
GEFL: Extended Filtration Learning for Graph Classification

Anonymous Author(s)

Anonymous Affiliation

Anonymous Email

Abstract

Extended persistence is a technique from topological data analysis to obtain global multiscale topological information from a graph. This includes information about connected components and cycles that are captured by the so-called persistence barcodes. We introduce extended persistence into a supervised learning framework for graph classification. Global topological information, in the form of a [barcode with four different types of bars](#) and their explicit cycle representatives, is combined into the model by the readout function which is computed by extended persistence. The entire model is end-to-end differentiable. We use a link-cut tree data structure and parallelism to lower the complexity of computing extended persistence, obtaining a speedup of more than 60x over the state-of-the-art [for extended persistence computation](#). This makes extended persistence feasible for machine learning. We show that, under certain conditions, extended persistence surpasses both the WL[1] graph isomorphism test and 0-dimensional barcodes in terms of expressivity because it adds more global (topological) information. In particular, arbitrarily long cycles can be represented, which is difficult for finite receptive field message passing graph neural networks. Furthermore, we show the effectiveness of our method on real world datasets compared to many existing recent graph representation learning methods.¹

1 Introduction

Graph classification is an important task in machine learning. Applications range from classifying social networks to chemical compounds. These applications require global as well as local topological information of a graph to achieve high performance. Message passing graph neural networks (GNNs) are an effective and popular method to achieve this task.

These existing methods crucially lack quantifiable information about the relative prominence of cycles and connected component to make predictions. Extended persistence is an unsupervised technique from topological data analysis that provides this information through a generalization of hierarchical clustering on graphs. It obtains both 1- and 0-dimensional multiscale global homological information.

Existing end-to-end filtration learning methods [1, 2] that use persistent homology do not compute extended persistence because of its high computational cost at scale. [A general matrix reduction approach \[3\] has time complexity of \$O\(\(n+m\)^\omega\)\$ for graphs with \$n\$ nodes and \$m\$ edges where \$\omega\$ is the exponent for matrix multiplication.](#) We address this by improving upon the work of [4] and introducing a link-cut tree data structure and a parallelism for computation. This allows for $O(\log n)$ update and query operations on a spanning forest with n nodes.

We consider the expressiveness of our model in terms of extended persistence barcodes and the cycle representatives. We characterize the barcodes in terms of size, what they measure, and their expressivity in comparison to WL[1] [2]. We show that it is possible to find a filtration where one of its cycle's length can be measured as well as a filtration where the size of each connected component can be measured. We also consider the case of barcodes when no learning of the filtration occurs.

¹code to be released if the paper is accepted, <https://anonymous.4open.science/r/GraphExtendedFiltrationLearning-34CB>




Figure 1: Lower and upper filtrations for extended persistence and the resulting barcode for a graph. The green bar comes from a pairing of a green edge with a vertex in the lower filtration. Similarly the blue bar in the upper filtration comes from a vertex-edge pairing in the upper filtration. The two dark blue bars count connected components and come from pairs of two vertices. The two red bars count cycles and come from pairs of edges. Both \mathcal{B}_0^{ext} and \mathcal{B}_1^{ext} bars cross from the lower filtration to the upper filtration. The multiset of bars forms the barcode. **Cycle reps. are shown in both filtrations.**

41 We consider several simple examples where our model can perfectly distinguish two classes of
 42 graphs that no GNN with expressivity at most that of WL[1] (henceforth called WL[1] bounded
 43 GNN) can. Furthermore, we present a case where experimentally 0-dimensional standard persistence
 44 [2, 5], the only kind of persistence considered in learning persistence so far, are insufficient for graph
 45 classification.

46 Our contributions are as follows:

- 47 1. We introduce extended persistence and its cycle representatives into the supervised learning
 48 framework in an end-to-end differentiable manner, for graph classification.
- 49 2. For a graph with m edges and n vertices, we introduce the link-cut tree data structure into the
 50 computation of extended persistence, resulting in an $O(m \log n)$ depth and $O(mn)$ work parallel
 51 algorithm, achieving more than 60x speedup over the state-of-the-art for [extended persistence](#)
 52 [computation](#), making extended persistence amenable for machine learning tasks.
- 53 3. We analyze conditions and examples upon which extended persistence can surpass the WL[1]
 54 graph isomorphism test [6] and 0-dimensional standard persistence and characterize what extended
 55 persistence can measure from additional topological information.
- 56 4. We perform experiments to demonstrate the feasibility of our approach against standard baseline
 57 models and datasets as well as an ablation study on the readout function for a learned filtration.

58 2 Background

59 2.1 Computational Topology for Graphs

60 Let $G = (V, E)$ be a graph where V is the set of vertices and $E \subset V \times V$ is the set of edges. Let
 61 $n = |V|$ and $m = |E|$ be the number of nodes and edges of G , respectively. Graphs in our case are
 62 undirected and simple, containing at most a single edge between any two vertices. Define a filtration
 63 function $F : G \rightarrow \mathbb{R}$ where F has a value in \mathbb{R} on each vertex and edge, denoted by $F(u)$ or $F(e)$ for
 64 $u \in V$ or $e \in E$. Given such a graph $G = (V, E)$, we define the λ -sublevel graph as $G_\lambda = (V_\lambda, E_\lambda)$
 65 w.r.t. F and a $\lambda \in \mathbb{R}$ where $V_\lambda = \{v \in V : F(v) \leq \lambda\}$ and $E_\lambda = \{e \in E : F(e) \leq \lambda\}$. Sublevel
 66 graphs of G are subgraphs of G . If we change λ from $-\infty$ to $+\infty$ we obtain an increasing sequence
 67 of sublevel graphs $\{G_\lambda\}_{\lambda \in \mathbb{R}}$ which we call a sublevel set filtration. Such a filtration can always be
 68 converted into a sequence of subgraphs of G : $\emptyset = G_0 \subset G_1 \subset \dots \subset G_{n+m} = G$ (See [7, Page
 69 102]) s.t. $\sigma_i = G_{i+1} \setminus G_i$ is a single edge or vertex and $F_i := F(\sigma_i)$. The sequence of vertices
 70 and edges $\sigma_0, \sigma_1, \dots, \sigma_{n+m-1}$ thus obtained is called the index filtration. Define a vertex-induced
 71 lower filtration for a vertex function $f_G : V \rightarrow \mathbb{R}$ as an index filtration where a vertex v has a value
 72 $F(v) := f_G(v)$ and any edge (u, v) has the value $F(u, v) := \max(F(u), F(v))$ and $F_i \leq F_{i+1}$.
 73 Similarly define an upper filtration for f_G as an index filtration where $F(v) := f_G(v)$ and the edge
 74 (u, v) has value $F(u, v) := \min(f_G(u), f_G(v))$ and $F_i \geq F_{i+1}$.

75 **Persistent homology (PH)** tracks changes in homological features of a topological space as the
 76 sublevel set for a given function grows; see books [7, 8]. For graphs, these features are given by
 77 evolution of components and cycles over the intervals determined by pairs of vertices and edges.
 78 A vertex $v_i = G_{i+1} \setminus G_i$ begins a connected component (CC) signalling a birth at filtration value
 79 $F(v_i)$ in zeroth homology group H_0 . An edge $e_j = G_{j+1} \setminus G_j$ may join two components signalling

80 a death of a class in H_0 at filtration value $F(e_j)$, or it may create a cycle signalling a birth in the 1st
 81 homology group H_1 at filtration value $F(e_j)$. When a death occurs in H_0 by an edge e_j , the youngest
 82 of the two components being merged is said to die giving a birth-death pair $(b, d) = (F(v_i), F(e_j))$
 83 if the dying component was created by vertex v_i . For cycles, there is no death and thus they have
 84 death at ∞ . The multiset of birth death pairs $\mathcal{B} = \{(b, d)\}$ given by the persistent homology
 85 is called the barcode. Each pair (b, d) provides a closed-open interval $[b, d)$, which is called a bar.
 86 The persistence of each bar $[b, d)$ in a barcode is defined as $|d - b|$. Notice that, both in 0- and
 87 1-dimensional persistence, some bars may have infinite persistence since some components (H_0
 88 features) and cycles (H_1 features) never die, equivalently, have death at ∞ .

89 **Extended persistence**(PH_{ext}) takes an extended filtration F_{f_G} as input, which is obtained by
 90 concatenating lower filtration of the graph G and an upper filtration of the coned space of G induced
 91 by a vertex filtration function f_G . Concatenation here simply means concatenating two index filtration
 92 sequences. More specifically, let α be an additional vertex for the graph G . Define an extended
 93 function $f_{G \cup \{\alpha\}}$ whose value is equal to f_G on all vertices except α on which it has a value larger
 94 than any other vertices. The cone of a vertex u is given by the edge (α, u) and the cone of an edge
 95 (u, v) is given by the triangle (α, u, v) . As a result, in extended persistence all 0- and 1-dimensional
 96 features die (bars are finite; see [3] for details). Four different persistence pairings or bars result from
 97 PH_{ext} . The barcode $\mathcal{B}_0^{\text{low}}$ results from the vertex-edge pairs within the lower filtration, the barcode
 98 $\mathcal{B}_0^{\text{up}}$ results from the vertex-edge pairs within the upper filtration, the barcode $\mathcal{B}_0^{\text{ext}}$ results from the
 99 vertex-vertex pairs that represent the persistence of connected components born in the lower filtration
 100 and die in the upper filtration, and the barcode $\mathcal{B}_1^{\text{ext}}$ results from edge-edge pairs that represent the
 101 persistence of cycles that are born in the lower filtration and die in the upper filtration. The barcodes
 102 $\mathcal{B}_0^{\text{low}}$, $\mathcal{B}_0^{\text{up}}$, and $\mathcal{B}_0^{\text{ext}}$ represent persistence in the 0th homology H_0 . The barcode $\mathcal{B}_1^{\text{ext}}$ represents
 103 persistence in the 1st homology H_1 . In the TDA literature, $\mathcal{B}_0^{\text{low}}$, $\mathcal{B}_0^{\text{up}}$, $\mathcal{B}_0^{\text{ext}}$, and $\mathcal{B}_1^{\text{ext}}$ also go by the
 104 names of Ord_0 , Rel_1 , Ext_0 , Ext_1 respectively.

105 See Figure 1 for an illustration of the filtration and barcode one obtains for a simple graph with
 106 vertices taking on values from 0...6 denoted by the variable t . In particular, at each t , we have the
 107 filtration subgraph G_t of all vertices and edges of filtration function value less than or equal to t .
 108 Each line indicates the values 0...6 from the bottom to top. Repetition in the bar endpoints across all
 109 bars which appear on the right of Figure 1 is highly likely in general due to the fact that there are
 110 only $O(n)$ filtration values but $O(m)$ possible bars.

111 2.2 Message Passing Graph Neural Networks (MPGNN)

112 A message passing GNN (MPGNN) convolutional layer takes a vertex embedding \mathbf{h}_u in a finite
 113 dimensional Euclidean space and an adjacency matrix A_G as input and outputs a vertex embedding
 114 \mathbf{h}'_u for some $u \in V$. The k th layer is defined generally as

$$\mathbf{h}_u^{k+1} \leftarrow \text{AGG}(\{\text{MSG}(\mathbf{h}_v^k) | v \in N_{A_G}(u)\}, \mathbf{h}_u^k), u \in V$$

115 where $N_{A_G}(u)$ is the neighborhood of u . The functions MSG and AGG have different implementa-
 116 tions and depend on the type of GNN.

