

Generalizing Downsampling from Regular Data to Graphs

Davide Bacciu, Alessio Conte, Francesco Landolfi*

Department of Computer Science, Università di Pisa
 davide.bacciu@unipi.it, alessio.conte@unipi.it, francesco.landolfi@phd.unipi.it

Abstract

Downsampling produces coarsened, multi-resolution representations of data and it is used, for example, to produce lossy compression and visualization of large images, reduce computational costs, and boost deep neural representation learning. Unfortunately, due to their lack of a regular structure, there is still no consensus on how downsampling should apply to graphs and linked data. Indeed reductions in graph data are still needed for the goals described above, but reduction mechanisms do not have the same focus on preserving topological structures and properties, while allowing for resolution-tuning, as is the case in regular data downsampling.

In this paper, we take a step in this direction, introducing a unifying interpretation of downsampling in regular and graph data. In particular, we define a graph coarsening mechanism which is a graph-structured counterpart of controllable equispaced coarsening mechanisms in regular data. We prove theoretical guarantees for distortion bounds on path lengths, as well as the ability to preserve key topological properties in the coarsened graphs. We leverage these concepts to define a graph pooling mechanism that we empirically assess in graph classification tasks, providing a greedy algorithm that allows efficient parallel implementation on GPUs, and showing that it compares favorably against pooling methods in literature.

1 Introduction

The concept of information coarsening is fundamental in the adaptive processing of data, as it provides a simple, yet effective, means to obtain multi-resolution representations of information at different levels of abstraction. In large scale problems coarsening also serves to provide computational speed-ups by solving tasks on the reduced representation, ideally with a contained loss in precision with respect to solving the original problem.

Coarsening is key in Convolutional Neural Networks (CNNs, Fukushima 1980; LeCun et al. 1989), where pooling is often used to repeatedly subsample an image to extract visual feature detectors at increasing levels of abstraction (e.g., blobs, edges, parts, objects, etc). Downsampling is also popular in the adaptive processing of timeseries where,

for instance, it is used in clockwork-type Recurrent Neural Networks (Koutník et al. 2014; Carta, Sperduti, and Bacciu 2021) to store information extracted at different frequencies and timescales. More recently, the Graph Convolutional Networks (GCNs, Micheli 2009; Gori, Monfardini, and Scarselli 2005; Bacciu et al. 2020) community popularized graph reduction mechanisms as a structured counterpart of the image pooling mechanism in *classical* CNNs.

The definition of a reduction mechanism that downsamples information at regular intervals between data points (e.g., a sample, a pixel, a timestamped observation, etc) is straightforward when working with images and time series. It can be achieved simply by picking up a data point every k ones, where k is a given reduction factor defining the distance between the sampled points in the original data, possibly aggregating the properties of non-selected point with appropriate functions. The same approach cannot be straightforwardly applied to graphs, which lack regularity and a consistent ordering among their constituent data points, i.e., the nodes. Therefore, defining a well-formed notion of downsampling for graphs becomes non-trivial. The research community has been tackling this issue by a number of approaches, including differentiable clustering of node embeddings (Ying et al. 2018; Bianchi, Grattarola, and Alippi 2020), graph reductions (Shuman, Faraji, and Vandergheynst 2016; Loukas 2019), and node ranking (Cangea et al. 2018; Gao and Ji 2019). Notably, approaches like the latter select important nodes in a graph and simply discard the rest without protecting the linked structure of the network, while reduction methods typically focus on preserving structure without accounting for the role or relevance of nodes involved.

What is yet an open problem is how to define a controllable graph coarsening method, which reduces the size while preserving the overall structure by sampling *representative* yet *evenly* spaced elements, similarly to the approaches discussed above for image and time series reduction.

This paper provides a first approach introducing such a topology-preserving graph coarsening and its use in graph pooling. We provide mechanisms which are the graph equivalent of pooling and striding operators on regular data, accompanying our intuition with formal proofs (in the Supplementary Material) of the equivalence of such operators on graphs which model regular data.

Central to our contribution is the definition of a mechanism

*Corresponding author. Supplementary Material available at <https://arxiv.org/abs/2208.03523>.
 Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

to find a set of nodes that are approximately equally spaced (at distance no less than k) in the original graph. We build on the graph-theoretic concept of Maximal k -Independent Sets (k -MIS), that also comes with the ability to pin-point important nodes in each area of the graph. The selected nodes are then used as vertices of the reduced graph whose topology is defined in such a way that key structural properties of the original graph are well preserved. To this end, we provide theoretical guarantees regarding distance distortions between a graph and its reduction. Additionally, we prove the reduced graph has the same number of connected components as the original. The latter point is particularly relevant for a graph pooling mechanism as it guarantees that the structure is not broken in disconnected fragments, which can hinder the performance of neural message passing in the GCN layers.

Such properties are fundamental to ensure that the original graph is downsampled evenly throughout its structure, preserving distances and sparseness of the key focal points in the graph. By this means, the reduced graph can be used as an accurate fast estimator of the distances between nodes in the original graph, where the amount of compression can be easily regulated through the choice of the k reduction factor.

Concurrently, we borrow from node-ranking methods (Gao and Ji 2019) to produce k -MISs that maximize the total weights associated to the selected nodes, in order to preserve relevant nodes without compromising structure.

In summary, our contributions are the following:

- We introduce a graph coarsening method leveraging k -MIS that is the graph-structured counterpart of equispaced sampling in flat data. We provide a greedy parallel algorithm to efficiently compute the k -MIS reduction, which is well suited to use in GPU accelerators (Section 3).
- We give formal proof of equivalence of our approach to regular downsampling in convolutional neural networks, when applied to diagonal grid graphs (Section 4 and Supplementary Material).
- We prove theoretical guarantees on the distance distortions between a graph and its reduction. We provide also a formal complexity analysis of the introduced algorithms, proving, both theoretically and experimentally, their scalability on large real-world graphs (Section 4 and Supplementary Material).
- We integrate k -MIS reduction both as a pooling layer and as a downsampling operator for GCNs, providing an empirical confirmation of its advantages over literature approaches on graph classification benchmarks (Section 6).

