Extended Abstract Track

Topological Blindspots: Understanding and Extending Topological Deep Learning Through the Lens of Expressivity

Editors: List of editors' names

Abstract

Topological deep learning (TDL) is a rapidly growing field that seeks to leverage topological structure in data and facilitate learning from data supported on topological objects. Most TDL architectures can be unified under the framework of higher-order message-passing (HOMP), which generalizes graph message-passing to higher order domains. In the first part of the paper, we explore HOMP's expressive power from a topological perspective, demonstrating the framework's inability to express fundamental topological and metric invariants such as diameter, orientability, planarity, and homology. In the second part of the paper, we develop two new classes of TDL architectures – multi-cellular networks (MCN) and scalable MCN (SMCN) – which draw inspiration from expressive graph architectures. MCN can reach full expressivity, but scaling it to large data objects can be computationally expansive. Therefore, SMCN is designed as a more scalable alternative that still mitigates many of HOMP's expressivity limitations. In the third part of the paper, we design benchmarks for evaluating TDL models on their ability to learn topological properties of complexes. We then evaluate SMCN on these benchmarks as well as real-world graph datasets, demonstrating improvements both over HOMP baselines and expressive graph methods, highlighting the value of expressively leveraging topological information.

Keywords: Topological Deep Learning, Higher-Order Message-Passing, Expressivity.

1. Introduction

Topological deep learning (TDL) is an emerging field focused on learning from data supported on topological objects. Higher-order message-passing (HOMP) (Hajij et al., 2022a,b) has emerged as a key framework in TDL, unifying architectures designed for various topological data types. Originally introduced for simplicial complexes (Bodnar et al., 2021b), HOMP has been successively adapted to cellular complexes (Bodnar et al., 2021a; Hajij et al.), and more recently to combinatorial complexes (Hajij et al., 2022a,b). The HOMP framework extends traditional message-passing neural networks (MPNNs) (Gilmer et al., 2017), widely used in graph learning, to higher-order topological domains.

Despite their widespread adoption, MPNNs are known to struggle with expressivity limitations, often failing to distinguish even simple non-isomorphic graphs (Morris et al., 2019; Xu et al., 2018). This realization has led to a substantial body of work dedicated to developing more expressive graph architectures (Morris et al., 2023). Given HOMP's similarity to MPNNs, a natural question arises: what are the limitations of higher-order message-passing architectures in distinguishing topological objects? This question, highlighted in a recent position paper (Papamarkou et al., 2024), is the main focus of this paper. In the first part of the paper, we demonstrate HOMP's inability to differentiate between complexes based

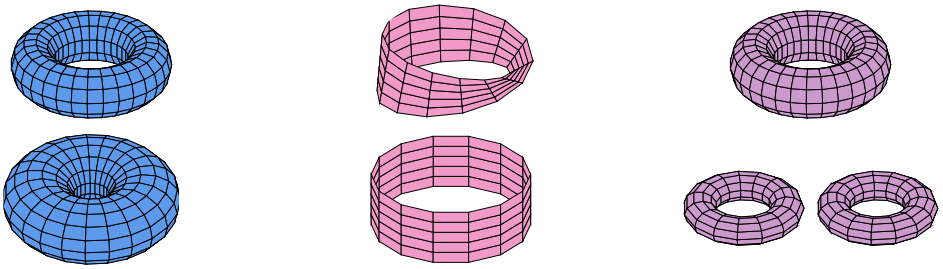

(a) Diameter (b) Orientability/planarity (c) Homology

Figure 1: Pairs of CCs differing in fundamental metric/topological properties that are HOMP-indistinguishable . 1(a) tori with different diameters (20 and 22); 1(b) a Möbius strip and a cylinder differing in both orientability and planarity; 1(c) a torus and a pair of disconnected tori, differing in all non-trivial homology groups.

on several fundamental topological and metric invariants, including diameter, orientability, planarity, and homology groups.

In the second part of the paper, we introduce two new classes of TDL architectures, multi-cellular networks (MCN) and scalable MCN (SMCN), designed to address HOMP's expressivity limitations. MCN draws inspiration from invariant graph networks (IGNs) (Maron et al., 2018, 2019; Keriven and Peyré, 2019) and can reach *full expressivity*, although scaling it to large structures can be computationally expensive. SMCN is a more scalable alternative to MCN that applies expressive GNNs to graph structures defined over the cells of the complex. SMCN still mitigates many of HOMP's expressivity limitations and can be easily applied to large scale real-world datasets.

We create two datasets for assessing the ability of TDL architectures to capture topological information. We empirically evaluate SMCN on these datasets and on real-world graph benchmarks. SMCN outperforms both standard HOMP architectures and expressive GNNs, highlighting the value of expressively leveraging topological information.

2. Preliminaries

Combinatorial complexes. Combinatorial complexes (CCs) are a class of higher-order objects that can flexibly model many types of hierarchical data. Most TDL domains, including simplicial complexes, cellular complexes, and hypergraphs, are subclasses of combinatorial complexes. Therefore, throughout the paper data is assumed to be in the form of a combinatorial complex.

Definition 2.1 (Combinatorial complex) *A combinatorial complex (CC) is a* 3*-tuple* $(S, \mathcal{X}, \mathrm{rk})$ *comprised of a node set* S*, a cell set* $\mathcal{X} \subseteq \mathcal{P}(S) \backslash \varnothing$*, and a rank function* $\mathrm{rk} : \mathcal{X} \to \mathbb{Z}_{\geqslant 0}$ *such that* $\forall s \in S$*,* $\{s\} \in \mathcal{X}$*,* $\mathrm{rk}(\{s\}) = 0$*, and* $\forall x, y \in \mathcal{X}\ x \subseteq y \Rightarrow \mathrm{rk}(x) \leqslant \mathrm{rk}(y)$*.*

The set of cells of rank r (r-cells) is called the r-skeleton and is denoted by $\mathcal{X}_r = \mathrm{rk}^{-1}(r)$. By slight abuse of notation, we sometimes use $\mathcal{X}$ to refer to the entire CC.

Higher-order message-passing. Data supported on a CC can be viewed as a collection of functions $\{\mathbf{h} : \mathcal{X}_r \to \mathbb{R}^{d_r}\}_{r=0}^{\ell}$. Higher-order message passing (HOMP) (Hajij et al., 2022b) is a general computational framework, encompassing many TDL architectures, that

processes cellular information by exchanging messages among cells. HOMP architectures are defined via the following update rule:

$$\mathbf{h}_x^{(t+1)} = \beta\left(\bigotimes_{i=1}^{k} \bigoplus_{y\in\mathcal{N}_i(x)} \mathsf{MLP}^{(t)}_{i,\mathrm{rk}(x)}\left(\mathbf{h}_x^{(t)}, \mathbf{h}_y^{(t)}\right)\right), \tag{1}$$

where $\bigotimes$ and $\bigoplus$ are aggregation functions, $\bigoplus$ is permutation invariant, $\{\mathcal{N}_i : \mathcal{X} \to \mathcal{P}(\mathcal{X})\}_{i=1}^{k}$ are a fixed set of *neighborhood functions* controlling cell connectivity, and β is a non-linearity. The neighborhood functions used in HOMP are the upper/lower incidence and (co)adjacency defined in Appendix A.

3. Expressive Power of Higher-Order Message-Passing

As with GNN expressivity, we study HOMP's ability to distinguish non-isomorphic CCs. We develop a topological criterion that identifies when a pair of CCs is indistinguishable by HOMP, extending the results of Bamberger (2022) from graphs to combinatorial complexes.

Theorem 3.1 (HOMP indistinguishability criterion) *Let $\mathcal{X}$ and $\mathcal{X}'$ be ℓ-dimensional CCs such that $|\mathcal{X}_0| = |\mathcal{X}'_0|$. If there exists a CC $\tilde{\mathcal{X}}$ that covers each of the connected components of both $\mathcal{X}$ and $\mathcal{X}'$, then for every HOMP model M, $M(\mathcal{X}) = M(\mathcal{X}')$.*

We use Theorem 3.1 to demonstrate HOMP's inability to distinguish between complexes based on the following common topological and metric invariants: (1) diameter, which measures how "spread out" the complex is; (2) orientability, which captures whether a "side" or "direction" can be consistently defined across the entire space; (3) planarity which captures whether the complex can be embedded in the plane; (4) homology groups, which encode the structure of "r-dimensional holes".

Theorem 3.2 (Topological blindspots) *For any invariant $\boldsymbol{I} \in \{$diameter, orientability, planarity, homology$\}$ there exists a pair of HOMP-indistinguishable CCs that differ in $\boldsymbol{I}$.*

Figure 1 depicts pairs of HOMP-indistinguishable CCs which differ in each of the above invariants. In addition, Appendix B provides a detailed discussion regarding each invariant, as well as complete proofs of Theorem 3.2 and Theorem 3.1.

4. (Scalable) Multi-cellular Networks

MCN and SMCN. To mitigate the expressivity limitations outlined in the previous section, we design two classes of TDL architectures, inspired by expressive GNN architectures. The first class, termed Multi-cellular Networks (MCN) adapts the k-IGN framework (Maron et al., 2018) to CCs by introducing updates over multi-cellular cochain spaces. Similarly to k-IGNs, MCN can reach full expressivity but is expensive to implement in full generality. Therefore, we construct scalable MCN (SMCN) by replacing the most computationally expensive updates with expressive GNNs often used as a practical alternative to k-IGNs. Full descriptions of both methods are provided in Appendix C.

MCN and SMCN expressive power. As mentioned, MCN is fully expressive – for every pair of non-isomorphic CCs, $\mathcal{X}$ and $\mathcal{X}'$, there exists an MCN model M such that $M(\mathcal{X}) \neq M(\mathcal{X}')$. For SMCN, every pair of HOMP-indistinguishable CCs listed in Figure 1

can be separated by an SMCN model. Additionally, SMCN can compute the diameter of any CC, and the Betti numbers of cellular complexes representing surfaces embeddable in $\mathbb{R}^3$. A detailed discussion of MCN and SMCN's expressive power appears in Appendix D.

5. Experiments

We empirically evaluate SMCN with two sets of experiments. First, we evaluate SMCN's ability to learn topological and metric properties of cellular complexes on two novel benchmarks. Then, we test SMCN on graph classification and regression tasks[1].

Model	ZINC	MOLHIV	MOLESOL
	MAE (↓)	ROC-AUC (↑)	RMSE (↓)
GCN (Kipf and Welling, 2016)	0.321 ± 0.009	76.06 ± 0.97	1.114 ± 0.036
GIN (Xu et al., 2018)	0.163 ± 0.004	75.58 ± 1.40	1.173 ± 0.057
CIN (Bodnar et al., 2021a)	0.079 ± 0.006	80.94 ± 0.57	1.288 ± 0.026
CIN++ (Giusti et al., 2023)	0.077 ± 0.004	80.63 ± 0.94	–
Cellular Transformer (Ballester et al., 2024)	0.080^*	79.46^*	–
PPGN (Maron et al., 2019)	0.079 ± 0.005	–	–
PPGN++ (6) (Puny et al., 2023)	0.071 ± 0.001	–	–
DSS-GNN (Bevilacqua et al., 2021)	0.102 ± 0.003	76.78 ± 1.66	–
SUN (Frasca et al., 2022)	0.083 ± 0.003	80.03 ± 0.55	–
GNN-SSWL+ (Zhang et al., 2023)	0.070 ± 0.005	79.58 ± 0.35	–
Subgraphormer (Bar-Shalom et al., 2023)	0.067 ± 0.007	80.38 ± 1.92	0.832 ± 0.043
Subgraphormer + PE (Bar-Shalom et al., 2023)	0.063 ± 0.001	79.48 ± 1.28	0.826 ± 0.010
SMCN (ours)	$\mathbf{0.060} \pm 0.004$	$\mathbf{81.16} \pm 0.90$	$\mathbf{0.809} \pm 0.037$

Table 1: SMCN outperforms MPNNs, HOMP and expressive GNNs on various graph regression and classification tasks.

Topological/metric property prediction. We design two benchmarks based on the ZINC dataset (Sterling and Irwin, 2015). Graphs are lifted into CCs by adding cycles of length $\leqslant 18$ as 2-cells. We test SMCN's ability to classify the resulting CCs by their cross diameter (18 possible values) and 2nd Betti number (6 possible values). We compare CIN (Bodnar et al., 2021a), a custom HOMP architecture designed and tuned for each task, and SMCN. The results, presented in Table 2, demonstrate SMCN's superior ability in learning topological and metric properties of CCs.

Model	Cross-diameter	2-nd Betti number
	Accuracy (↑)	Accuracy (↑)
CIN	$34.78 \pm 3.00\%$	$42.15 \pm 25.22\%$
Custom HOMP	$67.87 \pm 12.26\%$	$81.76 \pm 10.06\%$
SMCN	$\mathbf{92.76} \pm 0.53\%$	$\mathbf{99.61} \pm 0.12\%$

Table 2: Accuracy of predicting cross-diameter and the 2-nd Betti number of lifted ZINC graphs.

Graph benchmarks. On ZINC (Sterling and Irwin, 2015; Dwivedi et al., 2023), MOLHIV and MOLESOL (Hu et al., 2020) SMCN outperforms *both* standard HOMP baselines as well as expressive graph methods, highlighting the value of expressively leveraging higher-order topological information.

See Appendix E for additional experiments and further details.

6. Conclusion

We analyze the topological expressivity of HOMP — the dominant framework in TDL. Despite the fact that HOMP operates on topological data, we prove its inability to express key topological and metric invariants (diameter, orientability, planarity and homology). To mitigate this limitation we extend HOMP, constructing two expressive TDL frameworks. MCN generalizes k-IGNs to CCs and SMCN utilizes efficient and expressive GNNs to increase HOMPs expressivity. MCN can reach full expressivity and that SMCN is strictly more expressive than HOMP. Experimentally, SMCN outperforms both expressive graph methods and HOMP architectures on several benchmarks.

1. All graphs lifted to CCs by adding cycles as 2-cells.

References

Waiss Azizian and Marc Lelarge. Expressive power of invariant and equivariant graph neural networks. *arXiv preprint arXiv:2006.15646*, 2020.