117 Since there should not be a canonical ordering to the nodes of a GNN in graph classification, a GNN
 118 for graph classification should be permutation invariant with respect to node indices. To achieve
 119 permutation invariance [9], as well as achieve a global view of the graph, there must exist a readout
 120 function or pooling layer in a GNN. The readout function is crucial to achieving power for graph
 121 classification. With a sufficiently powerful readout function, a simple 1-layer MPGNN with $O(\Delta)$
 122 number of attributes [10] can compute any Turing computable function, Δ being the max degree of
 123 the graph. Examples of simple readout functions include aggregating the node embeddings, or taking
 124 the element-wise maximum of node embeddings [11]. See Section 3 for various message passing
 125 GNNs and readout functions from the literature.

126 3 Related Work

127 Graph Neural Networks (GNN)s have achieved state of the art performance on graph classification
 128 tasks in recent years. For a comprehensive introduction to GNNs, see the survey [12]. In terms of the
 129 Weisfeler Lehman (WL) hierarchy, there has been much success and efficiency in GNNs [11, 13, 14]
 130 bounded by the WL[1] [15] graph isomorphism test. In recent years, the WL[1] bound has been

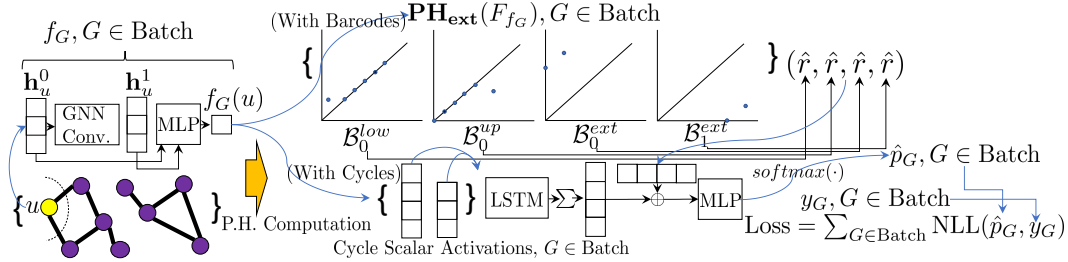


Figure 2: The extended persistence architecture (bars+cycles) for graph representation learning. The negative log likelihood (NLL) loss is used for supervised classification. The yellow arrow denotes extended persistence computation, which can compute both barcodes and cycle representatives.

131 broken by heterogenous message passing [16], high order GNNs [17], and put into the framework of
 132 cellular message passing networks [18]. Furthermore, a sampling based pooling layer is designed in
 133 [19]. It has no theoretical guarantees and its code is not publicly available for comparison. Other
 134 readout functions include [20], [21] [22]. For a full survey on global pooling, see [23].

135 Topological Data Analysis (TDA) based methods [2, 5, 24–28] that use learning with persistent
 136 homology have achieved favorable performance with many conventional GNNs in recent years. All
 137 existing methods have been based on 0-dimensional standard persistent homology on separated lower
 138 and upper filtrations [5]. We sidestep these known limitations by introducing extended persistence
 139 into supervised learning while keeping computation efficient.

140 A TDA inspired cycle representation learning method in [29] learns the task of knowledge graph
 141 completion. It keeps track of cycle bases from shortest path trees and has a $O(|V| \cdot |E| \cdot k)$, k
 142 a constant, computational complexity per graph. This high computational cost is addressed in our
 143 method by a more efficient algorithm for keeping track of a cycle basis.

144 On the computational side, fast methods to compute higher dimensional PH using GPUs, a necessity
 145 for modern deep learning, have been introduced in [30]. In [27, 31] neural networks have been
 146 shown to successfully approximate the persistence diagrams with learning based approach. However,
 147 differentiability and parallel extended persistence computation has not been implemented. Given
 148 the expected future use of extended persistence in graph data, a parallel differentiable extended
 149 persistence algorithm is an advance on its own.

150 4 Method

151 Our method as illustrated in Figure 2 introduces extended persistence as the readout function for
 152 graph classification. In our method, an upper and lower filtration, represented by a filtration function,
 153 coincides with a set of scalar vertex representations from standard message passing GNNs. This
 154 filtration function is thus learnable by MPGNN convolutional layers. Learning filtrations was
 155 originally introduced in [5] with standard persistence. As we show in Section 6 and Section 5
 156 arbitrary cycle lengths are hard to distinguish by both standard GNN readout functions [32] as well as
 157 standard persistence due to the lack of explicitly tracking paths or cycles. Extended persistence, on the
 158 other hand, explicitly computes learned displacements on cycles of some cycle basis as determined
 159 by the filtration function as well as explicit cycle representatives.

160 We represent the map from graphs to learnable filtrations by any message passing GNN layer such as
 161 GIN, GCN or GraphSAGE followed by a multi layer perceptron (MLP) as a Jumping Knowledge
 162 (JK) [33] layer. The JK layer with concatenation is used since we want to preserve the higher
 163 frequencies from the earlier layers [34]. Our experiments demonstrate that fewer MPGNN layers
 164 perform better than more MPGNN layers. This prevents oversmoothing [35, 36], which is exacerbated
 165 by the necessity of scalar representations.

166 The readout function, the function that consolidates a filtration into a global graph representation,
 167 is determined by computing four types of bars for the extended persistence on the concatenation of
 168 the lower and upper filtrations followed by compositions with four rational hat functions \hat{r} as used
 169 in [1, 2, 5]. To each of the four types of bars in barcode \mathcal{B} , we apply the hat function \hat{r} to obtain a

170 k -dimensional vector. The function \hat{r} is defined as:

$$\hat{r}(\mathcal{B}) := \left\{ \sum_{\mathbf{p} \in \mathcal{B}} \frac{1}{1 + |\mathbf{p} - \mathbf{c}_i|_1} - \frac{1}{1 + ||r_i| - |\mathbf{p} - \mathbf{c}_i|_1|} \right\}_{i=1}^k \quad (1)$$

171 where $r_i \in \mathbb{R}$ and $\mathbf{c}_i \in \mathbb{R}^2$ are learnable parameters. The intent of Equation 1 is to have controlled
 172 gradients. It is derived from a monotonic function, see [1]. This representation is then passed through
 173 MLP layers followed by a softmax to obtain prediction probability vector \hat{p}_G for each graph G . The
 174 negative log likelihood loss from standard graph classification is then used on these vectors \hat{p}_G .

175 **If the filtration values on the nodes and edges are distinct**, the extended persistence barcode repre-
 176 sentation is permutation invariant **with respect to node indices**. **Isomorphic graphs with permuted**
 177 **indices and an index filtration with distinct filtration values will have a unique sorted index filtration**.
 178 **Node filtration values are usually distinct since computed floating points rarely coincide**. However to
 179 **break ties and eliminate any dependence on node indices for edges**, implement edge filtration values
 180 **for lower filtration as $F(u, v) = \max(F(u), F(v)) + \epsilon \cdot \min(F(u), F(v))$ and for upper filtration as**
 181 **$F(u, v) = \min(F(u), F(v)) + \epsilon \cdot \max(F(u), F(v))$, ϵ very small.**

182 **Cycle Representatives:** Because computing extended persistence results in computing a cycle basis,
 183 we can explicitly store the cycle representatives, or sequences of filtration scalars, along with the
 184 barcode on graph data. This slightly improves the performance in practice and guarantees cycle
 185 length classification for arbitrary lengths. After the cycle representatives are stored, we pass them
 186 through a bidirectional LSTM then aggregate these LSTM representation per graph and then sum this
 187 graph representation by cycles with the vectorization of the graph barcode by the rational hat function
 188 of Equation 1, see Figure 2. The aggregation of the cycle representations is permutation invariant
 189 due to the composition of aggregations [9]. In particular, the sum of the barcode vectorization and
 190 the mean of cycle representatives, our method’s graph representation, must be permutation invariant.
 191 What makes keeping track of cycle representatives unique to standard message passing GNNs is that
 192 a finite receptive field message passing GNN would never be able to obtain such cycle representations
 193 and certainly not from a well formed cycle basis.

194 4.1 Efficient Computation of Extended Persistence

195 The computation for extended persistence can be reduced to applying a matrix reduction algorithm
 196 to a coned matrix as detailed in [8]. In [4], this computation was found to be equivalent to a graph
 197 algorithm, which we improve upon.

198 4.1.1 Algorithm

199 Our algorithm is as follows and written in Algorithm 1. We perform the 0-dimensional persistence
 200 algorithm, PH_0 , using the union find data structure in $O(m \log n)$ time and $O(n)$ memory for the
 201 upper and lower filtrations **in lines 1 and 2**. **See the Appendix Section D.1 for a description of this**
 202 **algorithm**. These two lines generate the vertex-edge pairs for $\mathcal{B}_0^{\text{low}}$ and $\mathcal{B}_0^{\text{up}}$. We then measure the
 203 minimum lower filtration value and maximum upper filtration value of each vertex in the union-find
 204 data structure found from the PH_0 algorithm as in **lines 3 and 4 using the roots of the union-find data**
 205 **structure U_{up} formed by the algorithm**. These produce the vertex-vertex pairs in $\mathcal{B}_0^{\text{ext}}$.

206 For computing edge-edge pairs in $\mathcal{B}_1^{\text{ext}}$ with cycle representatives, we implement the algorithm in [4]
 207 with a link-cut tree data structure that facilitates deleting and inserting edges in a spanning tree and
 208 employ a parallel algorithm to enumerate the edges in a cycle. **See the Appendix Section D.2 for a**
 209 **more thorough explanation of the link-cut tree implementation and the operations we use on it**. We
 210 collect the max spanning forest T of negative edges, edges that join components, from the upper
 211 filtration by repeatedly applying the link operation $n - 1$ times **in lines 6-8 in decreasing order of F_{up}**
 212 **values** and sort the list of the remaining positive edges, which create cycles **in line 9**. Then, for each
 213 positive edge $e = (u, v)$, in order of the upper filtration **(line 10)**, we find the least common ancestor
 214 (lca) of u and v in the spanning forest T we are maintaining as in **line 11**. Next, we apply the parallel
 215 primitive [37] of *list ranking* twice, once on the path u to lca and the other on the path v to lca in
 216 **line 12**. List ranking allows a list to populate an array in parallel in logarithmic time. The tensor
 217 concatenation of the two arrays is appended to a list of cycle representatives as in **line 13**. This is so
 218 that the cycle maintains order from u to v . We then apply an $\text{ARGMAXREDUCECYCLE}(T, u, v, \text{lca})$
 219 which finds the edge having a maximum filtration value on it over the cycle formed by u, v and lca .

