complexes,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, pp. 15382–15391, 2024.

APPENDIX

The expressive power of Topological Neural Networks (TNNs) is usually (but not only [33]) measured by evaluating their ability to distinguish isomorphic combinatorial topological spaces (CTS) [5]. The Weisfeiler-Leman (WL) test, an isomorphism test for graphs [34], has been proven to be an upper bound on GNNs discriminative power [6], i.e. no vanilla GNN can distinguish isomorphic graphs better than the WL. Graph liftings, such as the directed flag complex lifting (introduced in Sec. II), canonically map graphs to Combinatorial Topological Spaces (CTS), enriching their structure with higher-order relations while preserving the original graph’s nodes (0-simplices) and edges (1-simplices). These liftings ensure that isomorphic graphs are mapped to isomorphic CTSs and non-isomorphic graphs to non-isomorphic CTSs. For this reason, topological variants of the WL test have been developed to exploit the additional higher-order structure introduced by graph lifting [5], [26], [35]. These extensions, as their graph counterpart, have been proven to be an upper bound on TNNs ability to distinguish isomorphic CTS, and, more importantly, they have been used to prove that TNNs operating on lifted graphs are more expressive than traditional GNNs. In other words, TNNs have demonstrated a greater capacity to distinguish between non-isomorphic graphs that GNNs may otherwise classify as identical. In this appendix, we prove Theorem 1, i.e. we formally show that Directed Simplicial Neural Networks (Dir-SNNs) leveraging higher-order structures and directional information through directed flag complex liftings are strictly more expressive than Directed Graph Neural Networks (Dir-GNNs) [20]. Our proof is a non-trivial adaptation of the Simplicial Weisfeiler-Leman (SWL) test from [5] to the directed setting (defined in Sec. III). In particular, we first introduce the Directed Simplicial Weisfeiler-Leman (D-SWL) test, the first isomorphism test for directed simplicial complexes. We then rigorously show that D-SWL is a more powerful isomorphism test than D-WL [29], being an upper bound on the expressiveness of Dir-GNNs. Finally, the proof is concluded by showing that the D-SWL is an upper bound on the expressiveness of Dir-SNNs. The structure of our proofs, theorems, and notation follows closely the work of

[5], and we encourage readers to consult it for further context and foundational concepts.

Directed Simplicial Weisfeiler-Leman (D-SWL) Test. We assume $k = 1$ and omit this specification from the notation of the up- and down-adjacencies in (4)-(3). Let \mathcal{K} be a directed simplicial complex, and let c^t denote the colouring of simplices at iteration t in the D-SWL test. Let $\sigma \in \mathcal{K}$ be an n -simplex, we define the following arguments at step t :

1. *Boundary colouring:* Represents the ordered tuple of colours assigned to each element of the boundary of σ , $d_i(\sigma)$.

$$c_B^t(\sigma) = (c_{d_0(\sigma)}, \dots, c_{d_n(\sigma)}).$$

2. *Coboundary colouring:* Let $c_{d_i^{-1}(\sigma)} = \{\{c_\tau \mid \tau \in d_i^{-1}(\sigma)\}\}$ be the multiset of colours for the preimage of the i -th face map of σ . Then the coboundary colouring of σ is given by:

$$c_C^t(\sigma) = \bigcup_{i=0}^{\dim(\sigma)+1} c_{d_i^{-1}(\sigma)}.$$

3. *Down adjacency colouring:* Captures the pairs of colours associated with simplices τ that are downward adjacent to σ , where κ represents a shared face.

$$(c_\downarrow^{ij})^t(\sigma) = \{\{(c_\tau^t, c_\kappa) \mid \tau \in \mathcal{A}_\downarrow^{ij}(\sigma), d_i(\sigma) = \kappa = d_j(\tau)\}\}.$$

4. *Up adjacency colouring:* Represents the pairs of colours of simplices τ that are upward adjacent to σ , where κ represents a shared higher-dimensional face.

$$(c_\uparrow^{ij})^t(\sigma) = \{\{(c_\tau^t, c_\kappa) \mid \tau \in \mathcal{A}_\uparrow^{ij}(\sigma), d_i(\kappa) = \sigma, d_j(\kappa) = \tau)\}.$$

The D-SWL test distinguishes non-isomorphic directed simplicial complexes through iterative color refinement of simplices. The procedure consists of the following steps:

1) *Initialization:* At $t = 0$, assign an initial colouring to all simplices.
2) *Color Propagation and Refinement:* For each simplex σ , propagate and refine the colours using the injective update rule:

$$c_\sigma^{t+1} = \text{HASH} \left(c_\sigma^t, c_B^t(\sigma), c_C^t(\sigma), ((c_\downarrow^{ij})^t(\sigma))_{i,j}, ((c_\uparrow^{i,j})^t(\sigma))_{i,j} \right),$$

where $i, j \in \{0, \dots, \dim(\sigma)\}$ and $i, j \in \{0, \dots, \dim(\sigma)\}$ for the down and upper adjacencies, respectively.