Rubén Ballester, Pablo Hern'andez-Garc'ia, Mathilde Papillon, Claudio Battiloro, Nina Miolane, Tolga Birdal, Carles Casacuberta, Sergio Escalera, and Mustafa Hajij. Attending to topological spaces: The cellular transformer. *ArXiv*, abs/2405.14094, 2024. URL https://api.semanticscholar.org/CorpusID:269982011.

Jacob Bamberger. A topological characterisation of weisfeiler-leman equivalence classes. In *Topological, Algebraic and Geometric Learning Workshops 2022*, pages 17–27. PMLR, 2022.

Guy Bar-Shalom, Beatrice Bevilacqua, and Haggai Maron. Subgraphormer: Subgraph GNNs meet graph transformers. In *NeurIPS 2023 Workshop: New Frontiers in Graph Learning*, 2023. URL https://openreview.net/forum?id=e8ba9Hu1mM.

Guy Bar-Shalom, Yam Eitan, Fabrizio Frasca, and Haggai Maron. A flexible, equivariant framework for subgraph gnns via graph products and graph coarsening. *arXiv preprint arXiv:2406.09291*, 2024.

Beatrice Bevilacqua, Fabrizio Frasca, Derek Lim, Balasubramaniam Srinivasan, Chen Cai, Gopinath Balamurugan, Michael M Bronstein, and Haggai Maron. Equivariant subgraph aggregation networks. *arXiv preprint arXiv:2110.02910*, 2021.

Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yu Guang Wang, Pietro Liò, Guido Montufar, and Michael M. Bronstein. Weisfeiler and lehman go cellular: CW networks. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021a. URL https://openreview.net/forum?id=uVPZCMVtsSG.

Cristian Bodnar, Fabrizio Frasca, Yuguang Wang, Nina Otter, Guido F Montufar, Pietro Lió, and Michael Bronstein. Weisfeiler and lehman go topological: Message passing simplicial networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1026–1037. PMLR, 18–24 Jul 2021b. URL https://proceedings.mlr.press/v139/bodnar21a.html.

Michael Bronstein. Using subgraphs for more expressive gnns. https://towardsdatascience.com/using-subgraphs-for-more-expressive-gnns-8d06418d5ab, December 2021.

Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

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

Fabrizio Frasca, Beatrice Bevilacqua, Michael Bronstein, and Haggai Maron. Understanding and extending subgraph gnns by rethinking their symmetries. *Advances in Neural Information Processing Systems*, 35:31376–31390, 2022.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.

Lorenzo Giusti, Teodora Reu, Francesco Ceccarelli, Cristian Bodnar, and Pietro Liò. Cin++: Enhancing topological message passing. *arXiv preprint arXiv:2306.03561*, 2023.

Mustafa Hajij, Kyle Istvan, and Ghada Zamzmi. Cell complex neural networks.

Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, and Karthikeyan Natesan Ramamurthy. Higher-order attention networks. *arXiv preprint arXiv:2206.00606*, 2(3):4, 2022a.

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

Jason Hartford, Devon Graham, Kevin Leyton-Brown, and Siamak Ravanbakhsh. Deep models of interactions across sets. In *International Conference on Machine Learning*, pages 1909–1918. PMLR, 2018.

Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002. ISBN 0-521-79160-X; 0-521-79540-0.

Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.

Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

Thomas Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *ArXiv*, abs/1609.02907, 2016. URL https://api.semanticscholar.org/CorpusID:3144218.

Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. *arXiv preprint arXiv:1812.09902*, 2018.

Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *Advances in neural information processing systems*, 32, 2019.

Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4602–4609, 2019.

Christopher Morris, Yaron Lipman, Haggai Maron, Bastian Rieck, Nils M Kriege, Martin Grohe, Matthias Fey, and Karsten Borgwardt. Weisfeiler and leman go machine learning: The story so far. *The Journal of Machine Learning Research*, 24(1):15865–15923, 2023.

Theodore Papamarkou, Tolga Birdal, Michael M Bronstein, Gunnar E Carlsson, Justin Curry, Yue Gao, Mustafa Hajij, Roland Kwitt, Pietro Lio, Paolo Di Lorenzo, et al. Position: Topological deep learning is the new frontier for relational learning. In *Forty-first International Conference on Machine Learning*, 2024.

Omri Puny, Derek Lim, Bobak Kiani, Haggai Maron, and Yaron Lipman. Equivariant polynomials for graph neural networks. In *International Conference on Machine Learning*, pages 28191–28222. PMLR, 2023.

Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.

Yanbo Wang and Muhan Zhang. An empirical study of realized gnn expressiveness. In *Forty-first International Conference on Machine Learning*, 2024.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.

Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small relu networks are powerful memorizers: a tight analysis of memorization capacity. *Advances in Neural Information Processing Systems*, 32, 2019.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.

Bohang Zhang, Guhao Feng, Yiheng Du, Di He, and Liwei Wang. A complete expressiveness hierarchy for subgraph gnns via subgraph weisfeiler-lehman tests. In *International Conference on Machine Learning*, pages 41019–41077. PMLR, 2023.

Appendix A. Extended Preliminaries

Notation. We use $[n] = \{1, \dots, n\}$. $\bigoplus$ and $\bigotimes$ denote aggregation functions (i.e. functions $\mathcal{S} \times \cdots \times \mathcal{S} \to \mathcal{S}$). $\bigoplus$ is used for *permutation invariant* aggregation (i.e. $\forall (s_1, \dots, s_n) \in \mathcal{S} \times \cdots \times \mathcal{S}, \forall \sigma \in S_n$, $\bigoplus_{i=1}^{n} s_{\sigma^{-1}(i)} = \bigoplus_{i=1}^{n} s_i$), and $\bigotimes$ denotes aggregations that are not necessarily permutation invariant. Multisets are denoted by $\{\!\{a_1, \dots, a_n\}\!\}$. The size of a set $\mathcal{S}$ is denoted by $|\mathcal{S}|$. We use bold lower case letters (e.g. $\boldsymbol{k}$) for tuples of integers. $\boldsymbol{e}_i \in \mathbb{N}^{\ell+1}$, 0-indexed, represents the tuple with one at the ith positions and zeros elsewhere.

Neighborhood functions. Neighborhood functions are a key component in HOMP, facilitating dynamic aggregation of information across cells. Formally, a neighborhood function is a function $\mathcal{N} : \mathcal{X} \to \mathcal{P}(\mathcal{X})$. The most common neighborhood functions are:

1. The (r_1, r_2)-adjacency and co-adjacency, defined by:

$$\begin{aligned} \mathcal{A}_{r_1,r_2}(x) &= \{y \in \mathcal{X}_{r_1} \mid \exists z \in \mathcal{X}_{r_2} \text{ s.t. } x, y \subseteq z\}, \\ \operatorname{co}\mathcal{A}_{r_1,r_2}(x) &= \{y \in \mathcal{X}_{r_1} \mid \exists z \in \mathcal{X}_{r_2} \text{ s.t. } z \subseteq x, y\}, \end{aligned} \tag{2}$$

   for $x \in \mathcal{X}_{r_1}$, and $\mathcal{A}_{r_1,r_2}(x) = \operatorname{co}\mathcal{A}_{r_1,r_2}(x) = \varnothing$ for $x \notin \mathcal{X}_{r_1}$.

2. The (r_1, r_2)-upper and lower incidence, defined by:

$$\mathcal{B}_{r_1,r_2}(x) = \{y \in \mathcal{X}_{r_2} \mid x \subseteq y\}, \quad \mathcal{B}^{\top}_{r_1,r_2}(x) = \{y \in \mathcal{X}_{r_2} \mid y \subseteq x\}, \tag{3}$$

   for $x \in \mathcal{X}_{r_1}$, and $\mathcal{B}_{r_1,r_2}(x) = \mathcal{B}^{\top}_{r_1,r_2}(x) = \varnothing$ for $x \notin \mathcal{X}_{r_1}$.

We call the above *natural* neighborhood functions, a collection denoted by $\boldsymbol{\mathcal{N}}_{\text{nat}}$. Given an enumeration of the cells, a neighborhood functions can be represented in matrix form. E.g. we can represent a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ as a one dimensional combinatorial complex $(\mathcal{V}, \mathcal{X}_0 \cup \mathcal{X}_1, \text{rk})$ by setting $\mathcal{X}_0 = \{\{v\} \mid v \in \mathcal{V}\}$, $\mathcal{X}_1 = \mathcal{E}$, and $\text{rk}(x) = 0$ for nodes and 1 for edges. In this case, the matrix form of the adjacency neighborhood function $\mathcal{A}_{0,1}$ is the graph adjacency matrix, while the incidence neighborhood function $\mathcal{B}_{0,1}$ corresponds to the graph incidence matrix.

Cochain spaces. Data defined over an ℓ-dimensional CC can be viewed as a collection of functions $\{\mathbf{h}_r : \mathcal{X}_r \to \mathbb{R}^{d_r}\}_{r=0}^{\ell}$ (in general, $d_{r_1} \neq d_{r_2}$, e.g. for a molecule atoms might have a different number of features than bonds). Each of these functions is called a *cochain* or a *cell feature map*. The vector space of all cochains on cells of rank r is denoted by $\mathcal{C}^r(\mathcal{X}, \mathbb{R}^{d_r})$, often abbreviated as $\mathcal{C}^r$. The feature associated with a cell $x \in \mathcal{X}_r$ is denoted by $\mathbf{h}_r(x)$, $(\mathbf{h}_r)_x$, or simply $\mathbf{h}_x$ when the rank is clear from context.

Higher-order message-passing. Higher-order message passing (HOMP) (Hajij et al., 2022b) is a general computational framework that processes cellular information defined on a higher-order domain by exchanging messages among cells using a set of neighborhood functions. Let $\boldsymbol{\mathcal{N}} = \{\mathcal{N}_1, \dots, \mathcal{N}_k\}$ be a collection of neighborhood functions, given an initial cochain $\mathbf{h}^{(0)} = \mathbf{h}$, HOMP is recursively defined via the following update rule:

$$\mathbf{h}_x^{(t+1)} = \beta\left(\bigotimes_{i=1}^{k} \bigoplus_{y \in \mathcal{N}_i(x)} \mathsf{MLP}^{(t)}_{i,\text{rk}(x)}\left(\mathbf{h}_x^{(t)}, \mathbf{h}_y^{(t)}\right)\right), \tag{4}$$

where $\mathbf{h}_x^{(t)}$ is the feature associated with cell $x \in \mathcal{X}$ at layer t, β is a non-linear activation, and $\{\mathsf{MLP}_{i,r}^{(t)}\}_{i=1,\ldots,k,r=0,\ldots,\ell,t=0,\ldots T-1}$ are MLPs, one for each neighborhood function, rank and layer. In the following, unless specified otherwise we assume $\boldsymbol{\mathcal{N}} \subseteq \boldsymbol{\mathcal{N}}_{\text{nat}}$. Similar to graph message-passing, the HOMP framework encompasses many TDL architectures, including simplicial complex architectures (Bodnar et al., 2021b), cellular complex architectures (Hajij et al.; Bodnar et al., 2021a; Giusti et al., 2023) and combinatorial complex architectures (Hajij et al., 2022a,b).

Tensor diagrams. To navigate the space of HOMP architectures Hajij et al. (2022a) introduce tensor diagrams, a notation scheme for specifying HOMP architectures using DAGs. Tensor diagrams allow for selective aggregation over different neighborhood functions for different cochain spaces in different layers of the network. The nodes of a tensor diagram correspond to cochain spaces, and the edges to neighborhood functions. At each level of the diagram the signal flows from the source nodes to the target nodes using the update rule specified in Equation 4, where aggregation is performed only over neighborhood functions associated with the incoming edges (this is equivalent to setting $\mathsf{MLP}_{i,r}^{(t)} \equiv \mathbf{0}$ for neighborhood functions $\mathcal{N}_i$ that are not associated with an incoming edge). See Hajij et al. (2022b) for a detailed overview.

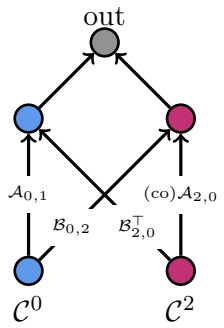

Figure 2: HOMP tensor diagram.

Appendix B. Expressive Limitations of HOMP

The expressivity of graph models is often evaluated based on their ability to distinguish between pairs of non-isomorphic graphs. Generalizing to other TDL domains, we develop topological tools to understand when a pair of CCs is indistinguishable by HOMP. Although CCs are combinatorial objects, they can induce various metric and topological structures. The shortest path distance with respect to any neighborhood function defines a metric on the cells of the CC. In addition, if the complex is cellular or simplicial, it can be canonically associated with a topological space. The topological properties of these spaces are invariants of the underlying CC, therefore the question of their usefulness in distinguishing CCs is well defined.

B.1. A Topological Criterion for HOMP-Indistinguishability

We extend the topological graph WL indistinguishability characterization introduced by Bamberger (2022) to combinatorial complexes. As a first step, we introduce the notion of a CC cover.

Definition B.1 (CC covering) *A $\tilde{\mathcal{X}}$ covers $\mathcal{X}$ if there exists a surjective rank and inclusion preserving map $\rho : \tilde{\mathcal{X}} \to \mathcal{X}$ which is a local isomorphism with respect to natural neighborhood functions (i.e. ρ bijectivly maps $\mathcal{N}(x') \to \mathcal{N}(\rho(x'))$ for all $x' \in \tilde{\mathcal{X}}$ and $\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\text{nat}}$).*

The next theorem, which is a formal restatement of Theorem 3.1, outlines the relationship between CC covering and HOMP-indistinguishability:

Theorem B.2 (HOMP indistinguishability criterion) *Let $\mathcal{X}$, $\mathcal{X}'$ be two CCs of dimension ℓ with no cell features such that $|\mathcal{X}_0| = |\mathcal{X}'_0|$. If $\mathcal{X}$ and $\mathcal{X}'$ admit decompositions into connected components*

$$\mathcal{X} = \bigsqcup_{\mathcal{Z} \in C(\mathcal{X})} \mathcal{Z}, \quad \mathcal{X}' = \bigsqcup_{\mathcal{Z}' \in C(\mathcal{X}')} \mathcal{Z}', \tag{5}$$

such that $\exists \tilde{\mathcal{X}}$ that is covers each of the connected components $\mathcal{Z} \in C(\mathcal{X}), \mathcal{Z}' \in C(\mathcal{X}')$, then for every HOMP model M, $M(\mathcal{X}) = M(\mathcal{X}')$.