Algorithm 1 Efficient Computation of PH_{ext}

Input: $G = (V, E)$, F_{low} : lower filtration function, F_{up} : upper filtration function
Output: $\mathcal{B}_0^{\text{low}}, \mathcal{B}_0^{\text{up}}, \mathcal{B}_0^{\text{ext}}, \mathcal{B}_1^{\text{ext}}, \mathcal{C}$: cycle reps.

- 1: $\mathcal{B}_0^{\text{low}}, E_{\text{pos}}^{\text{low}}, E_{\text{neg}}^{\text{low}}, U_{\text{low}} \leftarrow \text{PH}_0(G, F_{\text{low}}, \text{lower})$
- 2: $\mathcal{B}_0^{\text{up}}, E_{\text{pos}}^{\text{up}}, E_{\text{neg}}^{\text{up}}, U_{\text{up}} \leftarrow \text{PH}_0(G, F_{\text{up}}, \text{upper})$
- 3: $\text{roots} \leftarrow \{\text{GET_UNION-FIND_ROOTS}(U_{\text{up}}, v), v \in V\}$
- 4: $\mathcal{B}_0^{\text{ext}} \leftarrow \{\min(\text{roots}[v]), \max(\text{roots}[v]), v \in V\}$
- 5: $\mathbf{T} \leftarrow \{\}$ empty link-cut tree; $\mathcal{B}_1^{\text{ext}} \leftarrow \{\{\}\}$; $\mathcal{C} \leftarrow \{\}$ empty list of cycle representatives
/ $E_{\text{neg}}^{\text{up}}$ is sorted by PH_0 in decreasing order of F_{up} values (desc. filtr. values)*/*
- 6: **for** $e = (u, v) \in E_{\text{neg}}^{\text{up}}$ **do**
- 7: $\mathbf{T} \leftarrow \text{LINK}(\mathbf{T}, e, \{w\})$ */* $w \notin \mathbf{T}, w = u$ or v */*
- 8: **end for**
- 9: */* $E_{\text{pos}}^{\text{up}}$ is sorted by PH_0 with respect to F_{up} (descending filtration values)*/*
- 10: **for** $e = (u, v) \in E_{\text{pos}}^{\text{up}}$ **do**
- 11: $\text{lca} \leftarrow \text{LCA}(\mathbf{T}, u, v)$ (Get the least common ancestor of u and v to form a cycle)
- 12: $P_1 \leftarrow \text{LISTRANK}(\text{PATH}(u, \text{lca})); P_2 \leftarrow \text{LISTRANK}(\text{PATH}(v, \text{lca}))$
- 13: $\mathcal{C} \leftarrow \mathcal{C} \sqcup \{F_{\text{up}}(P_1) \sqcup F_{\text{up}}(\text{Reverse}(P_2))\}$ (Keep track of the scalar activations on the cycle)
- 14: $(u', v') \leftarrow \text{ARGMAXREDUCECYCLE}(\mathbf{T}, u, v, \text{lca})$
- 15: $\mathbf{T}_1, \mathbf{T}_2 \leftarrow \text{CUT}(\mathbf{T}, (u', v'))$; $\mathbf{T} \leftarrow \text{LINK}(\mathbf{T}_1, (u, v), \mathbf{T}_2)$
- 16: $\mathcal{B}_1^{\text{ext}} \leftarrow \mathcal{B}_1^{\text{ext}} \cup \{(F_{\text{low}}(u', v'), F_{\text{up}}(u, v))\}$
- 17: **end for**
- 18: **return** $(\mathcal{B}_0^{\text{low}}, \mathcal{B}_0^{\text{up}}, \mathcal{B}_0^{\text{ext}}, \mathcal{B}_1^{\text{ext}}, \mathcal{C})$

220 We then cut the spanning forest at the edge (u', v') , forming two forests as in [line 15](#). These two
 221 forests are then linked together at (u, v) as in [line 15](#). The bar $(F_{\text{low}}(u', v'), F_{\text{up}}(u, v))$ is now found
 222 and added to the multiset $\mathcal{B}_1^{\text{ext}}$. The final output of the algorithm is [four types of bars](#) and a list of
 223 cycle representatives: $((\mathcal{B}_0^{\text{low}}, \mathcal{B}_0^{\text{up}}, \mathcal{B}_0^{\text{ext}}, \mathcal{B}_1^{\text{ext}}), \mathcal{C})$.

224 4.1.2 Complexity

225 We improve upon the complexity of [4] by obtaining a $O(mn)$ work $O(m \log n)$ depth algorithm
 226 on $O(n)$ processors using $O(n)$ memory. Here m and n are the number of edges and vertices in
 227 the input graph. We introduce two ingredients for lowering the complexity, the first is the link-cut
 228 dynamic connectivity data structure and the second is the parallel primitives of list ranking. The
 229 link-cut tree data structure is a dynamic connectivity data structure that can keep track of the spanning
 230 forest with $O(\log n)$ amortized time for LINK, CUT, PATH, LCA, ARGMAXREDUCE. Furthermore,
 231 list ranking [38] is an $O(\log n)$ depth and $O(n)$ work parallel algorithm on $O(\frac{n}{\log n})$ processors that
 232 determines the distance of each vertex from the start of the path or linked list it is on. In other words,
 233 list ranking turns a linked list into an array in parallel. Sorting can be performed in parallel using
 234 $O(n \log n)$ work and $O(\log n)$ depth.

235 Notice that if we do not keep track of cycle representatives (remove lines 12 and 13 from Algorithm
 236 1), then we have an $O(m \log n)$ time sequential algorithm. The repeated calling of the supporting
 237 operation EXPOSE() dominates the complexity, see Appendix Section D.2.

238 5 Expressivity of Extended Persistence

239 We prove some properties of extended persistence barcodes. We also find a case where extended
 240 persistence with supervised learning can give high performance for graph classification. WL[1]
 241 bounded GNNs, on the other hand, are guaranteed to not perform well. *Certainly all such results also*
 242 *apply for the explicit cycle representatives since the min and max on the scalar activations on the*
 243 *cycle form the corresponding bar.*

244 5.1 Some Properties

245 The following Theorem 5.1 states some properties of extended persistence. This should be compared
 246 with the 0- and 1-dimensional persistence barcodes in the standard persistence. Every vertex and

247 edge is associated with some bar in the standard persistence though they can be both finite or infinite.
 248 However, in extended persistence all bars are finite and we form barcodes from an extended filtration
 249 of $2m + 2n$ edges and vertices instead of the standard $(m + n)$ -lengthed filtration.

250 **Theorem 5.1.** (*Extended Barcode Properties*)

251 $\mathbf{PH}_{\text{ext}}(G)$ produces four *multisets of bars*: $\mathcal{B}_1^{\text{ext}}, \mathcal{B}_0^{\text{ext}}, \mathcal{B}_0^{\text{low}}, \mathcal{B}_0^{\text{up}}$, s.t.

252 $|\mathcal{B}_1^{\text{ext}}| = \dim H_1 = m - n + C,$

253 $|\mathcal{B}_0^{\text{ext}}| = \dim H_0 = C,$

254 $|\mathcal{B}_0^{\text{low}}| = |\mathcal{B}_0^{\text{upper}}| = n - C,$

255 where there are C connected components and $\dim H_k$ is the dimension of the k th homology group
 256 s.t.:

- 257 1. the H_1 bars comes from a cycle basis of G which also constitutes a basis of its fundamental group,
- 258 2. $\dim H_1$ counts the number of chordless cycles when G is outer-planar, and
- 259 3. there exists an injective filtration function where the union of the resulting barcodes is strictly more
 260 expressive than the histogram produced by the WL[1] graph isomorphism test.

261 The barcodes found by extended persistence thus have more degrees of freedom than those obtained
 262 from standard persistence. For example, a cycle is now represented by two filtration values rather than
 263 just one. Furthermore, the persistence $|d - b|$ of a pair $(b, d) \in \mathcal{B}_1^{\text{ext}}$ or $\mathcal{B}_0^{\text{ext}}$ can measure topological
 264 significance of a cycle or a connected component respectively through persistence. Thus, extended
 265 persistence encodes more information than standard persistence. In Theorem 5.1, property 1 says that
 266 extended persistence actually computes pairs of edges of cycles in a cycle basis. A modification of
 267 the extended persistence algorithm could generate all or count certain kinds of important cycles, see
 268 [39]. Property 2 characterizes what extended persistence can count.

269 We makes some observations on the expressivity of \mathbf{PH}_{ext} .

270 **Observation 5.2.** (Cycle Lengths) For any graph G and a cycle $\mathbf{C} \subset G$, there exists an injective
 271 filtration function where \mathbf{PH}_{ext} of that filtration function can measure the number of edges along \mathbf{C} .

272 Such a result cannot hold for learning of the filtration by local message passing from constant node
 273 attributes. Thus, for the challenging 2CYCLE graphs dataset in Section B.2, it is a necessity to use
 274 the cycle representatives \mathcal{C} for each graph to distinguish pairs of cycles of arbitrary length. This
 275 should be compared with Top- K methods, K being a constant hyper parameter such as in [19, 40].
 276 The constant hyper parameter K prevents learning an arbitrarily long cycle length when the node
 277 attributes are all the same. Furthermore, a readout function like SUM is agnostic to graph topology
 278 and also struggles with learning when presented with an arbitrarily long cycle. This struggle for
 279 distinguishing cycles in standard MPGNNs is also reported in [41]. An observation similar to the
 280 previous Observation 5.2 can also be made for paths measured by $\mathcal{B}_0^{\text{ext}}$.

281 **Observation 5.3.** (Connected Component Sizes) For any graph G and all connected components
 282 $\mathbf{CC} \subset G$, there exists an injective filtration function where \mathbf{PH}_{ext} of that filtration can measure the
 283 number of vertices in \mathbf{CC} .

284 We investigate the case where no learning takes place, namely when the filtration values come from
 285 a random noise. We observe that even in such a situation some information is still encoded in the
 286 extended persistence barcodes with a probability that depends on the graph.

287 **Observation 5.4.** For any graph G where every edge belongs to some cycle and an extended
 288 filtration on it induced by randomly sampled vertex values $x_i \sim U([0, 1])$, \mathbf{PH}_{ext} has a H_1 bar
 289 $[\max_i(x_i), \min_i(x_i)]$ with probability $\sum_{v \in V} \frac{1}{n} \frac{\deg(v)}{n-1}$.

290 Notice that for a clique, the probability of finding the bar with maximum possible persistence is 1. It
 291 becomes lower for sparser graphs.

292 **Corollary 5.5.** In Observation 5.4, the expected persistence $\mathbb{E}[|\max_i(x_i) - \min_i(x_i)|]$ of bar
 293 $[\max_i(x_i), \min_i(x_i)]$ goes to 1 as $n \rightarrow \infty$.

294 What Corollary 5.5 implies is that, for certain graphs, even when nothing is learned by the GNN
 295 filtration learning layers, the longest $\mathcal{B}_1^{\text{ext}}$ bar indicates that n is large. This happens for graphs that
 296 are randomly initialized with vertex labels from the unit interval and occurs with high probability

297 for dense graphs by Observation 5.4. For large n , the empirical mean of the longest bar will have
 298 persistence near 1. Notice that \mathcal{B}_1^{ext} can measure this even though the number of H_1 bars, $m - n + C$,
 299 could tell us nothing about n .