3) *Termination:* Repeat this process until the colouring stabilizes. Two directed simplicial complexes are considered non-isomorphic if their stable color histograms differ.

We now demonstrate that, for distinguishing non-isomorphic directed simplicial complexes, the multisets of colours associated with the coboundaries and down-adjacencies can be removed without sacrificing the expressiveness of the D-SWL test. This will simplify the rest of the proof, and it also demonstrates that D-SWL (and, eventually, Dir-SNNs) can be made computationally more efficient.

Lemma 1. *D-SWL with $\text{HASH} \left(c_\sigma^t, c_B^t(\sigma), ((c_\uparrow^{ij})^t(\sigma))_{i,j} \right)$ is as powerful as D-SWL with the generalised update rule $c_\sigma^{t+1} = \text{HASH} \left(c_\sigma^t, c_B^t(\sigma), c_C^t(\sigma), ((c_\downarrow^{ij})^t(\sigma))_{i,j}, ((c_\uparrow^{i,j})^t(\sigma))_{i,j} \right)$.*

Proof. First, we show that coboundary multisets can be removed, following Lemma 25 of [5].

A. *Omitting $c_C^t(\sigma)$*

Let a^t and b^t denote the colourings at iteration t for the generalized update rule (which includes the coboundary color multisets) and the restricted update rule (which excludes them), respectively. To establish that these two colourings are equivalent, we need to show

that a^t refines b^t and vice versa. Since the refinement a^t to b^t is straightforward, we focus on proving that b^t refines a^t . We proceed by induction on t .

1) *Base Case:* At $t = 0$, all simplices are assigned the same initial color, so the base case holds trivially.

2) *Inductive Hypothesis:* Assume that for some iteration t , the restricted colouring b^t refines the generalized colouring a^t . That is, if $b_\sigma^t = b_\tau^t$ for two simplices $\sigma \in K_1$ and $\tau \in K_2$ with $\dim(\sigma) = \dim(\tau)$, then $a_\sigma^t = a_\tau^t$. We aim to show that this holds for the next iteration, $t + 1$, i.e., $b_\sigma^{t+1} = b_\tau^{t+1}$ implies $a_\sigma^{t+1} = a_\tau^{t+1}$.

3) *Inductive Step:* Suppose that $b_\sigma^{t+1} = b_\tau^{t+1}$. This implies equivalence in the arguments of the update rule at time t . In particular, the upward adjacency colouring multisets $(b_\uparrow^{ij})^t(\sigma)$ include the colours of all higher-dimensional simplices δ_σ whose i -th face is σ . This accounts for the colours of the elements in the preimage of the i -th face map $d_i^{-1}(\sigma)$. Thus, we have:

$$b_{d_i^{-1}(\sigma)} = b_{d_i^{-1}(\tau)}$$

for all i . Therefore, by definition, $b_C^t(\sigma) = b_C^t(\tau)$. Finally, by the inductive hypothesis, the arguments for the generalized hash update rule must also be equivalent at time t , so $a_\sigma^t = a_\tau^t$, leading to $a_\sigma^{t+1} = a_\tau^{t+1}$, as required, completing the proof.

Second, we prove that the update rule can be further refined by omitting the downward adjacencies.

B. *Omitting $((c_\downarrow^{ij})^t(\sigma))_{i,j}$*

Let a^t and b^t denote the colourings at iteration t for the generalized update rule (which includes down-adjacencies but excludes coboundary color multisets) and the restricted update rule (which excludes both down-adjacencies and coboundary color multisets), respectively. Our goal is to show that b^{2t} refines a^t by induction.

1) *Base Case:* At $t = 0$, all simplices are assigned the same initial color, so the base case trivially holds.

2) *Inductive Hypothesis:* Assume that for some iteration t , the restricted colouring b^{2t} refines the generalized colouring a^t . This means that if $b_\sigma^{2t} = b_\tau^{2t}$ for two simplices $\sigma \in K_1$ and $\tau \in K_2$ with $\dim(\sigma) = \dim(\tau)$, then $a_\sigma^t = a_\tau^t$. We aim to show that this holds for the next iteration, $t + 2$.