A combinatorial complex $\mathcal{X}$ is said to be *connected* if the Hasse graph, defined by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \mathcal{X}$ and $\mathcal{E} = \{(x, y) \mid x \subseteq y, \mathrm{rk}(x) = \mathrm{rk}(y) - 1\}$, is connected. To prove Theorem B.2, we first prove two lemmas.

Lemma B.3 *Let $\mathcal{X}$ and $\tilde{\mathcal{X}}$ be CCs of dimension ℓ such that $\tilde{\mathcal{X}}$ covers $\mathcal{X}$ via the covering map $\rho : \tilde{\mathcal{X}} \to \mathcal{X}$. In addition, let M be a HOMP model with T layers, and let $\mathbf{h}_x^{(t)}$ and $\tilde{\mathbf{h}}_{x'}^{(t)}$ denote the cell feature maps of $\mathcal{X}$ and $\tilde{\mathcal{X}}$ at layer t evaluated on cells $x \in \mathcal{X}$ and $x' \in \tilde{\mathcal{X}}$ respectively. Under these conditions, $\tilde{\mathbf{h}}_{x'}^{(t)} = \mathbf{h}_{\rho(x')}^{(t)}$, for $t = 0, \ldots, T$, $x' \in \tilde{\mathcal{X}}$.*

Proof We use induction on t. For $t = 0$, as both CCs have no initial cellular feature maps, HOMP initializes $\mathbf{h}_x^{(0)}, \tilde{\mathbf{h}}_{x'}^{(0)}$ by assigning a constant feature to all cells and the claim holds trivially. Assume the claim holds for some $t \in \{0, \ldots, T\}$. The HOMP update rule reads:

$$\begin{aligned} \mathbf{h}_x^{(t+1)} &= \beta \left(\bigotimes_{\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\mathrm{nat}}} \bigoplus_{y \in \mathcal{N}(x)} \mathsf{MLP}^{(t)}_{\mathcal{N}, \mathrm{rk}(x)}(\mathbf{h}_x^{(t)}, \mathbf{h}_y^{(t)}) \right), \\ \tilde{\mathbf{h}}_{x'}^{(t+1)} &= \beta \left(\bigotimes_{\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\mathrm{nat}}} \bigoplus_{y' \in \mathcal{N}(x')} \mathsf{MLP}^{(t)}_{\mathcal{N}, \mathrm{rk}(x')}(\tilde{\mathbf{h}}_{x'}^{(t)}, \tilde{\mathbf{h}}_{y'}^{(t)}) \right). \end{aligned} \tag{6}$$

Since ρ is a covering map, $\mathcal{N}(x')$ is bjectively mapped to $\mathcal{N}(\rho(x'))$ for every $x' \in \tilde{\mathcal{X}}$ and every neighborhood function $\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\mathrm{nat}}$. Additionally, $\mathrm{rk}(\rho(x')) = \mathrm{rk}(x')$. This, along with the fact that $\bigoplus$ is permutation invariant, and the induction hypothesis ($\tilde{\mathbf{h}}_{x'}^{(t)} = \mathbf{h}_{\rho(x')}^{(t)}$) implies that:

$$\bigoplus_{y' \in \mathcal{N}(x')} \mathsf{MLP}^{(t)}_{\mathcal{N}, \mathrm{rk}(x')}(\tilde{\mathbf{h}}_{x'}^{(t)}, \tilde{\mathbf{h}}_{y'}^{(t)}) = \bigoplus_{y \in \mathcal{N}(\rho(x'))} \mathsf{MLP}^{(t)}_{\mathcal{N}, \mathrm{rk}(\rho(x'))}(\mathbf{h}_{\rho(x')}^{(t)}, \mathbf{h}_y^{(t)}). \tag{7}$$

Thus, combining Equation (6) and Equation (7), we get $\tilde{\mathbf{h}}_{x'}^{(t+1)} = \mathbf{h}_{\rho(x')}^{(t+1)}$. ■

Lemma B.4 *If $\mathcal{X}$ is connected and $\rho : \tilde{\mathcal{X}} \to \mathcal{X}$ is a covering map, $\forall x \in \mathcal{X}, |\rho^{-1}(x)| = \frac{|\tilde{\mathcal{X}}_0|}{|\mathcal{X}_0|}$.*

Proof Since ρ is surjective and rank preserving, the above is equivalent to $\forall x, y \in \mathcal{X}$, $|\rho^{-1}(y)| = |\rho^{-1}(x)|$. Since $\mathcal{X}$ is connected, it suffices to show that this equality holds for any $x, y \in \mathcal{X}$ such that $y \in \mathcal{N}(x)$ for some function $\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\mathrm{nat}}$. We first show that for any natural neighborhood function $\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\mathrm{nat}}$ and cell $x \in \mathcal{X}$ the sets $\{\mathcal{N}(x') \mid x' \in \rho^{-1}(x)\}$ are pairwise disjoint. To see this, assume by contradiction that for a pair of cells $x'_1, x'_2 \in \rho^{-1}(x)$

we have $\mathcal{N}(x'_1) \cap \mathcal{N}(x'_2) \neq \varnothing$. If $z' \in \mathcal{N}(x'_1) \cap \mathcal{N}(x'_2)$, then there is a neighborhood function $\mathcal{N}^* \in \boldsymbol{\mathcal{N}}_{\text{nat}}$ such that $x'_1, x'_2 \in \mathcal{N}^*(z')$. Given that $\rho(x'_1) = \rho(x'_2)$, this would imply that ρ is not injective on $\mathcal{N}^*(z')$, contradicting the definition of a covering map. Now, since $y \in \mathcal{N}(x)$, for any $x' \in \rho^{-1}(x)$ there exists a $y' \in \mathcal{N}(x')$ such that $\rho(y') = y$. Since the set $\{\mathcal{N}(x') \mid x' \in \rho^{-1}(x)\}$ is pairwise disjoint this implies that $|\rho^{-1}(y)| \geqslant |\rho^{-1}(x)|$. Since $y \in \mathcal{N}(x)$, there exists a neighborhood function $\mathcal{N}^* \in \boldsymbol{\mathcal{N}}_{\text{nat}}$ such that $x \in \mathcal{N}^*(y)$, implying by the same reasoning above that $|\rho^{-1}(y)| \leqslant |\rho^{-1}(x)|$. We thus have $|\rho^{-1}(y)| = |\rho^{-1}(x)|$ which concludes the proof. ■

We are now ready to prove Theorem B.2.

Proof Let $\tilde{\mathcal{X}}$ be a combinatorial complex that covers all connected components $\mathcal{Z} \in C(\mathcal{X})$ and $\mathcal{Z}' \in C(\mathcal{X}')$ via maps the maps $\{\rho_{\mathcal{Z}}\}_{\mathcal{Z} \in C(\mathcal{X})}$ and $\{\rho_{\mathcal{Z}'}\}_{\mathcal{Z}' \in C(\mathcal{X}')}$ respectively. Let M be a HOMP model with T layers and let $\mathbf{h}^{(t)}$, $\mathbf{h}'^{(t)}$, and $\tilde{\mathbf{h}}^{(t)}$ denote the cell feature maps of $\mathcal{X}$, $\mathcal{X}'$, and $\tilde{\mathcal{X}}$ respectively at layer t. Lemma B.3 implies that for every $\mathcal{Z} \in C(\mathcal{X})$, $\mathcal{Z}' \in C(\mathcal{X}')$ and every $z \in \mathcal{Z}$, $z' \in \mathcal{Z}'$ we have

$$\begin{aligned} \mathbf{h}_z^{(T)} &= \tilde{\mathbf{h}}_y^{(T)} \quad \forall y \in \rho_{\mathcal{Z}}^{-1}(z), \\ \mathbf{h}_{z'}'^{(T)} &= \tilde{\mathbf{h}}_y^{(T)} \quad \forall y \in \rho_{\mathcal{Z}'}^{-1}(z'). \end{aligned} \tag{8}$$

This implies that the sets of unique values corresponding to the multisets $\{\!\{\mathbf{h}_x^{(T)} \mid x \in \mathcal{X}\}\!\}$, $\{\!\{\mathbf{h}_{x'}'^{(T)} \mid x' \in \mathcal{X}'\}\!\}$ and $\{\!\{\tilde{\mathbf{h}}_y^{(T)} \mid y \in \tilde{\mathcal{X}}\}\!\}$ are the same. Let $n_y, n'_y, \tilde{n}_y$ be the number of times the value $\tilde{\mathbf{h}}_y^{(T)}$ appear in the multisets $\{\!\{\mathbf{h}_x^{(T)} \mid x \in \mathcal{X}\}\!\}$, $\{\!\{\mathbf{h}_{x'}'^{(T)} \mid x' \in \mathcal{X}'\}\!\}$ and $\{\!\{\tilde{\mathbf{h}}_y^{(T)} \mid y \in \tilde{\mathcal{X}}\}\!\}$ respectively. Since each $\mathcal{Z}, \mathcal{Z}'$ are connected, we can use Lemma B.4 to get that $\forall z \in \mathcal{Z}$, $\forall z' \in \mathcal{Z}'$, $|\rho_{\mathcal{Z}}^{-1}(z)| = \frac{|\tilde{\mathcal{X}}_0|}{|\mathcal{Z}_0|}$ and $|\rho_{\mathcal{Z}'}^{-1}(z')| = \frac{|\tilde{\mathcal{X}}_0|}{|\mathcal{Z}'_0|}$. This implies that $\forall y \in \tilde{\mathcal{X}}$

$$n_y = \tilde{n}_y \cdot \left(\sum_{\mathcal{Z} \in C(\mathcal{X})} \frac{|\mathcal{Z}_0|}{|\tilde{\mathcal{X}}_0|} \right), \quad n'_y = \tilde{n}_y \cdot \left(\sum_{\mathcal{Z}' \in C(\mathcal{X}')} \frac{|\mathcal{Z}'_0|}{|\tilde{\mathcal{X}}_0|} \right). \tag{9}$$

Since $\sum_{\mathcal{Z} \in C(\mathcal{X})} |\mathcal{Z}_0| = |\mathcal{X}_0| = |\mathcal{X}'_0| = \sum_{\mathcal{Z}' \in C(\mathcal{X}')} |\mathcal{Z}'_0|$, this implies that $\forall y \in \tilde{\mathcal{X}}$, $n_y = n'_y$. We have shown the set of unique values corresponding to multisets $\{\!\{\mathbf{h}_x^{(T)} \mid x \in \mathcal{X}\}\!\}$ and $\{\!\{\mathbf{h}_x'^{(T)} \mid x' \in \mathcal{X}'\}\!\}$ is the same, and that the number of times each value appears in the multisets is the same, thus the two multisets are equal. Since the readout of a HOMP model can is a function this multiset, $\mathcal{X}$ and $\mathcal{X}'$ are indistinguishable by HOMP. ■

B.2. Topological and Metric Limitations of HOMP

In this section, we rigorously state and proof all results presented in Section 3 regrading HOMPs inability to distinguish CCs based topological/metric properties. We begin by defining the ℓ-dimensional torus CCs. This class contains examples of indistinguishable CCs which differ in both the diameter and homology groups.

Torus CCs. An ℓ-dimensional torus is a Cartesian product of ℓ cycles. More formally:

Definition B.5 (ℓ-dimensional torus CCs) *For a sequence of integers $p_1, \ldots, p_\ell$, the torus $T_{p_1,\ldots,p_\ell}$ is a combinatorial complex $(S, \mathcal{X}, \mathrm{rk})$ defined by:*

$$S = [p_1] \times \cdots \times [p_\ell], \tag{10}$$

$$\mathcal{X}_r = \{\boldsymbol{s}_{\boldsymbol{k}} \mid \boldsymbol{s} \in S, \boldsymbol{k} \in \{0,1\}^\ell, k_1 + \cdots + k_\ell = r\}, \tag{11}$$

where $\boldsymbol{s}_{\boldsymbol{k}}$ is defined by:

$$\boldsymbol{s}_{\boldsymbol{k}} = \{\boldsymbol{s} + \boldsymbol{k}' \mid \boldsymbol{k}' \in \{0,1\}^\ell, \boldsymbol{k}' \leqslant \boldsymbol{k}\}. \tag{12}$$

The sum $\boldsymbol{s} + \boldsymbol{k}'$ is coordinate-wise, where at coordinate j result is taken modulo p_j; we say that $\boldsymbol{k}' \leqslant \boldsymbol{k}$ if $k'_j \leqslant k_j, \forall j \in \{1, \ldots, \ell\}$. See Figure 1(*a*) for an illustration of a 2-dimensional torus. Note that Definition B.5 is only one possible realization of the an ℓ-dimensional torus as a combinatorial complex. As the next lemma shows, all ℓ-dimensional tori are locally isometric.

