

000 DIFFERENTIABLE LIFTING FOR TOPOLOGICAL NEURAL 001 NETWORKS 002

003 **Anonymous authors**
004

005 Paper under double-blind review
006

007 ABSTRACT 008

009 Topological neural networks (TNNs) enable leveraging high-order structures on
010 graphs (e.g., cycles and cliques) to boost the expressive power of message-passing
011 neural networks. In turn, however, these structures are typically identified *a priori*
012 through an unsupervised graph lifting operation. Notwithstanding, this choice is
013 crucial and may have a drastic impact on a TNN’s performance on downstream
014 tasks. To circumvent this issue, we propose ∂ lift (DiffLift), a general framework
015 for learning graph liftings to hypergraphs and cellular- and simplicial complexes in
016 an end-to-end fashion. In particular, our approach leverages learned vertex-level
017 latent representations to identify and parameterize distributions over candidate
018 higher-order cells for inclusion. This results in a scalable model which can be
019 readily integrated into any TNN. Our experiments show that ∂ lift outperforms
020 existing lifting methods on multiple benchmarks for graph and node classification
021 across different TNN architectures. Notably, our approach leads to gains of up to
022 45% over static liftings, including both connectivity- and feature-based ones.
023

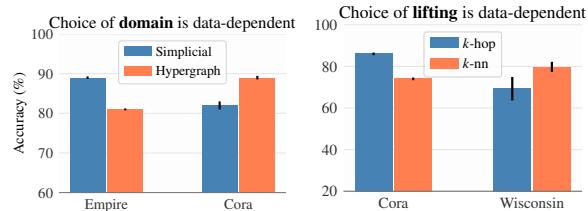
025 1 INTRODUCTION 026

027 Topological neural networks (TNNs) (Papillon et al., 2023b; Bodnar et al., 2021a; Verma et al., 2024)
028 have recently emerged as a prominent class of models for learning on topological domains, such as hy-
029 pergraphs and simplicial complexes, with many researchers arguing they represent the new frontier for
030 relational learning (Papamarkou et al., 2024). Akin to graph neural networks (GNNs) (Scarselli et al.,
031 2009; Gilmer et al., 2017), typical TNNs employ message-passing layers where each element of the
032 input (e.g., nodes or cells) updates its representation (features) based on those of its topological neigh-
033 bors. Thus, these models generalize convolution-like operations on graphs to higher-order relational
034 objects. Importantly, the primary application of TNNs has been to enhance the capabilities of graph-
035 based models, particularly in terms of expressivity (Bodnar et al., 2021a;b). In this context, the input
036 graphs must first be transformed to the domain on which a TNN operates — a process known as *lifting*.
037

038 Lifting methods explore graph connec-
039 tivity and features to create higher-
040 order relational structures. For in-
041 stance, *clique lifting* (Bodnar et al.,
042 2021b) produces a simplicial complex
043 by leveraging cliques in the input graph
044 while *cycle lifting* (Hajij et al., 2022)
045 detects cycles to create a cell complex.
046 In general, there are many lifting pro-
047 cedures for each topological domain —
048 c.f. Tab. 3 in Telyatnikov et al. (2024).
049

050 Not surprisingly, the optimal choice of topological domain and lifting procedure for each task is non-
051 obvious, and its impact on performance is highly data-dependent. Figure 1 compares TNNs on differ-
052 ent domains, showing opposite behaviors depending on data, even within the same topological domain.
053

054 Strikingly, despite the high impact of the lifting operation on TNNs, most lifting methods are not
055 supervised and thus not informed by the task at hand (Hajij et al., 2022; Telyatnikov et al., 2024),
056 which may lead to suboptimal architectures. To date, differentiable lifting has only been explored
057 in the context of cell complexes (Battiloro et al., 2023).
058



059 Figure 1: [Left] Lifting to different domains can lead to disparate
060 performances. Accuracies taken from the best TNNs in (Telyat-
061 nikov et al., 2024). [Right] Performances of liftings to the same
062 domain (hypergraph) vary greatly. Values taken from Table 4.
063

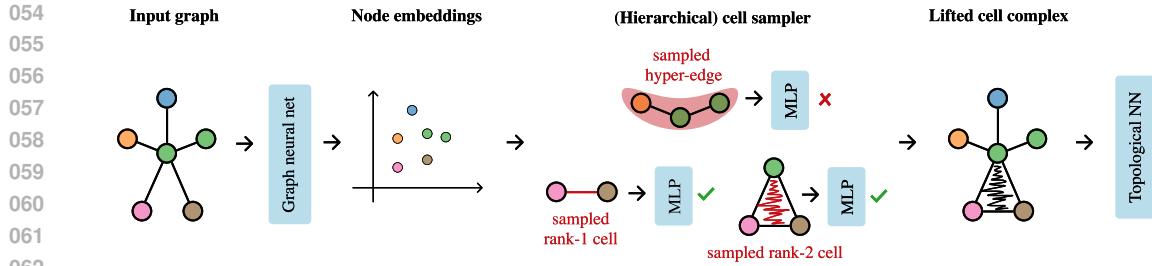


Figure 2: Overview of ∂ lift. For a given input graph, we first compute node embeddings using GNNs. Then, we use these embeddings to select cells/hyperedges. Cell-level embeddings run through MLPs responsible for returning acceptance probabilities. For hierarchical domains (e.g., cell complexes), cells are generated in increasing dimensionality. From the accepted cells, we form a relational object that is sent to an off-the-shelf TNN for graph/node-level predictions. The model is trained end-to-end.

This work proposes ∂ lift (DiffLift) – a general, differentiable lifting framework applicable to various domains, including hypergraphs and simplicial/cell complexes. Our method uses a probabilistic approach to sample candidate cells of adaptive sizes. Specifically, we parameterize distributions over cells using node embeddings derived from arbitrary graph models (e.g., GNNs or Graph Transformers (Rampášek et al., 2022)). For each candidate cell, we compute its embedding and use a multilayer perceptron (MLP) to estimate the probability of accepting or rejecting the cell — that is, determining whether it should be included in the output structure. Figure 2 provides a schematic overview of ∂ lift.

To model the typical hierarchical structure of topological objects, we propose an iterative sampling procedure, where cells are generated in increasing order of dimensionality: samples of dimension i are used to inform the sampling of dimension $(i + 1)$ -cells. Notably, our approach generalizes across multiple topological domains and can be seamlessly integrated into standard TNN pipelines.

We evaluate ∂ lift on 12 datasets spanning graph and node classification tasks using four different TNN models. Our results show that ∂ lift consistently outperforms unsupervised lifting methods in nearly all graph-level classification benchmarks — achieving superior performance in 22 out of 24 experiments, often by a substantial margin. These gains are robust across all TNN architectures. For example, when using CW Networks (Bodnar et al., 2021a), ∂ lift yields performance gains of up to 45%. For node classification, ∂ lift achieves competitive performance relative to static lifting methods and outperforms DCM (Battiloro et al., 2023) (a differentiable lifting baseline) overall. Additionally, we analyze the sensitivity of ∂ lift to the choice of its GNN component, highlighting that while this choice often impacts the overall performance, our design is robust and produces strong empirical results even when adopting simple GNNs (e.g., graph isomorphism networks (Xu et al., 2019)).

2 BACKGROUND

This section overviews the main types of relational structures and respective neighborhood notions, message-passing networks for relational data, and graph lifting methods. In the following, we assume readers are familiar with basic notions in topology; see (Munkres, 2000) for reference.

Graphs and hypergraphs. We denote an *undirected graph* as a tuple $G = (V, E)$ where V is a set of vertices (or nodes) and E is a set of unordered vertex pairs, i.e., edges. The set of neighbors of a node v in G is denoted by $\mathcal{N}^G(v) = \{u \in V : \{v, u\} \in E\}$. Hypergraphs generalize graphs by allowing edges to connect multiple nodes. Formally, a *hypergraph* on a nonempty set V is a pair (V, K) , where $K \subseteq 2^V \setminus \emptyset$ and its elements are called hyperedges.

Simplicial complexes are topological spaces comprised of simple mathematical objects called simplices (points (0-simplices), line segments (1-simplices), triangles (2-simplices), and their higher-dimensional analogues). In particular, an *abstract simplicial complex* (ASC) over a vertex set V is a set K of subsets of V (the *simplices*) such that, for every $\sigma \in K$ and every non-empty $\tau \subset \sigma$, we have that $\tau \in K$. Thus, we can define ASCs as a family of subsets $K \subseteq 2^V$ of V that is closed under taking subsets. The dimension of a simplex is equal to its cardinality minus 1, and the dimension of an ASC is the maximal dimension of its simplices. We say τ is on the boundary

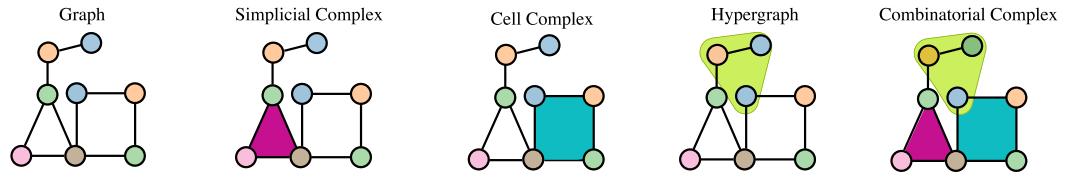


Figure 3: Examples of topological domains.

of a simplex σ , denoted by $\tau \prec \sigma$, iff $\tau \subset \sigma$ and there is no δ such that $\tau \subset \delta \subset \sigma$, i.e., \prec defines the boundary relation of K . We note that undirected graphs correspond to 1-dimensional ASCs.

Cell complexes. A *regular cell complex* (Hansen and Ghrist, 2019) is a topological space X with a partition $\{X_\sigma\}_{\sigma \in P_X}$ of subspaces X_σ of X called *cells* such that

1. For each $x \in X$, there is an open neighborhood of x that intersects finitely many cells;
2. For all $\sigma, \tau \in P_X$, $X_\tau \cap \overline{X_\sigma} \neq \emptyset$ only if $X_\tau \subseteq \overline{X_\sigma}$, where $\overline{X_\sigma}$ denotes the closure of X_σ (the intersection of all closed sets containing X_σ);
3. Every cell X_σ is homeomorphic to \mathbb{R}^{D_σ} for some D_σ which we call X_σ 's dimension;
4. For all $\sigma \in P_X$, there is a homeomorphism ϕ of a closed ball in \mathbb{R}^{D_σ} to $\overline{X_\sigma}$ such that the restriction of ϕ to the interior of the ball is a homeomorphism onto X_σ .

Importantly, the conditions (2) and (4) impose a poset structure $\tau \leq \sigma \iff X_\tau \subseteq \overline{X_\sigma}$ which fully characterizes the topology of the underlying cell complex X . This topological information can be described by the *boundary relation* \prec between two cells: $\sigma \prec \tau$ iff $\sigma < \tau$ and there is no cell δ such that $\sigma < \delta < \tau$, where $<$ denotes the strict version of the partial order \leq above. We note that the class of cell complexes subsumes simplicial complexes. For more details on cell complexes, we refer to Hatcher (2002); Bodnar et al. (2021a).

Combinatorial complexes. A *combinatorial complex* (CC) (Hajij et al., 2022) is a tuple (V, K, rk) where V is a finite set, $K \subseteq 2^V \setminus \emptyset$ comprises a set of cells, and $\text{rk} : K \rightarrow \mathbb{Z}_{\geq 0}$ is a ranking function s.t.