3) *Inductive Step:* Suppose that $b_\sigma^{2t+2} = b_\tau^{2t+2}$. This implies that $b_\sigma^{2t} = b_\tau^{2t}$, meaning the arguments for the restricted update rule at step $2t$ are equivalent. We now need to prove that $(b_\downarrow^{ij})^{2t}(\sigma) = (b_\downarrow^{ij})^{2t}(\tau)$ holds for all ordered pairs (i, j) . Assume, by contradiction, that there exists a pair of indices (i, j) such that $(b_\downarrow^{ij}(\sigma))^{2t} \neq (b_\downarrow^{ij}(\tau))^{2t}$. Without loss of generality, suppose there is a pair of colours (c_0, c_1) that appears more frequently in $(b_\downarrow^{ij}(\sigma))^{2t}$ than in $(b_\downarrow^{ij}(\tau))^{2t}$. For simplicity, as in [5], assume $b_\sigma^{2t} \neq c_0 \neq b_\tau^{2t}$. Define $c_\delta^j = b_{d_j^{-1}(\delta)}(c_0)$ to denote the multiplicity of c_0 in the set $b_{d_j^{-1}(\delta)}$, i.e., the number of elements in $d_j^{-1}(\delta)$ coloured with c_0 . Now, assume there exist two simplices δ_1 and δ_2 , such that $c_{\delta_1}^j \neq c_{\delta_2}^j$ (without loss of generality, assume $c_{\delta_1}^j > c_{\delta_2}^j$). Thus, for all $s \in \{0, \dots, \dim(\delta_1) + 1\}$, c_0 appears more frequently in $(b_\downarrow^{ij}(\delta_1))^{2t}$ than in $(b_\downarrow^{ij}(\delta_2))^{2t}$. This implies that $b_{\delta_1}^{2t+1} \neq b_{\delta_2}^{2t+1}$. Since by assumption, the number of tuples (c_0, c_1) is greater in $(b_\downarrow^{ij}(\sigma))^{2t}$ than in $(b_\downarrow^{ij}(\tau))^{2t}$, we can substitute $\delta_1 = d_i(\sigma)$, and $\delta_2 = d_i(\tau)$. This yields $c_{d_i(\sigma)}^j = b_{d_j^{-1}(d_i(\sigma))}(c_0) \neq c_{d_i(\tau)}^j = b_{d_j^{-1}(d_i(\tau))}(c_0)$. Therefore, we conclude:

$$b_{d_i(\sigma)}^{2t+1} \neq b_{d_i(\tau)}^{2t+1},$$

which implies that the boundary colouring tuples $b_{\mathcal{B}}^{2t+1}(\sigma) \neq b_{\mathcal{B}}^{2t+1}(\tau)$. Hence, $b_{\sigma}^{2t+2} \neq b_{\tau}^{2t+2}$, leading to a contradiction. Therefore, $(b_{\downarrow}^{ij})^{2t}(\sigma) = (b_{\downarrow}^{ij})^{2t}(\tau)$ for all ordered pairs (i, j) . Finally, by the inductive hypothesis, b^{2t} refines a^t , completing the proof. ■

D-SWL vs D-WL. We now prove the following result.

Lemma 2. *D-SWL with a directed flag complex lifting is strictly more powerful than D-WL.*

Proof. We begin by proving that D-SWL is at least as expressive as D-WL in distinguishing non-isomorphic directed simplicial complexes. The proof follows closely the structure of Lemma 27 in [5], with key differences in the treatment of upper adjacencies: D-SWL and D-WL employ distinct colouring multisets to capture asymmetric directional dependencies.

A. At Least As Expressive

Let \mathcal{K} be a directed simplicial complex. Denote by a^t and b^t the colourings of the same vertices in \mathcal{K} at iteration t of D-WL and D-SWL, respectively. We aim to show that the D-SWL colouring b^t refines the D-WL colouring a^t .

1) *Base Case:* At $t = 0$, all simplices are assigned the same initial color, so the base case trivially holds.

2) *Inductive Step:* Suppose that for some iteration t , $b_v^{t+1} = b_w^{t+1}$ for two vertices v and w in two arbitrary directed simplicial complexes \mathcal{K}_1 and \mathcal{K}_2 , meaning that the D-SWL colourings of v and w at step $t + 1$ are identical. Since vertices have no boundary simplices and are only upper adjacent, this implies:

$$b_v^t = b_w^t, \quad (b_{\uparrow}^{01})^t(v) = (b_{\uparrow}^{01})^t(w), \quad (b_{\uparrow}^{10})^t(v) = (b_{\uparrow}^{10})^t(w).$$

Expanding the definition of these multisets, we get:

$$\begin{aligned} \{\{b_z^t \mid (b_z^t, \cdot) \in (b_{\uparrow}^{01})^t(v)\}\} &= \{\{b_u^t \mid (b_u^t, \cdot) \in (b_{\uparrow}^{01})^t(w)\}\}, \\ \{\{b_z^t \mid (b_z^t, \cdot) \in (b_{\uparrow}^{10})^t(v)\}\} &= \{\{b_u^t \mid (b_u^t, \cdot) \in (b_{\uparrow}^{10})^t(w)\}\}. \end{aligned}$$