2 Notation and Definitions

We represent a graph G as a pair of disjoint sets (V, E) , where $V = \{1, \dots, n\}$ is its node set and $E \subset V \times V$ its edge set, with $|E| = m$. A graph can also be represented as a symmetric matrix $\mathbf{A} \in \mathbb{R}_+^{n \times n}$, such that $\mathbf{A}_{uv} = \mathbf{A}_{vu}$ is equal to a weight associated to the edge $uv \in E$ or zero if $uv \notin E$. The neighborhood $N(v)$ of v is the set of nodes adjacent to it (denoted $N[v]$ if includes v itself), and the degree $\deg(v)$ of v is defined as the number of its neighbors, i.e., $\deg(v) = |N(v)|$. The *unweighted* distance between two

nodes $u, v \in V$, denoted as $d(u, v)$, is defined as the length of the shortest path between the two nodes. If there is no path between the two nodes, then $d(u, v) = \infty$. The k -hop neighborhood $N_k(v)$ of v ($N_k[v]$ if inclusive) is the set of nodes that can be reached by a path in G of length at most k . The k -th power of a graph G^k is the graph where each node of G is connected to its k -hop neighbors. To avoid confusion, any function may be denoted with a subscript to specify the graph on which is defined (e.g., d_G). An *independent set*, is a set of nodes $S \subseteq V$ such that no two of which are adjacent in G . An independent set is *maximal* if is not a subset of another one in G . A (maximal) *k -independent set* is a (maximal) independent set of G^k .

3 Graph Coarsening with k -MWIS

When dealing with signals, images, or other kinds of Euclidean data, *downsampling* often amounts to *keeping every k -th* data point, where k is a given reduction factor. This means, for a generic discrete n -dimensional Euclidean datum, keeping a subset of its points such that every two of them are *exactly* k points far from each other on every of its dimensions. On graph-structured data, we lose this regularity along with the concept of dimensionality, and hence defining a new notion of downsampling that applies to graph becomes non-trivial.

Here we define a graph coarsening method that, similarly to *classical* downsampling, reduces the size of a graph G by a given “factor”, by finding a set of *almost* evenly spaced nodes within G . These nodes will form the node set of the reduced graph, while its topology will be constructed starting from G in a way in which some of its key properties will be preserved, such as connectivity, or approximated, such as pairwise node distances.

Coarsening algorithm. Given a graph $G = (V, E)$ and a distance k , we want to obtain a coarsen representation of G by first selecting a set of nodes $S \subseteq V$, that we refer to as *centroids*, such that every two centroids are more than k hops distant from each other, and such that no area of the graph remains unsampled; in other words, a *maximal k -independent sets* (k -MIS) of G : this way, each centroid will be more than k hops from every other, while the *maximality* ensures every node of G is within k hops from a centroid.

Any MIS of a graph G^k is a k -MIS of G (Agnarsson, Damaschke, and Halldórsson 2003), thus a k -MIS could be naïvely computed by known MIS algorithms, such as Luby (1985) or Blelloch, Fineman, and Shun (2012), on the k -th power of the adjacency matrix of G . Using this approach will require $O(n^2)$ space since the density of G^k increases rapidly with k , becoming rapidly impractical for real world graphs with millions or billions of nodes. To overcome this problem, we introduce Algorithm 1 that efficiently computes a k -MIS of G without explicitly computing its k -th power. Once the k -MIS $S \subseteq V$ is computed with Algorithm 1, we construct the coarsened graph $H = (S, E')$ as follows:

1. using Algorithm 2, we compute a partition \mathcal{P} of V of size $|S|$, such that
 - (a) every $P \in \mathcal{P}$ contains exactly one centroid and (a subset of) its k -hop neighbors, and

Algorithm 1 Parallel Greedy k -MIS algorithm, adapted from Blelloch, Fineman, and Shun (2012). Given a graph G , a subset of its nodes $U \subseteq V$, and a node ranking π , returns a maximal k -independent set in G , with $k \in \mathbb{N}$.

```

1: function  $k$ -MIS( $G, U, \pi$ )
2:   if  $|U| = 0$  then return  $\emptyset$ 
3:    $\pi_0 \leftarrow \pi$ 
4:   for  $i = 1, \dots, k$  do
5:     for  $v \in U$  do in parallel
6:        $\pi_i(v) \leftarrow \min_{u \in N[v] \cap U} \pi_{i-1}(u)$ 
7:    $S_0 \leftarrow \{v \in U \mid \pi(v) = \pi_k(v)\}$ 
8:   for  $i = 1, \dots, k$  do
9:      $S_i \leftarrow \bigcup_{v \in S_{i-1}} N[v]$ 
10:   $R \leftarrow U \setminus S_k$ 
11:  return  $S_0 \cup k$ -MIS( $G, R, \pi$ )

```

Algorithm 2 Parallel k -MIS partitioning algorithm. Given a graph G , $k \in \mathbb{N}$, and a node ranking π , returns a partition of G .

```

1: function CLUSTER( $G = (V, E), k, \pi$ )
2:   $S \leftarrow k$ -MIS( $G, V, \pi$ )
3:   $\pi_0 \leftarrow \pi$ 
4:  for  $v \in V \setminus S$  do in parallel
5:     $\pi_0(v) \leftarrow +\infty$ 
6:  for  $i = 1, \dots, k$  do
7:    for  $v \in V$  do in parallel
8:       $\pi_i(v) \leftarrow \min_{u \in N[v]} \pi_{i-1}(u)$ 
9:  return  $\{\{u \in V \mid \pi_k(u) = \pi(v)\}\}_{v \in S}$ 

```

- (b) for every node in P there is a centroid in P at distance at most k -hops;
2. for every edge in E we add an edge in E' joining the two nearest centroids in the partitions containing the source and destination nodes. If this generates multiple edges, we coalesce them into a single one, and we aggregate their weights according to a predefined aggregation function (e.g., sum);
3. (*pooling, optional*) in case of weights/labels associated to the nodes, these can also be aggregated according to the partitioning \mathcal{P} .

A detailed discussion of Algorithms 1 and 2 will be provided later in Section 4.