300 6 Experiments

301 We perform experiments of our method on standard GNN datasets. We also perform timing experi-
 302 ments for our extended persistence algorithm, showing impressive scaling. Finally, we investigate
 303 cases where experimentally our method distinguishes graphs that other methods cannot, demonstrating
 304 how our method learns to surpass the WL[1] bound.

305 6.1 Experimental Setup

306 We perform experiments on a 48 core Intel Xeon Gold CPU machine with 1 TB DRAM equipped
 307 with a Quadro RTX 6000 NVIDIA GPU with 24 GB of GPU DRAM.

308 Hyper parameter information can be found in Table 3. For all baseline comparisons, the hyperpa-
 309 rameters were set to their repository’s standard values. In particular, all training were stopped at 100
 310 epochs using a learning rate of 0.01 with the Adam optimizer. Vertex attributes were used along
 311 with vertex degree information as initial vertex labels if offered by the dataset. We perform a fair
 312 performance evaluation by performing standard 10-fold cross validation on our datasets. The lowest
 313 validation loss is used to determined a test score on a test partition. An average±standard deviation
 314 test score over all partitions determines the final evaluation score.

315 The specific layers of our architecture for the neural network for our filtration function f_G is given by
 316 one or two GIN convolutional layers, with the number of layers as determined by an ablation study.

Experimental Evaluation						
avg. acc. ± std.	DD	PROTEINS	IMDB-MULTI	MUTAG	PINWHEELS	2CYCLES
GFL	75.2 ± 3.5	73.0 ± 3.0	46.7 ± 5.0	87.2 ± 4.6	100 ± 0.0	50.0 ± 0.0
Ours+Bars	75.5 ± 2.9	74.9 ± 4.1	50.3 ± 4.7	88.3 ± 7.1	100 ± 0.0	50 ± 0.0
Ours+Bars+Cycles	75.9 ± 2.0	75.2 ± 4.1	51.0 ± 4.6	86.8 ± 7.1	100 ± 0.0	100 ± 0.0
GIN	72.6 ± 4.2	66.5 ± 3.8	49.8 ± 3.0	84.6 ± 7.9	50.0 ± 0.0	50.0 ± 0.0
GIN0	72.3 ± 3.6	67.5 ± 4.7	48.7 ± 3.7	83.5 ± 7.4	50.0 ± 0.0	50.0 ± 0.0
GraphSAGE	72.6 ± 3.7	59.6 ± 0.2	50.0 ± 3.0	72.4 ± 8.1	50.0 ± 0.0	50.0 ± 0.0
GCN	72.7 ± 1.6	59.6 ± 0.2	50.0 ± 2.0	73.9 ± 9.3	50.0 ± 0.0	50.0 ± 0.0
GraphCL	65.4 ± 12	62.5 ± 1.5	49.6 ± 0.4	76.6 ± 26	49.0 ± 8.0	50.5 ± 10
InfoGraph	61.5 ± 10	65.5 ± 12	40.0 ± 8.9	89.1 ± 1.0	50.0 ± 0.0	50.0 ± 0.0
ADGCL	74.8 ± 0.7	73.2 ± 0.3	47.4 ± 0.8	63.3 ± 31	42.5 ± 19	52.5 ± 21
TOGL	74.7 ± 2.4	66.5 ± 2.5	44.7 ± 6.5	-	47.0 ± 3	54.4 ± 5.8
Filt.+SUM	75.0 ± 3.2	73.5 ± 2.8	48.0 ± 2.9	86.7 ± 8.0	51.0 ± 11	50.0 ± 0.0
Filt.+MAX	67.6 ± 3.9	68.6 ± 4.3	45.5 ± 3.1	70.3 ± 5.4	48.0 ± 4.2	50.0 ± 0.0
Filt.+AVG	69.5 ± 2.9	67.2 ± 4.2	46.7 ± 3.8	81.4 ± 7.9	50.0 ± 13	50.0 ± 0.0
Filt.+SORT	76.9 ± 2.6	72.6 ± 4.6	49.0 ± 3.6	85.6 ± 9.2	51.0 ± 16	50.0 ± 0.0
Filt.+S2S	69.0 ± 3.3	67.8 ± 4.6	48.7 ± 4.2	86.8 ± 7.1	51.0 ± 13	50.0 ± 0.0

Table 1: Average accuracy ± std. dev. of our approach (EGFL) with and without explicit cycle representations, Graph Filtration Learning (GFL), GIN0, GIN, GraphSAGE, GCN, ADGCL, GraphCL and TOGL and a readout ablation study on the four TUDatasets: DD, PROTEINS, IMDB-MULTI, MUTAG as well as the two Synthetic WL[1] bound and Cycle length distinguishing datasets. Numbers in bold are highest in performance; bold-gray numbers show the second highest. The symbol – denotes that the dataset was not compatible with software at the time.

317 6.2 Performance on Real World and Synthetic Datasets

318 We perform experiments with the TUDatasets [42], a standard GNN benchmark. We compare with
 319 WL[1] bounded GNNs (GIN, GIN0, GraphSAGE, GCN) from the PyTorch Geometric [43, 44]
 320 benchmark baseline commonly used in practice as well as GFL[5], ADGCL [45], and InfoGraph [46],

10-fold cross validation ablation study on OGBG-MOL datasets by ROC-AUC							
avg. score ± std.	Ours+Bars	Ours+Bars +Cycles	Filt.+SUM	Filt.+MAX	Filt.+AVG	Filt.+SORT	Filt.+Set2Set
molbace	80.0 ± 3.6	81.6 ± 3.9	79.7 ± 4.6	71.9 ± 4.8	78.0 ± 3.0	78.4 ± 3.3	78.2 ± 3.6
molbbbp	78.0 ± 4.3	81.9 ± 3.3	76.7 ± 4.9	69.8 ± 8.7	78.5 ± 4.6	76.3 ± 4.3	78.0 ± 5.0

Table 2: Ablation study on readout functions. The average ROC-AUC ± std. dev. on the ogbg-mol datasets is shown for each readout function. Number coloring is as in Table 1

self-supervised methods. Self supervised methods are promising but should not surpass the performance of supervised methods since they do not use the label during representation learning. We also compare with existing topology based methods TOGL [2] and GFL [5]. We also perform an ablation study on the readout function, comparing extended persistence as the readout function with the SUM, AVERAGE, MAX, SORT, and SET2SET [47] readout functions. The hyper parameter k is set to the 10th percentile of all datasets when sorting for the top- k nodes activations. We do not compare with [19] since its code is not available online. The performance numbers are listed in Table 1. We are able to improve upon other approaches for almost all cases. The real world datasets include DD, MUTAG, PROTEINS and IMDB-MULTI. DD, PROTEINS, and MUTAG are molecular biology datasets, which emphasize cycles, while IMDB-MULTI is a social network, which emphasize cliques and their connections. We use accuracy as our performance score since it is the standard for the TU datasets.

We also verify that our method surpasses the WL[1] bound, a theoretical property which can be proven, as well as can count cycle lengths when the graph is sparse enough, e.g. when the set of cycles is equal to the cycle basis. This is achieved by the two datasets PINWHEELS and 2CYCLES. See the Appendix Sections B for the related experimental and dataset details. Both datasets are particularly hard to classify since they contain spurious constant node attributes, with the labels depending completely on the graph connectivity. This removal of node attributes is in simulation of the WL[1] graph isomorphism test, see [6]. Furthermore, doing so is a case considered in [48]. It is known that WL[1], in particular WL[2], cannot determine the existence of cycles of length greater than seven [49, 50].

Table 2 shows the ablation study of extended filtration learning on the ogbg datasets [51] OGBG MOLBACE and MOLBBBP. We perform a 10 fold cross validation with the test ROC-AUC score of the lowest validation loss used as the test score. This is performed instead of using the train/val/test split offered by the OGBG dataset in order to keep our evaluation methods consistent with the evaluation of the TUDATASETS and synthetic datasets.

From Section B, we know that there are special cases where extended persistence can distinguish graphs where WL[1] bounded GNNs cannot. We perform experiments to show that our method can surpass random guessing whereas other methods achieve only $\sim 50\%$ accuracy on average, which is no better than random guessing. Our high accuracy is guaranteed on PINWHEELS since such graphs are distinguished by counting bars through 0-dim standard persistence. Similarly, 2CYCLES is guaranteed high accuracy when keeping track of cycles and comparing the variance of cycle representations since cycle lengths can be distinguished by a LSTM on different lengthed cycle inputs. Of course, a barcode representation alone will not distinguish cycle lengths.

7 Conclusion

We introduce extended persistence into the supervised learning framework, bringing in crucial global connected component and cycle measurement information into the graph representations. We address a fundamental limitation of MPGNNs, which is their inability to measure cycles lengths. Our method hinges on an efficient algorithm for computing extended persistence. This is a parallel differentiable algorithm with an $O(m \log n)$ depth $O(mn)$ work complexity and scales impressively over the state-of-the-art. The speed with which we can compute extended persistence makes it feasible for machine learning. Our end-to-end model obtains favorable performance on real world datasets. We also construct cases where our method can distinguish graphs that existing methods struggle with.

References

- 364
- 365 [1] Christoph D Hofer, Roland Kwitt, and Marc Niethammer. Learning representations of persis-
366 tence barcodes. *J. Mach. Learn. Res.*, 20(126):1–45, 2019. 1, 4, 5, 21
- 367 [2] Max Horn, Edward De Brouwer, Michael Moor, Yves Moreau, Bastian Rieck, and Karsten
368 Borgwardt. Topological graph neural networks. *arXiv preprint arXiv:2102.07835*, 2021. 1, 2, 4,
369 9, 13, 16
- 370 [3] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Extending persistence using
371 poincaré and lefschetz duality. *Foundations of Computational Mathematics*, 9(1):79–103, 2009.
372 1, 3
- 373 [4] Zuoyu Yan, Tengfei Ma, Liangcai Gao, Zhi Tang, and Chao Chen. Link prediction with
374 persistent homology: An interactive view. In *International Conference on Machine Learning*,
375 pages 11659–11669. PMLR, 2021. 1, 5, 6
- 376 [5] Christoph Hofer, Florian Graf, Bastian Rieck, Marc Niethammer, and Roland Kwitt. Graph
377 filtration learning. In *International Conference on Machine Learning*, pages 4314–4323. PMLR,
378 2020. 2, 4, 8, 9
- 379 [6] Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra
380 which appears therein. *NTI, Series*, 2(9):12–16, 1968. 2, 9
- 381 [7] Tamal K. Dey and Yusu Wang. *Computational Topology for Data Analysis*. Cambridge Uni-
382 versity Press, 2022. [https://www.cs.purdue.edu/homes/tamaldey/book/CTDAbook/
383 CTDAbook.pdf](https://www.cs.purdue.edu/homes/tamaldey/book/CTDAbook/CTDAbook.pdf). 2
- 384 [8] Herbert Edelsbrunner and John Harer. *Computational topology: an introduction*. American
385 Mathematical Soc., 2010. 2, 5
- 386 [9] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov,
387 and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30,
388 2017. 3, 5
- 389 [10] Andreas Loukas. What graph neural networks cannot learn: depth vs width. *arXiv preprint
390 arXiv:1907.03199*, 2019. 3
- 391 [11] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural
392 networks? *arXiv preprint arXiv:1810.00826*, 2018. 3
- 393 [12] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A
394 comprehensive survey on graph neural networks. *IEEE transactions on neural networks and
395 learning systems*, 32(1):4–24, 2020. 3
- 396 [13] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large
397 graphs. In *Proceedings of the 31st International Conference on Neural Information Processing
398 Systems*, pages 1025–1035, 2017. 3
- 399 [14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional
400 networks. *arXiv preprint arXiv:1609.02907*, 2016. 3
- 401 [15] Sandra Kiefer, Neil Immerman, Pascal Schweitzer, and Martin Grohe. Power and limits of the
402 weisfeiler-leman algorithm. Technical report, Fachgruppe Informatik, 2020. 3
- 403 [16] Jiaxuan You, Jonathan Gomes-Selman, Rex Ying, and Jure Leskovec. Identity-aware graph
404 neural networks. *arXiv preprint arXiv:2101.10320*, 2021. 4
- 405 [17] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen,
406 Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural
407 networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages
408 4602–4609, 2019. 4
- 409 [18] Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yu Guang Wang, Pietro Liò, Guido F Montufar,
410 and Michael Bronstein. Weisfeiler and leman go cellular: Cw networks. *Advances in Neural
411 Information Processing Systems*, 34, 2021. 4
- 412 [19] Hongyang Gao, Yi Liu, and Shuiwang Ji. Topology-aware graph pooling networks. *IEEE
413 Transactions on Pattern Analysis and Machine Intelligence*, 43(12):4512–4518, 2021. 4, 7, 9
- 414 [20] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. Self-attention graph pooling. In *International
415 conference on machine learning*, pages 3734–3743. PMLR, 2019. 4