Lemma B.6 (Joint covering) *Let $T_{p_1,\ldots,p_\ell}$ and $T_{p'_1,\ldots,p'_\ell}$ be two ℓ-dimensional tori such that $\forall j \in \{1, \ldots, \ell\}$, $p_j, p'_j \geqslant 3$. The torus $T_{p_1 \cdot p'_1,\ldots,p_\ell \cdot p'_\ell}$ covers both $T_{p_1,\ldots,p_\ell}$ and $T_{p'_1,\ldots,p'_\ell}$.*

Proof Denote $\boldsymbol{p} = (p_1, \ldots, p_\ell)$, $\boldsymbol{p}' = (p'_1, \ldots, p'_\ell)$, $\tilde{\boldsymbol{p}} = (\tilde{p}_1, \ldots, \tilde{p}_\ell) = (p_1 \cdot p'_1, \ldots, p_\ell \cdot p'_\ell)$. Additionally, denote by S, S', $\tilde{S}$, and $\mathcal{X}$, $\mathcal{X}'$, $\tilde{\mathcal{X}}$ the nodes and cell sets of $T_{\boldsymbol{p}}$, $T_{\boldsymbol{p}'}$ and $T_{\tilde{\boldsymbol{p}}}$ respectively. Define $\rho : \tilde{S} \to S$, $\rho' : \tilde{S} \to S'$ by:

$$\begin{aligned} \rho(\boldsymbol{s}) &= \boldsymbol{s} \mod \boldsymbol{p}, \\ \rho'(\boldsymbol{s}) &= \boldsymbol{s} \mod \boldsymbol{p}', \end{aligned} \tag{13}$$

where $\boldsymbol{s} \mod \boldsymbol{p} = (s_1 \mod p_1, \ldots, s_\ell \mod p_\ell)$. We extend ρ and ρ' to $\tilde{\mathcal{X}}$ by $\rho(x) = \{\rho(s) \mid s \in x\}$. To prove that ρ is a covering map, we start by showing that $\forall r \in \{0, \ldots \ell\}$, $\rho(\tilde{\mathcal{X}}_r) = \mathcal{X}_r$ (i.e. ρ is rank preserving). Recall that all elements of $\tilde{\mathcal{X}}_r$ are of the form $\boldsymbol{s}_{\boldsymbol{k}}$ for some $\boldsymbol{s} \in \tilde{S}$ and $\boldsymbol{k} \in \{0,1\}^\ell$ such that $k_1 + \cdots + k_\ell = r$. Since $\boldsymbol{p} < \tilde{\boldsymbol{p}}$, for every $\boldsymbol{k}' \leqslant \boldsymbol{k}$:

$$(\boldsymbol{s} + \boldsymbol{k}' \mod \tilde{\boldsymbol{p}}) \mod \boldsymbol{p} = (\boldsymbol{s} \mod \boldsymbol{p}) + \boldsymbol{k}' \mod \boldsymbol{p} = \rho(\boldsymbol{s}) + \boldsymbol{k}' \mod \boldsymbol{p}. \tag{14}$$

Therefore, $\rho(\boldsymbol{s}_{\boldsymbol{k}}) = \rho(\boldsymbol{s})_{\boldsymbol{k}} \in \mathcal{X}_r$, and ρ is rank preserving. Notice that since ρ is defined on the node set $\tilde{S}$, for every $x, y, z \in \tilde{\mathcal{X}}$ we have:

- $x \subseteq y \Rightarrow \rho(x) \subseteq \rho(y)$.
- $x, y \subseteq z \Rightarrow \rho(x), \rho(y) \subseteq \rho(z)$.
- $z \subseteq x, y \Rightarrow \rho(z) \subseteq \rho(x), \rho(y)$.

Thus, ρ preserves all natural neighborhood functions. Finally, since $p_1, \ldots, p_\ell \geqslant 3$ it is easy to check that for any $x' \neq y' \in \tilde{\mathcal{X}}$ and $\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\mathrm{nat}}$:

$$y' \in \mathcal{N}(x') \Rightarrow \rho(x') \neq \rho(y'), \tag{15}$$

so ρ bijectively maps $\mathcal{N}(x') \to \mathcal{N}(\rho(x'))$. This implies that ρ is a covering map. An equivalent argument shows that ρ' is also a covering map, completing the proof. ■

Lemma B.6 gives rise to the following useful corollary.

Corollary B.7 *If $T_{p_1,\dots,p_\ell}$ and $T_{p'_1,\dots,p'_\ell}$ are ℓ-dimensional tori such that $p_1 \cdots p_\ell = p'_1 \cdots p'_\ell$ (i.e. $T_{p_1,\dots,p_\ell}$ and $T_{p'_1,\dots,p'_\ell}$ have the same number of 0-cells) and $\forall j \in \{1,\dots,\ell\}$, $p_j, p'_j \geqslant 3$, then for every HOMP model M, $M(T_{p_1,\dots,p_\ell}) = M(T_{p'_1,\dots,p'_\ell})$.*

Proof Both $T_{p_1,\dots,p_\ell}$ and $T_{p'_1,\dots,p'_\ell}$ are connected, have the same number of 0-cells ($(T_{p_1,\dots,p_\ell})_0 = p_1 \cdots p_\ell = p'_1 \cdots p'_\ell = (T_{p'_1,\dots,p'_\ell})_0$), and are covered by $T_{p_1 \cdot p'_i,\dots,p_\ell \cdot p'_\ell}$. Therefore, Theorem B.2 implies that $T_{p_1,\dots,p_\ell}$ and $T_{p'_1,\dots,p'_\ell}$ are indistinguishable by HOMP. ■

Note, that tori with the same number of nodes can still differ on a number of topological and metric properties. In the following we use the family of ℓ dimensional tori to produces examples of topologically/metrically distinct CCs that are indistinguishable by HOMP.

Diameter. For a given (co)adjacency neighborhood function (co)$\mathcal{A}_{r_1,r_2}$, the (r_1, r_2)-diameter of a combinatorial complex $\mathcal{X}$ is defined by

$$\operatorname{diam}_{(\text{co})\mathcal{A}_{r_1,r_2}}(\mathcal{X}) = \max_{x,x' \in \mathcal{X}_{r_1}} d_{(co)\mathcal{A}_{r_1,r_2}}(x, x'), \tag{16}$$

where $d_{(co)\mathcal{A}_{r_1,r_2}}$ is the shortest path distance with respect to neighborhood function (co)$\mathcal{A}_{r_1,r_2}$. Additionally, for $k \in \{0,\dots,\ell\}$, the (r_1, r_2, k) cross diameter is defined by

$$\operatorname{diam}^k_{(\text{co})\mathcal{A}_{r_1,r_2}}(\mathcal{X}) = \max_{\substack{x \in \mathcal{X}_{r_1} \\ y \in \mathcal{X}_k}} \min_{x' \subseteq y} d_{(co)\mathcal{A}_{r_1,r_2}}(x, x'). \tag{17}$$

In this section we show that HOMP is unable to compute diameters of CCs, using ℓ-dimensional tori as a counter example. Corollary B.7 implies that any pair of ℓ-dimensional tori with the same number of nodes (0-cells) is indistinguishable by HOMP, therefore it is enough to construct such tori with different diameters. E.g. the tori $T_{4,4,32}$ and $T_{8,8,8}$ have the same number of 0-cells but different diameters and cross-diameters for any (co)adjacency function and $k = 1, 2, 3$. This can be extended to tori of any dimensions. More formally we have the following proposition for the $(0, 1)$-diameter.

Proposition B.8 *If $T_{p1,\dots,p_\ell}$ and $T_{p'_1,\dots,p'_\ell}$ are ℓ-dimensional tori satisfying*

1. $p_1 \cdots p_\ell = p'_1 \cdots p'_\ell$,
2. $\forall j \in \{1,\dots,\ell\}$, $p_j, p'_j \geqslant 3$, *and*
3. $\sum_{j=1}^{\ell} \lfloor \frac{p_j}{2} \rfloor \neq \sum_{j=1}^{\ell} \lfloor \frac{p'_j}{2} \rfloor$,

then

$$diam_{\mathcal{A}_{0,1}}(T_{p_1,\dots,p_\ell}) \neq diam_{\mathcal{A}_{0,1}}(T_{p'_1,\dots,p'_\ell}) \tag{18}$$

but for any HOMP model M,

$$M(T_{p_1,\dots,p_\ell}) = M(T_{p'_1,\dots,p'_\ell}). \tag{19}$$

Proof Conditions 1 and 2 imply that $T_{p_1,\dots,p_\ell}$ and $T_{p'_1,\dots,p'_\ell}$ are indistinguishable by HOMP. To see that they have different diameters, observe that the graph induced on the nodes of $T_{p_1,\dots,p_\ell}$ by the adjacency neighborhood $\mathcal{A}_{0,1}$ is the Cartesian product of the cyclic graphs

$\mathrm{Cyc}(p_1), \ldots, \mathrm{Cyc}(p_\ell)$. Consequently, since the diameter of a Cartesian product is equal to the sum of diameters over the factors of the product, we have:

$$\begin{aligned}\mathrm{diam}_{\mathcal{A}_{0,1}}(T_{p_1,\ldots,p_\ell}) = \sum_{j=1}^{\ell} \mathrm{diam}(\mathrm{Cyc}(p_j)) = \sum_{j=1}^{\ell} \lfloor \frac{p_j}{2} \rfloor \neq \sum_{j=1}^{\ell} \lfloor \frac{p'_j}{2} \rfloor \\ = \sum_{j=1}^{\ell} \mathrm{diam}(\mathrm{Cyc}(p'_j)) \\ = \mathrm{diam}_{\mathcal{A}_{0,1}}(T_{p'_1,\ldots,p'_\ell}).\end{aligned} \tag{20}$$

■

Homology and Betti numbers. The r-th homology group of a cellular complex [2] encodes the structure of "r-dimensional holes" in the space (e.g. a circle has a singe 1-dimensional hole, a sphere has a single 2-dimensional hole etc). We denote the r-th homology of a CC $\mathcal{X}$ by $H_r(\mathcal{X})$. The rank of the r-th homology group (the size of its minimal generating set) is called the r-th *Betti number*, denoted by $b_r(\mathcal{X})$.

Proposition B.9 (HOMP cannot distinguish complexes based on homology) *Let $T = T_{p_1,\ldots,p_\ell}$ be an ℓ-dimensional torus and $T' = T_{p_1^1,\ldots,p_\ell^1} \sqcup T_{p_1^2,\ldots,p_\ell^2}$ be a disjoint union of two disconnected tori. If $p_1 \cdots p_\ell = p_1^1 \cdots p_\ell^1 + p_1^2 \cdots p_\ell^2$ and $\forall j \in \{1, \ldots, \ell\}$, $p_j, p_j^1, p_j^2 \geqslant 3$, then T and T' are HOMP-indistinguishable but have different homology groups and Betti number of all orders: $\forall r \in \{0, \ldots, \ell\}$, $H_r(T) \neq H_r(T')$, $b_r(T) \neq b_r(T')$.*

Proof First, Lemma B.6 implies that the T, T_1, and T_2 have a common cover. Thus, since T and T' have the same number of cells, Theorem B.2 implies they are HOMP-indistinguishable. Additionally, for every $H_r(T) = \mathbb{Z}^{\binom{\ell}{r}}$ (see e.g. Hatcher (2002)) and since, T' is a disjoint union of T_1 and T_2, $H_r(T') = H_r(T_1) \times H_r(T_2) = \mathbb{Z}^{\binom{\ell}{r}} \times \mathbb{Z}^{\binom{\ell}{r}} = \mathbb{Z}^{2\binom{\ell}{r}}$. Therefore, $\forall r \in \{0, \ldots, \ell\}$, $H_r(T) \neq H_r(T')$ and $b_r(T) = \binom{\ell}{r} \neq 2\binom{\ell}{r} = b_r(T')$. ■

Orientability. We now turn our attention to HOMPs capability to to detect another common topological property: orientability. Loosely speaking, a surface is orientable if one can distinguish between an "inner" and an "outer" side of the surface. A common example of two locally isomorphic surfaces where one is orientable and the other is not is the Möbius strip and a cylinder. For an in depth discussion about orientability and the Möbius strip see Hatcher (2002). We now realize both of these surfaces as cellular complexes. A visualization of the construction can be seen in Figure 1(*b*). We begin by defining two auxiliary functions.

Definition B.10 *For $h, p \in \mathbb{N}$ define the following two functions $\rho^{h,p}_{cyl}, \rho^{h,p}_{möb} : \mathbb{Z}^2 \to \mathbb{Z}^2$:*

$$\rho^{h,p}_{cyl}(\boldsymbol{s}) = (s_1, s_2 \mod p) \tag{21}$$

$$\rho^{h,p}_{möb}(\boldsymbol{s}) = \begin{cases} s_1, s_2 \mod r & s_2 \mod 2p \leqslant p \\ (h+1-s_1, s_2 \mod r) & s_2 \mod 2p > p. \end{cases} \tag{22}$$

2. Homology is not defined for general combinatorial complexes, only for simplicial/cellular complexes.

Using $\rho^{h,p}_{\text{cyl}}$ and $\rho^{h,p}_{\text{möb}}$ we can construct the cylinder and the Möbius strip.