Node ordering. A key property of our k -MIS algorithm (similarly to the one of Blelloch, Fineman, and Shun (2012)) is that it is *deterministic*: given a graph G and a *ranking* of its nodes $\pi : V \rightarrow \{1, \dots, n\}$, that defines the position of the nodes in a given ordering, Algorithm 1 will always produce the same k -MIS, for any $k \geq 0$. This property has some interesting consequences:

- The ranking π can be used to lead Algorithm 1 to greedily include nodes having a higher rank under a given order of importance, such as a centrality measure, a task-dependent relevance, or a (possibly learned) *scoring value*. (Note that the computation of the ranking can impact the complexity of the algorithm.)
- If the ranking can be *uniquely* determined by the nodes themselves (e.g., in function of their attributes or their neighbors), Algorithms 1 and 2 become injective and hence,




Figure 1: (*first column*) Average pooling and (*second to fourth columns*) our method using different ranking and aggregation functions, for varying values of k .

permutation invariant.¹ This can be obtained by ranking the nodes with respect to a score computed by means of a (sufficiently expressive) GCN, as learning injective functions over the nodes in a graph is a problem strictly related to the one of graph isomorphism, a topic that is gaining a lot of traction in the graph learning community (Morris et al. 2019; Xu et al. 2019; Maron et al. 2019; Loukas 2020; Geerts and Reutter 2021; Papp and Wattenhofer 2022).

- A properly chosen ranking can produce a marginally greater total score of the selected nodes with respect to the one that we would get by greedily selecting the top scoring ones. This aspect will be discussed more in detail in Section 4.

We now provide two examples on how we can change the ranking of the nodes to prioritize salient aspects according to a specific preference. Examples are conducted on the graph defined by the first sample of the MNIST dataset (LeCun, Cortes, and Burges 2010), a 28×28 monochromatic image (first row of Fig. 1) where every pixel is connected to the ones in the same pixel row, column or diagonal.

First, we simulate the typical downsampling on images (also known as *average pooling* (Fukushima 1980)), where squared partitions of $p \times p$ pixels are averaged together (first

¹Notice that, in our setting, if $\pi : V \rightarrow \{1, \dots, n\}$ is injective, then it is also bijective and, hence, a permutation.

column of Fig. 1). To do this, we set the ranking π of Algorithm 2 as the *lexicographic* ordering: given (i, j) the coordinate of a pixel, we rank the nodes in decreasing order of $28i + j$. The resulting reduction is in the second column of Fig. 1: averaging intensities of pixels in the same partition produces a coarsened graph which is identical to classical downsampling. Note that this result is partly due to the fact that Algorithm 2 also makes use of π to define the clustering, such that the nodes in a partition have always a lower rank with respect to the centroid in the same partition.

Secondly, we rank nodes in decreasing order of *intensity*, thus prioritizing the pixels (i.e., the nodes) belonging to the drawn digit. Here we show two different results: the first, where we average the lightness and coordinates of the nodes in the same clusters (third column of Fig. 1), and a second one, where we just keep the ones belonging to the nodes in the k -MIS (fourth column). We see that the reduced graphs indeed prioritized the digit against other pixels, producing a coarsened representation where the digit is also remarkably recognizable.

4 Theoretical Analysis and Results

Regular downsampling. Downsampling plays a key role in Convolutional Neural Networks (CNNs, Goodfellow, Bengio, and Courville 2016), where it is adopted, for instance, in *strided* convolutions and *pooling* layers. In strided convolutions, an input tensor (e.g., a time series, an image, or a voxel grid) is reduced by applying the convolved filter every s -th of its entries, along every dimension, while skipping the others. In pooling layers, instead, tensors are reduced by summarizing every p -sided sub-tensors, taken at regular intervals. (More specific reductions are also possible, where distinct intervals are used for every dimension.)

We can show that, on n -dimensional diagonal grid graphs (i.e., grids where nodes are also diagonally adjacent), Algorithms 1 and 2 behave *exactly* as the aforementioned downsampling strategies, if we rank their nodes by their position in lexicographic order. This is of particular interest as the adjacencies in these graphs can represent the receptive fields of a single convolutional layer when applied to a some regular data of the same shape, like images (2-dimensional) or voxel grids (3-dimensional). Specifically, if $G = (V, E)$ is a diagonal grid constructed using the entries of a given tensor as nodes, and π is the ranking of these entries in lexicographic order of their position indices, we have that

1. k -MIS(G, V, π) selects the same entries of a strided convolution with $s = k + 1$,
2. CLUSTER(G, k, π) partitions the tensor as a pooling layer with $p = k + 1$, and
3. the reduced graph obtained by contracting the resulting partition is again a diagonal grid of the same dimensions of their output tensor.

A formal restatement and proof of these properties are provided in the Supplementary Material, while in Fig. 1 we show an example of the equivalence between pooling (first column) and our reduction method (second column).

Connectivity of the reduced graph. For the sake of conciseness, hereafter we denote with $(H, \rho) = \mathcal{R}(G, k)$ the function reducing a graph G by contracting the clusters obtained with Algorithm 2, as described in Section 3. The term $H = (S, E')$ denotes the reduced graph, where S is the k -MIS of G , while $\rho : V \rightarrow S$ is the function mapping every node to the (exactly one) centroid in its cluster. The following results are invariant with respect to the ranking parameter and the aggregation function used to reduce the edges or the nodes.

We follow a simple observation: for every edge in $uv \in E'$ with $u \neq v$, the nodes u and v are within $2k + 1$ hops in G , since two nodes in S are connected in H only if an edge in G crosses their two clusters. This property, combined with the lower bound implicitly defined by the k -MIS, yields the following bounds.

Remark 1. For any $uv \in E(H)$ such that $u \neq v$, we have that $k + 1 \leq d_G(u, v) \leq 2k + 1$.

An example of this property is shown in Fig. 2, where bounds in Remark 1 apply for the Minnesota road network (Davis and Hu 2011) reduced with different values of k . From the above observation, we can obtain the two following properties.

Proposition 1. Let G be a connected graph and $(H, \rho) = \mathcal{R}(G, k)$, with $k \geq 0$. Then, $\forall u, v \in V(G)$,

$$d_H(\rho(u), \rho(v)) \leq d_G(u, v) \leq (2k+1) d_H(\rho(u), \rho(v)) + 2k.$$

Corollary 1. For any $k \geq 0$, G and $H = \mathcal{R}(G, k)$ have the same number of connected components.

The full proofs are provided in the Supplementary Material. Both Proposition 1 and Corollary 1 are fundamental in our proposal of using k -MIS reduction as a pooling method in Graph Neural Networks. In particular: (i) differently from several other pooling techniques (Cangea et al. 2018; Gao and Ji 2019; Knyazev, Taylor, and Amer 2019; Lee, Lee, and Kang 2019; Zhang et al. 2020; Ranjan, Sanyal, and Talukdar 2020; Ma et al. 2020), we can guarantee that the input graph is not divided in multiple components, and that, if applied repeatedly, our method will eventually produce a single representation node for the whole graph; (ii) when training with batches of graphs at a time, our method guarantees also that different graphs are not joined together.