- 416 [21] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural
417 networks. *arXiv preprint arXiv:1511.05493*, 2015. 4
- 418 [22] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec.
419 Hierarchical graph representation learning with differentiable pooling. *Advances in neural*
420 *information processing systems*, 31, 2018. 4
- 421 [23] Chuang Liu, Yibing Zhan, Chang Li, Bo Du, Jia Wu, Wenbin Hu, Tongliang Liu, and Dacheng
422 Tao. Graph pooling for graph neural networks: Progress, challenges, and opportunities. *arXiv*
423 *preprint arXiv:2204.07321*, 2022. 4
- 424 [24] Alexandros D Keros, Vedit Nanda, and Kartic Subr. Dist2cycle: A simplicial neural network
425 for homology localization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36
426 (7):7133–7142, Jun. 2022. doi: 10.1609/aaai.v36i7.20673. URL [https://ojs.aaai.org/
427 index.php/AAAI/article/view/20673](https://ojs.aaai.org/index.php/AAAI/article/view/20673). 4
- 428 [25] Joshua Levy, Christian Haudenschild, Clark Barwick, Brock Christensen, and Louis Vaickus.
429 Topological feature extraction and visualization of whole slide images using graph neural
430 networks. In *BIOCOMPUTING 2021: Proceedings of the Pacific Symposium*, pages 285–296.
431 World Scientific, 2020.
- 432 [26] Guido Montúfar, Nina Otter, and Yuguang Wang. Can neural networks learn persistent homology
433 features? *arXiv preprint arXiv:2011.14688*, 2020.
- 434 [27] Zuoyu Yan, Tengfei Ma, Liangcai Gao, Zhi Tang, Yusu Wang, and Chao Chen. Neural
435 approximation of extended persistent homology on graphs. *CoRR*, abs/2201.12032, 2022. URL
436 <https://arxiv.org/abs/2201.12032>. 4
- 437 [28] Qi Zhao, Ze Ye, Chao Chen, and Yusu Wang. Persistence enhanced graph neural network. In
438 Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International*
439 *Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine*
440 *Learning Research*, pages 2896–2906. PMLR, 26–28 Aug 2020. URL [https://proceedings.
441 mlr.press/v108/zhao20d.html](https://proceedings.mlr.press/v108/zhao20d.html). 4
- 442 [29] Zuoyu Yan, Tengfei Ma, Liangcai Gao, Zhi Tang, and Chao Chen. Cycle representation learning
443 for inductive relation prediction. In *ICLR 2022 Workshop on Geometrical and Topological*
444 *Representation Learning*, 2022. 4
- 445 [30] Simon Zhang, Mengbai Xiao, and Hao Wang. Gpu-accelerated computation of vietoris-rips
446 persistence barcodes. In *36th International Symposium on Computational Geometry (SoCG*
447 *2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. 4
- 448 [31] Thibault de Surrel, Felix Hensel, Mathieu Carrière, Théo Lacombe, Yuichi Ike, Hiroaki Kurihara,
449 Marc Glisse, and Frédéric Chazal. Ripsnet: a general architecture for fast and robust estimation
450 of the persistent homology of point clouds, 2022. URL [https://arxiv.org/abs/2202.
451 01725](https://arxiv.org/abs/2202.01725). 4
- 452 [32] Clement Vignac, Andreas Loukas, and Pascal Frossard. Building powerful and equivariant graph
453 neural networks with structural message-passing. *Advances in Neural Information Processing*
454 *Systems*, 33:14143–14155, 2020. 4
- 455 [33] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and
456 Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In
457 *International conference on machine learning*, pages 5453–5462. PMLR, 2018. 4
- 458 [34] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Be-
459 yond homophily in graph neural networks: Current limitations and effective designs. *Advances*
460 *in Neural Information Processing Systems*, 33:7793–7804, 2020. 4
- 461 [35] Wenbing Huang, Yu Rong, Tingyang Xu, Fuchun Sun, and Junzhou Huang. Tackling over-
462 smoothing for general graph convolutional networks. *arXiv preprint arXiv:2008.09864*, 2020.
463 4
- 464 [36] Kenta Oono and Taiji Suzuki. Optimization and generalization analysis of transduction through
465 gradient boosting and application to multi-scale graph neural networks. *Advances in Neural*
466 *Information Processing Systems*, 33:18917–18930, 2020. 4
- 467 [37] Guy E Blelloch and Bruce M Maggs. Parallel algorithms. In *Algorithms and theory of*
468 *computation handbook: special topics and techniques*, pages 25–25. 2010. 5

- 469 [38] Richard J Anderson and Gary L Miller. Deterministic parallel list ranking. *Algorithmica*, 6(1):
470 859–868, 1991. 6
- 471 [39] Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can graph neural networks count
472 substructures? *arXiv preprint arXiv:2002.04025*, 2020. 7
- 473 [40] Muhan Zhang, Zhicheng Cui, Marion Neumann, and Yixin Chen. An end-to-end deep learning
474 architecture for graph classification. In *Proceedings of the AAAI conference on artificial
475 intelligence*, volume 32, 2018. 7
- 476 [41] Vikas Garg, Stefanie Jegelka, and Tommi Jaakkola. Generalization and representational limits
477 of graph neural networks. In *International Conference on Machine Learning*, pages 3419–3430.
478 PMLR, 2020. 7
- 479 [42] Christopher Morris, Nils M Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion
480 Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. *arXiv
481 preprint arXiv:2007.08663*, 2020. 8
- 482 [43] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric.
483 *arXiv preprint arXiv:1903.02428*, 2019. 8
- 484 [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,
485 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative
486 style, high-performance deep learning library. *Advances in neural information processing
487 systems*, 32, 2019. 8
- 488 [45] Susheel Suresh, Pan Li, Cong Hao, and Jennifer Neville. Adversarial graph augmentation to
489 improve graph contrastive learning. *Advances in Neural Information Processing Systems*, 34,
490 2021. 8
- 491 [46] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. Infograph: Unsupervised and
492 semi-supervised graph-level representation learning via mutual information maximization. *arXiv
493 preprint arXiv:1908.01000*, 2019. 8
- 494 [47] Haoji Hu and Xiangnan He. Sets2sets: Learning from sequential sets with neural networks. In
495 *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery &
496 Data Mining*, pages 1491–1499, 2019. 9
- 497 [48] Pan Li, Yanbang Wang, Hongwei Wang, and Jure Leskovec. Distance encoding: Design
498 provably more powerful neural networks for graph representation learning. *Advances in Neural
499 Information Processing Systems*, 33:4465–4478, 2020. 9
- 500 [49] Vikraman Arvind, Frank Fuhlbrück, Johannes Köbler, and Oleg Verbitsky. On weisfeiler-leman
501 invariance: Subgraph counts and related graph properties. *Journal of Computer and System
502 Sciences*, 113:42–59, 2020. 9
- 503 [50] Martin Fürer. On the combinatorial power of the weisfeiler-lehman algorithm. In *International
504 Conference on Algorithms and Complexity*, pages 260–271. Springer, 2017. 9
- 505 [51] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
506 Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs.
507 *arXiv preprint arXiv:2005.00687*, 2020. 9
- 508 [52] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. Measuring and relieving the
509 over-smoothing problem for graph neural networks from the topological view. In *Proceedings
510 of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3438–3445, 2020. 16
- 511 [53] Robin Sibson. Slink: an optimally efficient algorithm for the single-link cluster method. *The
512 computer journal*, 16(1):30–34, 1973. 17
- 513 [54] Bernard A Galler and Michael J Fisher. An improved equivalence algorithm. *Communications
514 of the ACM*, 7(5):301–303, 1964. 17
- 515 [55] Daniel D Sleator and Robert Endre Tarjan. A data structure for dynamic trees. In *Proceedings
516 of the thirteenth annual ACM symposium on Theory of computing*, pages 114–122, 1981. 18, 19
- 517 [56] Daniel Dominic Sleator and Robert Endre Tarjan. Self-adjusting binary search trees. *Journal of
518 the ACM (JACM)*, 32(3):652–686, 1985. 18, 19

A Proofs

Theorem A.1. (Theorem 5.1)

$\mathbf{PH}_{\text{ext}}(G)$ produces four types of bars: $\mathcal{B}_1^{\text{ext}}, \mathcal{B}_0^{\text{ext}}, \mathcal{B}_0^{\text{low}}, \mathcal{B}_0^{\text{up}}$, s.t.

$$|\mathcal{B}_1^{\text{ext}}| = \dim H_1 = m - n + C,$$

$$|\mathcal{B}_0^{\text{ext}}| = \dim H_0 = C,$$

$$|\mathcal{B}_0^{\text{low}}| = |\mathcal{B}_0^{\text{upper}}| = n - C,$$

where there are C connected components and $\dim H_k$ is the dimension of the k th homology group s.t.:

1. the H_1 barcode comes from a cycle basis of G which also constitutes a basis of its fundamental group,

2. $\dim H_1$ counts the number of chordless cycles when G is outer-planar, and

3. there exists an injective filtration function where the union of the resulting barcodes is strictly more expressive than the histogram produced by the $\text{WL}[1]$ graph isomorphism test.

Proof. There are n bars with vertex births since every vertex creates exactly one connected component. The number of these bars which are in $\mathcal{B}_0^{\text{ext}}$ is C , which counts the number of global connected components. In other words, $|\mathcal{B}_0^{\text{ext}}| = \dim(H_0) = C$. Thus, we have $n - C = |\mathcal{B}_0^{\text{low}}| = |\mathcal{B}_0^{\text{upper}}|$.