1. For all $v \in V, \{v\} \in K$;
2. For all $\sigma, \sigma' \in K, \sigma \subseteq \sigma' \implies \text{rk}(\sigma) \leq \text{rk}(\sigma')$.

The idea of CCs is to generalize hierarchical structures (e.g., simplicial complexes) by imposing mild relationships between cells via ranking functions — CCs only require the order-preserving property in condition (2) — while being flexible to accommodate non-hierarchical structures such as hypergraphs. Figure 3 depicts the most popular relational structures in topological deep learning.

Neighborhood structures. We can exploit boundary relations (or rank functions) to specify local neighbors for each cell. In particular, Bodnar et al. (2021b) introduce four neighborhood structures:

- Boundary and co-boundary: $\mathcal{N}_B(\sigma) = \{\tau : \tau \prec \sigma\}$ and $\mathcal{N}_C(\sigma) = \{\tau : \sigma \prec \tau\}$, respectively
- Upper/lower adjacency: $\mathcal{N}_\uparrow(\sigma) = \{\tau : \exists \delta \text{ st } \tau \prec \delta, \sigma \prec \delta\}$ and $\mathcal{N}_\downarrow(\sigma) = \{\tau : \exists \delta \text{ st } \delta \prec \tau, \delta \prec \sigma\}$

Analogs of these neighborhoods can also be obtained via ranking functions (Hajij et al., 2022).

Features / signals. In this work, we consider relational structures equipped with features. Let K be a set of cells or hyperedges of a relational domain. Its attributed counterpart is a tuple (K, x) where $x : K \rightarrow \mathbb{R}^d$ assigns a feature vector $x(\sigma)$ to each cell σ . Hereafter, we denote the features of σ by x_σ .

Topological neural networks (TNNs). Most TNNs use message-passing mechanisms to obtain cell-level representations (Papillon et al., 2023b). In particular, let \mathcal{N}_i be a finite sequence of neighborhood structures, $\mathcal{N}_C(\sigma, \tau) = \mathcal{N}_C(\sigma) \cap \mathcal{N}_C(\tau)$, and $\mathcal{N}_B(\sigma, \tau) = \mathcal{N}_B(\sigma) \cap \mathcal{N}_B(\tau)$. In its general form, starting from $h_\sigma^0 = x_\sigma$ for all σ , a message-passing TNN (Bodnar et al., 2021a) recursively computes

$$m_{i,\sigma}^\ell = \begin{cases} \{\phi_{\ell,i}(h_\tau^\ell, h_\sigma^\ell, h_\delta^\ell) : \tau \in \mathcal{N}_i(\sigma), \delta \in \mathcal{N}_B(\sigma, \tau)\}, & \text{if } \mathcal{N}_i = \mathcal{N}_\downarrow \\ \{\phi_{\ell,i}(h_\tau^\ell, h_\sigma^\ell, h_\delta^\ell) : \tau \in \mathcal{N}_i(\sigma), \delta \in \mathcal{N}_C(\sigma, \tau)\}, & \text{if } \mathcal{N}_i = \mathcal{N}_\uparrow \\ \{\phi_{\ell,i}(h_\sigma^\ell) : \tau \in \mathcal{N}_i(\sigma)\}, & \text{otherwise.} \end{cases} \quad (1)$$

$$h_\sigma^{\ell+1} = \varphi \left(h_\sigma^\ell, \bigotimes_i \text{Agg}_\ell(m_{i,\sigma}^\ell) \right) \quad (2)$$

162 where h_σ^ℓ is the embedding of σ at layer ℓ , \otimes and Agg_ℓ are inter- and intra-neighborhood aggregation
 163 functions, respectively, and φ is an update function (e.g., MLP).

164 **Graph lifting.** A *graph lifting* is a map $\text{lift} : \mathbb{G} \rightarrow \mathbb{T}$ from the space of attributed graphs, \mathbb{G} , to a
 165 target domain, \mathbb{T} , such that $G \cong_{\mathbb{G}} G' \implies \text{lift}(G) \cong_{\mathbb{T}} \text{lift}(G')$, where $\cong_{\mathbb{T}}$ denotes the *isomorphism*
 166 relation in domain \mathbb{T} . One of the most widely used methods for lifting graphs to cell complexes is
 167 *cycle lifting* (Bodnar et al., 2021a). This is a static (non-learnable) approach that constructs 2-cells
 168 by identifying basic cycles (elements of a cycle basis) or chordless cycles (Bodnar et al., 2021a) in
 169 input graphs. Specifically, the vertices involved in a basic cycle are grouped to form a 2-cell in the
 170 resulting cell complex. A cycle basis of a graph G is a minimal set of cycles such that any other cycle
 171 in G can be expressed as a modulo-2 sum of cycles from this set.

172

173 3 DIFFERENTIABLE LIFTING

174

175 In this section, we introduce ∂lift (read DiffLift), a general framework for learning graph lifting
 176 functions. Section 3.1 provides an iterative description of our method, allowing for learning structures
 177 of increasingly higher order — when the target domain is hierarchically structured. Section 3.2 and
 178 Section 3.3 instantiate ∂lift for graph-to-hypergraph and graph-to-cell-complex liftings, respectively.
 179 Moreover, we formulate our approach for simplicial- and combinatorial complexes in Appendix A.

180

181

182 3.1 GENERAL FORMULATION

183

184 Lifting consists of determining which higher-order cells should be added to an input graph G ,
 185 satisfying the constraints of the target domain. To do so, we propose the following recipe.

186

187 $\partial\text{lift} : \text{general recipe for differentiable graph liftings}$

188

189 **Input:** Attributed graph $G = (V, E, x)$, target domain \mathbb{T} , and maximum dimension D_{\max} .

190

191 **Step 1: Compute node embeddings.** Use an arbitrary GNN to compute a vector representation
 192 (embedding) z_v for each node $v \in V$. This GNN component can be either a pre-trained
 193 model or learned end-to-end. Set the current domain dimension to $D = 1$.

194

195 **Step 2: Elicit candidate cells.** Given the node embeddings $\{z_v\}_{v \in V}$, define a set of
 196 candidate cells $\mathcal{C} \subseteq 2^V$ of dimension D . For each cell $C \in \mathcal{C}$, compute an embedding
 197 $z_C = \bigoplus_{v \in C} z_v$, where \bigoplus is an arbitrary permutation-invariant aggregation function. Note
 198 that the exact procedure for defining candidate cells depends on the target domain \mathbb{T} , as
 199 candidates must respect possible hierarchical constraints.

200

201 **Step 3: Accept/reject candidate cells.** Apply a neural network ϕ (e.g., an MLP) that defines
 202 an acceptance probability $\phi(z_C)$ for each candidate cell C . Finally, draw a sample y_C from
 203 a Bernoulli distribution with parameter $\phi(z_C)$ indicating whether cell C is accepted or not.
 204 The resulting domain is then given by $V \cup E \cup \{C \in \mathcal{C} : y_C = 1 \text{ with } y_C \sim \text{Ber}(\phi(z_C))\}$.

205

206 **Step 4: Termination check.** If $D = D_{\max}$, halt; otherwise, $D \leftarrow D + 1$ and return to Step 2.

207

208 Importantly, ∂lift is learned in an end-to-end fashion, using the straight-through estimator (Bengio
 209 et al., 2013) to propagate gradients through samples at Step 3. For hypergraphs, we assume
 210 hyperedges have dimension one, causing ∂lift to stop once it reaches Step 4. We note that [Steps 2 and 3 are the only domain-dependent ingredients of our algorithm](#). Next, we explain how these
 211 steps can be adapted to specific domains.

212

213 3.2 GRAPH-TO-HYPERGRAPH LIFTING

214

215 \Rightarrow [Step 2] For notational convenience, suppose we wish to learn up to one hyperedge per node.
 216 For each node v , we define a candidate hyperedge $C(v)$ using the k_v nearest neighbors of v in the
 217 embedding space:

218

$$C(v) = \{S \subset V : |S| = k_v \text{ and } w \notin S \implies \text{dist}(z_w, z_v) \geq \max_{u \in S} \text{dist}(z_u, z_v)\}, \quad (3)$$

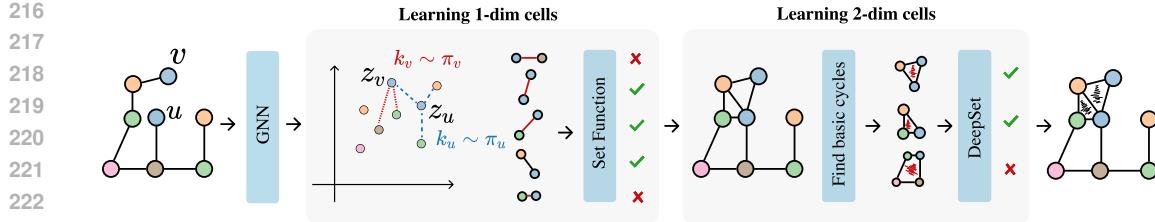


Figure 4: Two iterations of ∂ lift for cell complexes. At the first iteration, we leverage each node u 's GNN embedding z_u to delineate candidate 1-dim cells. Specifically, we consider cell-equivalent of edges linking u to each of its k_u NNs in embedding space, where k_u is a random variable parameterized by z_u . We use a set function over the embedding of nodes within each cell to compute their acceptance probabilities. At the second iteration onwards, we use cycle lifting in our augmented cell complex to elicit candidate cells, whose acceptance probabilities are computed similarly to the first step.

where $\text{dist}(\cdot, \cdot)$ denotes a dissimilarity metric. Here, we consider the Euclidean distance.

To allow for adaptive hyperedge sizes, we sample k_v according to a probability distribution parameterized by (a function of) v 's embedding z_v . More specifically, we define the $(k_{\max} - k_{\min} + 1)$ -dimensional probability vector $\pi_v \propto \exp \circ \text{MLP}(z_v)$ and draw $k_v \sim \text{Categorical}(\pi_v)$, where k_{\min} and k_{\max} are lower- and upper-bounds on k_v .

⇒ [Step 3] We define the probability of acceptance (i.e., of $b_v = 1$) for $C(v)$ as a function of the (multiset of embeddings of) nodes in $C(v)$. More specifically, we define b_v as

$$b_v \sim \text{Ber}(\Psi(\{\{z_u : u \in C(v)\}\})), \quad (4)$$

where Ψ is learned and maps from multisets (i.e., Ψ is order-invariant) of elements in \mathbb{R}^d to $(0, 1)$.

Feature lifting. Each accepted hyperedge $C(v)$ receives a feature vector $x_{C(v)}$ computed as a multiset operation over $\{x_u : u \in C(v)\}$. Specifically, we employ a *scaled sum projection*:

$$x_{C(v)} = \frac{1}{k_v} \sum_{u \in C(v)} x_u, \quad \forall v \text{ such that } b_v = 1. \quad (5)$$

3.3 GRAPH-TO-CELL-COMPLEX LIFTING

For computational reasons, we split the lifting procedure for cell complexes into two cases. We provide an overview of our proposed graph-to-cell-complex lifting in Figure 4.

Case $D = 1$: Learning edges

⇒ [Step 2] Similarly to Step 2 of graph-to-hypergraph lifting, for each node v , we sample a neighborhood size k_v and define a set $C(v) \subseteq V \setminus \{v\}$ containing the nodes associated with the k_v nearest neighbors of v in the embedding space, excluding v itself.

⇒ [Step 3] Next, we construct candidate edges (1-cells) by considering each pair (v, v') and define their probability of acceptance (i.e., of $b_{v,v'} = 1$) as a function of the embeddings of v and v' . Specifically, we set $b_{v,v'} \sim \text{Ber}(\Psi(\{z_v, z_{v'}\}))$, where Ψ is an order-invariant function.