Definition B.11 (Cylinder as CC) *Given a height h and perimeter p, the cylinder $Cyl_{h,p}$ is a 2-dimensional combinatorial complex $(S, \mathcal{X}, \text{rk})$ defined by:*

$$S = [h] \times [p], \tag{23}$$

$$\mathcal{X}_r = \{\boldsymbol{s_k} \mid \boldsymbol{s} \in S, \boldsymbol{k} \in \{0,1\}^2, k_1 + k_2 = r, \rho^{h,p}_{cyl}(\boldsymbol{s} + \boldsymbol{k}) \in S\}, \tag{24}$$

$$\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{X}_2, \tag{25}$$

where $\boldsymbol{s_k}$ is defined by:

$$\boldsymbol{s_k} = \{\rho^{h,p}_{cyl}(\boldsymbol{s} + \boldsymbol{k}') \mid \boldsymbol{k}' \in \{0,1\}^2, \boldsymbol{k}' \leqslant \boldsymbol{k}\}. \tag{26}$$

Definition B.12 (Möbius strip as a CC) *Given two integers h, p, the Möbius strip $Möb_{h,p}$ is a 2-dimensional combinatorial complex $(S, \mathcal{X}, \text{rk})$ defined by:*

$$S = [h] \times [p], \tag{27}$$

$$\mathcal{X}_r = \{\boldsymbol{s_k} \mid \boldsymbol{s} \in S, \boldsymbol{k} \in \{0,1\}^2, k_1 + k_2 = r, \rho^{h,p}_{möb}(\boldsymbol{s} + \boldsymbol{k}) \in S\}, \tag{28}$$

$$\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1 \cup \mathcal{X}_2, \tag{29}$$

where $\boldsymbol{s_k}$ is defined by:

$$\boldsymbol{s_k} = \{\rho^{h,p}_{möb}(\boldsymbol{s} + \boldsymbol{k}) \mid \boldsymbol{k}' \in \{0,1\}^2, \boldsymbol{k}' \leqslant \boldsymbol{k}\}. \tag{30}$$

We now show HOMP is unable to distinguish between CCs based on orientability:

Proposition B.13 (HOMP cannot detect orientability) *For any two integers $h, p \in \mathbb{N}$ such that $h, p \geqslant 3$, and for every HOMP model M, $Cyl_{h,p}$ and $Möb_{h,p}$ are HOMP-indistinguishable, but $Cyl_{h,p}$ is orientable as a topological space while $Möb_{h,r}$ is not.*

Proof First, the fact that the cylinder is orientable, whereas the Möbius strip is not is well known (see e.g. Hatcher (2002) for proof). As for HOMP-indistinguishably, consider the wide cylinder $\text{Cyl}_{h,2p}$ with height h and perimeter $2p$. We show that $\text{Cyl}_{h,2p}$ covers both $\text{Cyl}_{h,p}$ and $\text{Möb}_{h,p}$. Since the two CCs are connected and have the same number of nodes, Theorem B.2 implies that they are HOMP-indistinguishable. Denote by $\tilde{S}$, S^{cyl}, $S^{\text{möb}}$ and $\tilde{\mathcal{X}}$, $\mathcal{X}^{\text{cyl}}$, $\mathcal{X}^{\text{möb}}$ the sets of nodes and cells of $\text{Cyl}_{h,2p}$, $\text{Cyl}_{h,p}$ and $\text{Möb}_{h,p}$ respectively. Define $\rho : \tilde{S} \to S^{\text{cyl}}$ and $\rho' : \tilde{S} \to S^{\text{möb}}$ by $\rho = \rho^{h,p}_{\text{cyl}}\big|_{\tilde{S}}$ and $\rho' = \rho^{h,p}_{\text{möb}}\big|_{\tilde{S}}$. It's easy to verify that $\rho(\tilde{S}) = S^{\text{cyl}}$ and $\rho'(\tilde{S}) = S^{\text{möb}}$, thus ρ and ρ' are well defined and surjective. ρ, ρ' induce maps $\mathcal{P}(\tilde{S}) \to \mathcal{P}(S^{\text{cyl}})$ and $\mathcal{P}(\tilde{S}) \to \mathcal{P}(S^{\text{möb}})$; by abuse of notation we refer to these maps by ρ, ρ' as well. To show that ρ and ρ' are covering maps, we first show that they are rank preserving (i.e. that $\rho(\tilde{\mathcal{X}}_r) = \mathcal{X}^{\text{cyl}}_r$ and $\rho'(\tilde{\mathcal{X}}_r) = \mathcal{X}^{\text{möb}}_r$),

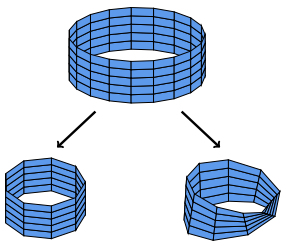

Figure 3: $\text{Cyl}_{h,2p}$ covers both $\text{Cyl}_{h,p}$ and $\text{Möb}_{h,p}$.

and then show that they are local isomorphisms. Recall that all elements of $\tilde{\mathcal{X}}_r$ are of the form $\tilde{\boldsymbol{s}}_{\boldsymbol{k}}$ for some $\tilde{\boldsymbol{s}} \in \tilde{S}$ and $\boldsymbol{k} \in \{0,1\}^2$ such that $k_1 + k_2 = r$. For every $\boldsymbol{k}' \leqslant \boldsymbol{k}$

$$\rho(\rho_{\text{cyl}^{h,2p}}(\tilde{\boldsymbol{s}} + \boldsymbol{k}')) = \rho_{\text{cyl}^{h,p}}(\rho(\tilde{\boldsymbol{s}}) + \boldsymbol{k}'), \tag{31}$$

so $\rho(\tilde{\boldsymbol{s}}_{\boldsymbol{k}}) = \rho(\tilde{\boldsymbol{s}})_{\boldsymbol{k}}$. Additionally,

$$\rho'(\rho^{h,2p}_{\text{cyl}}(\tilde{\boldsymbol{s}} + \boldsymbol{k}')) = \begin{cases} \rho^{h,p}_{\text{möb}}(\rho'(\tilde{\boldsymbol{s}}) + \boldsymbol{k}') & \tilde{s}_1 \leqslant p \\ \rho^{h,p}_{\text{möb}}(\rho'(\tilde{\boldsymbol{s}}) + (-k'_1, k'_2)) & \tilde{s}_1 > p. \end{cases} \tag{32}$$

so

$$\rho'(\tilde{\boldsymbol{s}}_{\boldsymbol{k}}) = \begin{cases} \rho'(\tilde{\boldsymbol{s}})_{\boldsymbol{k}} & \tilde{s}_1 \leqslant p \\ (\rho'(\tilde{\boldsymbol{s}}) + (-1,0))_{\boldsymbol{k}} & \tilde{s}_1 > p. \end{cases} \tag{33}$$

By the definitions $\tilde{\mathcal{X}}_r$, $\mathcal{X}^{\text{cyl}}$ and $\mathcal{X}^{\text{möb}}$ we now have $\rho(\tilde{\mathcal{X}}_r) = \mathcal{X}^{\text{cyl}}_r$ and $\rho'(\tilde{\mathcal{X}}_r) = \mathcal{X}^{\text{möb}}_r$ as needed. Since ρ and ρ' are extended to $\mathcal{P}(\tilde{S})$ from $\tilde{S}$, for every $x, y, z \in \tilde{\mathcal{X}}$

- $x \subseteq y \Rightarrow \rho(x) \subseteq \rho(y)$ and $\rho'(x) \subseteq \rho'(y)$.
- $x, y \subseteq z \Rightarrow \rho(x), \rho(y) \subseteq \rho(z)$ and $\rho'(x), \rho'(y) \subseteq \rho'(z)$
- $z \subseteq x, y \Rightarrow \rho(z) \subseteq \rho(x), \rho(y)$ and $\rho'(z) \subseteq \rho'(x), \rho'(y)$.

Therefore, ρ and ρ' preserves all natural neighborhood functions. Finally, since $h, p \geqslant 3$, for $x, y \in \tilde{\mathcal{X}}$ and $\mathcal{N} \in \boldsymbol{\mathcal{N}}_{\text{nat}}$, $y \in \mathcal{N}(x) \Rightarrow \rho(x) \neq \rho(y)$ and $\rho'(x) \neq \rho'(y)$. This implies that ρ and ρ' are local isomorphisms, completing the proof. ■

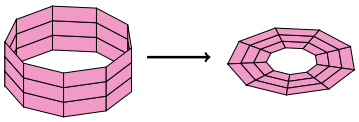

Figure 4: Cylinders are planar.

Planarity. A topological space is considered planar if it can be continuously embedded in $\mathbb{R}^2$. Proposition B.13 provides us with the following corollary.

Corollary B.14 (HOMP cannot detect planarity) *There exist pairs of cellular complexes $\mathcal{X}, \mathcal{X}'$ such that the induced topology of $\mathcal{X}$ is planar while the induced topology of $\mathcal{X}'$ is not, but $\mathcal{X}$ and $\mathcal{X}'$ are HOMP-indistinguishable.*

Proof The CCs $\text{Cyl}_{h,p}$ and $\text{Möb}_{h,p}$ for $p, h \geqslant$ are HOMP-indistinguishable according to Proposition B.13. The Möbius strip is not planar (see e.g., Hatcher (2002)), whereas the cylinder is. ■

Appendix C. (Scalable) Multi-cellular Networks

In this section we motivate and formally define the MCN and SMCN architectures.

C.1. Multi-Cellular Networks

In graph learning message-passing expresivity issues have been successfully addressed by invariant graph networks (IGNs) (Maron et al., 2018, 2019; Keriven and Peyré, 2019; Azizian and Lelarge, 2020). These models are constructed by stacking equivariant linear layers between tensor spaces. We propose a similar approach, incorporating equivariant linear

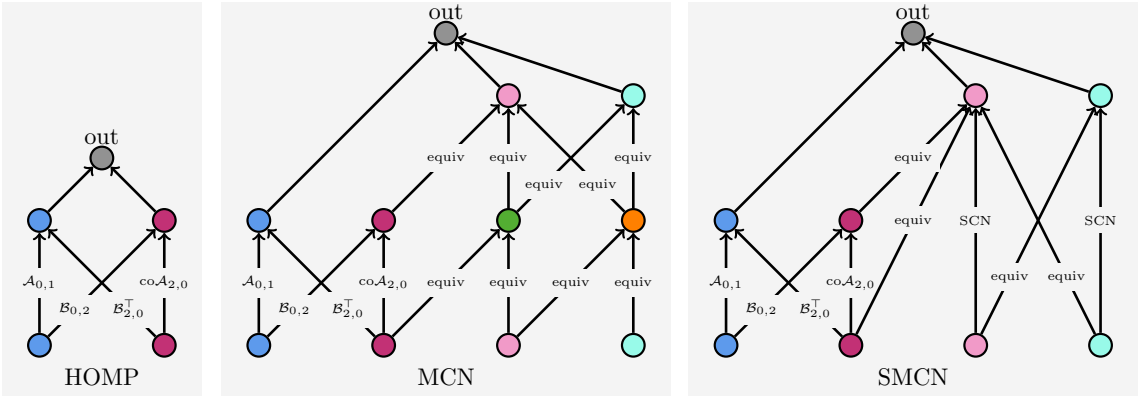

Node labels						Edge labels		
HOMP		MCN		MCN		HOMP	MCN	SMCN
●	●	●	●	●	●	$\mathcal{A}_{r_1,r_2}$ / (co)$\mathcal{A}_{r_1,r_2}$ / $\mathcal{B}_{r_1,r_2}$ / $\mathcal{B}_{r_1,r_2}^\top$	"equiv" label	"SCN" label
$\mathcal{C}^{\boldsymbol{e}_0}$	$\mathcal{C}^{\boldsymbol{e}_2}$	$\mathcal{C}^{\boldsymbol{e}_0+\boldsymbol{e}_1}$	$\mathcal{C}^{2\boldsymbol{e}_0}$	$\mathcal{C}^{2\boldsymbol{e}_0+\boldsymbol{e}_1}$	$\mathcal{C}^{3\boldsymbol{e}_0}$	Neighborhood function induced update	Equivariant linear layer	Sub complex network

Figure 5: Example tensor diagrams for HOMP, MCN, and SMCN. HOMP diagrams can only use nodes labeled with standard cochain spaces. MCN diagram can additionally use nodes labeled with multi-cellular cochain spaces and edges labeled with "equiv" updates. SMCN diagrams introduce edges labeled with "SCN".

layers into the HOMP framework. We integrate these layer by introducing new node and edge types into existing HOMP tensor diagrams. We first define multi-cellular cochain spaces which parallel the tensor spaces used in IGNs.

Multi-cellular cochain spaces. The multi-cellular cochain space associated with an ℓ-dimensional combinatorial complex $\mathcal{X}$ and an $\ell+1$-tuple of integers $\boldsymbol{k} = (k_0, \dots, k_\ell)$, denoted by $\mathcal{C}^{\boldsymbol{k}}(\mathcal{X}, \mathbb{R}^d)$, is the space of functions $\mathbf{h} : \mathcal{X}^{\boldsymbol{k}} := \mathcal{X}_0^{k_0} \times \cdots \times \mathcal{X}_\ell^{k_\ell} \to \mathbb{R}^d$. For brevity, we often denote $\mathcal{C}^{\boldsymbol{k}}(\mathcal{X}, \mathbb{R}^d)$ by $\mathcal{C}^{\boldsymbol{k}}$.

Multi-cellular cochain symmetry group. Given enumerations of the skeletons $\mathcal{X}_0 = \{x_0^0, \dots, x_{n_0}^0\}, \dots, \mathcal{X}_\ell = \{x_0^\ell, \dots, x_{n_\ell}^\ell\}$, a multi-cellular cochain $\mathbf{h} \in \mathcal{C}^{\boldsymbol{k}}(\mathcal{X}, \mathbb{R}^d)$ can be identified with the tensor $\mathbf{A}_{\mathbf{h}} \in \mathbb{R}^{n_0^{k_0} \times \cdots \times n_\ell^{k_\ell} \times d}$, defined by

$$(\mathbf{A}_{\mathbf{h}})_{\boldsymbol{i}_0,\dots,\boldsymbol{i}_\ell,:} = \mathbf{h}\left(x^0_{(\boldsymbol{i}_0)_1}, \dots, x^0_{(\boldsymbol{i}_0)_{k_0}}, \dots, x^\ell_{(\boldsymbol{i}_\ell)_1}, \dots, x^\ell_{(\boldsymbol{i}_\ell)_{k_\ell}}\right) \tag{34}$$

for multi-indices $\boldsymbol{i}_1 \in \{1, \dots, n_0\}^{k_0}, \dots, \boldsymbol{i}_\ell \in \{1, \dots, n_\ell\}^{k_\ell}$. Therefore, $\mathcal{C}^{\boldsymbol{k}}$ can be identified with the space $\mathbb{R}^{n_0^{k_0} \times \cdots \times n_\ell^{k_\ell} \times d}$. The group $G = S_{n_0} \times \cdots \times S_{n_\ell}$ acts on $\mathcal{C}^{\boldsymbol{k}}$ by

$$(\boldsymbol{\sigma} \cdot \mathbf{h})(\boldsymbol{x}_0, \dots, \boldsymbol{x}_\ell) = \mathbf{h}(\sigma_0 \cdot \boldsymbol{x}_0, \dots, \sigma_\ell \cdot \boldsymbol{x}_\ell), \tag{35}$$

where if $\boldsymbol{x}_r = \left(x^r_{j_1}, \ldots, x^r_{j_{k_r}}\right) \in \mathcal{X}_r^{k_r}$, $\sigma_r \cdot \boldsymbol{x} = \left(x^r_{\sigma_r^{-1}(j_1)}, \ldots, x^r_{\sigma_r^{-1}(j_{k_r})}\right)$, i.e. the action of S_{n_r} changes permutes each of the r-rank cells in the tuple $\boldsymbol{x}_r$. We aim to construct architectures invariant to the action of G.