Algorithm discussion and complexity. In order to avoid computing the k -th graph power of a possibly large-scale graph, Algorithm 1 modifies the one by Blelloch, Fineman, and Shun (2012, Algorithm 2, also restated in the Supplementary Material) to compute the k -MIS without explicitly generating every k -hop neighborhood. Given a graph $G = (V, E)$, a subset of its nodes $U \subseteq V$, and a (injective) node mapping $\pi : V \rightarrow \{1, \dots, n\}$ (that we can consider as a ranking of the nodes under a given permutation), Algorithm 1 works as follows:

1. if $U \subseteq V$ is not empty, in Lines 3 to 7 we find the set of nodes S_0 with minimum rank among their k -hop neighbors (i.e., their neighbors in G^k). This is done with k steps of label propagation such that, at each step, every node




Figure 2: Minnesota road network (Davis and Hu 2011) reduced with different values of k . For $k = 0$, the two bounds coincide, as the graph is not reduced at all. For $k = 1$, the real distance covered by an edge is polarized (is either 2 or 3). For greater values of k , the edges' real distance span over all the range $[k + 1, 2k + 1] \cap \mathbb{N}$.

takes the minimum label found within their (1-hop) neighbors. We only propagate labels belonging to nodes still in U ;

2. in Lines 8 to 10 we remove from U all the nodes that are at most k -hops from a node in S_0 (i.e., all their neighbors in G^k). This is also done with k steps of label propagation starting from the nodes in S_0 , where this time the propagated label is a flag signaling that the node shall be removed;
3. finally, the algorithm makes a recursive call in Line 11 using only the remaining nodes. The resulting set is merged with S_0 and returned.

It is easy to see that, if $k = 1$, steps 1 to 3 become exactly Blelloch's algorithm, whereas by taking a general k every step is extended to consider k -hop neighbors of G , thus efficiently emulating Blelloch's algorithm on G^k .

As for complexity, Blelloch, Fineman, and Shun (2012) propose several trade-offs between *work* and *depth* on a *concurrent-read/concurrent-write* PRAM model (CRCW, with minimum priority concurrent write). Here, we consider one version (Algorithm 2 from Blelloch, Fineman, and Shun (2012)) which allows an efficient parallel implementation with $O(m)$ work and $O(\log^3 n)$ depth with high probability (see Blelloch, Fineman, and Shun 2012, Lemma 4.2), and most closely resembles the structure of Algorithm 1. Our algorithm introduces a factor k (compared to the one of Blelloch, Fineman, and Shun (2012)) on the operations performed on lines Lines 3 to 7 and Lines 8 to 10 to compute the k -hop neighborhood. It follows that the work and depth of Algorithm 1 are bounded by k times that of Blelloch's algorithm, i.e., $O(k(n + m))$ work and $O(k \log^3 n)$ depth w.h.p., where an extra $O(n)$ work is needed to generate the additional vector of labels, which is modified every k iterations. Regarding Algorithm 2, after computing the k -MIS, the algorithm performs k steps of label propagation, which add $O(k(n + m))$ work and $O(k \log n)$ depth to the total computation. Total space consumption is $O(n + m)$, comprising input and $O(1)$ label vectors of size $O(n)$.

Proposition 2. *Given a graph G , an integer $k \in \mathbb{N}$, and a random ranking of the nodes π , both Algorithms 1 and 2 can be implemented to run on a CRCW PRAM using $O(k(n + m))$ work, $O(k \log^3 n)$ depth, and $O(n + m)$ space. The depth bound holds w.h.p.*

Bounds on the total weight. In any greedy MIS algorithm, whenever we add a node to the independent set we have to remove all of its neighbors from the graph. Having observed this, a typical heuristic to compute larger-weight independent sets is to select nodes with high weight and low degree (Caro 1979; Wei 1981). Following this intuition, Sakai, Togasaki, and Yamazaki (2003) proposed the following rules: given $\mathbf{x} \in \mathbb{R}_+^n$ a vector of positive weights associated to each node, add to the independent set the node v maximizing either (i) $\mathbf{x}_v / (\deg(v) + 1)$, or (ii) $\mathbf{x}_v / (\sum_{u \in N[v]} \mathbf{x}_u)$. Both rules can be trivially extended to k -hop neighborhoods by computing G^k , which would however require $O(n^2)$ space, unless done sequentially. Parallel computation of the *neighborhood function* $\deg_k(v) = |N_k(v)|$ in limited space can be achieved only by resorting to approximations, e.g. using Monte Carlo methods (Cohen 1997) or approximate sets representations (Palmer, Gibbons, and Faloutsos 2002; Boldi, Rosa, and Vigna 2011), and still this would not extend to approximate rule (ii).

To overcome these limitations, we overestimate the sum of the weights in the k -hop neighborhood of each node, by computing instead $\mathbf{c}_k = (\mathbf{A} + \mathbf{I})^k \mathbf{x} \in \mathbb{R}_+^n$, where $\mathbf{A}, \mathbf{I} \in \{0, 1\}^{n \times n}$ are, respectively, the adjacency and the identity matrices. The matrix $(\mathbf{A} + \mathbf{I})^k \in \mathbb{N}_0^{n \times n}$ represents the number of k -walks (i.e., sequences of adjacent nodes of length k) from every pair of nodes in the graph. Clearly, $[(\mathbf{A} + \mathbf{I})^k]_{uv} \geq 1$ if $v \in N_k[u]$, while the equality holds for every pair of nodes for $k = 1$. When $\mathbf{x} = \mathbf{1}$, \mathbf{c}_k is equal to the k -path centrality (Sade 1989; Borgatti and Everett 2006). Notice that we do not need to compute $(\mathbf{A} + \mathbf{I})^k$ explicitly, as \mathbf{c}_k can be obtained with a sequence of k matrix-vector products, that can be computed in $O(n + m)$ space, $O(k(n + m))$ work and $O(k \log n)$ depth.

In the following, we provide a generalization of the bounds of Sakai, Togasaki, and Yamazaki (2003) when a k -MIS is computed by Algorithm 1 with the ranking defined by rules (i)-(ii) approximated by the k -walk matrix \mathbf{c}_k . We remark that, for $k = 1$, the following theorems are providing the same bounds as the one given by Sakai, Togasaki, and Yamazaki (2003). The full proofs can be found in the Supplementary Material.