At end of this iteration, the obtained cell complex is given by:

$$K^1 = V \cup E \cup \{\{v, v'\} : v \in V, v' \in C(v) \text{ with } b_{v,v'} = 1\}. \quad (6)$$

Regarding *feature lifting*, we apply scaled sum projection, identically to the hypergraph case.

Case $D \geq 2$: Learning D -cells

To select candidate cells of arbitrary dimension, we need the notion of n -cycles of a cell complex. Let $C_n(K)$ denote the n -chains of the cell complex K equipped with $\mathbb{Z}/2\mathbb{Z}$ -vector space structure. Also let $\partial_n : C_n(K) \rightarrow C_{n-1}(K)$ be the boundary linear map on K . Then, the n -cycles of K are given by $Z_n(K) = \ker(\partial_n)$. We provide further details in the supplementary material.

270 \Rightarrow [Step 2] Let K^{D-1} be the cell complex at the end of iteration $D - 1$. We select a basis for
 271 $(D - 1)$ -set of cycles $Z_{D-1}(K^{D-1})$ in K^{D-1} to serve as candidate cells. Recall a basis for cycles is a
 272 minimal collection of cycles such that any cycle can be written as a module-2 sum of cycles in the basis.
 273 We note that we can also employ chordless cycles to define the 2-cells, as in (Bodnar et al., 2021a).

274 \Rightarrow [Step 3] Let \mathcal{C} be the set of candidate $(D - 1)$ -cycles from Step 2. We define the probability of
 275 accepting $C \in \mathcal{C}$ (i.e., setting $b_C = 1$) using a DeepSet model (Zaheer et al., 2017) over the multiset
 276 of embeddings $\{z_v\}_{v \in C}$ of all nodes v in C . The output complex at this iteration is then
 277

$$278 \quad K^D = K^{D-1} \cup \{C \in \mathcal{C} : b_C = 1\}. \quad (7)$$

280 For simplicity, again, the features of D -cells are obtained via sum projection lifting.
 281

282 **Remark 1** *Despite the generality of ∂ lift, in the experiments we only consider 2-dimensional cell
 283 complexes ($D_{\max} = 2$) and use the algorithm in (Paton, 1969) (available at the toolbox Net-
 284 workX (Hagberg et al., 2008)) to identify basic 1-cycles in graphs. This is mainly due to the
 285 fact that current implementations of TNNs for cell complexes only support 2-dimensional objects —
 286 for instance, see TopoBenchmarkX (Telyatnikov et al., 2024).*

288 **Remark 2** *We can obtain a deterministic version of ∂ lift using a probability threshold, i.e., we
 289 simply set $b_C = \mathbb{1}[\Psi(\cdot) > \gamma]$ with, e.g., $\gamma = 0.5$, for all candidate cells C .*

291 Experiments regarding the deterministic version can be found in the Appendix E.2.
 292

293 4 RELATED WORKS

296 **Topological deep learning.** Traditional graph deep learning methods are limited to modeling only
 297 pairwise interactions, making them unsuitable for capturing higher-order dependencies involving
 298 multiple nodes (Hajij et al., 2022; Papillon et al., 2023a). To address this limitation, a variety of
 299 deep topological learning methods have been developed for hypergraphs (Bai et al., 2021; Yadati
 300 et al., 2019), simplicial complexes (Hajij et al., 2021; Goh et al., 2022; Maggs et al., 2023; Yang
 301 et al., 2022), and cell complexes (Hajij et al., 2022; 2020), the works that are based on the topological
 302 signal processing field (Barbarossa and Sardellitti, 2020; Schaub et al., 2021; Roddenberry et al.,
 303 2022; Sardellitti et al., 2021). (Papillon et al., 2025) also use GNNs to enhance TDL, where the lifted
 304 topological domain is transformed into augmented Hasse graphs. These methods have demonstrated
 305 their effectiveness across several practical applications, including action recognition (Wang et al.,
 306 2022; Hao et al., 2021), bioinformatics (Liu et al., 2022), and neuroscience (Wang et al., 2023).

307 **Liftings to topological domains.** Most relational datasets and benchmarks are defined on discrete
 308 structures such as graphs. To apply topological deep learning methods to these datasets, a trans-
 309 formation process known as lifting is required, which maps discrete data into topological domains
 310 (Telyatnikov et al., 2024; Hajij et al., 2022; Bernárdez et al., 2024). This lifting process can be either
 311 predefined – e.g., based on structural features like node proximity or the presence of cycles – or learned
 312 directly from the data (Battiloro et al., 2023; Kazi et al., 2022). Graph structure learning methods
 313 (Qian et al., 2024; Kazi et al., 2022; Franceschi et al., 2019; Topping et al., 2021; Sun et al., 2023; Chen
 314 et al., 2020; Jin et al., 2020) are closely related to the graph lifting literature and can be interpreted as
 315 instances of graph lifting to graph domain. Our approach represents the most general form of learnable
 316 lifting proposed so far and empirically outperforms the aforementioned methods in many benchmarks.
 317

318 **Static liftings.** To the best of our knowledge, Bodnar et al. (2021b;a) were the first to combine
 319 static liftings and high-order message passing, focusing on simplicial- and cell complexes. These
 320 static liftings embed a graph into a topological domain by, e.g., aggregating each node’s n-hop
 321 neighborhood or by tracing its cycles. The repertoire of static liftings was later broadened by
 322 the ICML TDL challenge (Bernárdez et al., 2024), which added methods based on kNN, Voronoi
 323 decompositions, and random walks. Our work proposes a more flexible, data-driven approach to
 defining liftings, which offers benefits across a range of tasks.

324 **5 EXPERIMENTS**

325

326 In this section, we evaluate ∂ lift on two complementary tasks: graph classification and node classification.
 327 We compare it against broadly used lifting schemes for both hypergraphs and cell complexes. We
 328 also report results across different TNNs within each of these domains. We run experiments using Py-
 329 Torch (Paszke et al., 2017) and PyTorch Geometric (Fey and Lenssen, 2019); our code is anonymously
 330 available at https://anonymous.4open.science/r/tdl_knn_lifting_gnn-AD22.
 331

332 **5.1 GRAPH CLASSIFICATION**

333

334 **Datasets.** We evaluate model performance on six widely used graph-level benchmark datasets for
 335 molecular property prediction: NCI1, NCI109, MUTAG, MOLHIV, PROTEINS, and ZINC (Kersting
 336 et al., 2016; Dwivedi et al., 2023b; Hu et al., 2020). These datasets are standard benchmarks in the
 337 literature for assessing the effectiveness of graph-based models (Dwivedi et al., 2023a; Telyatnikov
 338 et al., 2024). All tasks are binary classification problems, with the exception of ZINC, which is a
 339 regression task. We provide more details regarding datasets in the Appendix.

340 **Baselines.** We compare ∂ lift with four existing graph lifting methods: cycle lifting, k -hop lifting,
 341 k -nearest-neighbor (k -NN) lifting, and kernel lifting. Among these, cycle lifting is the most widely
 342 adopted strategy for graph-to-cell-complex liftings and has become the *de facto* standard in most of
 343 TNNs operating on cell complexes (Telyatnikov et al., 2024). Similarly, k -hop lifting is the predomi-
 344 nant approach for constructing hypergraphs from graphs and is often the sole method considered in re-
 345 cent benchmarks such as (Telyatnikov et al., 2024). We also consider k -NN lifting as it shares similar-
 346 ities with our approach due to the use of k -NN. Finally, we consider kernel lifting, one of the most suc-
 347 cessful approaches in the ICML TDL challenge (Bernárdez et al., 2024). Notably, our choice of base-
 348 lines covers lifting methods based on connectivity (cycle and k -hop liftings), features (k -NN lifting),
 349 and both connectivity and features (kernel lifting). We provide formulations for the baseline liftings
 350 in Appendix B. We consider the following TNNs: CWN, CIN (Bodnar et al., 2021b) and CXN (Hajij
 351 et al., 2020) for cell complexes, and UniGCNII and UniGIN (Huang and Yang, 2021) for hypergraphs.
 352

353 **Evaluation setup.** For ZINC and MOLHIV, we use the publicly available train/val/test data splits; for
 354 the remaining datasets, we use a random 80/10/10% split. We optimize the hyper-parameters of the
 355 lifting methods and take the optimal hyperparameter values from (Telyatnikov et al., 2024) whenever
 356 available; otherwise, we select optimal values based on the optimal results using cycle or k -hop lifting.
 357 We provide further details on the choice of hyperparameters and model selection in the supplementary
 358 material. We compute the mean and standard deviation of the performance metrics (MAE \downarrow for ZINC,
 359 AUC \uparrow for MOLHIV, and accuracy \uparrow for all other datasets) over three independent runs.

360 **Table 1:** Graph classification: ∂ lift vs static liftings. We denote the best-performing model for each
 361 dataset/TNN in bold. For any fixed TNN and dataset, ∂ lift is better than static liftings in 90% of
 362 cases, offering a performance improvement of up to 45%.

363 Domain	364 TNN	365 Lifting	366 NCI1\uparrow	367 NCI109\uparrow	368 MOLHIV\uparrow	369 MUTAG\uparrow	370 Proteins\uparrow	371 ZINC\downarrow
372 Graph	373 GCN	374 -	375 74.45 ± 1.05	376 76.46 ± 1.03	377 74.99 ± 1.09	378 64.91 ± 4.96	379 70.18 ± 1.35	380 0.64 ± 0.04
	381 GIN	382 -	383 76.89 ± 1.75	384 76.90 ± 0.80	385 70.76 ± 2.46	386 80.70 ± 2.48	387 72.50 ± 2.31	388 0.59 ± 0.03
	389 CWN	390 Cycle	391 76.93 ± 1.18	392 76.71 ± 1.34	393 70.15 ± 3.98	394 66.67 ± 12.41	395 69.05 ± 2.95	396 0.46 ± 0.01
	397 ∂ lift	398 79.81 ± 0.40	399 80.55 ± 0.50	400 75.37 ± 0.80	401 85.96 ± 4.96	402 70.54 ± 3.34	403 0.17 ± 0.00	
404 Cellular	405 CXN	406 Cycle	407 72.02 ± 1.69	408 75.01 ± 0.62	409 69.17 ± 1.20	410 61.40 ± 2.48	411 70.83 ± 1.52	412 0.79 ± 0.02
	413 ∂ lift	414 82.08 ± 1.50	415 82.57 ± 0.40	416 74.83 ± 1.96	417 84.21 ± 4.30	418 69.94 ± 2.10	419 0.17 ± 0.01	
	420 CIN	421 Cycle	422 75.91 ± 1.11	423 76.11 ± 1.09	424 68.46 ± 2.16	425 66.96 ± 1.46	426 67.86 ± 0.89	427 0.42 ± 0.01
	428 ∂ lift	429 79.59 ± 1.50	430 81.06 ± 0.40	431 72.37 ± 1.65	432 88.72 ± 4.30	433 72.43 ± 2.10	434 0.20 ± 0.01	
435 Hypergraph	436 UniGCN2	437 k -hop	438 72.70 ± 0.52	439 72.01 ± 1.55	440 50.72 ± 1.06	441 61.40 ± 2.48	442 72.92 ± 1.11	443 0.66 ± 0.02
		444 k -NN	445 71.78 ± 0.20	446 68.60 ± 0.93	447 57.73 ± 6.84	448 64.91 ± 2.48	449 73.51 ± 0.42	450 1.10 ± 0.01
		451 kernel	452 73.80 ± 0.94	453 72.64 ± 0.40	454 57.07 ± 10.32	455 63.16 ± 8.59	456 73.21 ± 0.73	457 0.79 ± 0.02
		458 ∂ lift	459 77.45 ± 1.88	460 75.30 ± 1.10	461 69.32 ± 1.62	462 89.47 ± 4.30	463 73.51 ± 0.84	464 0.56 ± 0.03
465 Hypergraph	466 UniGIN	467 k -hop	468 65.50 ± 1.99	469 66.97 ± 7.25	470 63.49 ± 9.55	471 64.91 ± 2.48	472 71.43 ± 0.73	473 1.15 ± 0.01
		474 k -NN	475 72.83 ± 1.09	476 70.14 ± 1.48	477 52.34 ± 3.21	478 59.65 ± 4.96	479 72.62 ± 1.52	480 1.10 ± 0.02
		481 kernel	482 60.50 ± 1.26	483 66.59 ± 1.49	484 49.60 ± 0.07	485 57.89 ± 4.30	486 66.67 ± 1.83	487 1.45 ± 0.02
		488 ∂ lift	489 64.88 ± 1.09	490 79.74 ± 0.23	491 72.04 ± 0.88	492 66.67 ± 6.56	493 73.81 ± 1.52	494 0.92 ± 0.05