Considering all $2m$ edges on the extended filtration, every edge gets paired. Furthermore, $n - C$ of the edges in the lower filtration are negative edges paired with vertices that give birth to connected components. Similarly there are $n - C$ edges paired with vertices in the upper filtration. We thus have $\frac{2m - 2(n - C)}{2}$ edge-edge pairings in $\mathcal{B}_1^{\text{ext}}$ because every edge gets paired. Thus, $|\mathcal{B}_1^{\text{ext}}| = m - n + C$. Since each bar in $\mathcal{B}_1^{\text{ext}}$ counts a birth of a 1-dimensional homological class which together span the 1-dimensional homological classes in H_1 , we have that $\dim H_1 = |\mathcal{B}_1^{\text{ext}}|$.

1. This follows from the discussion above.

2. By Euler's formula, we have $n - m + F = C + 1$ for planar graphs where F is the number of faces of the planar graph as embedded in \mathbb{S}^2 . For outer planar graphs, since $F - 1$ interior faces lie on one hemisphere of \mathbb{S}^2 and one exterior face covers the opposite hemisphere, each interior face must be a chordless cycle.

3. This follows directly by the result in [2] stating that 0-dimensional barcodes are more expressive than the $\text{WL}[1]$ graph isomorphism test. In extended persistence, $\mathcal{B}_0^{\text{low}}$ and $\mathcal{B}_0^{\text{ext}}$ are computed. Since all bars in $\mathcal{B}_0^{\text{ext}}$ correspond to infinite bars denoted \mathcal{B}_0^∞ in the 0-dimensional standard persistence, we have that $\mathcal{B}_0^{\text{low}}$ and $\mathcal{B}_0^{\text{ext}}$ carry at least the same amount of information as a 0-dimensional barcode as determined by $\mathcal{B}_0^{\text{low}}$ and \mathcal{B}_0^∞ .

551

□

Observation A.2. (Observation 5.2) For any graph G and a cycle $\mathbf{C} \subset G$, there exists an injective filtration function where \mathbf{PH}_{ext} of the induced filtration can measure the number of edges along \mathbf{C} .

Proof. Number the vertices of the cycle \mathbf{C} of length k in descending order and counter clockwise as $n - 1 \dots n - k$. For each vertex $u \in \mathbf{C}$, set $f_G(u)$ to be the index of u . For the other vertices, assign arbitrary different values less than $n - k$ and then apply ε -perturbation to make the filtration injective. For example, for vertex u and all its incident edges of same filtration value, one can subtract different $\varepsilon \in \mathbb{R}^+$ from each edge to impose injectivity on the induced filtration of f_G . We then get that every edge on the cycle \mathbf{C} except one: $(n - 1, n - k)$ becomes negative and thus belongs to the negative spanning forest of the upper filtration. The positive edge of smallest value in the upper filtration is edge $(n - 1, n - k)$. The extended persistence algorithm, after computing $\mathcal{B}_0^{\text{low}}$ and $\mathcal{B}_0^{\text{up}}$, pairs the edge $e = (n - 1, n - k)$ with the edge having maximum value in the lower filtration in the cycle \mathbf{C} that e forms with the spanning forest. This paired edge is $(n - 1, n - 2)$ and has lower filtration value $n - 1$. We thus have the bar $[n - 1, n - k]$ which encodes the length k of the cycle \mathbf{C} .

565

□

566 **Observation A.3.** (Observation 5.3) For any graph G and all connected components $\mathbf{CC} \subset G$,
 567 there exists an injective filtration function where \mathbf{PH}_{ext} of that filtration can measure the number of
 568 vertices in \mathbf{CC} .

569 *Proof.* For each connected component \mathbf{CC} in G , index the vertices in \mathbf{CC} in consecutive order
 570 where indices in each connected component remain distinct. Then define $f_G(u)$ equal to the index of
 571 u in G . By ε -perturbation, we can make this an injective filtration function. Since $\mathcal{B}_0^{\text{ext}}$ has each bar
 572 $[\min_{u \in \mathbf{CC}} f_G(u), \max_{u \in \mathbf{CC}} f_G(u)]$ and since all indices are consecutive, each bar's persistence in
 573 $\mathcal{B}_0^{\text{ext}}$ measures how many vertices are in the connected component they constitute.

574 □

575 **Observation A.4.** (Observation 5.4) For any graph G where every edge belongs to some cycle and
 576 an extended filtration on it is induced by randomly sampling vertex values $x_i \sim U([0, 1])$, \mathbf{PH}_{ext}
 577 has the H_1 bar $[\max_i(x_i), \min_i(x_i)]$ with probability $\sum_{v \in V} \frac{1}{n} \frac{\text{deg}(v)}{n-1}$.

578 *Proof.* Since the probability of finding a given permutation on n vertices sampled uniformly at
 579 random without replacement is equivalent to the probability of a given order on the vertices sampled
 580 uniformly at randomly n times, it suffices to find the probability of sampling uniformly at random
 581 without replacement two vertices that are connected with an edge in G .

582 For a fixed $\sigma \in S_n$, a permutation from the group S_n of permutations on n vertices, we have:

$$\begin{aligned} \frac{1}{n!} &= P(x_n < x_{n-1} < \dots < x_1, x_i \sim U([0, 1])) \\ &= \int_0^1 \int_0^{x_1} \dots \int_0^{x_{n-1}} dx_n dx_{n-1} \dots dx_1 = P(\sigma \sim U(S_n)) \end{aligned}$$

583 Let $G = (V, E)$ be the graph with vertex values sampled from a uniform distribution. Let $G' =$
 584 (V', E') be the same graph with vertex values in $\{0, 1, \dots, n-1\}$ sampled uniformly without
 585 replacement. We know that the probability for a given order on these vertices is the same for both
 586 graphs. By the law of total probability:

$$\begin{aligned} &P((\min_i x_i, \max_i x_i) \in E, x_i \sim U([0, 1])) \\ &= \sum_{v \in V} (P(v = \max_i x_i, x_i \sim U([0, 1])) \cdot P(\min_i x_i \in \text{Nbr}(v) | v = \max_i x_i, x_i \sim U([0, 1]))) \\ &= \sum_{v \in V} (n-1)! \int_0^1 \int_0^{x_1} \dots \int_0^{x_{n-1}} dx_n dx_{n-1} \dots dx_1 \cdot \text{deg}(v)(n-2)! \int_0^1 \int_0^{x_2} \dots \int_0^{x_{n-1}} dx_n dx_{n-1} \dots dx_2 \\ &= P((n-1, 0) \in E') = \sum_{v \in V'} (P(v = n-1) \cdot P(0 \in \text{Nbr}(v) | v = n-1)) \\ &= \sum_{v \in V'} \frac{1}{n} \frac{\text{deg}(v)}{n-1} \end{aligned}$$

587 We now show that if $(\min_i x_i, \max_i x_i)$ occurs as an edge in $G = (V, E)$, where every edge belongs to
 588 some cycle, then the bar $[\max_i x_i, \min_i x_i]$ is guaranteed to occur.

589 The spanning tree comprised of negative edges that begins the computation for $\mathcal{B}_1^{\text{ext}}$ as in line 6
 590 of Algorithm 1 for the H_1 barcode computation is a maximum spanning tree. This is because
 591 the negative edges are just those found by the Kruskal's algorithm for the 0-dimensional standard
 592 persistence applied to an upper filtration. Since $e = (\min_i x_i, \max_i x_i)$ has value $\min_i x_i$ in the
 593 upper filtration and since every edge belongs to at least one cycle, it cannot be in the maximum
 594 spanning tree. Thus e is a positive edge.

595 Since e is positive in the upper filtration, it will be considered at some iteration of the for loop in line
 596 10 of Algorithm 1. When we consider it, it will form a cycle \mathbf{C} with the dynamically maintained
 597 spanning forest. To form a persistence H_1 bar for e , we pair it with the maximum edge in the cycle
 598 \mathbf{C} in the lower filtration. This forms a bar $[\max_i x_i, \min_i x_i]$.

599 □

600 **Corollary A.5.** *In Observation 5.4, the expected persistence of bar $[\max_i(x_i), \min_i(x_i)]$,*
 601 *$\mathbb{E}[\max_i(x_i) - \min_i(x_i)]$, goes to 1 as $n \rightarrow \infty$.*

602 *Proof.* Define the random variable $X_n = |\max_i x_i - \min_i x_i|$ for n random points drawn uniformly
 603 from $[0, 1]$. We find $\lim_{n \rightarrow \infty} \mathbb{E}[X_n]$. The following sequence of equations follow by repeated
 604 substitution.

$$\begin{aligned} \mathbb{E}[X_n] &= n! \int_0^1 \int_0^{x_1} \dots \int_0^{x_{n-1}} (x_1 - x_n) dx_n \dots dx_1 \\ &= n! \int_0^1 \left(\frac{x_1^n}{(n-1)!} - \frac{x_1^n}{n!} \right) dx_1 = \frac{n-1}{n+1} \end{aligned}$$

605 where the $n!$ comes from symmetry.

606 Therefore: $\lim_{n \rightarrow \infty} \mathbb{E}[X_n] = 1$. □

607 B Demonstrating the Expressivity of Learned Extended Persistence

608 We present some cases where the classification performance of our method excels. We look for
 609 graphs that cannot be distinguished by WL[1] bounded GNNs. We find that pinwheeled cycle graphs
 610 and varied length cycle graphs can be perfectly distinguished by learned extended persistence and, in
 611 practice, with much better performance than random guessing using our model. See the experiments
 612 Section 6 to see the empirical results for our method against other methods on this synthetic data.

613 B.1 Pinwheeled Cycle Graphs (The PINWHEELS Dataset)




Figure 3: Class 0: 2 triangles with pinwheel at each vertex.




Figure 4: Class 1: A hexagon with pinwheel at each vertex.

614 We consider pinwheeled cycle graphs. To form the base skeleton of these graphs, we take the standard
 615 counter example to the WL[1] test of 2 triangles and 1 hexagon. We then append pinwheels of a
 616 constant number of vertices to the vertices of these base skeletons. The node attributes are set to a
 617 spurious constant noise vector. They have no effect on the labels.

618 It is easy to check that both Class 0 and Class 1 graphs are indistinguishable by WL[1]; see Figures 3
 619 and 4. Notice that if there are 6 core vertices and edges in the base skeleton and if there are pinwheels
 620 of size k , then with edge deletions and vertex deletions composed, we have a $1 - (\frac{6}{6k+6})^2$ probability
 621 of only deleting a pinwheel edge or vertex and thus not affecting H_1 . This probability converges to 1
 622 as $k \rightarrow \infty$. According to Theorem 5.1, $\dim H_1$ measures the number of cycles and $\dim H_0$ measures
 623 the number of connected components. If neither of these counts are affected by training during
 624 supervised learning, our method is guaranteed to distinguish the two classes simply by counting
 625 according to Theorem 5.1.