Equivariant linear maps between multi-cellular cochain spaces. Following the construction of other popular permutation invariant architectures Maron et al. (2018); Zaheer et al. (2017); Hartford et al. (2018); Bronstein et al. (2021) we aim to characterize the space of linear equivariant maps $L : \mathcal{C}^{\boldsymbol{k}} \to \mathcal{C}^{\boldsymbol{k}'}$ for pairs of ℓ-tuples $\boldsymbol{k} = (k_0, \ldots, k_\ell)$, $\boldsymbol{k}' = (k'_0, \ldots, k'_\ell)$. Since $\mathcal{C}^{\boldsymbol{k}}$ can be identified with $\mathbb{R}^{n_0^{k_0} \times \cdots \times n_\ell^{k_\ell} \times d}$, we can use the basis for the space of equivariant linear layers $L : \mathbb{R}^{n_0^{k_0} \times \cdots \times n_\ell^{k_\ell} \times d} \to \mathbb{R}^{n_0^{k'_0} \times \cdots \times n_\ell^{k'_\ell} \times d'}$ that was constructed in Maron et al. (2018). Using this basis, denoted by $\{L_\gamma\}_{\gamma \in \Gamma(\boldsymbol{k}, \boldsymbol{k}', d, d')}$, we construct learnable equivariant layers of the form

$$L(\mathbf{h}) = \beta \left(\sum_{\gamma \in \Gamma(\boldsymbol{k}, \boldsymbol{k}', d, d')} w_\gamma L_\gamma(\mathbf{A}_{\mathbf{h}}) \right), \tag{36}$$

where $\{w_\gamma\}_{\gamma \in \Gamma(\boldsymbol{k}, \boldsymbol{k}', d, d')}$ are learnable parameters and β is a non-linearity. We expand the HOMP framework, adding equivariant layers to the tensor diagram scheme (as depicted in Figure 5), defining a new class of TDL architecture we call multi-cellular networks (MCNs). We now formally describe the components of the MCN scheme.

Diagram. Similar to HOMP tensor diagrams, MCN tensor diagrams are layered directed graphs with labeled nodes and edges. Each node is labeled by a multi-cellular cochain space, extending the class of node labels used in HOMP tensor diagrams. Directed edges with source and target nodes labeled by $\mathcal{C}^{\boldsymbol{e}_r}$ can be labeled by any neighborhood function, while edges between nodes labeled by other types of multi-cellular cochain spaces are labeled with the new label "equiv".

Input. The input to the MCN model is determined by the 0-th layer of the tensor diagram, whose nodes can be labeled by the following types of multi-cellular cochain spaces: (1) nodes labeled by $\mathcal{C}^{\boldsymbol{e}_r}$ which take the r-rank cell features as input; (2) Nodes labeled by $\mathcal{C}^{\boldsymbol{e}_{r_1} + \boldsymbol{e}_{r_2}}$ which take the matrix form of the incidence neighborhood $\mathcal{B}_{r_1, r_2}$ as input; (3) Nodes labeled by $\mathcal{C}^{2\boldsymbol{e}_r}$ which take the matrix form of the (co)adjacency matrices (co)$\mathcal{A}_{r,r'}$.

Update. The multi-cellular cochains computed at the t-th layer of the network are denoted by $\mathbf{h}^{(t)}$ (i.e. if $\boldsymbol{x} \in \mathcal{X}_0^{k_0} \times \cdots \times \mathcal{X}_\ell^{k_\ell}$ is a multi-cell and there is a node in the t-th layer of the diagram marked by $\mathcal{C}^{\boldsymbol{k}}$ we use $\mathbf{h}^{(t)}_{\boldsymbol{x}}$ to denote its multi-cellular feature). For each directed edge (v_t, v_{t+1}) in the diagram we compute a message $\boldsymbol{m}^{v_t, v_{t+1}}$ based on $\mathbf{h}^{(t)}$; the message computation depends on the label of (v_t, v_{t+1}). If the edge is labeled by a neighborhood function $\mathcal{N}$, (in which case v_t and v_{t+1} are labeled by standard cochain spaces), the message $\boldsymbol{m}^{v_t, v_{t+1}}$ is computed by $\bigoplus_{y \in \mathcal{N}(x)} M^{v_t, v_{t+1}}(\mathbf{h}^{(t)}_x, \mathbf{h}^{(t+1)}_y)$, where $M^{v_t, v_{t+1}}$ is parameterized by an MLP. Note that this is the exact message used in HOMP tensor diagrams. For an edge labeled by "equiv", $\boldsymbol{m}^{v_t, v_{t+1}}$ is computed using an equivariant layer as in Equation (36). After computing the messages for all incoming edges to a target node v_{t+1} labeled by $\mathcal{C}^{\boldsymbol{k}}$

we aggregate them to for form a new multi-cellular cochain defied by

$$\mathbf{h}_{\boldsymbol{x}}^{(t+1)} = \bigotimes_{\substack{(u,v_{t+1}) \text{ is an} \\ \text{edge in the diagram}}} \boldsymbol{m}_{\boldsymbol{x}}^{u,v_{t+1}}. \tag{37}$$

for $\boldsymbol{x} \in \mathcal{X}_0^{k_0} \times \cdots \times \mathcal{X}_\ell^{k_\ell}$. The last layer of the tensor diagram contains a single node representing a final readout layer.

C.2. Scalable Multi-cellular Networks

Despite its strong expressive power, implementing MCN in full generality is impractical. This is because both the computational complexity and the size of the basis $\{L_\gamma\}_{\gamma \in \Gamma(\boldsymbol{k},\boldsymbol{k}',d,d')}$ grow exponentially with $\boldsymbol{k}$ and $\boldsymbol{k}'$ — $|\Gamma(\boldsymbol{k},\boldsymbol{k}',d,d')| = d \cdot d' \cdot \prod_{i=1}^{\ell} b(k_i + k_i')$ where $b(k)$ is the k-th bell number. SMCN is a more scalable version of MCN that still addresses many of the expressivity issues observed in HOMP. First we restrict the SMCN framework to only use multi-cellular cochain space for which $\sum_{i=0}^{\ell} k_i \leqslant 3$. Of these, the two types of cochain spaces that introduce the largest computational overheads are $\mathcal{C}^{3\boldsymbol{e}_r}$ and $\mathcal{C}^{2\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$. We replace the equivariant linear updates induced by these multi-cellular cochain spaces with new updates inspired by expressive GNNs. This provide a middle ground between expressive power and scalability. The adaptation of these graph architectures to the domain of CCs is defined using the augmented Hasse graphs.

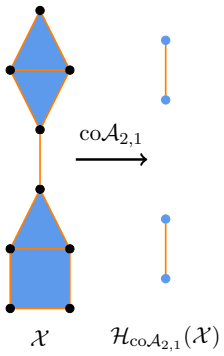

Figure 6: $\mathcal{H}_{\text{co}\mathcal{A}_{2,1}}$.

Definition C.1 (Augmented Hasse graphs) *Given a combinatorial complex $\mathcal{X}$, its augmented Hasse graph with respect to the (co)adjacency neighborhood* (co)$\mathcal{A}_{r,r'}$*, is defined by* $\mathcal{H}_{(\text{co})\mathcal{A}_{r,r'}} = (\mathcal{V}, \mathcal{E})$*, with* $\mathcal{V} = \mathcal{X}_{r_1}$ *and* $\mathcal{E} = \{(x,y) \mid y \in \mathcal{A}_{r,r'}(x)\}$.

See Figure 6 for an example. Intuitively, $\mathcal{H}_{(\text{co})\mathcal{A}_{r,r'}}$'s graph structure encodes the relational interactions between r-rank cells given by (co)$\mathcal{A}_{r,r'}$. Using augmented Hasse graphs, we define new multi-cellular updates and use them to create a more scalable version of MCN called scalable MCN (SMCN).

Replacing $\mathcal{C}^{2\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$ with $\mathcal{C}^{\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$. Recall that the space $\mathcal{C}^{2\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$ can be identified with $\mathbb{R}^{n_{r_1}^2 \times n_{r_2} \times d}$. Under the action of $S_{n_{r_1}} \times S_{n_{r_2}}$, the tensor $\mathbf{H}_{\mathbf{h}} \in \mathbb{R}^{n_{r_1} \times n_{r_1} \times n_{r_2} \times d}$ can be viewed as a bag of tensors $\{\mathbf{H}_{\mathbf{h}}^k \in \mathbb{R}^{n_{r_1} \times n_{r_1} \times d} \mid k \in [n_{r_2}]\}$, each of which is considered up-to row and column S_{n_r}-permutations. This is precisely the data structure processed by subgraph neural networks (Bevilacqua et al., 2021), which operate on a set of adjacency matrices corresponding to different subgraphs defined over a fixed set of nodes. Subgraph networks have stronger expressive power than MPNNs and demonstrated strong experimental performance. In addition, They have quadratic runtime complexity as opposed to $\mathcal{O}(n_{r_1}^2 \cdot n_{r_2})$ for $\mathcal{C}^{2\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}} \to \mathcal{C}^{2\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$ equivariant layers. Therefore, they are a good candidate for replacing the latter equivariant updates. For a comprehensive review of subgraph networks, see Bronstein (2021).

Following the above discussion we formally define subcomplex network (SCN) updates. Let v_t, v_{t+1} be tensor diagram nodes which are both labeled by $\mathcal{C}^{\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$ and whose tensor diagram edge (v_t, v_{t+1}) is labeled by "SCN". We compute the message $\boldsymbol{m}^{v_t,v_{t+1}} \in \mathcal{C}^{\boldsymbol{e}_{r_1}+\boldsymbol{e}_{r_2}}$:

$$\begin{aligned}\boldsymbol{m}_{x,y}^{v_t,v_{t+1},r,r'} &= \mathsf{MLP}^{r,r'}\left(\mathbf{h}_{x,y}^{(t)}, \mathbf{h}_{(\mathrm{co})\mathcal{A}_{r_1,r}(x),y}^{(t)}, \mathbf{h}_{x,(\mathrm{co})\mathcal{A}_{r_2,r'}(y)}^{(t)}, \mathbf{h}_{x,\mathcal{B}_{r_1,r_2}(x)}, \mathbf{h}_{\mathcal{B}^{\top}_{r_2,r_1}(y),y}\right),\\ \boldsymbol{m}_{x,y}^{v_t,v_{t+1}} &= \bigotimes_{\substack{r=0\\r'=0}}^{\ell} \boldsymbol{m}_{x,y}^{v_t,v_{t+1},r,r'},\end{aligned} \tag{38}$$

where if $\mathcal{Q}_1 \subseteq \mathcal{X}_{r_1}$ and $\mathcal{Q}_2 \subseteq \mathcal{X}_{r_2}$ are sets of cells, $\mathbf{h}_{\mathcal{Q}_1,y} := \sum_{x' \in \mathcal{Q}_1} \mathbf{h}_{x',y}$ and $\mathbf{h}_{x,\mathcal{Q}_2} := \sum_{y' \in \mathcal{Q}_2} \mathbf{h}_{x,y'}$. Note that looking at the case $r = r'$, the intermediate messages $\{\boldsymbol{m}^{v_t,v_{t+1},r,r}\}_{r=0}^{\ell}$ can be viewed as applying a subgraph network update to the bag of augmented Hasse graphs $\{\mathcal{H}^{x}_{(\mathrm{co})gA_{r_1,r}} \mid x \in \mathcal{X}_r\}$, where $\mathcal{H}^{x}_{\mathcal{A}_{r_1,r}}$ denotes the graph $\mathcal{H}_{\mathcal{A}_{r_1,r}}$ with the cells $\{y \in \mathcal{X}_{r_1} \mid y \subseteq x\}$ marked.

Replacing $\mathcal{C}^{3\boldsymbol{e}_r}$ with $\mathcal{C}^{2\boldsymbol{e}_r}$. Recall that the space $\mathcal{C}^{3\boldsymbol{e}_r}$ can be identified with $\mathbb{R}^{n_r^3 \times d}$. Thus, equivariant linear layers $\mathcal{C}^{3\boldsymbol{e}_r} \to \mathcal{C}^{3\boldsymbol{e}_r}$ can be identified with 3-IGN layers which take as input 3-tensors indexed by $\mathcal{X}_r$. The expressive GNN literature offers a few candidates for efficient 3-IGN substitutes. The first option we considered is PPGN (Maron et al., 2019), which matches 3-IGN's 3-WL expressive power with a runtime of $\mathcal{O}(|\mathcal{V}|^{2.p})$ (run time bottleneck is $|\mathcal{V}| \times |\mathcal{V}|$ matrix multiplication). The other option is using subgraph networks with node marking. These network have a runtime of $\mathcal{O}(|\mathcal{V}| \cdot |\mathcal{E}|)$, are strictly more expressive than MPNNs (> 2-WL), but are less powerful than 3-IGNs. We experimented with both versions and found no significant performance improvement in using PPGN. Therefore, for simplicity of the method we continue with the subgraph version, which is formally detailed next, but note that (since PPGN can implement subgraph networks) all theoretical results hold for the PPGN case as well. The $\mathcal{C}^{2\boldsymbol{e}_r} \to \mathcal{C}^{2\boldsymbol{e}_r}$ message is computed using the same update rule as in Equation (38) with $r_1 = r_2 = r$, which in this case simplifies to:

$$\begin{aligned}\boldsymbol{m}_{x,x'}^{v_t,v_{t+1},r',r''} &= \mathsf{MLP}^{r',r''}\left(\mathbf{h}_{x,x'}^{(t)}, \mathbf{h}_{x,x}^{(t)}, \mathbf{h}_{x',x'}^{(t)}, \sum_{x'' \in (\mathrm{co})\mathcal{A}_{r,r'}(x)} \mathbf{h}_{x'',x'}^{(t)}, \sum_{x'' \in (\mathrm{co})\mathcal{A}_{r,r''}(x')} \mathbf{h}_{x,x''}^{(t)}\right),\\ \boldsymbol{m}_{x,x'}^{v_t,v_{t+1}} &= \bigotimes_{\substack{r'=0\\r''=0}}^{\ell} \boldsymbol{m}_{x,x'}^{v_t,v_{t+1},r',r''}.\end{aligned} \tag{39}$$

Appendix D. MCN and SMCN Expressive Power

D.1. MCN.