378 **Results.** Table 1 shows that ∂ lift is the best-performing lifting method in over 90% of the TNN/dataset
 379 combinations, both for cell complexes and hypergraphs. Notably, ∂ lift resulted in an improvement in
 380 average accuracy of up to $\approx 45\%$ compared to static liftings using the same TNN. For CWN, CIN, and
 381 UniGCNII, our method outperforms static liftings on all datasets. On NCI109 and ZINC, ∂ lift is con-
 382 sistently better than the static liftings across all TNN backbones. We also observe that TNNs perform
 383 better than GNNs in this setup with ∂ lift. Overall, these results validate the effectiveness ∂ lift.

384 **Impact of GNN choice on performance.** We aim to assess how sensitive our approach is to the
 385 choice of GNN. We report results using GIN (Xu et al., 2019) and GPS (Rampášek et al., 2022).
 386

387 Table 2 indicates that choosing GNNs that are able to generate richer and more informative latent
 388 node representations leads to better results in ∂ lift. In particular, GPS performs better than GIN in
 389 most datasets. A possible explanation for this observation is the greater expressivity of GPS, which
 390 benefits from the incorporation of positional encodings. Notably, on cell complexes and ZINC dataset,
 391 GPS allows reducing the MAE from 0.46 to 0.17.

392 Table 2: Effect of GNN backbone on the performance of ∂ lift. The results suggest that the expressive
 393 power of backbone GNNs have a direct impact in ∂ lift’s performance. Except for MOLHIV and
 394 NCI1, GPS leads to better performance than GIN overall.

TNN	GNN	NCI1 \uparrow	NCI109 \uparrow	MOLHIV \uparrow	MUTAG \uparrow	Proteins \uparrow	ZINC \downarrow
CWN	GPS	79.81 \pm 0.40	80.55 \pm 0.50	64.31 \pm 5.32	87.72 \pm 2.48	70.54 \pm 3.34	0.17 \pm 0.00
	GIN	81.59 \pm 0.80	78.69 \pm 1.43	75.37 \pm 0.80	82.46 \pm 2.48	71.13 \pm 2.76	0.46 \pm 0.00
CWN	GPS	79.97 \pm 0.61	82.57 \pm 0.40	65.58 \pm 4.42	84.21 \pm 4.30	69.94 \pm 2.10	0.17 \pm 0.01
	GIN	81.35 \pm 2.29	79.98 \pm 1.01	72.25 \pm 3.23	77.19 \pm 2.48	67.86 \pm 2.19	0.43 \pm 0.01
UniGCN2	GPS	78.67 \pm 1.46	74.50 \pm 1.16	68.22 \pm 2.38	89.47 \pm 4.30	73.81 \pm 0.42	0.56 \pm 0.03
	GIN	75.38 \pm 1.39	74.98 \pm 1.12	68.73 \pm 2.05	64.91 \pm 6.56	73.51 \pm 0.84	0.63 \pm 0.01
UniGIN	GPS	66.42 \pm 1.79	79.74 \pm 0.23	68.32 \pm 3.12	66.67 \pm 6.56	72.32 \pm 0.73	1.01 \pm 0.05
	GIN	64.40 \pm 0.41	78.53 \pm 0.69	68.86 \pm 3.05	70.18 \pm 6.56	72.02 \pm 3.74	1.12 \pm 0.26

408 5.2 NODE CLASSIFICATION

409 **Datasets.** For node classification, We evaluate ∂ lift on four datasets: Cora , Citeseer Yang et al.
 410 (2016), Texas, and Winsconsin Rozemberczki et al. (2020). Within these datasets, two are knowingly
 411 homophilic (Cora and Citeseer) and two are heterophilic datasets (Texas and Wisconsin). Dataset
 412 statistics can be found in Appendix D.

413 **Baselines.** We also compare our method (for cell domains) against the learnable approach in
 414 (Battiloro et al., 2023), called Differentiable Cell Complex Module (DCM), which was originally
 415 evaluated on node classification tasks. To do so, we consider ∂ lift combined with CWN and
 416 TopoTune (Papillon et al., 2025). We also include results of them with cycle lifting. We consider
 417 the same hypergraph TNN baselines as in Section 5.1.

418 **Evaluation setup.** For all datasets, we use random train/val/test data split with 60/20/20% split.
 419 Similarly to the experiments for graph classification, we optimize the hyper-parameters of the lifting
 420 methods and take the optimal TNN hyperparameters from (Telyatnikov et al., 2024) when available.
 421 Otherwise, we choose them to maximize the validation accuracy using k -hop lifting. For more
 422 details, please refer to the supplementary material. We report the average accuracy and standard
 423 deviation over three independent runs.

424 **Results.** Table 3 compares ∂ lift against DCM and cycle lifting. Notably, ∂ lift is the best-performing
 425 method in all datasets except for Wisconsin, in which it achieves the second-best performance. It is
 426 also worth mentioning that ∂ lift (with either CWN or TopoTune) outperforms DCM for all datasets,
 427 sometimes by a large margin — c.f., Texas and Cora.

428 Table 4 reports results of lifting methods for hypergraph neural networks. Compared to k -hop, our
 429 approach is better on heterophilic datasets but worse on homophilic ones for UniGCN2. Additionally,
 430 Table 4 shows that ∂ lift leads to better average accuracy than other static liftings.

432 Table 3: Comparison of DCM and ∂ lift on node classification. ∂ lift achieves the highest average
 433 performance across all datasets (rhs) and significantly outperforms DCM on heterophilic datasets.
 434

TNN	Lifting	Cora	Citeseer	Texas	Wisconsin	Avg
GCN	-	85.64 \pm 0.51	70.43 \pm 0.71	58.91 \pm 0.76	49.14 \pm 0.66	66.03
GAT	-	86.17 \pm 0.33	73.82 \pm 0.45	58.38 \pm 1.05	49.41 \pm 0.95	66.95
GIN	-	85.50 \pm 0.54	72.20 \pm 0.60	59.10 \pm 0.80	48.50 \pm 0.70	66.33
DCM	-	80.73 \pm 0.33	77.90 \pm 0.80	56.76 \pm 6.62	73.86 \pm 1.85	72.31
CWN	Cycle	74.80 \pm 0.08	75.83 \pm 0.90	63.06 \pm 7.75	80.39 \pm 4.24	73.52
CWN	∂ Lift	80.17 \pm 1.59	72.83 \pm 2.15	80.18 \pm 3.37	77.78 \pm 3.70	77.74
TopoTune	Cycle	69.03 \pm 0.88	72.90 \pm 0.85	71.56 \pm 1.27	70.16 \pm 1.85	70.91
TopoTune	∂ Lift	86.82 \pm 0.75	78.23 \pm 1.08	72.97 \pm 0.00	65.36 \pm 2.45	75.84

445 Table 4: Comparison of ∂ lift and static lifting baselines for hypergraphs on node classification. Our
 446 method outperforms all static liftings on average across the selected node classification datasets.
 447

TNN	Lifting	Cora	Citeseer	Texas	Wisconsin	Avg
UniGCN2	k -hop	86.03 \pm 0.63	78.40 \pm 0.36	66.67 \pm 6.74	69.28 \pm 5.62	75.09
	k -NN	74.00 \pm 0.65	18.17 \pm 0.05	70.27 \pm 3.82	79.74 \pm 2.45	60.54
	kernel	29.93 \pm 1.33	18.07 \pm 0.21	57.66 \pm 7.75	57.52 \pm 10.42	40.79
	∂ Lift	81.93 \pm 1.11	78.03 \pm 0.91	69.37 \pm 2.55	73.20 \pm 5.62	75.63
UniGIN	k -hop	78.73 \pm 0.66	74.47 \pm 1.72	65.77 \pm 9.19	58.82 \pm 5.77	69.44
	k -NN	62.00 \pm 1.08	19.33 \pm 0.48	65.77 \pm 1.27	73.20 \pm 4.03	55.07
	kernel	40.93 \pm 2.52	18.53 \pm 0.61	58.56 \pm 7.09	51.63 \pm 4.03	42.41
	∂ Lift	84.23 \pm 0.53	77.97 \pm 0.45	63.96 \pm 6.37	63.40 \pm 4.03	72.39

458 6 CONCLUSION

461 Topological neural networks (TNNs) are receiving increasing attention in the graph machine learning
 462 community. Yet, their effectiveness depends crucially on the choice of graph lifting procedure.
 463 Despite its central role, lifting has remained largely unsupervised and task-agnostic, which can lead
 464 to the construction of suboptimal topological representations for downstream learning.

465 To address this limitation, we introduced ∂ lift, a general-purpose, differentiable lifting framework
 466 that is compatible with multiple topological domains. Across a broad set of benchmarks and TNN
 467 architectures, ∂ lift consistently outperformed traditional unsupervised lifting methods, demonstrating
 468 the benefit of making the lifting process learnable and task-informed.

469 **Limitations.** For hypergraph domains, ∂ lift can create candidate hyper-edges and decide whether
 470 to keep them in embarrassingly parallel fashion — rendering ∂ lift especially compute-efficient for
 471 this domain. However, for cell complexes, we need to compute a cycle basis to elicit candidate
 472 cells, which may come at a cubic with respect to the number of nodes in the input graph. In this
 473 case, we may reduce the number of candidate cells by, for instance, regularizing the k_v variables or
 474 shifting their distribution towards zero. Nonetheless, devising more efficient algorithms for candidate
 475 identification in hierarchical domains is a clear direction of improvement for future works.

476 **Future work.** Our method can be extended to other topological domains, such as point clouds,
 477 making it applicable to 3D mapping tasks. Additionally, future work could focus on addressing
 478 the computational challenges of DiffLift and scaling it to handle larger graphs. Another promising
 479 direction is to explore differentiable lifting in dynamic or temporal graphs, where topological
 480 structures evolve over time. Moreover, integrating ∂ lift with pretraining strategies could yield
 481 generalizable topological priors across tasks.

482 We also believe that formally analyzing the impact of enriching topological structures with learnable
 483 liftings on mitigating oversmoothing and oversquashing in TNNs is an interesting research direction.
 484 For instance, in scenarios where long and narrow paths connect dense substructures, static liftings
 485 are limited to adding cycles within each local region, leaving information between communities to
 486 traverse the original bottleneck. In contrast, ∂ lift can learn to introduce 2-cells that effectively create

486
 487 shortcuts across such bridges, reducing the effective distance between distant nodes and improving
 488 information flow. While a full theoretical treatment is left for future work, this adaptive ability
 489 provides a plausible mechanism for mitigating oversquashing.