626 Certainly the pinwheeled cycle graphs, are distinguishable by counts of bars. We check this experi-
 627 mentally by constructing a dataset of 1000 graphs of two classes of graph evenly split. Class 0 is as
 628 in Figure 3 and involves two triangles with pinwheels of random sizes. Class 1 is as in Figure 4 with

629 a hexagon and pinwheels of random sizes attached. We obtain on average 100% accuracy. This is
 630 confirmed experimentally in Table 1. This matches the performance of GFL, since counting bars, or
 631 Betti numbers, can also be done through 0-dim. standard persistence. Interestingly TOGL does not
 632 achieve a score of 100 accuracy on this dataset. We conjecture this is because their layers are not able
 633 to ignore the spurious and in fact misleading constant node attributes.

634 B.2 Regular Varied Length Cycle Graphs (The 2CYCLES dataset)




Figure 5: Class 0: A 15 node cycle and an 85 node cycle.




Figure 6: Class 1: A 50 node cycle with a 50 node cycle.

635 We further consider varied length cycle graphs. These are graphs that involve two cycles. Class 0 has
 636 one short and one long cycle while Class 1 has two near even lengthed cycles. The node attributes are
 637 all the same and spurious in this dataset. Extended persistence should do well to distinguish these
 638 two classes. We conjecture this based on Observation 5.2, which states that there is some filtration
 639 that can measure the length of certain cycles.

640 It is the path length, coming from Observation 5.3, which is being measured. The 0-dimensional
 641 standard persistence is insufficient for this purpose. The infinite bars of 0-dimensional standard
 642 persistence are determined only by a birth time. Furthermore, extended persistence without cycle
 643 representatives is also insufficient since a message passing GNN learns a constant filtration function
 644 over the nodes. However, with cycle representatives, or a list of scalar node activations per cycle for
 645 each graph, we can easily distinguish the average sequence representation since the pair of sequence
 646 lengths are different. In class 0, a short cycle and a long cycle are paired while in class 1, two cycles
 647 of medium lengths are paired.

648 A similar but more challenging dataset to the PINWHEELS dataset, the 2CYCLES dataset, is similar
 649 to the necklaces dataset from [2] and is illustrated in Section B.2 but with more misleading node
 650 attributes and simplified to two cycles. It involves 400 graphs consisting of two cycles. There are two
 651 classes as shown in Figures 5 and 6.

652 The experimental performance on 2CYCLES surpasses random guessing while all other methods
 653 just randomly guess as stated in Section B.2. Certainly WL[1] bounded GNNs cannot distinguish
 654 the two classes in 2CYCLES since they are all regular. As discussed, because GFL and TOGL use
 655 learned 0-dimensional standard persistence, these approaches do no better than random guessing on
 656 this dataset.

657 B.2.1 Number of Convolutional Layers Experiment

658 We also perform an experiment to determine the number of layers in the MPGNN of the filtration
 659 function that has the highest performance. Due to oversmoothing [52], which is exacerbated by the
 660 required scalar-dimensional vertex embeddings, as we increase the number of layers for the filtration
 661 function the performance drops. See Figure 7 for an illustration of this phenomenon on the PROTEINS
 662 and MUTAG dataset. For these two datasets, two layers perform the best.

663 C Timing of Extended Persistence Algorithm (without storing cycle 664 representations)

665 Since the persistence computation, especially extended persistence computation, is the bottleneck to
 666 any machine learning algorithm that uses it, it is imperative to have a fast algorithm to compute it.
 667 We perform timing experiments with a C++ torch implementation of our fast extended persistence
 668 algorithm. In our implementation each graph in the batch has a single thread assigned to it.




Figure 7: An exhibit of oversmoothing in the filtration convolutional layers. Plot of the average accuracy with std. dev. as a function of the number of convolutional layers before the Jumping Knowledge MLP and the extended persistence readout. The Proteins and Mutag datasets were used in (a) and (b) respectively.

669 Our experiment involves two parameters, the sparsity, or probability, p for the edges of an Erdos-
 670 Renyi graph and the number of vertices of such a graph n . We plot our speedup over GUDHI, the
 671 state of the art software for computing extended persistence, as a function of p with n held fixed. We
 672 run GUDHI and our algorithm 5 times and take the average and standard deviation of each run's
 673 speedup. Since our algorithm has lower complexity, our speedup is theoretically unbounded. We
 674 obtain up to 62x speedup before surpassing 12 hours of computation time for experimentation. The
 675 plot is shown in Figure 8. The speedup is up to 2.8x, 9x, 24x, and 62x for $n = 200, 500, 1000, 2000$
 676 respectively.




Figure 8: Average speedup with std. dev. as a function of sparsity p and number of vertices n on Erdos Renyi graphs.

677 D Algorithm and Data Structure Details

678 Here we detail the algorithmic details of computing extended persistence.

679 D.1 The PH_0 Algorithm

680 Here is the union-find algorithm that computes 0 dimensional persistent homology. The algorithm
 681 is a single-linkage clustering algorithm [53]. It starts with n nodes, 0 edges, and a union-find data
 682 structure [54] on n nodes. The edges are sorted in ascending order if a lower filtration function
 683 is given. Otherwise, the edges are sorted in descending order. It then proceeds to connect nearest
 684 neighbor clusters, or connected components, in a sequential fashion by introducing edges in order one
 685 at a time. Two connected components are nearest to each other if they have two nodes closer to each

Algorithm 2 PH₀ Algorithm

Input: $G = (V, E)$, F : filtration function, $order$: flag to denote an upper or lower filtration
Output: $\mathcal{B}_0, E_{pos}, E_{neg}, U$: H_0 bars, pos. edges, neg. edges, and union-find data structure
 1: $U \leftarrow V$ /* a union-find data structure populated by n unlinked nodes*/
 2: $\mathcal{B}_0 \leftarrow \{\}$ /*A multiset*/
 3: **if** $order = lower$ **then**
 4: $SORT_{incr}(E)$ /*increasing w.r.t. F;*/
 5: **else**
 6: $SORT_{decr}(E)$ /*decreasing w.r.t F;*/
 7: **end if**
 8: **for** $e = (u, v) \in E$ **do**
 9: $root_u \leftarrow U.FIND(u)$
 10: $root_v \leftarrow U.FIND(v)$
 11: **if** $root_u = root_v$ **then**
 12: $E_{pos} \leftarrow E_{pos} \cup \{e\}$
 13: **else**
 14: $E_{neg} \leftarrow E_{neg} \cup \{e\}$
 15: **end if**
 16: **if** $order = lower$ **then**
 17: $b \leftarrow \max(F(root_u), F(root_v))$
 18: **else**
 19: $b \leftarrow \min(F(root_u), F(root_v))$
 20: **end if**
 21: $d \leftarrow F(e)$
 22: $\mathcal{B}_0 \leftarrow \mathcal{B}_0 \cup \{(b, d)\}$
 23: $U.LINK(root_u, root_v)$
 24: **end for**
 25: **return** $(\mathcal{B}_0, E_{pos}, E_{neg}, U)$

686 other than any other pair of connected components. This is achieved by iterating through the edges in
 687 sorted order and merging the connected components that they connect. When given a lower filtration
 688 function, when a connected component merges with another connected component, the connected
 689 component with the larger connected component root value has its root filtration function value a
 690 birth time. This birth time is paired with the current edge's filtration value and form a birth death
 691 pair. The smaller of the two connected component root values is used as birth time when an upper
 692 filtration function is given. The two connected components are subsequently merged in a union-find
 693 data structure by the LINK operation.

694 D.2 A Brief Overview of the Link-Cut Tree Data Structure

695 The link-cut tree data structure [55] is a well known dynamic connectivity data structure. For
 696 modifying the tree of n nodes, it takes $O(\log n)$ amortized time for deleting an edge (cut) and joining
 697 two trees (link). Furthermore, it takes $O(\log n)$ amortized time for the composition of associative
 698 reductions, such as max, min, sum, on some path from any node to its root. We may view the
 699 link-cut tree data structure as a collection of trees and thus as a forest as well. Details of this forest
 700 implementation are omitted.

701 The link-cut tree decomposes a tree T into a disjoint union of preferred paths, or sequences of nodes
 702 that strictly decreasing in depth (distance from the root of T) on T . A path has each consecutive node
 703 connected by a single edge. In particular, each node in T has a preferred child, forming a preferred
 704 edge. The maximally connected sequence of preferred edges forms a preferred path. The preferred
 705 path decomposition will change as the link-cut tree gets operated on. Each preferred path is in one to
 706 one correspondence with a splay tree [56] called an auxiliary tree on the set of nodes in the preferred
 707 path. For any node v in a preferred path's auxiliary tree, its left subtree is made up of nodes higher up
 708 (closer to the root in T) than v and its right subtree is made up of nodes lower (farther from the root
 709 in T) than v . Each auxiliary tree contains a pointer, termed the auxiliary tree's parent-pointer, from
 710 its root to the parent of the highest (closest to the root) node in the preferred path associated with the
 711 auxiliary tree.

712 The most important supporting operation to a link-cut tree T is the $\text{EXPOSE}()$ operation. The result
 713 of $\text{EXPOSE}(v)$ for $v \in T$ is the formation of a unique preferred path from the root of T to v with this
 714 preferred path's set of nodes forming an auxiliary tree. Furthermore, it results in v to be the root
 715 of the auxiliary tree it belongs to. The complexity of $\text{EXPOSE}(v)$ is $O(\log n)$. For implementation
 716 details, see [55].

717 Let T_1, T_2 be two link-cut trees and $u \in T_1, v \in T_2$. Define the operation $\text{LINK}(T_1, (u, v), T_2)$ as
 718 the operation that joins T_1 to T_2 by connecting u with v by an edge and outputs the resulting tree.
 719 This is achieved by simply calling $\text{EXPOSE}(u)$ then $\text{EXPOSE}(v)$, which makes u and v the roots of
 720 their respective auxiliary trees, then in the auxiliary tree of u , set the left child of u to v .

721 Let T be a link-cut tree and $u, v \in T$ connected by an edge. Define the operation $\text{CUT}(T, (u, v))$ as
 722 the operation that disconnects T by deleting the edge between u and v . This is achieved by simply
 723 calling $\text{EXPOSE}(u)$ and then making u a root by making the left child of v point to *null*.

724 Let T be a link-cut tree and $u, v \in T$. Define the operation $\text{LCA}(T, (u, v))$ as the operation that finds
 725 the least common ancestor of u and v in T . This is achieved by calling $\text{EXPOSE}(u)$ then $\text{EXPOSE}(v)$
 726 and then taking the node pointed to by the parent-pointer of the auxiliary tree of which u is root.

727 Let T be a link-cut tree and $u, v \in T$. Define the operation $\text{PATH}(u, v)$ as the operation that returns a
 728 linked list of the path from u to v in T . This is achieved by obtaining the parent v' of v first. The
 729 parent of v can be obtained by calling $\text{EXPOSE}(v)$ then traversing the splay tree it is a root of for its
 730 parent in T . Call $\text{EXPOSE}(u)$ to form a preferred path from u to the root of T then $\text{EXPOSE}(v')$ to
 731 detach v' from this preferred path. Let $\text{SPLAY}(u)$ be the operation that rotates the unique splay tree,
 732 or preferred path, containing u so that u becomes the root of its splay tree. After calling $\text{SPLAY}(u)$,
 733 u becomes the root of a linked-list splay tree. It is a linked-list since u is the lowest (farthest from the
 734 root) node in its splay tree and the rest of the preferred path is made up of a path of strictly decreasing
 735 distance to the root. Return this linked-list splay tree as the resulting path from u to v .