In this section, we analyze the expressive power of MCN defined in Section 4. We begin by formally defining CC isomorphism, as described in Hajij et al. (2022b).

Definition D.1 (CC isomorphism) *A pair of CCs $(S, \mathcal{X}, \mathrm{rk})$, $(S', \mathcal{X}', \mathrm{rk}')$ are isomorphic if there exists a bijective map $\rho : \mathcal{X} \to \mathcal{X}'$ such that:*

1. $\mathrm{rk}(x) = \mathrm{rk}'(\rho(x)) \quad \forall x \in \mathcal{X}$,
2. $x \subseteq y \Rightarrow \rho(x) \subseteq \rho(y) \quad \forall x, y \in \mathcal{X}$.

Is there exists an isomorphism $\rho : \mathcal{X} \to \mathcal{X}'$ we say that $\mathcal{X}$ and $\mathcal{X}'$ are *isomorphic*; if such an isomorphism does not exist, we say taht $\mathcal{X}$ and $\mathcal{X}'$ are *non-isomorphic*.

Proposition D.2 (MCN expressive power) *If $\mathcal{X}$ and $\mathcal{X}'$ are non-isomorphic there exists an MCN model M such that*

$$M(\mathcal{X}) \neq M(\mathcal{X}'). \tag{40}$$

Proof First, let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ and $\mathcal{H}' = (\mathcal{V}', \mathcal{E}')$ be the Hasse graphs of $\mathcal{X}$ and $\mathcal{X}^*$ respectively, defined by

$$\mathcal{V} = \mathcal{X}, \tag{41}$$
$$\mathcal{V}' = \mathcal{X}', \tag{42}$$
$$\mathcal{E} = \{(x, y) \in \mathcal{X} \times \mathcal{X} \mid x \subseteq y, \mathrm{rk}(x) = \mathrm{rk}(y) - 1\}, \tag{43}$$
$$\mathcal{E}' = \{(x', y') \in \mathcal{X}' \times \mathcal{X}' \mid x' \subseteq y', \mathrm{rk}'(x') = \mathrm{rk}'(y') - 1\}. \tag{44}$$

It was shown in Hajij et al. (2022b) that a pair of CCs is isomorphic if and only if their corresponding Hasse graphs are isomorphic. Therefore, in our case $\mathcal{H}$ and $\mathcal{H}'$ are non-isomorphic graphs. Since any pair of non-isomorphic graphs of size n are n-WL distinguishable, and k-IGN networks can distinguish between any pair of k-WL indistinguishable graphs (see Maron et al. (2019)), it is enough to prove that there exists a MCN model which is able to simulate any k-IGN network on the Hasse graphs. Let $\boldsymbol{A}$ be the adjacency matrix of $\mathcal{H}$ and define $n = |\mathcal{X}|$, $n_r = |\mathcal{X}_r|$ for all $r \in \{0, \dots, \ell\}$. $\boldsymbol{A}$ can be decomposed into block matrices $\boldsymbol{A}^{r_1, r_2}$ for $r_1, r_2 \in \{0, \dots, \ell\}$ defined by:

$$\boldsymbol{A}^{r_1, r_2} = \begin{cases} \mathbf{0}_{n_i \times n_j} & r_1 \neq r_2 + 1 \\ \boldsymbol{B}_{r_1, r_2} & r_1 = r_2 + 1, \end{cases} \tag{45}$$

where $\boldsymbol{B}_{r_1, r_2}$ is the matrix form of neighborhood function $\mathcal{B}_{r_1, r_2}$. matrices $\boldsymbol{A}^{r_1, r_2}$ can be view as a multi-cellular cochains $\mathbf{h}^{r_1, r_2} \in \mathcal{C}^{\boldsymbol{e}_{r_1} + \boldsymbol{e}_{r_2}}(\mathcal{X}, \mathbb{R})$ so $\boldsymbol{A}$ can be realized as an element of $\mathcal{Q} := \times_{r_1=0, r_2=0}^{\ell} \mathcal{C}^{\boldsymbol{e}_{r_1} + \boldsymbol{e}_{r_2}}(\mathcal{X}, \mathbb{R})$. Recall that all neighborhood matrices $\boldsymbol{B}_{r_1, r_2}$ are given as input to the MCN model and so we can recover $\boldsymbol{A}$. To show that MCN can simulate any k-IGN update on $\boldsymbol{A}$, we need to show that it can compute $L(\boldsymbol{A})$ for any S_n-equivariant linear function $L : \mathcal{Q}^{\otimes k} \to \mathcal{Q}^{\otimes k'}$, where $\mathcal{Q}^{\otimes k}$ represents taking the tensor product of $\mathcal{Q}$ with itself k times. Let $G < S_n$ be the subgroup of permutations preserving the subsets $\{1, \dots, n_0\}$, $\{n_0 + 1, \dots, n_0 + n_1\}$, $\dots$, $\{n_0 + \dots + n_{\ell-1} + 1, \dots, n_0 + \dots + n_\ell\}$; $G \cong S_{n_0} \times \dots \times S_{n_\ell} \subseteq [n]$. Since G is a subgroup of S_n, all S_n equivariant linear maps are also G-equivariant. Thus it is enough to show that we are can compute $L(\mathbf{h})$ for all G-equivariant linear maps $L : \mathcal{Q}^{\otimes k} \to \mathcal{Q}^{\otimes k'}$.

The space $\mathcal{Q} = \times_{r_1=0, r_2=0}^{\ell} \mathcal{C}^{\boldsymbol{e}_{r_1} + \boldsymbol{e}_{r_2}}(\mathcal{X}, \mathbb{R})$ can be embedded into the multi-cellular cochain space $\mathcal{C}^{\mathbf{1}_{\ell+1}}(\mathcal{X}, \mathbb{R}^{(\ell+1)^2})$ via the following map:

$$T(\mathbf{h})(x_0, \dots x_\ell) = \Big\Vert_{r_1=0, r_2=0}^{\ell} \mathbf{h}^{r_1, r_2}(x_{r_1}, x_{r_2}), \tag{46}$$

where $\Vert$ stands for concatenation, $\mathbf{1}_{\ell+1} = (1, \dots, 1) \in \mathbb{R}^{\ell+1}$ is the all ones vector, $x_r \in \mathcal{X}_r$ is a cell of rank r and $\mathbf{h} \in \mathcal{Q}$ composed of the multi-cellular cochains $\mathbf{h}^{r_1, r_2} \in \mathcal{C}^{\boldsymbol{e}_{r_1} + \boldsymbol{e}_{r_2}}(\mathcal{X}, \mathbb{R})$.

MCN can use any linear function $L : \mathcal{C}^{k\cdot\mathbf{1}_{\ell+1}}(\mathcal{X}, \mathbb{R}^{(\ell+1)^2}) \to \mathcal{C}^{k'\cdot\mathbf{1}_{\ell+1}}(\mathcal{X}, \mathbb{R}^{(\ell+1)^2})$ which is G-equivariant, and so it can compute $L(\mathbf{h})$ for all linear maps as define above, concluding the proof. ■

D.2. SMCN

In this section we formally demonstrate the SMCN's ability to mitigate many of the expressive limitations demonstrated in Appendix B. We begin by providing a useful lemma which allows us to leverage several expressivity results from the subgraph GNN literature in our setting. We then provide an in depth discussion on the ability of SMCN to express each one of the four aforementioned metric/topological properties: diameter, orientability, planarity and homology.

Lemma D.3 *For any CS-GNN (Bar-Shalom et al., 2024) model M operating on the Hasse graph $\mathcal{H}_{(co)\mathcal{A}_{r_1,r_2}}$ using cells of rank $r \geqslant r_1$ as super-nodes, there exits an SMCN model M', such that for any CC $\mathcal{X}$ of dimension $\geqslant r_1, r_2, r$, $M(\mathcal{H}_{(co)\mathcal{A}_{r_1,r_2}}) = M'(\mathcal{X})$.*

Proof First, note that the incidence matrix $\mathcal{B}_{r_1,r} \in \mathcal{C}^{\boldsymbol{e}_{r_1}+\boldsymbol{e}_r}$ is equivalent to the "simple node marking" defined in Bar-Shalom et al. (2024), so SMCN can recover the input to the CS-GNN architecture. Second, by taking

$$\mathsf{MLP}^{r,r'}(x,y) = \begin{cases} \mathsf{MLP}(x,y) & \text{if } r = r_2 \text{ and } r' = r_1, \\ \mathbf{0} & \text{otherwise} \end{cases} \tag{47}$$

for a fixed MLP, Equation 38 becomes identical to the CS-GNN update. ■

Remark D.4 *For the case where $r = r_1$ (i.e. super-nodes are regular Hasse graph nodes) the CS-GNN architecture becomes equivalent to GNN-SSWL+ (Zhang et al., 2023).*

Diameter. We first show SMCN is capable to fully leverage the information provided by the (cross) diameters of an input CC. see Appendix B for a definition.

Proposition D.5 (SMCN can compute diameters) *If $\mathcal{X}, \mathcal{X}'$ are CCs such that*

$$diam^r_{\mathcal{A}_{r_1,r_2}}(\mathcal{X}) \neq diam^r_{\mathcal{A}_{r_1,r_2}}(\mathcal{X}'), \tag{48}$$

for $r_1, r_2, r \in \mathbb{N}$ with $r_1 \leqslant r$, then there exists an SMCN model M such that $M(\mathcal{X}) \neq M(\mathcal{X}')$

Proof In Zhang et al. (2023), it was shown that GNN-SSWL+, with standard node marking applied to a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, can compute a final feature representation:

$$\mathbf{h}^{(T)}_{u,v} = d_{\mathcal{G}}(u,v) \quad \text{for } u, v \in \mathcal{V}. \tag{49}$$

By taking the maximum over $\mathbf{h}^{(T)}_{u,v}$, GNN-SSWL+ can distinguish between graphs with different diameters. Similarly, It was shown in Bar-Shalom et al. (2024) that CS-GNN with standard node marking applied to a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and super-node set $\mathcal{V}^*$ can compute a final feature representation

$$\mathbf{h}^{(T)}_{S,v} = d_{\mathcal{G}}(S,v) \quad \text{for } v \in \mathcal{V} \text{ and } S \in \mathcal{V}^*. \tag{50}$$

By taking the maximum over $\mathbf{h}_{S,v}^{(T)}$, CS-GNN with standard node marking can distinguish between graphs with different cross diameters. Thus, applying Lemma D.3 and Remark D.4 on the Hasse graph $\mathcal{H}_{\mathcal{A}_{r_1,r_2}}$ with $\mathcal{X}_r$ as "super-nodes" we get that SMCN can distinguish between CCs with different (cross) diameters. ■

Orientability and planarity. We now show SMCN is able to separate the cylinder and the Möbius strip. This implies that SMCN is strictly better than HOMP at detecting planarity and orientability. Understanding SMCN's ability to fully detect orientability or planarity is still and open question and is left for future work.

Proposition D.6 (SMCN can separate a cylinder and a Möbius strip) *For any two integers $h, p \in \mathbb{N}$ such that $h, p \geqslant 3$, there exists an SMCN model M, such that:*

$$M(Cyl_{h,p}) \neq M(Möb_{h,p}). \tag{51}$$

Proof First, using the terms "edge" and "1-cell" interchangeably, we define two types of edges on $\mathrm{Cyl}_{h,p}$ and $\mathrm{Möb}_{h,p}$. An edge $x \in \mathcal{X}_1$ is called an *interior edge* if $|\mathcal{B}_{1,2}(x)| > 1$, otherwise it's called a *boundary edge*. We denote the boundary edge graph (node set are the nodes contained in the boundary edges and edge set is the boundary edges themselves) of a CC $\mathcal{X}$ by $\partial\mathcal{X}$. We construct the model M by first using a $\mathcal{B}_{1,2}$ aggregation to get the cochain $\mathbf{h}^{(1)} \in \mathcal{C}^{\boldsymbol{e}_1}(\mathcal{X}, \mathbb{R})$

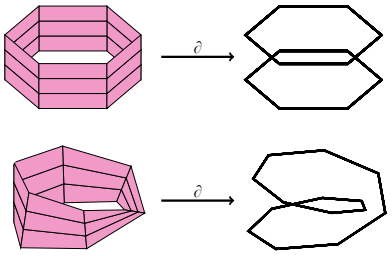

Figure 7: Boundary 1-cells.

$$\mathbf{h}^{(1)}(x) = \deg_{\mathcal{B}_{1,2}}(x). \tag{52}$$

Next, we use an equivariant linear update to construct the multi-cellular cochain $\mathbf{h}^{(2)} \in \mathcal{C}^{2\boldsymbol{e}_1}(\mathcal{X}, \mathbb{R}^2)$ defined by:

$$\mathbf{h}^{(2)}_{x_1,x_2} = \deg_{\mathcal{B}_{1,2}}(x_1) \,\|\, \deg_{\mathcal{B}_{1,2}}(x_2), \tag{53}$$

where, $\|$ denotes concatenation. Recall that the matrix form of $\mathrm{co}\mathcal{A}_{1,0}$ defines a cochain $\mathbf{h}_{\mathrm{co}\mathcal{A}_{1,0}} \in \mathcal{C}^{2\boldsymbol{e}_1}$ which can be used as input to SMCN. Using $\mathbf{h}_{co\mathcal{A}_{1,0}}$ can now construct

$$\mathbf{h}^{(3)}_{x_1,x_2} = (\mathbf{h}_{\mathrm{co}\mathcal{A}_{1,0}})_{x_1,x_2} \,\|\, \deg_{\mathcal{B}_{1,2}}(x_1) \,\|\, \deg_{\mathcal{B}_{1,2}}(x_2). \tag{54}$$

Finally, using a stack of equivariant linear layers, we can construct a fourth cochain $\mathbf{h}^{(4)}_{x_1,x_2} = m(\mathbf{h}^{(3)}_{x_1,x_2})$ where m is parameterized by an MLP. We use the Memorization Theorem (Yun et al., 2019), and to construct an MLP MLP that satisfies

$$\mathsf{MLP}(a,b,c) = \begin{cases} 1 & a = b = c = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{55}$$

$\mathbf{h}^{(4)}$ represents the adjacency matrix of $\partial\mathcal{X}$. $\partial\mathrm{Cyl}_{h,p}$ is composed of two disconnected cycles of length p; $\partial\mathrm{Möb}_{h,p}$ is composed of a single cycle of length $2p$. These two graphs are distinguishable by subgraph architectures like GNN-SSWL+. Thus, using Lemma D.3 and Remark D.4 we can continue the construction of M so that it will be able to differentiate between $\mathrm{Cyl}_{h,p}$ and $\mathrm{Möb}_{h,p}$. ■

Homology. We first show that SMCN is able to count the number of connected components i.e. the 0-th homology.