490 **Ethics Statement.** We do not identify any immediate, direct societal harms from the technical
 491 contributions presented in this work. Our method operates on standard, non-sensitive benchmarks
 492 and does not require or expose personally identifiable information.

493 **Reproducibility Statement.** An anonymized repository containing the code is available at https://anonymous.4open.science/r/tdl_knn_lifting_gnn-AD22. We provide further
 494 495 details on used datasets and implementation details (e.g., parameter selection) in Appendices D and C.

496
497 REFERENCES

498 Song Bai, Feihu Zhang, and Philip HS Torr. Hypergraph convolution and hypergraph attention. *Pattern*
 499 Recognition, 110:107637, 2021.

500
 501 Sergio Barbarossa and Stefania Sardellitti. Topological signal processing over simplicial complexes. *IEEE*
 502 Transactions on Signal Processing, 68:2992–3007, 2020.

503 Claudio Battiloro, Indro Spinelli, Lev Telyatnikov, Michael Bronstein, Simone Scardapane, and Paolo Di Lorenzo.
 504 From latent graph to latent topology inference: Differentiable cell complex module. *arXiv preprint*
 505 *arXiv:2305.16174*, 2023.

506 Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic
 507 neurons for conditional computation. *ArXiv e-prints*, 2013.

509 Guillermo Bernárdez, Lev Telyatnikov, Marco Montagna, Federica Baccini, Mathilde Papillon, Miquel Ferriol-
 510 Galmés, Mustafa Hajij, Theodore Papamarkou, Maria Sofia Bucarelli, Olga Zaghen, et al. Icml topological
 511 deep learning challenge 2024: Beyond the graph domain. *arXiv preprint arXiv:2409.05211*, 2024.

512 Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yuguang Wang, Pietro Liò, Guido F Montufar, and Michael
 513 Bronstein. Weisfeiler and lehman go cellular: CW networks. In *Advances in Neural Information Processing*
 514 Systems (NeurIPS), 2021a.

515 Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lio, and Michael
 516 Bronstein. Weisfeiler and lehman go topological: Message passing simplicial networks. In *International*
 517 *Conference on Machine Learning (ICML)*, 2021b.

518 Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural networks: Better and
 519 robust node embeddings. *Advances in neural information processing systems*, 33:19314–19326, 2020.

521 V. P. Dwivedi, C. K. Joshi, A. T. Luu, T. Laurent, Y. Bengio, and X. Bresson. Benchmarking graph neural
 522 networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023a.

523 Vijay Prakash Dwivedi, Chaitanya K Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson.
 524 Benchmarking graph neural networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023b.

526 M. Fey and J. E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *Workshop track of the*
 527 *International Conference on Representation Learning (ICLR)*, 2019.

528 Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiao He. Learning discrete structures for graph
 529 neural networks. In *International conference on machine learning*, pages 1972–1982. PMLR, 2019.

531 J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum
 532 chemistry. In *International Conference on Machine Learning (ICML)*, 2017.

533 Christopher Wei Jin Goh, Cristian Bodnar, and Pietro Lio. Simplicial attention networks. *arXiv preprint*
 534 *arXiv:2204.09455*, 2022.

536 Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function
 537 using networkx. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th*
 538 *Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.

539 Mustafa Hajij, Kyle Istvan, and Ghada Zamzmi. Cell complex neural networks. *arXiv preprint arXiv:2010.00743*,
 2020.

540 Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Vasileios Maroulas, and Xuanting Cai. Simplicial
 541 complex representation learning. *arXiv preprint arXiv:2103.04046*, 2021.

542

543 Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, Karthikeyan Natesan
 544 Ramamurthy, Tolga Birdal, Tamal K Dey, Soham Mukherjee, Shreyas N Samaga, et al. Topological deep
 545 learning: Going beyond graph data. *arXiv preprint arXiv:2206.00606*, 2022.

546 J. Hansen and R. Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational
 547 Topology*, 2019.

548 XIAOKE HAO, JIE LI, YINGCHUN GUO, TAO JIANG, AND MING YU. Hypergraph neural network for skeleton-based
 549 action recognition. *IEEE Transactions on Image Processing*, 30:2263–2275, 2021.

550

551 Allen Hatcher. *Algebraic topology*. Cambridge University Press, 2002.

552 Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure
 553 Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*,
 554 2020.

555 Jing Huang and Jie Yang. UniGNN: a unified framework for graph and hypergraph neural networks. *arXiv
 556 preprint arXiv:2105.00956*, 2021.

557

558 Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. Graph structure learning
 559 for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on
 560 knowledge discovery & data mining*, pages 66–74, 2020.

561 Anees Kazi, Luca Cosmo, Seyed-Ahmad Ahmadi, Nassir Navab, and Michael M Bronstein. Differentiable
 562 graph module (dgm) for graph convolutional networks. *IEEE Transactions on Pattern Analysis and Machine
 563 Intelligence*, 45(2):1606–1617, 2022.

564 Kristian Kersting, Nils M. Kriege, Christopher Morris, Petra Mutzel, and Marion Neumann. Benchmark data
 565 sets for graph kernels, 2016. URL <http://graphkernels.cs.tu-dortmund.de>.

566

567 Xuan Liu, Congzhi Song, Shichao Liu, Menglu Li, Xionghui Zhou, and Wen Zhang. Multi-way relation-
 568 enhanced hypergraph representation learning for anti-cancer drug synergy prediction. *Bioinformatics*, 38(20):
 569 4782–4789, 2022.

570 Kelly Maggs, Celia Hacker, and Bastian Rieck. Simplicial representation learning with neural k -forms. *arXiv
 571 preprint arXiv:2312.08515*, 2023.

572 J.R. Munkres. *Topology*. Featured Titles for Topology. Prentice Hall, Incorporated, 2000. ISBN 9780131816299.
 573 URL <https://books.google.fi/books?id=XjoZAQAAIAAJ>.

574

575 Theodore Papamarkou, Tolga Birdal, Michael M. Bronstein, Gunnar E. Carlsson, Justin Curry, Yue Gao, Mustafa
 576 Hajij, Roland Kwitt, Pietro Lio, Paolo Di Lorenzo, Vasileios Maroulas, Nina Miolane, Farzana Nasrin,
 577 Karthikeyan Natesan Ramamurthy, Bastian Rieck, Simone Scardapane, Michael T Schaub, Petar Veličković,
 578 Bei Wang, Yusu Wang, Guowei Wei, and Ghada Zamzmi. Position: Topological deep learning is the new
 579 frontier for relational learning. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria
 580 Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on
 581 Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 39529–39555. PMLR,
 582 21–27 Jul 2024.

583 Mathilde Papillon, Mustafa Hajij, Audun Myers, Helen Jenne, Johan Mathe, Theodore Papamarkou, Aldo
 584 Guzmán-Sáenz, Neal Livesay, Tamal Dey, Abraham Rabinowitz, et al. Icml 2023 topological deep learning
 585 challenge: Design and results. In *Topological, Algebraic and Geometric Learning Workshops 2023*, pages
 3–8. PMLR, 2023a.

586 Mathilde Papillon, Sophia Sanborn, Mustafa Hajij, and Nina Miolane. Architectures of topological deep learning:
 587 A survey on topological neural networks. *ArXiv e-prints*, 2023b.

588

589 Mathilde Papillon, Guillermo Bernárdez, Claudio Battiloro, and Nina Miolane. Topotune : A framework for
 590 generalized combinatorial complex neural networks, 2025. URL <https://arxiv.org/abs/2410.06530>.

591

592 A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer.
 593 Automatic differentiation in pytorch. In *Advances in Neural Information Processing Systems (NeurIPS -
 594 Workshop)*, 2017.

594 Keith Paton. An algorithm for finding a fundamental set of cycles of a graph. *Commun. ACM*, 12(9):514–518,
 595 September 1969. ISSN 0001-0782. doi: 10.1145/363219.363232. URL <https://doi.org/10.1145/363219.363232>.

597 Chendi Qian, Andrei Manolache, Christopher Morris, and Mathias Niepert. Probabilistic graph rewiring via
 598 virtual nodes. *arXiv preprint arXiv:2405.17311*, 2024.

599 Ladislav Rampášek, Mikhail Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique
 600 Beaini. Recipe for a general, powerful, scalable graph transformer. In *Proceedings of the 36th International
 601 Conference on Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran
 602 Associates Inc. ISBN 9781713871088.

603 T Mitchell Roddenberry, Michael T Schaub, and Mustafa Hajij. Signal processing on cell complexes. In *ICASSP
 604 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages
 605 8852–8856. IEEE, 2022.

606 Benedek Rozemberczki, Carl Allen, and Rik Sarkar. Multi-scale attributed node embedding, 2020. URL
 607 <https://openreview.net/forum?id=HJxiMAVtPH>.

608 Stefania Sardellitti, Sergio Barbarossa, and Lucia Testa. Topological signal processing over cell complexes. In
 609 *2021 55th Asilomar Conference on Signals, Systems, and Computers*, pages 1558–1562. IEEE, 2021.

610 F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE
 611 Transactions on Neural Networks*, 20(1):61–80, 2009.

612 Michael T Schaub, Yu Zhu, Jean-Baptiste Seby, T Mitchell Roddenberry, and Santiago Segarra. Signal processing
 613 on higher-order networks: Livin’ on the edge... and beyond. *Signal Processing*, 187:108149, 2021.

614 Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective
 615 classification in network data. *AI Magazine*, 29(3):93, Sep. 2008. doi: 10.1609/aimag.v29i3.2157. URL
 616 <https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/2157>.

617 Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. Pitfalls of graph
 618 neural network evaluation. *Relational Representation Learning Workshop, NeurIPS 2018*, 2018.

619 Qingyun Sun, Jianxin Li, Beining Yang, Xingcheng Fu, Hao Peng, and Philip S Yu. Self-organization preserved
 620 graph structure learning with principle of relevant information. In *Proceedings of the AAAI Conference on
 621 Artificial Intelligence*, volume 37, pages 4643–4651, 2023.

622 Lev Telyatnikov, Guillermo Bernardez, Marco Montagna, Pavlo Vasylenko, Ghada Zamzmi, Mustafa Hajij,
 623 Michael T Schaub, Nina Miolane, Simone Scardapane, and Theodore Papamarkou. Topobenchmarkx: A
 624 framework for benchmarking topological deep learning. *arXiv preprint arXiv:2406.06642*, 2024.

625 Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M Bronstein.
 626 Understanding over-squashing and bottlenecks on graphs via curvature. *arXiv preprint arXiv:2111.14522*,
 627 2021.

628 Yogesh Verma, Amauri H Souza, and Vikas Garg. Topological neural networks go persistent, equivariant, and
 629 continuous. In *International Conference on Machine Learning (ICML)*, 2024.

630 Cheng Wang, Nan Ma, Zhixuan Wu, Jin Zhang, and Yongqiang Yao. Survey of hypergraph neural networks
 631 and its application to action recognition. In *CAAI International Conference on Artificial Intelligence*, pages
 632 387–398. Springer, 2022.

633 Junqi Wang, Hailong Li, Gang Qu, Kim M Cecil, Jonathan R Dillman, Nehal A Parikh, and Lili He. Dynamic
 634 weighted hypergraph convolutional network for brain functional connectome analysis. *Medical image analysis*,
 635 87:102828, 2023.