Theorem 1. *Let $G = (V, E)$ be a graph, with (unweighted) adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$ and with $\mathbf{x} \in \mathbb{R}_+^n$ representing a vector of positive node weights. Let $k \in \mathbb{N}$ be an*

integer, then define $w : V \rightarrow \mathbb{R}_+$ as

$$w(v) = \frac{\mathbf{x}_v}{[(\mathbf{A} + \mathbf{I})^k \mathbf{1}]_v}, \quad (1)$$

and π_w as the ranking of the nodes in decreasing order of w . Then, k -MIS(G, V, π_w) outputs a maximal k -independent set S such that $\sum_{u \in S} \mathbf{x}_u \geq \sum_{v \in V} w(v)$.

Theorem 2. Let $G = (V, E)$ be a graph, with (unweighted) adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$ and with $\mathbf{x} \in \mathbb{R}_+^n$ representing a vector of positive node weights. Let $k \in \mathbb{N}$ be an integer, then define $w : V \rightarrow \mathbb{R}_+$ as

$$w(v) = \frac{\mathbf{x}_v}{[(\mathbf{A} + \mathbf{I})^k \mathbf{x}]_v}, \quad (2)$$

and π_w as the ranking of the nodes in decreasing order of w . Then, k -MIS(G, V, π_w) outputs a maximal k -independent set S such that $\sum_{u \in S} \mathbf{x}_u \geq \sum_{v \in V} w(v) \cdot \mathbf{x}_v$.

Theorem 3. Let $G = (V, E)$ be a non-empty graph with positive node weights $\mathbf{x} \in \mathbb{R}_+^n$, and let π_w be a ranking defined as in Theorem 1 or 2 for any given $k \in \mathbb{N}$. Then $\sum_{v \in S} \mathbf{x}_v \geq \alpha(G^k) / \Delta_k$, where $S = k$ -MIS(G, V, π_w) and $\Delta_k = \max_{v \in V} [(\mathbf{A} + \mathbf{I})^k \mathbf{1}]_v$.

Recalling that $\alpha(G^k)$ is the optimal solution, Theorem 3 shows that our heuristics guarantee a Δ_k^{-1} approximation. This bound degrades very quickly as the value of k increases, since the number of k -walks may exceed the total number of nodes in the graph. In the Supplementary Material we show that, in practice, the total weight produced by Algorithm 1 is on par with respect to the one obtained using the exact neighborhood function for low values of k . This aspect is of practical value as in general the k values used for graph pooling are on the low-end.

5 Related Works

Maximal k -Independent Sets. Computing a k -MIS can be trivially done in (superlinear) polynomial time and space using matrix powers (Agnarsson, Damaschke, and Halldórsson 2003) and any greedy MIS algorithm (e.g., Luby 1985; Blelloch, Fineman, and Shun 2012). Koerts (2021) proposed a formulation of the problem both as an integer linear program and as a semi-definite program, but still relying on the k -th power of the input graph. Several papers propose efficient algorithms to solve the maximum (weighted or unweighted) k -IS problem on specific classes of graphs (Agnarsson, Greenlaw, and Halldórsson 2000; Agnarsson, Damaschke, and Halldórsson 2003; Eto, Guo, and Miyano 2014; Bhattacharya and Houle 1999; Duckworth and Zito 2003; Pal and Bhattacharjee 1996; Hsiao, Tang, and Chang 1992; Hota, Pal, and Pal 2001; Saha and Pal 2003), which fall beyond the scope of this article. To the best of our knowledge, the only other parallel algorithm for computing a maximal k -independent set was proposed by Bell, Dalton, and Olson (2012) as a generalization of the one of Luby (1985) for $k \geq 1$. This algorithm is essentially the same as Algorithms 1 and 2, but without the ranking argument, making the algorithm non-deterministic, as the nodes are always extracted in a random order.

Graph Coarsening and Reduction. MISs (i.e., with $k = 1$) were adopted as a first sampling step in Barnard and Simon (1994), although their final reduction step may not preserve the connectivity of the graph. Using MIS was also suggested by Shuman, Faraji, and Vandergheynst (2016) as an alternative sampling step for their graph reduction method. The spectral reduction proposed by Loukas (2019, neighborhood variant) does not use sampling as a first reduction step, but sequentially contracts node neighborhoods until a halting condition is reached, performing similar steps to the classical greedy algorithm for maximum-weight independent sets.

Graph Pooling. In a contemporary and independent work, Stanovic, Gaüzère, and Brun (2022) introduced a pooling mechanism based on maximal independent (vertex) sets, named MIVSPool. Their method is analogous to ours, but restricted to the case of 1-MIS, that they compute using the parallel algorithm of Meer (1989). Another related model is EDGEPOOL (Diehl et al. 2019), which computes instead a maximal matching, i.e., a maximal independent set of edges, selecting the edges depending on a learned scoring function. Nouranizadeh et al. (2021) also proposed a pooling method that constructs an independent set maximizing the mutual information between the original and the reduced graph. To do so, the authors leverage on a sequential algorithm with cubic time complexity and also no guarantees that the resulting set is maximal. Apart from a few other cases (Dhillon, Guan, and Kulis 2007; Luzhnica, Day, and Lio' 2019; Ma et al. 2019; Wang et al. 2020; Bacciu and Di Sotto 2019; Bacciu et al. 2021; Bianchi et al. 2020), pooling in Graph Neural Networks (GNNs) usually entails an adaptive approach, typically realized by means of another neural network. These pooling methods can be divided in two types: *dense* and *sparse*. Dense methods, such as DIFFPOOL (Ying et al. 2018), MINCUTPOOL (Bianchi, Grattarola, and Alippi 2020; Bianchi et al. 2020), MEMPOOL (Khasahmadi et al. 2019), STRUCTPOOL (Yuan and Ji 2019), and DMON (Tsitsulin et al. 2022), compute for each node a soft-assignment to a fixed number of clusters defined by a reduction factor $r \in (0, 1)$, thus generating a matrix requiring $O(rn^2)$ space. Sparse methods, such as GPOOL/TOPKPOOL (Gao and Ji 2019; Cangea et al. 2018), SAGPOOL (Lee, Lee, and Kang 2019; Knyazev, Taylor, and Amer 2019), GSAPool (Zhang et al. 2020), ASAPool (Ranjan, Sanyal, and Talukdar 2020), PANPOOL (Ma et al. 2020), IPOOL (Gao et al. 2021), and TAGPOOL (Gao, Liu, and Ji 2021), instead, compute a score for each node (requiring $O(n)$ space), and reduce the graph by keeping only the top $\lceil rn \rceil$ scoring ones and dropping the rest. Although scalable, these methods provide no theoretical guarantees regarding the preservation of connectivity of the reduced graph, as the $n - \lceil rn \rceil$ dropped nodes may disconnect the graph.