Proposition D.7 (SMCN can count connected components) *Let $\mathcal{X}, \mathcal{X}'$ be CCs. If the number of connected components of the augmented Hasse graphs $\mathcal{H}_{\mathcal{A}_{r_1,r_2}}$ and $\mathcal{H}'_{\mathcal{A}_{r_1,r_2}}$ is different for some $r_1, r_2 \in \mathbb{N}$ then there exists an SMCN model M such that $M(\mathcal{X}) \neq M(\mathcal{X}')$.*

Proof Using Lemma D.3 and Remark D.4, it suffices to show that GNN-SSWL+ can distinguish graphs with different numbers of connected components. It was shown in Zhang et al. (2023) that adding an additional aggregation of the form $\mathbf{h}_{u,v} \mapsto \sum_{v' \in V} \mathbf{h}_{u,v'}$ to GNN-SSWL+ does not effect its capacity to separate graphs. Therefore, for the remainder of this proof, we include this aggregation in GNN-SSWL+. As previously demonstrated, GNN-SSWL+ can compute a feature vector of the form:

$$\mathbf{h}_{u,v}^{(t)} = d_{\mathcal{G}}(u, v) \quad \text{for } u, v \in \mathcal{V}. \tag{56}$$

If u and v are in different connected components, their distance is encoded as -1. Let $g_1 : [-1, |\mathcal{V}|] \to \mathbb{R}$ be a continuous function such that:

$$g_1(x) = \begin{cases} 0 & \text{if } x = -1, \\ 1 & \text{if } x \geqslant -\frac{1}{2}. \end{cases} \tag{57}$$

We can approximate g_1 using an MLP and apply it to $\mathbf{h}_{u,v}^{(t)}$ to obtain:

$$\mathbf{h}_{u,v}^{(t+1)} = \begin{cases} 0 & \text{if } v \notin \mathcal{G}_u, \\ 1 & \text{if } v \in \mathcal{G}_u. \end{cases} \tag{58}$$

We now take $\mathbf{h}_{u,v}^{(t+2)} = \sum_{v' \in V} \mathbf{h}_{u,v'}^{(t+1)}$, to get

$$\mathbf{h}_{u,v}^{(t+2)} = |\mathcal{G}_u|. \tag{59}$$

Define $g_2 : [1, |\mathcal{V}|] \to \mathbb{R}$ to be

$$g_2(x) = \frac{1}{x}. \tag{60}$$

We can approximate g_2 using an MLP and apply it to $\mathbf{h}_{u,v}^{(t+2)}$ to obtain

$$\mathbf{h}_{u,v}^{(t+3)} = \frac{1}{|\mathcal{G}_u|}. \tag{61}$$

We can now perform a readout of the form $h_{\text{out}} = \sum_{u,v \in \mathcal{V}} \mathbf{h}_{u,v}^{(t+3)}$ to obtain

$$h_{\text{out}} = \sum_{u,v \in \mathcal{V}} \frac{1}{|\mathcal{G}_u|} = \sum_{\mathcal{G}^* \in C(\mathcal{G})} \sum_{u \in \mathcal{G}^*} \frac{|\mathcal{V}|}{|\mathcal{G}^*|} = \sum_{\mathcal{G}^* \in C(\mathcal{G})} |\mathcal{V}| = |\mathcal{V}||C(\mathcal{G})|. \tag{62}$$

Now for let $\mathcal{G}, \mathcal{G}'$ be a pair of graphs with a different number of connected components. If these two graphs have a different number of nodes, they can be easily distinguished by GNN-SSWL+. On the other hand, if they have the same number of nodes they can

be distinguished by GNN-SSWL+ based on Equation 62. Thus, we have shown that an augmented GNN-SSWL+ model can distinguish between $\mathcal{H}_{\mathcal{A}_{r_1,r_2}}$ and $\mathcal{H}'_{\mathcal{A}_{r_1,r_2}}$, and therefore, there exists an SMCN model M that can separate $\mathcal{X}$ and $\mathcal{X}'$. ■

Since the 0-th homology satisfies $H_0(\mathcal{X}) = \mathbb{Z}^{|C(\mathcal{X})|}$ we additionally get the following corollary.

Corollary D.8 (SMCN can compute the 0-th homology) *If $\mathcal{X}, \mathcal{X}'$ are CCs such that the 0-th homology group of their induced topological spaces are different, then there exists an SMCN model M such that $M(\mathcal{X}) \neq M(\mathcal{X}^*)$.*

Exploring SMCN's capacity to differentiate between CCs based on their higher-order homology groups is left for future work. As a first step, we show that SMCN can successfully separate a natural family of CCs — two-dimensional surfaces embeddable in $\mathbb{R}^3$ — based on any homology group/Betti number.

Proposition D.9 (SMCN can compute homology groups of surfaces) *Let $\mathcal{X}, \mathcal{X}'$ be two cellular complexes that are realizations of 2-dimensional manifolds (with or without boundary) $\mathcal{M}, \mathcal{M}'$ which are embeddable in $\mathbb{R}^3$. If $\exists r \in \mathbb{N}$ such that $H_r(\mathcal{M}) \neq H_r(\mathcal{M}')$ then there is an SMCN model M such that $M(\mathcal{X}) \neq M(\mathcal{X}')$.*

Proof First, since $\mathcal{M}$ is 2-dimensional, the only non-trivial homology groups it may have are of order $0 \leqslant r \leqslant 2$. The 0-th homology group of $\mathcal{M}$, is of the form $H_0(\mathcal{M}) = \mathbb{Z}^{k_0}$ where k_0 is the number of $\mathcal{M}$'s connected components. Furthermore, as each connected component of $\mathcal{M}$ is a connected 2-dimensional manifold with boundary that can be embedded in $\mathbb{R}^3$, it must either be orientable or have a non-empty boundary. If such a component is orientable, then by the Poincaré duality, its 2-nd homology group is $\mathbb{Z}$. On the other hand, if it has a boundary, it is homotopic to a 1-dimensional cellular complex, and thus its 2-nd homology group is trivial. Therefore, $H_2(\mathcal{M}) = \mathbb{Z}^{k_2}$, where k_2 is the number of connected components of $\mathcal{M}$ with no boundary. Finally, since $\mathcal{M}$ is embeddable in $\mathbb{R}^3$, its 1-st homology groups is $H_1(\mathcal{M}) = \mathbb{Z}^{k_1}$ for some integer k_1. The Euler characteristic of the manifold $\mathcal{M}$ defined by $k_0 - k_1 + k_2$ can be computed using the number of cells of $\mathcal{X}$ using the following formula:

$$\chi(\mathcal{M}) = k_2 - k_1 + k_0 = |\mathcal{X}_2| - |\mathcal{X}_1| + |\mathcal{X}_0|. \tag{63}$$

Thus in order to separate $\mathcal{X}$ from $\mathcal{X}'$ we need to be able to construct a SMCN model which is able to separate CCs which are different in either one of the following three quantities:

1. The Euler characteristic.
2. The number of connected components.
3. The number of connected components with no boundary.

Computing the Euler characteristic is be computed by standard HOMP updates, as it is a function of the sizes of $\mathcal{X}_0, \mathcal{X}_1$, and $\mathcal{X}_2$. For the second quantity, we have seen SMCN models can separate CCs with a different number of connected components in Proposition D.7. As for the third quantity, a connected component of $\mathcal{X}$ has a boundary if and only if it contains 1-cells whose degree with respect to the neighborhood function $\mathcal{B}_{1,2}$ is exactly

1. We can use a stack of standard HOMP layers to compute the 1-cells features

$$\mathbf{h}_x = \begin{cases} 1 & x \text{ is in the same connected component as a a boundary edge} \\ 0 & \text{otherwise.} \end{cases} \tag{64}$$

Using $\mathbf{h}$, we can adjust the proof of Proposition D.7 by summing in Equation 62 only over 1-cells for which $\mathbf{h}_x = 0$, resulting in the the number of connected components of $\mathcal{X}$ with no boundary. This shows that SMCN can distinguish CCs based on *either one* of the aforementioned three properties, concluding the proof. ■

Appendix E. Experimental Details

The first set of experiments were designed to benchmark TDL architecture on their ability to learn topological and metric properties of objects. First, we design a synthetic topological expressivity benchmark where models are tasked with distinguishing pairs of cellular complexes representing disjoint unions of tori. We then task models with predicting topological and metric properties of cellular complexes constructed by applying cyclic lifting (Bodnar et al., 2021a) to molecular graphs.

In the second set of experiments we test SMCN on a variety of graph classification and regression tasks, demonstrating performance increase over *both* higher-order message-passing architectures (e.g. CW-networks) and expressive graph methods (e.g. subgraph neural networks).

E.1. Topological and Metric Properties

Torus dataset. To construct the torus dataset we first select three parameters: m which specifies the number of nodes in the smallest CC in the dataset, M which specifies the number of nodes in the largest CC, and n, which specifies the maximum number of connected components in any CC within the dataset. The dataset is then constructed by iterating over all possible choices for the the number of nodes and connected components, generating all possible disjoint unions of 2-dimensional tori with the specified parameters. We then select all the pairs that have the same size (number of nodes). As mentioned in the paper, each such pair is indistinguishable by HOMP despite differing in basic metric/topological properties: they either have distinct homology , or they differ in the diameters of some of the components. In our experiments, we selected $m = 18$ (the smallest size that admits indistinguishable pairs), $M = 40$, and $n = 3$, resulting in 223 pairs.

To evaluate the ability of both HOMP and SMCN to distinguish between each pair, we employ the training and evaluation schemes proposed in Wang and Zhang (2024). We require a statistically significant difference between the outputs of the model on each the CC in the pair. Our experiments show that while HOMP is unable to distinguish any of the pairs, SMCN is able to distinguish all of them as depicted in Table 3.

Model	Pairs distinguished (↑)	Accuracy (↑)
CIN	0	0%
SMCN	223	100%

Table 3: Pair distinguishing accuracy on the torus dataset.

Lifted ZINC cross-diameter. We construct a CC dataset by adding cycles of length $\leqslant 18$ as 2-cells to graphs taken from the ZINC dataset (Sterling and Irwin, 2015). We remove edge and node features, and predict the (0, 1, 2) cross diameter, computed by

$$\max_{\substack{x\in\mathcal{X}_0,\ x'\in y\\ y\in\mathcal{X}_2}} \min d_{\mathcal{A}_{0,1}}(x,x') \tag{65}$$

where $d(x,x')$ is the shortest path distance. Training targets are normalized to have mean 0 and standard deviation 1. At test time we evaluate the accuracy of predicting the cross-diameter value (there are 18 possible values). We compare 3 architectures: CIN (Bodnar et al., 2021a), custom HOMP, and SMCN. Custom HOMP is a CIN variant optimized for cross-diameter computation and SMCN is implemented using 0-2 subcomplex layers. Results are presented in Table 2. Results are reported as mean $\pm$ std across 5 seeds.

Lifted ZINC Betti numbers. For the second topological property prediction task we tested our model's ability to learn to predict the 2-nd order Betti numbers- the ranks of the 2-nd homology group. To this end we constructed our benchmark dataset the following way: We started with the ZINC-FULL datasets (containing $250k$ molecular graphs), lifting all graphs to CCs as in the cross-diameter task. We then computed the 2-nd Betti number for each of the lifted graphs and randomly selected 850 samples from each of the 6 most common values (which were $0, 1, 2, 3, 4$ and 6), resulting in a balanced dataset of size 5,100. We used a 60%, 20%, 20% random split for training, validation, and test sets. We normalized training targets to have mean 0 and standard deviation 1 and train using MSE loss. At test time we evaluated the accuracy of predicting the 2-nd Betti number. Results are presented in Table 2. Results are reported as mean $\pm$ std across 5 seeds.

The results of the above three synthetic experiments further solidify SMCN's superior capability in capturing topological properties of CCs compared to existing HOMP architectures. This is demonstrated both for synthetically generated inputs as well as inputs constructed by applying common lifting procedures on real world graphs. This reinforces our theoretical findings which also suggest SMCN is more capable in learning both cross diameters and homology groups presented in Appendix D.

E.2. Graph Benchmarks

ZINC (Sterling and Irwin, 2015; Dwivedi et al., 2023). The ZINC dataset comprises 12,000 molecular graphs, extracted from the ZINC database, which is a collection of commercially available chemical compounds. These molecular graphs vary in size, ranging from 9 to 37 nodes each. In these graphs, nodes correspond to heavy atoms, encompassing 28 distinct atom types. Edges in the graphs represent chemical bonds, with three possible bond types. We perform regression on the constrained solubility (logP) of the molecules. The dataset is pre-partitioned into training, validation, and test sets, containing 10,000, 1,000, and 1,000 molecular graphs, respectively.

MOLHIV and MOLESOL (Hu et al., 2020). MOLHIV and MOLESOL are molecular property prediction datasets, adapted by the Open Graph Benchmark (OGB) from MoleculeNet. These datasets employ a unified featurization for nodes (atoms) and edges (bonds), encapsulating various chemophysical properties. The task in MOLHIV is to predict

the capacity of compounds to inhibit HIV replication. The task in MOLESOL is regression on water solubility (log solubility in mols per liter) for common organic small molecules.