636 Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks?
 637 In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ryGs6iA5Km>.

638 Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar.
 639 Hypergcn: A new method for training graph convolutional networks on hypergraphs. *Advances in neural
 640 information processing systems*, 32, 2019.

641 Maosheng Yang, Elvin Isufi, and Geert Leus. Simplicial convolutional neural networks. In *ICASSP 2022-2022
 642 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8847–8851.
 643 IEEE, 2022.

648 Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph
 649 embeddings. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International*
 650 *Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 40–48,
 651 New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/yang16.html>.

653 M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. Salakhutdinov, and A. Smola. Deep sets. In *Advances in*
 654 *Neural Information Processing Systems (NeurIPS)*, 2017.

656 A ADDITIONAL BACKGROUND AND FORMULATIONS

658 A.1 CHAINS, BOUNDARY OPERATORS, AND CYCLES

660 Here, we introduce some basic notions in algebraic topology. For simplicity, our exposition considers
 661 abstract simplicial complexes (ASCs) equipped with coefficients in the finite field $\mathbb{Z}/2\mathbb{Z} = \{0, 1\}$.

663 The space of n -chains is the vector space of all formal sums of n -dimensional simplices of an ASC
 664 K . Formally, let $n \geq 0$ and $K_{(n)} = \{\sigma \in K : \dim(\sigma) = n\}$ be the n -skeleton of K . The n -chains
 665 of K is the set $C_n(K)$ whose elements take the form

$$666 \quad \sum_{\sigma \in K_{(n)}} \epsilon_{\sigma} \sigma \quad (8)$$

669 where for all $\sigma \in K_{(n)}$, $\epsilon_{\sigma} \in \mathbb{Z}/2\mathbb{Z}$.

670 Let $c = \sum_{\sigma \in K_{(n)}} \epsilon_{\sigma} \sigma$ and $c' = \sum_{\sigma \in K_{(n)}} \epsilon'_{\sigma} \sigma$ be two n -chains. The sum of two chains $(c + c')$ and
 671 the product of a chain by a scalar (λc) are respectively defined by

$$673 \quad c + c' = \sum_{\sigma} (\epsilon_{\sigma} + \epsilon'_{\sigma}) \sigma \quad (9)$$

$$675 \quad \lambda c = \sum_{\sigma} (\lambda \epsilon_{\sigma}) \sigma \quad (10)$$

677 where sums and products are module 2.

679 We define the boundary of a n -simplex σ , denoted by $\partial_n \sigma$ as the sum of its constituents $(n - 1)$ -
 680 simplices, i.e.,

$$681 \quad \partial_n \sigma = \sum_{\tau \subset \sigma : |\tau| = |\sigma| - 1} \tau \quad (11)$$

684 This boundary extends linearly to chain spaces. In particular, the boundary operator ∂_n is a linear
 685 map $\partial_n : C_n(K) \rightarrow C_{n-1}(K)$ defined by

$$687 \quad \partial_n c = \partial_n \sum_{\sigma \in K_{(n)}} \epsilon_{\sigma} \sigma = \sum_{\sigma \in K_{(n)}} \epsilon_{\sigma} \partial_n \sigma. \quad (12)$$

690 Finally, we can define n -cycles. For $n \geq 0$, the n -cycles of K is the set $Z_n(K)$ given by the kernel
 691 of ∂_n , that is

$$692 \quad Z_n(K) = \{c \in C_n(K) : \partial_n c = 0\}. \quad (13)$$

694 A.2 ∂ lift FOR SIMPLICIAL COMPLEXES

696 Note that for when $D = 1$ — i.e., when we must decide which edges to add — Steps 2 and 3 of
 697 ∂ lift for cell-complexes naturally result in a simplicial complex. To fully specify ∂ lift for simplicial
 698 complexes, we are left with defining these steps when $D > 1$.

699 Case $D \geq 2$: Learning D -simplices

701 \Rightarrow [Step 2] When creating simplices of dimension D , we must ensure they respect the hierarchical
 702 structure of simplicial complexes. Let K^{ℓ} be the cell complex at the end of iteration $\ell \leq D - 1$. To

702 identify a preliminary set of candidates \mathcal{C}' , we run static D -clique lifting on K^1 . For $D > 2$, it is
 703 possible that a lower-order clique within some $C \in \mathcal{C}'$ does not belong to K^{D-1} . Therefore, we must
 704 filter out these elements, defining a refined set of candidates:
 705

$$706 \quad \mathcal{C} = \{C \in \mathcal{C}' \mid S \in K^{D-1} \text{ for all } S \subset C\} \quad (14)$$

708 \Rightarrow [Step 3] Similarly to this respective step for cell complexes, we define the probability of accepting
 709 $C \in \mathcal{C}$ (i.e., setting $b_C = 1$) applying a DeepSet over the embeddings $\{z_v\}_{v \in C}$, subsequently
 710 sampling the Bernoulli variables $\{b_C\}_{C \in \mathcal{C}}$. The output complex at this iteration is then
 711

$$712 \quad K^D = K^{D-1} \cup \{C \in \mathcal{C} : b_C = 1\}. \quad (15)$$

713 We define the features for D -simplices using sum projection lifting.
 714

715 A.3 ∂ lift FOR COMBINATORIAL COMPLEXES

717 There are multiple ways to combine cell complexes with hypergraphs to obtain valid combinatorial
 718 complexes (CC). Here, we would like to preserve the property that hyperedges exchange messages
 719 with nodes via boundary (or lower incidences) neighborhoods. Thus, we propose first running ∂ lift to
 720 either cell or simplicial complexes — where ranking functions are given by cell/simplex dimensions.
 721 Let K be the resulting complex. Then, we employ (in parallel) ∂ lift to a hypergraph H , where
 722 edges/hyperedges have rank 1. To ensure a valid combinatorial complex, we prune the sampled
 723 hyperedges to include only those that are not supersets of any cell of rank greater than 1 in K .
 724 Formally, the resulting CC is given by $\{h \in H : \nexists \sigma \in K \text{ s.t. } \sigma \subseteq h\} \cup K$.
 725

726 B TOPOLOGICAL LIFTINGS

728 **Clique lifting.** The set of cliques in a graph G is given by $Cl(G) = \{c \subseteq V(G) : u \neq v \in c \implies$
 729 $\{u, v\} \in E(G)\}$, i.e., each element of $Cl(G)$ is a complete subgraph of G . The k -cliques of G are
 730 the elements of $Cl(G)$ of size k , for $k > 1$, and we denote them as $Cl_k(G)$. Formally, the k -clique
 731 lifting operation is given by
 732

$$733 \quad \text{lift}_{\text{clique}, k}(G) = V(G) \cup \bigcup_{i=2}^k Cl_i(G). \quad (16)$$

735 Note that the inclusion of all cliques of size smaller than k ensures the function returns a valid abstract
 736 simplicial complex.
 737

738 **Cycle lifting.** The idea of cycle lifting is to identify basic cycles in the input graph and use the tuple
 739 of vertices in a cycle as a 2-rank cell of the output complex.
 740

741 Let us consider modulo-2 sum operations for vertices and edges. Also, let ∂_1 be the edge boundary
 742 map for a graph G , i.e., $\partial_1(\{u, v\}) = \{u\} + \{v\}$ for any edge $\{u, v\} \in E(G)$. The cycles of G are
 743 $L(G) = \{l \subseteq E(G) : \sum_{e \in l} \partial_1(e) = 0\}$.
 744

745 A basis for cycles of G is a minimal collection of cycles such that any cycle in G can be written
 746 as a sum of cycles in the basis — i.e., the smallest set $B \subseteq L$ such that $\forall l \in L, \exists B' \subseteq B$ with
 747 $l = \sum_{b \in B'} b$. The cycle lifting map is
 748

$$749 \quad \text{lift}_{\text{cycle}}(G) = V(G) \cup E(G) \cup \{V(b) : b \in B(G)\}, \quad (17)$$

750 where $V(b)$ denotes the set of vertices in the cycle b , and $\{V(b) : b \in B(G)\}$ is the set of 2-dim cells.
 751

752 **DCM.** Battiloro et al. (2023) proposed a novel layer composed of several modules, with the
 753 Differentiable Cell Complex Module (DCM) being central to latent topology inference. The DCM
 754 first samples the 1-skeleton of the latent cell complex using the α -Differentiable Graph Module
 755 (α -DGM). It then selects polygons—representing higher-order interactions—formed by cycles in the
 756 sampled graph using the Polygon Inference Module (PIM). For a detailed description of α -DGM and
 757 PIM, we refer the reader to Section 3 of Battiloro et al. (2023).
 758

756 ***k*-hop lifting.** The k -hop neighborhood of a node $v \in V(G)$ is defined as
 757

$$758 \quad N_k[v] = \{u \in V(G) : \text{dist}(u, v) \leq k\}, \quad (18)$$

759 where $\text{dist}(u, v)$ is the shortest-path distance in the graph G , measured by the number of edges in
 760 the path.

761 To construct the k -hop hypergraph H from G , a hyperedge is formed for each node $v \in V(G)$ based
 762 on its k -hop neighborhood:

$$763 \quad \text{lift}_{k\text{-hop}}(G) = \{N_k[v] : v \in V(G)\}.$$

764 One can note that when $k = 1$, k -hop is equal to neighborhood lifting. The parameter k controls the
 765 extent of the neighborhoods included as hyperedges, with larger k values progressively incorporating
 766 nodes farther away in terms of shortest-path distance.

767 ***k*-NN lifting.** k -NN lifting constructs hyperedges by identifying the k nearest neighbors based on
 768 their node features (feature space). For every node, a separate hyperedge is formed that includes the
 769 node itself and its k closest neighbors.

770 **Kernel lifting.** Kernel lifting is a procedure that constructs hyperedges based on similarity measures
 771 derived from kernels over graph nodes. These kernels can be defined in three ways: (i) over the graph
 772 structure itself, (ii) over the node features, or (iii) as a composition that jointly incorporates both graph
 773 and feature information. For a given reference node v , the method computes similarities between
 774 v and all other nodes v^* using a kernel function. A hyperedge is then formed by selecting a fixed
 775 fraction (typically 0.5) of the nodes that are most similar to v according to the chosen kernel. This
 776 process is repeated for each node to construct a set of hyperedges. The kernels can be defined several
 777 forms: over nodes $K_g(v, v^*)$, features $K_x(x, x^*)$, or over nodes and features $C(K(x, x^*, v, v^*))$,
 778 where C is a valid composition function. Kernels over features are calculated as standard RBF or
 779 exponential kernels Duvenaud (2014), whereas kernels over graphs can be calculated as heat or
 780 Matérn kernels Schölkopf and Smola (2002); Borovitskiy et al. (2021); Nikitin et al. (2022).

781 C IMPLEMENTATION DETAILS

782 C.1 MODELS

783 Our implementation relies mainly on the Pytorch (Paszke et al., 2017) and Pytorch Geometric (Fey
 784 and Lenssen, 2019) libraries. For TNN models and static lifting we used TopoX (Hajij et al., 2024)
 785 and TopoBenchmarkX (Telyatnikov et al., 2024).

786 Regarding the base TNNs, we use the hyperparameters (including learning rate, optimizer, batch size,
 787 width, depth, and so on) reported in TopoBenchmarkX for CWN, CXN, and UniGCNII on NCI1,
 788 NCI109, MOLHIV, MUTAG, Proteins, ZINC, Cora and Citeseer (Telyatnikov et al., 2024). Since
 789 TopoBenchmarkX does not report optimal hyperparameters for UniGIN, we use the same used for
 790 UniGCNII. While for CIN, we used part of the hyperparameters reported in (Bodnar et al., 2021b)
 791 and for the TopoTune, since they do not report the hyperparameters, we used a grid search similar to
 792 the available in their repo.