6 Experimental Analysis

Table 1 summarizes the average classification accuracy obtained on selected classification benchmarks using the same underlying Graph Neural Networks (GNNs) and different kinds of pooling mechanisms. For classification tasks, we chose those datasets having the highest num-

Model	DD	REDDIT-B	REDDIT-5K	REDDIT-12K	GITHUB
BASELINE	75.51 ± 1.07	78.40 ± 8.68	48.32 ± 2.38	45.04 ± 6.63	69.89 ± 0.28
BDO	76.69 ± 1.79	85.63 ± 1.43	45.95 ± 5.49	41.89 ± 7.14	65.64 ± 0.90
GRACLUS	75.17 ± 2.11	84.05 ± 5.81	43.22 ± 12.24	43.08 ± 9.32	67.64 ± 0.57
EDGEPOOL	74.70 ± 1.57	85.98 ± 1.57	52.44 ± 1.11	47.58 ± 0.78	68.72 ± 0.52
TOPKPOOL	74.92 ± 2.03	81.10 ± 3.82	45.28 ± 3.88	38.55 ± 2.35	65.93 ± 0.45
SAGPOOL	73.26 ± 2.26	84.90 ± 3.94	46.29 ± 5.61	42.30 ± 3.70	64.29 ± 5.70
ASAPool	73.73 ± 2.18	78.37 ± 5.22	39.53 ± 7.76	39.14 ± 3.58	66.98 ± 0.96
PANPOOL	73.26 ± 1.94	77.44 ± 4.95	46.04 ± 3.78	40.97 ± 3.02	62.48 ± 2.84
<i>k</i> -MIS (<i>strided</i>)	76.44 ± 1.50	86.32 ± 1.90	54.30 ± 0.53	46.06 ± 0.58	67.87 ± 0.48
<i>k</i> -MIS (<i>max-pool</i>)	76.91 ± 1.06	87.57 ± 1.96	53.44 ± 1.52	47.51 ± 0.99	68.24 ± 0.94
<i>k</i> -MIS (<i>mean-pool</i>)	73.56 ± 1.19	86.98 ± 1.13	54.00 ± 0.94	46.73 ± 0.85	68.60 ± 0.67

Table 1: Classification accuracy on selected benchmark datasets (*mean ± std*)

ber of nodes from in the *TUDataset* (Morris et al. 2020) collection (i.e., DD (Dobson and Doig 2003), GITHUB-STARGAZERS (Rozemberczki, Kiss, and Sarkar 2020), REDDIT-BINARY and -MULTI-5K/12K (Yanardag and Vishwanathan 2015)), where pooling layers may prove more useful. All datasets were divided in *training* (70%), *validation* (10%), and *test* (20%) sets using a randomized stratified split with fixed seed. All models have the same general architecture: 3 GNNs (optionally) interleaved by 2 layers of pooling, a global pooling method (*sum* and *max*), and a final MLP with dropout (Srivastava et al. 2014) as classifier. All models were trained using Adam optimizer (Kingma and Ba 2017). We performed a model selection using the training and validation split, and then we computed the average test set classification accuracy obtained by the best configuration, on 10 inference runs using different seed values. The hyperparameter concerning the reduction factor (*k* in our case, or *r* for other methods) has been chosen among the other parameters during the model selection phase. All models have been implemented and trained using PyTorch (Paszke et al. 2019) and PyTorch Geometric (Fey and Lenssen 2019). A detailed description of the datasets, models, and experimental settings are provided in the Supplementary Material, together with additional experiments regarding the controlled scaling and efficiency of our method, showing that it can reduce graphs with over 100 million edges in less than 10 seconds.

We compared our reduction method against different kinds of pooling layers readily available on the PyG library (we avoided DIFFPOOL-like dense methods as they do not scale well on the selected datasets), and also against our own method using random rankings of the nodes, as done in the aggregation scheme of Bell, Dalton, and Olson (BDO, 2012). This method, along with the baseline (with no pooling) and GRACLUS, are the only compared architectures that require no additional parameters. For our method, the node scoring function is computed by means of a sigmoid-activated linear layer having as input the features of the nodes. As in the other parametric methods, the feature vectors of the nodes are multiplied by the final score beforehand, to make the scoring function end-to-end trainable. The computed scores constitute the node weights and, as described in Section 4, the resulting ranking is obtained according to Eq. (2). The reduction is performed in two settings: *strided*, in which we

perform no aggregation, and *pool*, in which we aggregate the feature vectors in each partition using, respectively, *max* and *mean* aggregations.

Looking at Table 1 (where the top two accuracy scores for each dataset are in boldface) it is immediately evident how the proposed *k*-MIS-based approaches obtain consistently high accuracy, suggesting that the *evenly-spaced* centroid selection is indeed able to capture essential properties of each graph. On the other hand, the random permutation variant of Bell, Dalton, and Olson (2012), seem to overall perform worse than the other *k*-MIS-based strategies, while still obtaining a considerable result on DD and REDDIT-B. This suggests that exploiting the ranking function of Algorithm 1 to select relevant nodes is indeed able to improve the representativeness of the downsampled graph. It is also particularly noteworthy how one of the best performing model, EDGEPOOL, is also the only other parametric pooling method to preserve the connectivity of the graph, as its reduction step consists of contracting a maximal matching. This highlights the importance of preserving the connectivity of the network when performing pooling in GNNs, while also suggesting that evenly-spaced reductions can benefit graph representation learning tasks. Finally, we observe a remarkable performance of the baseline algorithm (no pooling) on the GITHUB dataset: we may speculate that the graphs are simple enough to not require pooling, yet at the same time *k*-MIS approaches obtains competitive accuracy, suggesting it is a reliable and versatile choice.

7 Conclusions

We introduced a new general graph coarsening approach that aims to preserve fundamental topological properties of the original graph, acting like a structured counterpart of down-sampling methods for regular data. The coarsening reduction can be regulated by the parameter *k*, going from the original graph, when *k* = 0, to up to a single node as *k* approaches the graph’s diameter, shrinking graphs uniformly as pairwise distances maintain a stretch controlled by *k*. Furthermore, we showed how this parameter generalizes to the pooling and stride intervals when applied to diagonal grid graphs.