These are equivalent to the upper in- and out-neighborhood sets of the vertices v and w , defined in [20] as $N_{\leftarrow}(v)$ and $N_{\rightarrow}(v)$, respectively. Thus, we can rewrite the expressions as:

$$\begin{aligned} \{\{b_z^t \mid z \in N_{\leftarrow}(v)\}\} &= \{\{b_u^t \mid u \in N_{\leftarrow}(w)\}\}, \\ \{\{b_z^t \mid z \in N_{\rightarrow}(v)\}\} &= \{\{b_u^t \mid u \in N_{\rightarrow}(w)\}\}. \end{aligned}$$

By the induction hypothesis, we know that $a_v^t = a_w^t$, $(a_{\uparrow}^{01})^t(v) = (a_{\uparrow}^{01})^t(w)$, and $(a_{\uparrow}^{10})^t(v) = (a_{\uparrow}^{10})^t(w)$. These are the arguments that the D-WL hash function uses to compute the colours of v and w at the next iteration, which implies that $a_v^{t+1} = a_w^{t+1}$. Hence, b^t refines a^t , proving that D-SWL is at least as expressive as D-WL.

B. Strictly More Expressive

To demonstrate that D-SWL is strictly more powerful than D-WL, we present a counterexample: a pair of digraphs that cannot be distinguished by D-WL but whose directed flag complexes can be distinguished by D-SWL. In Fig. 8, we show such a pair of digraphs. While the D-WL test produces identical colourings for both digraphs, their directed flag complexes reveal a crucial difference: one complex contains directed triangles, while the other contains none. Since D-SWL detects these higher-order simplicial structures, it successfully distinguishes the two digraphs, completing the proof. ■

Dir-SNNs vs Dir-GNNs. We can finally prove Thm. 1 from Sec. III.

Theorem 1. *There exist Dir-SNNs that are more powerful than Dir-GNNs [20] at distinguishing non-isomorphic digraphs when using a directed flag complex lifting.*

Proof. First, we demonstrate that Dir-SNNs with sufficient layers and injective neighborhood aggregators are as powerful as the D-SWL test in distinguishing non-isomorphic directed simplicial complexes. The proof parallels the structure of Lemma 9 in [5].

A. Dir-SNNs As Powerful As D-SWL

Let c^t and h^t represent the colouring at iteration t of D-SWL and the t -th layer of a Dir-SNN, respectively. We consider a Dir-SNN with L layers, and for $t > L$, we assume that $h^t = h^L$. First, we show by induction that the colouring c^t of D-SWL refines the colouring h^t of Dir-SNN.

1) *Base Case:* At $t = 0$, all simplices are assigned the same initial color in both D-SWL and Dir-SNN, so the base case trivially holds.

2) *Inductive Step:* For the inductive step, suppose that for some iteration t , the colouring $c_{\sigma}^{t+1} = c_{\tau}^{t+1}$ for two simplices σ and τ in a directed simplicial complex. Because the D-SWL colouring is injective, the arguments to the HASH function must also be equal. This implies the following equalities:

$$\begin{aligned} c_{\sigma}^t = c_{\tau}^t, \quad c_{\mathcal{B}}^t(\sigma) = c_{\mathcal{B}}^t(\tau), \quad c_{\mathcal{C}}^t(\sigma) = c_{\mathcal{C}}^t(\tau), \\ ((c_{\downarrow}^{ij})^t(\sigma))_{i,j} = ((c_{\downarrow}^{ij})^t(\tau))_{i,j}, \quad \text{and} \quad ((c_{\uparrow}^{ij})^t(\sigma))_{i,j} = ((c_{\uparrow}^{ij})^t(\tau))_{i,j}. \end{aligned}$$

By the induction hypothesis, these arguments will also be equal under the Dir-SNN colouring h^t . Since the same arguments are supplied as input to the message-passing, aggregate, and update functions of Dir-SNN, their outputs will be identical for σ and τ . Therefore, $h_{t+1}(\sigma) = h_{t+1}(\tau)$. Following the discussion in Theorem 9 of [5], and assuming that the boundary aggregation function is injective and *non-permutation invariant*, we can derive the reverse implication: h^t refines the colouring of c^t , proving that Dir-SNNs with sufficient layers and injective neighborhood aggregators are at least as powerful as D-SWL.

B. Strictly More Expressive

Finally, applying Lemma 2, we conclude that Dir-SNNs, when utilizing a directed flag complex lifting, are strictly more powerful than Dir-GNNs, as D-SWL is more expressive than D-WL in distinguishing non-isomorphic digraphs. ■