793 We note that MOLHIV, Texas and Wisconsin datasets were not present in TopoBenchmarkX. We use
 794 two TNN layers for MOLHIV and one for Texas and Wisconsin with respective learning rates 10^{-2} ,
 795 5×10^{-3} and 5×10^{-3} . For these datasets, we fix the embedding size in 64 for all layers. And for
 796 Texas and Wisconsin we used weight decay 5×10^{-6} .

797 We are left with the task of optimizing the hyperparameters for the lifting operations (∂ lift, k -NN,
 798 kernel). For ∂ lift, we consider using both GPS and GIN as backbone GNNs, with embedding
 799 dimensions in $\{32, 64, 128\}$, network depth in $\{2, 3\}$, and $k_{\max} = \{3, 5, 7, 9, 11\}$. For k -NN lifting
 800 we choose k in $\{3, 5, 7, 9\}$. For kernel lifting, we consider equally-spaced temperature values within
 801 0.1 and 9.6, with 0.5 increments.

802 All models were trained for 200 epochs and with early stopping after 50 epochs without improvement
 803 on validation accuracy. We run three independent trials for computing mean and standard deviation
 804 of the performance metrics. We select the optimal hyperparameters based on validation accuracy.

810 C.2 HARDWARE
811812 For all experiments, we use a cluster with Nvidia V100 GPUs — details regarding the compute
813 infrastructure are omitted for anonymity.
814815 D DATASETS
816817 **Graph-level tasks.** The datasets NCI1, NCI109, PROTEINS, and MUTAG are part of the
818 TUDatasets (Kersting et al., 2016) — a dataset collection broadly used for benchmarking GNNs.
819 We also use ZINC-12K and MOLHIV (Hu et al., 2020), popular benchmarks for molecular property
820 prediction. Statistics for each dataset are given in Table 5.
821822 Table 5: Statistics of datasets for graph-level tasks.
823

Dataset	#graphs	#classes	Avg #nodes	Avg #edges	Train%	Val%	Test%
NCI1	4110	2	29.87	32.30	80	10	10
NCI109	4127	2	29.68	32.13	80	10	10
MUTAG	188	2	17.93	19.79	80	10	10
PROTEINS	1113	2	39.06	72.82	80	10	10
MOLHIV	41127	2	25.5	27.5	Public Split		
ZINC	12000	-	23.16	49.83	Public Split		

830 **Node-level tasks.** For node classification, we use six popular benchmarks: Cora, Citeseer (Sen et al.,
831 2008; Shchur et al., 2018), Texas, and Wisconsin (Rozemberczki et al., 2020). Cora and Citeseer
832 are citation networks where nodes represent papers and edges denote citation between them. Node
833 features are given by bag-of-word vectors and node labels comprise the academic topics of the
834 underlying articles. Texas and Wisconsin are datasets of webpages from university departments.
835 Nodes represent webpages and edges are hyperlinks between them.
836837 For citation networks, we use the same data split as in (Chen et al., 2018), and for the remaining ones
838 we use the split in (Pei et al., 2020). These are the standard and most used splits. Table 6 provides
839 more details about the datasets.
840Table 6: Statistics of datasets for node classification.
841

Dataset	#Nodes	#Edges	#Features	#Classes	#Train	#Val	#Test
Cora	2708	5429	1,433	7	1,208	500	1,000
Citeseer	3327	4732	3,703	6	1,827	500	1,000
Texas	183	309	1703	5	87	59	37
Wisconsin	251	499	1703	5	120	80	51

847 E ADDITIONAL RESULTS
848849 E.1 RUNTIME AND COMPLEXITY COST
850852 **Wall-clock time.** Table 7 reports per-epoch training and test times (seconds) for the different lifting
853 methods and target neural networks (TNNs) used in our experiments. Overall, ∂lift incurs only a
854 modest runtime overhead compared to static liftings while offering the flexibility of task-adaptive
855 topology.
856857 **Memory usage.** Table 8 reports GPU and RAM consumption (GB) for the different architectures
858 and lifting schemes. Reported values are mean \pm standard deviation over runs. Overall, ∂lift
859 demonstrates moderate memory requirements across both cellular and hypergraph domains, with
860 **GPU usage scaling proportionally to the complexity of the lifted topology.**
861862 **Complexity analysis.** Hypergraph ∂lift runs in $\mathcal{O}(N^2D + N^2 \log k_{\max} + Nk_{\max}D)$, where the N^2D
863 term comes from computing all pairwise node distances in a D -dimensional embedding space, the
864 $N^2 \log k_{\max}$ term for selecting each node’s top- k neighbors, and the $Nk_{\max}D$ term for aggregating
865 across the sampled hyperedges. By contrast, a static kernel-based lifting requires $\mathcal{O}(N^2D + N^3)$
866

864

865

866 Table 7: Per-epoch training and test times (seconds) for the different lifting methods. Reported values
867 are mean \pm standard deviation over runs.

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

Table 7: Per-epoch training and test times (seconds) for the different lifting methods. Reported values
are mean \pm standard deviation over runs.

Dataset/Phase	CWN Cycle	CWN ∂ lift	CXN Cycle	CXN ∂ lift
Cellular Domain				
NCI1 Train	47.41 \pm 2.59	97.37 \pm 5.02	37.71 \pm 2.31	66.77 \pm 3.00
NCI1 Test	1.42 \pm 0.05	4.33 \pm 0.62	1.05 \pm 0.03	2.67 \pm 0.06
NCI109 Train	52.11 \pm 2.58	97.05 \pm 3.80	45.76 \pm 0.97	53.50 \pm 2.50
NCI109 Test	1.67 \pm 0.02	5.19 \pm 0.41	1.56 \pm 0.02	2.15 \pm 0.08
MUTAG Train	2.55 \pm 0.12	4.96 \pm 3.29	1.57 \pm 0.12	3.60 \pm 2.83
MUTAG Test	0.09 \pm 0.00	0.23 \pm 0.02	0.05 \pm 0.00	0.12 \pm 0.00
Proteins Train	18.07 \pm 1.75	20.12 \pm 2.10	13.63 \pm 2.22	15.18 \pm 1.50
Proteins Test	0.65 \pm 0.00	0.82 \pm 0.10	0.46 \pm 0.00	0.72 \pm 0.05
ZINC Train	58.24 \pm 3.10	118.45 \pm 6.20	46.18 \pm 2.40	82.30 \pm 4.50
ZINC Test	2.15 \pm 0.08	5.92 \pm 0.35	1.68 \pm 0.06	3.84 \pm 0.18
Dataset/Phase				
UniGCNII k-hop				
UniGCNII ∂lift				
UniGIN k-hop				
UniGIN ∂lift				
Hypergraph Domain				
NCI1 Train	42.26 \pm 0.46	69.34 \pm 2.55	37.46 \pm 0.31	61.32 \pm 2.83
NCI1 Test	1.23 \pm 0.01	1.94 \pm 0.07	0.91 \pm 0.01	2.03 \pm 0.06
NCI109 Train	31.66 \pm 2.26	88.22 \pm 3.28	28.62 \pm 2.35	69.61 \pm 3.70
NCI109 Test	0.89 \pm 0.05	2.72 \pm 0.03	0.63 \pm 0.02	2.30 \pm 0.05
MUTAG Train	1.94 \pm 0.14	5.12 \pm 3.41	2.04 \pm 1.57	3.06 \pm 0.13
MUTAG Test	0.06 \pm 0.00	0.14 \pm 0.00	0.05 \pm 0.00	0.09 \pm 0.00
Proteins Train	11.93 \pm 0.89	12.45 \pm 1.00	10.12 \pm 2.33	11.50 \pm 1.20
Proteins Test	0.36 \pm 0.00	0.60 \pm 0.05	0.24 \pm 0.00	0.50 \pm 0.03
ZINC Train	51.80 \pm 1.50	76.70 \pm 3.80	45.30 \pm 1.20	66.40 \pm 3.50
ZINC Test	1.58 \pm 0.04	2.68 \pm 0.12	1.28 \pm 0.03	2.85 \pm 0.11

891

892

893

894

895

Table 8: Memory usage (GB) for the different lifting methods and architectures. Reported values are
mean \pm standard deviation over runs.

Memory Type	TNN	MUTAG	NCI1	NCI109	PROTEINS	ZINC
GPU Memory (GB)						
UniGCNII k-hop						
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.03 \pm 0.00	0.02 \pm 0.00
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.05 \pm 0.01	0.02 \pm 0.00
UniGIN k-hop						
	0.02 \pm 0.00					
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.04 \pm 0.00	0.02 \pm 0.00
CWN Cycle						
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.03 \pm 0.00	0.02 \pm 0.00
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.07 \pm 0.02	0.02 \pm 0.00
CWN ∂lift						
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.03 \pm 0.00	0.02 \pm 0.00
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.04 \pm 0.00	0.02 \pm 0.00
CXN Cycle						
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.03 \pm 0.00	0.02 \pm 0.00
	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.02 \pm 0.00	0.04 \pm 0.00	0.02 \pm 0.00
RAM (GB)						
UniGCNII k-hop						
	1.22 \pm 0.02	1.31 \pm 0.01	1.31 \pm 0.02	1.24 \pm 0.01	1.34 \pm 0.00	
	1.45 \pm 0.02	1.46 \pm 0.01	1.45 \pm 0.00	1.46 \pm 0.01	1.47 \pm 0.01	
UniGIN k-hop						
	1.21 \pm 0.01	1.31 \pm 0.01	1.31 \pm 0.01	1.24 \pm 0.01	1.34 \pm 0.01	
	1.42 \pm 0.01	1.45 \pm 0.01	1.45 \pm 0.02	1.42 \pm 0.00	1.47 \pm 0.01	
CWN Cycle						
	1.20 \pm 0.01	1.46 \pm 0.02	1.45 \pm 0.00	1.35 \pm 0.01	1.71 \pm 0.01	
	1.38 \pm 0.01	1.42 \pm 0.01	1.43 \pm 0.01	1.39 \pm 0.01	1.42 \pm 0.01	
CWN ∂lift						
	1.21 \pm 0.00	1.47 \pm 0.01	1.46 \pm 0.01	1.37 \pm 0.02	1.70 \pm 0.00	
	1.40 \pm 0.01	1.43 \pm 0.00	1.44 \pm 0.00	1.43 \pm 0.01	1.45 \pm 0.02	

917

time, while a k -hop static lifting scales as $\mathcal{O}(Nk_{\max}\bar{d})$ with \bar{d} the average node degree. Cellular ∂ lift runs in $\mathcal{O}(N^2D + (E + Nk_{\max})^2 + C\ell D)$, where $(E + Nk_{\max})^2$ captures pairwise adjacency among lifted edges or candidate cells and $C\ell D$ covers the embedding aggregation over C candidate cells of average size ℓ . A static cycle-basis lifting costs $\mathcal{O}(N(N + E) + C\ell)$, whose dominant component is the cycle-basis computation. In both variants, the modest additional runtime is justified by consistent accuracy gains of approximately 5–10 percentage points compared to static liftings.

Choice of Hyperparameters. Although ∂ lift includes additional learnable components — such as the GNN used for neighborhood scoring, the networks computing acceptance probabilities and adaptive neighborhood sizes, and the maximum neighbor budget k_{\max} — the resulting hyperparameter overhead remains modest. Our search space comprised fewer than 30 configurations overall (Appendix C), and we found that tuning was somewhat stable across datasets. Crucially, the adaptivity introduced by these components leads to consistent performance gains while adding only minor computational and memory costs, as reflected in Tables 7 and 8.