The algorithm is designed to provide such guarantees while at the same time allowing a scalable parallel implementation,

which processes graphs with up to 100 million edges in just a few seconds on a single GPU.

The empirical analysis provided evidence of effectiveness of our k -MIS pooling in several graph classification benchmarks, showing superior performance with respect to related parametric and non-parametric methods from the literature.

This approach fills a methodological gap between reduction techniques for structured data and their rigorous counterparts on regular data. Given its generality and scalability, it has potential of positively impacting a plethora of computationally-intense applications for large scale networks, such as graph visualization, 3D mesh simplification, and classification.

Acknowledgments

This research was partially supported by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215. We would also like to thank Federico Poloni and Federico Errica for their most useful suggestions on earlier versions of this paper.

References

- Agnarsson, G.; Damaschke, P.; and Halldórsson, M. M. 2003. Powers of geometric intersection graphs and dispersion algorithms. *Discrete Applied Mathematics*, 132(1): 3–16.
- Agnarsson, G.; Greenlaw, R.; and Halldórsson, M. M. 2000. On Powers of Chordal Graphs And Their Colorings. In *Powers of Chordal Graphs and Their Colorings, Congressus Numerantium*, 144: 41–65.
- Bacciu, D.; Conte, A.; Grossi, R.; Landolfi, F.; and Marino, A. 2021. K -plex cover pooling for graph neural networks. *Data Mining and Knowledge Discovery*, 35(5): 1–21.
- Bacciu, D.; and Di Sotto, L. 2019. A Non-negative Factorization Approach to Node Pooling in Graph Convolutional Neural Networks. In Alviano, M.; Greco, G.; and Scarcello, F., eds., *AI*IA 2019 – Advances in Artificial Intelligence*, Lecture Notes in Computer Science, 294–306. Cham: Springer International Publishing. ISBN 978-3-030-35166-3.
- Bacciu, D.; Errica, F.; Micheli, A.; and Podda, M. 2020. A gentle introduction to deep learning for graphs. *Neural Networks*, 129: 203–221.
- Barnard, S. T.; and Simon, H. D. 1994. Fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems. *Concurrency: Practice and Experience*, 6(2): 101–117. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.4330060203>.
- Bell, N.; Dalton, S.; and Olson, L. N. 2012. Exposing Fine-Grained Parallelism in Algebraic Multigrid Methods. *SIAM Journal on Scientific Computing*, 34(4): C123–C152. Publisher: Society for Industrial and Applied Mathematics.
- Bhattacharya, B. K.; and Houle, M. E. 1999. Generalized Maximum Independent Sets for Trees in Subquadratic Time. In *Algorithms and Computation*, 435–445. Springer, Berlin, Heidelberg.
- Bianchi, F. M.; Grattarola, D.; and Alippi, C. 2020. Spectral Clustering with Graph Neural Networks for Graph Pooling. *International Conference on Machine Learning*, 1.
- Bianchi, F. M.; Grattarola, D.; Livi, L.; and Alippi, C. 2020. Hierarchical Representation Learning in Graph Neural Networks With Node Decimation Pooling. *IEEE Transactions on Neural Networks and Learning Systems*, 1–13. Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- Blelloch, G. E.; Fineman, J. T.; and Shun, J. 2012. Greedy sequential maximal independent set and matching are parallel on average. In *Proceedings of the twenty-fourth annual ACM symposium on Parallelism in algorithms and architectures*, SPAA '12, 308–317. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-1213-4.
- Boldi, P.; Rosa, M.; and Vigna, S. 2011. HyperANF: approximating the neighbourhood function of very large graphs on a budget. In *Proceedings of the 20th international conference on World wide web*, WWW '11, 625–634. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-0632-4.
- Borgatti, S. P.; and Everett, M. G. 2006. A Graph-theoretic perspective on centrality. *Social Networks*, 28(4): 466–484.
- Cangea, C.; Veličković, P.; Jovanović, N.; Kipf, T.; and Liò, P. 2018. Towards Sparse Hierarchical Graph Classifiers. *arXiv:1811.01287 [cs, stat]*. ArXiv: 1811.01287.
- Caro, Y. 1979. New results on the independence number. Technical report, Technical Report, Tel-Aviv University.
- Carta, A.; Sperduti, A.; and Bacciu, D. 2021. Incremental Training of a Recurrent Neural Network Exploiting a Multi-scale Dynamic Memory. In Hutter, F.; Kersting, K.; Lijffijt, J.; and Valera, I., eds., *Machine Learning and Knowledge Discovery in Databases*, Lecture Notes in Computer Science, 677–693. Cham: Springer International Publishing. ISBN 978-3-030-67658-2.
- Cohen, E. 1997. Size-Estimation Framework with Applications to Transitive Closure and Reachability. *Journal of Computer and System Sciences*, 55(3): 441–453.
- Davis, T. A.; and Hu, Y. 2011. The university of Florida sparse matrix collection. *ACM Transactions on Mathematical Software*, 38(1): 1:1–1:25.
- Dhillon, I. S.; Guan, Y.; and Kulis, B. 2007. Weighted Graph Cuts without Eigenvectors A Multilevel Approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11): 1944–1957.
- Diehl, F.; Brunner, T.; Le, M. T.; and Knoll, A. 2019. Towards Graph Pooling by Edge Contraction. In *ICML 2019 Workshop on Learning and Reasoning with Graph-Structured Data*.
- Dobson, P. D.; and Doig, A. J. 2003. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of Molecular Biology*, 330(4): 771–783.
- Duckworth, W.; and Zito, M. 2003. Large 2-Independent Sets of Regular Graphs. *Electronic Notes in Theoretical Computer Science*, 78: 223–235.
- Eto, H.; Guo, F.; and Miyano, E. 2014. Distance- $\$d\$$ independent set problems for bipartite and chordal graphs. *Journal of Combinatorial Optimization*, 27(1): 88–99.
- Fey, M.; and Lenssen, J. E. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Fukushima, K. 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4): 193–202.
- Gao, H.; and Ji, S. 2019. Graph U-Nets. In *International Conference on Machine Learning*, 2083–2092. ISSN: 1938-7228 Section: Machine Learning.
- Gao, H.; Liu, Y.; and Ji, S. 2021. Topology-Aware Graph Pooling Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12): 4512–4518. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Gao, X.; Dai, W.; Li, C.; Xiong, H.; and Frossard, P. 2021. iPool–Information-Based Pooling in Hierarchical Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning*