736 Let T be a link-cut tree, $u, v \in T$ with v higher up in the tree than u (it is closer to the root of T than
 737 u). Define $\text{REDUCE}(T, u, v, op)$ to be an associative reduction on the path from u to v . To do this,
 738 apply $\text{EXPOSE}(u)$ then $\text{EXPOSE}(v)$, then apply the associative operation on the whole auxiliary tree,
 739 as implemented on a splay tree in [56]. The associative reduction takes $O(\log n)$ time. This splay
 740 tree corresponds to the preferred path from u to v formed from the two EXPOSE operations. Notice
 741 that $\text{EXPOSE}(u)$ results in a preferred path from u to the root while the second call $\text{EXPOSE}(v)$
 742 detaches the path from v to the root of T from the preferred path of u to the root.

743 Let T be a link-cut tree, $u, v \in T$ and lca the least common ancestor of $u, v \in T$. Assume the
 744 nodes are labeled by a pair of their value and index. Define $\text{ARGMAXREDUCECYCLE}(T, u, v, lca)$
 745 as the operation that finds the edge with one of its nodes containing the maximum value on the
 746 cycle formed by u, v and lca . There are many ways to implement this. We describe a method that
 747 maintains the $O(\log n)$ complexity of link-cut tree operations. We first compute $(value(w_1), w_1) :=$
 748 $\text{REDUCE}(T, u, lca, max)$ to find the maximum value node along the path from u to lca , then compute
 749 $(value(w_2), w_2) := \text{REDUCE}(T, v, lca, max)$ to find the maximum value node along the path from
 750 v to lca . Let w to be the maximum valued vertex between w_1 and w_2 . If $w \neq lca(u, v)$, then find the
 751 parent z of w otherwise apply $\text{EXPOSE}(u)$ then $\text{EXPOSE}(v)$ and keep track of the child z of w that gets
 752 detached during $\text{EXPOSE}(v)$. Parent of w can be found by $\text{EXPOSE}(w)$ then traversing its splay tree
 753 to find the parent of $w \in T$. The edge (z, w) is returned by $\text{ARGMAXREDUCECYCLE}(T, u, v, lca)$

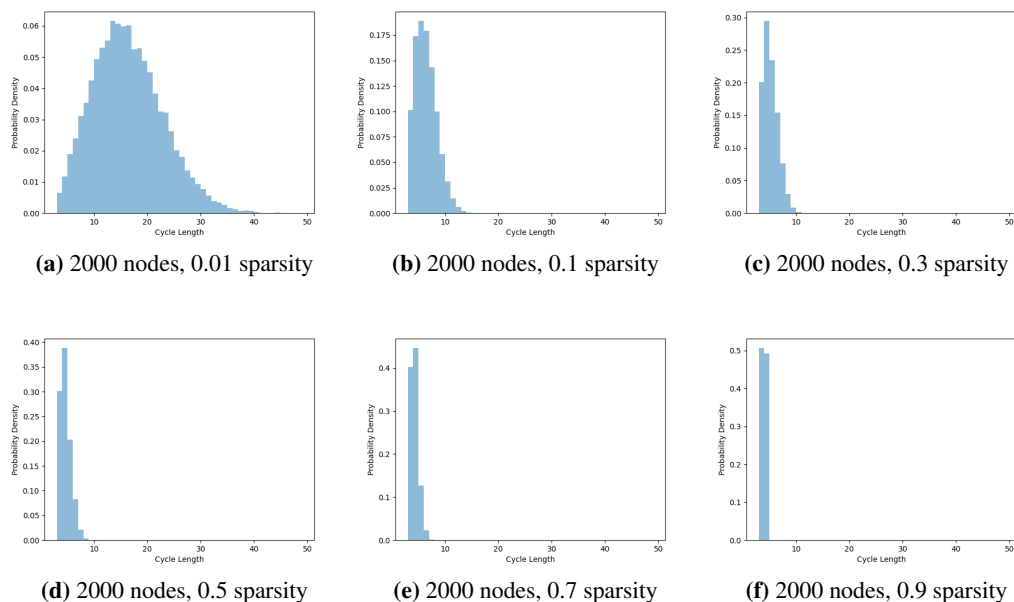


Figure 9: Cycle length histograms of the cycle representatives output by the extended persistence algorithm on sampled Erdos-Renyi graphs

E Cycle Length Distribution of the Cycle Basis found by Extended Persistence Algorithm for Erdos-Renyi Graphs

754
755

756 We perform an experiment to determine the cycle length distribution of cycle representatives output
757 by our algorithm on random Erdos-Renyi graphs. We observe that, as the graph becomes more dense,
758 the distribution of cycles shifts towards very short cycles. We also find that the cycle lengths for most
759 Erdos-Renyi sparsity hyperparameters rarely become very long.

760 For a given node count n , edge count m , and sparsity hyperparameter, $0 \leq s \leq 1$ which we define
761 as the Erdos-Renyi probability for keeping an edge from a clique on n nodes, we sample three
762 Erdos-Renyi graphs. We collect the multiset of $m - n + 1$ cycle lengths in the cycle basis found by
763 the algorithm. This multiset can be visualized as a histogram. Each histogram is a relative frequency
764 mixture of the three cycle length histograms for each graph. See Figure 9 for the histograms we
765 obtained from sampled Erdos-Renyi graphs. Notice that, even for 0.01 sparsity, Erdos-Renyi samples
766 of graphs on 2000 nodes have the average cycle length of 15, which is 0.75% of $n = 2000$.

767 To put this in perspective, assume that we can relate the Erdos-Renyi sparsity s by $\hat{s} := \frac{m}{n^2}$. For
768 the datasets of our experiments, we have $\hat{s} \approx 0.009, 0.0048, 0.39, 0.062, 0.084, 0.0018, 0.032,$ and
769 0.045 for DD, PROTEINS, IMDB-MULTI, MUTAG, PINWHEELS, 2CYCLES, MOLBACE,
770 and MOLBBBP, respectively. The sparsity estimator is in the range of $0.0018 \leq \hat{s} \leq 0.39$, which
771 tells us that most of the cycle lengths found by our algorithm are short.

772 F Rational Hat Function Visualization

773 Figure 10 and Figure 11 visualize the rational hat function for fixed r value and varying x and y
 774 values. Notice the boundedness of the plot as $(x, y) \rightarrow \infty$. For the theory behind the rational hat
 function, see [1].




Figure 10: The function \hat{r} , output sliced at one dimension, as a function of $|(x, y)|_1$ with $r = 0.5$ from Equation 1. The point (x, y) is given by $(x, y) = \mathbf{p} - \mathbf{c}$.




Figure 11: The function \hat{r} , output sliced at one dimension, as a function of $|(x, y)|_1$ with $r = 1.0$ from Equation 1. The point (x, y) is given by $(x, y) = \mathbf{p} - \mathbf{c}$.

775

776 G Datasets and Hyperparameter Information

777 Here are the datasets, both synthetic and real world, used in all of our experiments along with training
 778 hyperparameter information.

779 The barcode vectorization layer, or concatenation of four-rational hat functions, is set to a dimension
 780 of 256. The LSTM used on the explicit cycle representatives was set to a 2-layer bidirectional LSTM
 781 with single channel inputs and 256 dimensional vector representations. Due to the fact that our
 782 algorithm on random Erdos-Renyi graphs rarely encounters long cycles, we set the LSTM layers to a
 783 small number like 2 to avoid overfitting.

Dataset and Hyperparameter Information								
Dataset	Graphs	Classes	Avg. Vertices	Avg. Edges	lr	Node At- trs.(Y/N)	num. layers	Class ratio
DD	1178	2	284.32	715.66	0.01	Yes	2	691/487
PROTEINS	1113	2	39.06	72.82	0.01	Yes	2	663/422
IMDB-MULTI	1500	3	13.00	65.94	0.01	No	2	500/500/500
MUTAG	188	2	17.93	19.79	0.01	Yes	1	63/125
PINWHEELS	100	2	71.934	437.604	0.01	No	2	50/50
2CYCLES	400	2	551.26	551.26	0.01	No	2	200/200
MOLBACE	1513	2	34.09	36.9	0.001	Yes	2	822/691
MOLBBBP	2039	2	24.06	25.95	0.001	Yes	2	479/1560

Table 3: Dataset statistics and training hyperparameters used for all datasets in scoring experiments of Table 1 and Table 2

784 H Implementation Dependencies

785 Our experiments have the following dependencies: python 3.9.1, torch 1.10.1, torch_geometric 2.0.5,
 786 torch_scatter 2.0.9, torch_sparse 0.6.13, scipy 1.6.3, numpy 1.21.2, CUDA 11.2, GCC 7.5.0.

787 **I Visualization of Graph Filtrations**

788 We visualize the filtration functions f_G learned on graphs G for the datasets: IMDB-MULTI, MUTAG, and REDDIT-BINARY. The value of $f_G(v)$ for each $v \in V$ is shown in each figure.

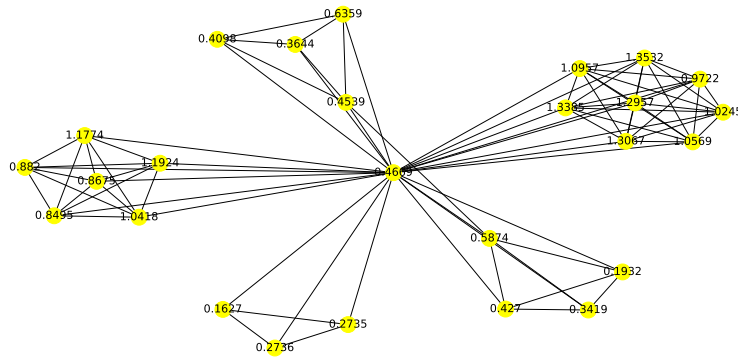


Figure 12: IMDB-MULTI learned filtration function

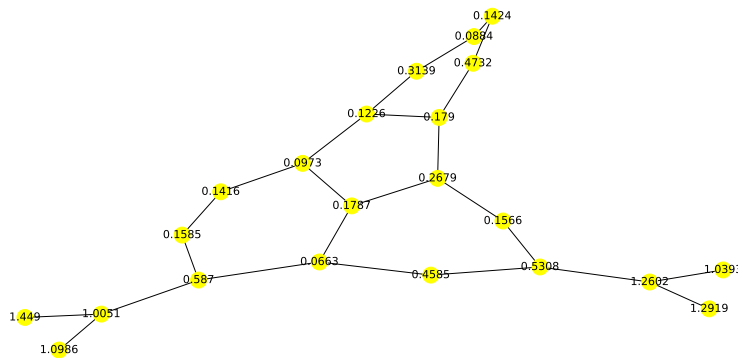


Figure 13: MUTAG learned filtration function

789