E.2 DETERMINISTIC VERSION

As we mentioned in the main text, we can derive a deterministic version of ∂ lift by thresholding probabilities. Table 9 shows compares the ∂ lift with its deterministic variant trained using $\text{threshold} = 0.5$. Note that (random) ∂ lift outperforms its deterministic counterpart in most cases. Nonetheless, while we fixed the threshold at 0.5 here, we leave open the possibility that tuning it as a hyper-parameter could yield performance improvements.

Table 9: Comparison between ∂ lift (random) and its deterministic variant (thresholding at 0.5).

Domain	TNN	∂ lift	NCI1	NCI109	ZINC (MAE)	PROTEINS	MUTAG
Cellular	CWN	deterministic	80.62 ± 0.75	77.64 ± 0.30	1.33 ± 5.33	70.54 ± 0.00	82.46 ± 2.48
Cellular	CWN	random	79.81 ± 0.40	80.55 ± 0.50	0.17 ± 0.00	70.54 ± 3.34	85.96 ± 4.96
Hypergraph	UniGCNII	deterministic	71.53 ± 1.92	69.76 ± 6.03	0.70 ± 0.01	72.62 ± 0.42	75.44 ± 2.48
Hypergraph	UniGCNII	random	77.45 ± 1.88	75.30 ± 1.10	0.56 ± 0.03	73.51 ± 0.84	89.47 ± 4.30

E.3 FURTHER COMPARISON AGAINST DCM

DCM versus ∂ lift. DCM (Battiloro et al., 2023) introduces a learnable lifting approach specifically for cell complexes through a two-step procedure: first, the α -Differentiable Graph Module (α -DGM) learns the 1-skeleton (edges) using α -entmax sampling to generate sparse, non-regular graphs; then, the Polygon Inference Module (PIM) samples polygons from induced cycles of the learned graph. DCM uses α -entmax for both edge and polygon sampling, and trains end-to-end using auxiliary reward-based losses (Eqs. 11 and 13 in their paper) that encourage edges/polymers involved in correct predictions. The method is evaluated exclusively on node classification tasks (homophilic and heterophilic datasets) and is limited to 2-dimensional cell complexes.

Some key differences between DCM and ∂ Lift include:

1. *Domain generality*: our framework applies to hypergraphs, simplicial complexes, and combinatorial complexes, not just cell complexes. Extending DCM to hypergraphs is non-trivial because it relies on cycle-based candidate generation, which is inherently tied to the graph structure and does not naturally translate to hyperedge formation;
2. *Sampling mechanism and training*: we use Bernoulli sampling with straight-through estimators rather than α -entmax, and critically, we do not require auxiliary reward-based losses — our framework is trained purely with the task loss, making it simpler and better suited for end-to-end learning. Additionally, our hypergraph variant is embarrassingly parallel whereas DCM’s two-step procedure is sequential;
3. *Adaptive cell sizes*: we learn distributions over k_v (neighborhood sizes) allowing adaptive hyperedge/cell cardinalities, while DCM’s polygon sizes are constrained by the induced cycles in the learned graph;
4. *Evaluation scope*: we assess both node and graph classification across 12 datasets with multiple TNN architectures, demonstrating broader applicability.

972 **Experiments on point clouds.** We compared ∂ lift against DCM on point-cloud node classification
 973 (graphs with no edges) to provide a direct comparison. Table 10 shows results across four datasets.
 974 ∂ lift combined with different TNNs achieves competitive or superior performance: UniGCNII+ ∂ lift
 975 achieves 84.97% on Wisconsin (vs. 71.24% for DCM), and UniGIN+ ∂ lift achieves 83.78% on
 976 Texas (vs. 62.16% for DCM), demonstrating substantial gains on heterophilic datasets. On the
 977 homophilic dataset Cora, DCM achieves 73.07% compared to 71.47% for the best ∂ lift variant,
 978 showing comparable performance. As expected, the performance in heterophilic datasets are higher
 979 in this setting compared to homophilic when referring Table 3 and Table 4. These results demonstrate
 980 that ∂ lift’s framework provides competitive performance while offering greater flexibility across
 981 topological domains and TNN architectures.

982
983 Table 10: Point cloud performance comparison of ∂ lift across TNNs versus DCM.
984

TNN + ∂ lift	Citeseer	Texas	Wisconsin	Cora
UniGCNII	74.73 \pm 0.12	81.98 \pm 1.27	84.97 \pm 2.45	70.83 \pm 0.37
CWN	51.80 \pm 1.67	71.17 \pm 1.27	75.82 \pm 0.92	55.90 \pm 2.45
CXN	62.03 \pm 1.59	81.08 \pm 0.00	78.43 \pm 0.00	57.43 \pm 1.19
UniGIN	72.73 \pm 0.61	83.78 \pm 2.21	83.66 \pm 0.92	71.47 \pm 1.77
DCM	74.40 \pm 0.45	62.16 \pm 5.84	71.24 \pm 3.33	73.07 \pm 0.92

992
993
994
995 E.4 COMPARISON AGAINST PROBABILISTIC REWIRING
996
997

998 **IPR-MPNN versus ∂ lift.** IPR-MPNN (Qian et al., 2024) introduces probabilistic graph rewiring
 999 by connecting original graph nodes to a small set of virtual nodes in an end-to-end differentiable
 1000 manner. The method uses an upstream MPNN to compute priors θ for assigning each original node
 1001 to k virtual nodes (from m total virtual nodes, where $m \ll n$), sampling assignment matrices via
 1002 differentiable k -subset sampling. A downstream MPNN then operates on the augmented graph with
 1003 message passing among: (1) original nodes to virtual nodes, (2) among virtual nodes (forming a
 1004 complete subgraph), and (3) among original nodes.

1005 There are relevant distinctions between ∂ lift and IPR-MPNN: (i) *Explicit vs. implicit structure*:
 1006 IPR-MPNN implicitly routes long-range information through virtual nodes, while ∂ lift explicitly
 1007 constructs higher-order cells (hyperedges, simplices, polygons) that directly encode multi-way
 1008 interactions; (ii) *Domain flexibility*: IPR-MPNN operates within the graph domain augmented with
 1009 virtual nodes, whereas ∂ lift learns liftings to diverse topological domains (hypergraphs, simplicial
 1010 complexes, cell complexes, combinatorial complexes); (iii) *Sampling strategy*: both use differentiable
 1011 k -subset sampling, but IPR-MPNN samples node-to-virtual-node assignments while ∂ lift samples
 1012 which higher-order cells to include in the lifted structure; (iv) *Computational approach*: IPR-MPNN
 1013 requires managing virtual node features and specialized message-passing between hierarchies, while
 1014 ∂ lift integrates directly with standard TNN architectures designed for each topological domain.

1015 **Results on graph classification.** Table 11 compares IPR-MPNN to ∂ lift combined with various
 1016 TNNs on five molecular datasets. To ensure a fair comparison, we evaluated IPR-MPNN using our
 1017 experimental setup, with the data splits described in Appendix D and a hyperparameter grid similar
 1018 to the one in Appendix C, with embedding dimensions in {32, 64, 128} and network depths in {2, 3}.
 1019 All other hyperparameters remained fixed, taken from the configurations in the official IPR-MPNN
 1020 repository. For instance, on ZINC we do not use edge features.

1021 ∂ lift achieves substantial improvements on multiple benchmarks: CXN+ ∂ lift reaches 82.08% on
 1022 NCI1 and 82.57% on NCI109 compared to IPR-MPNN’s 77.44% and 77.08%. On ZINC, both
 1023 CWN+ ∂ lift and CXN+ ∂ lift achieve 0.17 MAE versus IPR-MPNN’s 0.39, representing a 56% error
 1024 reduction. UniGCNII+ ∂ lift obtains the highest MUTAG accuracy at 89.47%. While IPR-MPNN
 1025 demonstrates the advantages of learnable graph augmentation through virtual nodes, our results show
 1026 that explicitly learning higher-order topological structures through ∂ lift can provide complementary
 1027 benefits, particularly when combined with TNNs designed to exploit these structures.

1026
1027 Table 11: Graph classification performance comparison between IPR-MPNN and ∂ lift combined
1028 with various TNNs

Method	MUTAG \uparrow	NCI1 \uparrow	NCI109 \uparrow	PROTEINS \uparrow	ZINC \downarrow
IPR-MPNN	70.18 \pm 2.48	77.44 \pm 1.13	77.08 \pm 0.30	73.73 \pm 1.93	0.39 \pm 0.05
CWN + Cycle	66.67 \pm 12.41	76.93 \pm 1.18	76.71 \pm 1.34	69.05 \pm 2.95	0.46 \pm 0.01
CWN + ∂ lift	85.96 \pm 4.96	79.81 \pm 0.40	80.55 \pm 0.50	70.54 \pm 3.34	0.17 \pm 0.00
CXN + Cycle	61.40 \pm 2.48	72.02 \pm 1.69	75.01 \pm 0.62	70.83 \pm 1.52	0.79 \pm 0.02
CXN + ∂ lift	84.21 \pm 4.30	82.08 \pm 1.50	82.57 \pm 0.40	69.94 \pm 2.10	0.17 \pm 0.01
UniGCNII + k -hop	61.40 \pm 2.48	72.70 \pm 0.52	72.01 \pm 1.55	72.92 \pm 1.11	0.66 \pm 0.02
UniGCNII + ∂ lift	89.47 \pm 4.30	77.45 \pm 1.88	75.30 \pm 1.10	73.51 \pm 0.84	0.56 \pm 0.03
UniGIN + k -hop	64.91 \pm 2.48	65.50 \pm 1.99	66.97 \pm 7.25	71.43 \pm 0.73	1.15 \pm 0.01
UniGIN + ∂ lift	66.67 \pm 6.56	64.88 \pm 1.09	79.74 \pm 0.23	73.81 \pm 1.52	0.92 \pm 0.05

1041
1042 SUPPLEMENTARY REFERENCES
10431044 Viacheslav Borovitskiy, Iskander Azangulov, Alexander Terenin, Peter Mostowsky, Marc Deisenroth, and Nicolas
1045 Durrande. Matérn gaussian processes on graphs. In *International Conference on Artificial Intelligence and*
1046 *Statistics*, pages 2593–2601. PMLR, 2021.1047 Jie Chen, Tengfei Ma, and Cao Xiao. Fastgcn: Fast learning with graph convolutional networks via importance
1048 sampling. In *International Conference on Learning Representations*, 2018.1049 David Duvenaud. *Automatic model construction with Gaussian processes*. PhD thesis, 2014.1050 Mustafa Hajij, Mathilde Papillon, Florian Frantzen, Jens Agerberg, Ibrahem AlJabea, Ruben Ballester, Claudio
1051 Battiloro, Guillermo Bernárdez, Tolga Birdal, Aiden Brent, et al. Topox: a suite of python packages for
1052 machine learning on topological domains. *Journal of Machine Learning Research*, 25(374):1–8, 2024.1053 Alexander V Nikitin, ST John, Arno Solin, and Samuel Kaski. Non-separable spatio-temporal graph kernels via
1054 spdes. In *International Conference on Artificial Intelligence and Statistics*, pages 10640–10660. PMLR, 2022.1055 Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph
1056 convolutional networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=S1e2agrFvS>.1057 Bernhard Schölkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization,*
1058 *optimization, and beyond*. MIT press, 2002.1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079