- Systems*, 1–13. Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- Geerts, F.; and Reutter, J. L. 2021. Expressiveness and Approximation Properties of Graph Neural Networks. In *International Conference on Learning Representations*.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press.
- Gori, M.; Monfardini, G.; and Scarselli, F. 2005. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, 729–734 vol. 2. ISSN: 2161-4407.
- Hota, M.; Pal, M.; and Pal, T. K. 2001. An Efficient Algorithm for Finding a Maximum Weight k -Independent Set on Trapezoid Graphs. *Computational Optimization and Applications*, 18(1): 49–62.
- Hsiao, J. Y.; Tang, C. Y.; and Chang, R. S. 1992. An efficient algorithm for finding a maximum weight 2-independent set on interval graphs. *Information Processing Letters*, 43(5): 229–235.
- Khasahmadi, A. H.; Hassani, K.; Moradi, P.; Lee, L.; and Morris, Q. 2019. Memory-Based Graph Networks. In *International Conference on Learning Representations*.
- Kingma, D. P.; and Ba, J. 2017. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*. ArXiv: 1412.6980.
- Knyazev, B.; Taylor, G. W.; and Amer, M. 2019. Understanding Attention and Generalization in Graph Neural Networks. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; Alché-Buc, F. d.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32*, 4202–4212. Curran Associates, Inc.
- Koerts, H. O. 2021. *On the k -Independent Set Problem*. Bachelor Thesis, Eindhoven University of Technology.
- Koutník, J.; Greff, K.; Gomez, F.; and Schmidhuber, J. 2014. A clockwork RNN. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14, II-1863-II-1871*. Beijing, China: JMLR.org.
- LeCun, Y.; Boser, B.; Denker, J. S.; Henderson, D.; Howard, R. E.; Hubbard, W.; and Jackel, L. D. 1989. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, 1(4): 541–551. Publisher: MIT Press.
- LeCun, Y.; Cortes, C.; and Burges, C. 2010. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2.
- Lee, J.; Lee, I.; and Kang, J. 2019. Self-Attention Graph Pooling. In *International Conference on Machine Learning, 3734–3743*. ISSN: 1938-7228 Section: Machine Learning.
- Loukas, A. 2019. Graph Reduction with Spectral and Cut Guarantees. *Journal of Machine Learning Research*, 20(116): 1–42.
- Loukas, A. 2020. How hard is to distinguish graphs with graph neural networks? *Advances in Neural Information Processing Systems*, 33.
- Luby, M. 1985. A simple parallel algorithm for the maximal independent set problem. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing, STOC '85, 1–10*. Providence, Rhode Island, USA: Association for Computing Machinery. ISBN 978-0-89791-151-1.
- Luzhnica, E.; Day, B.; and Lio' , P. 2019. Clique pooling for graph classification. *arXiv:1904.00374 [cs, stat]*. ArXiv: 1904.00374.
- Ma, Y.; Wang, S.; Aggarwal, C. C.; and Tang, J. 2019. Graph Convolutional Networks with EigenPooling. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19, 723–731*. Anchorage, AK, USA: Association for Computing Machinery. ISBN 978-1-4503-6201-6.
- Ma, Z.; Xuan, J.; Wang, Y. G.; Li, M.; and Lio, P. 2020. Path Integral Based Convolution and Pooling for Graph Neural Networks. In *Advances in Neural Information Processing Systems (NeurIPS) 33*. Curran Associates, Inc.
- Maron, H.; Ben-Hamu, H.; Serviansky, H.; and Lipman, Y. 2019. Provably Powerful Graph Networks. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; Alché-Buc, F. d.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32*, 2156–2167. Curran Associates, Inc.
- Meer, P. 1989. Stochastic image pyramids. *Computer Vision, Graphics, and Image Processing*, 45(3): 269–294.
- Micheli, A. 2009. Neural Network for Graphs: A Contextual Constructive Approach. *IEEE Transactions on Neural Networks*, 20(3): 498–511. Conference Name: IEEE Transactions on Neural Networks.
- Morris, C.; Kriege, N. M.; Bause, F.; Kersting, K.; Mutzel, P.; and Neumann, M. 2020. TUDataset: A collection of benchmark datasets for learning with graphs. In *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*. ArXiv: 2007.08663.
- Morris, C.; Ritzert, M.; Fey, M.; Hamilton, W. L.; Lenssen, J. E.; Rattan, G.; and Grohe, M. 2019. Weisfeiler and Leman Go Neural: Higher-Order Graph Neural Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 4602–4609. Number: 01.
- Nouranizadeh, A.; Matinkia, M.; Rahmati, M.; and Safabakhsh, R. 2021. Maximum Entropy Weighted Independent Set Pooling for Graph Neural Networks. ArXiv:2107.01410 [cs, math].
- Pal, M.; and Bhattacharjee, G. P. 1996. A sequential algorithm for finding a maximum weight K -independent set on interval graphs. *International Journal of Computer Mathematics*, 60(3-4): 205–214. Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/00207169608804486>.
- Palmer, C. R.; Gibbons, P. B.; and Faloutsos, C. 2002. ANF: a fast and scalable tool for data mining in massive graphs. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '02, 81–90*. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-58113-567-1.
- Papp, P. A.; and Wattenhofer, R. 2022. A Theoretical Comparison of Graph Neural Network Extensions. ArXiv:2201.12884 [cs].
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. 8026–8037.
- Ranjan, E.; Sanyal, S.; and Talukdar, P. 2020. ASAP: Adaptive Structure Aware Pooling for Learning Hierarchical Graph Representations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 5470–5477. Number: 04.
- Rozemberczki, B.; Kiss, O.; and Sarkar, R. 2020. Karate Club: An API Oriented Open-Source Python Framework for Unsupervised Learning on Graphs. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20, 3125–3132*. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-6859-9.
- Sade, D. S. 1989. Sociometrics of Macaca Mulatta III: n -path centrality in grooming networks. *Social Networks*, 11(3): 273–292.
- Saha, A.; and Pal, M. 2003. Maximum weight k -independent set problem on permutation graphs. *International Journal of Computer Mathematics*, 80(12): 1477–1487.

- Sakai, S.; Togasaki, M.; and Yamazaki, K. 2003. A note on greedy algorithms for the maximum weighted independent set problem. *Discrete Applied Mathematics*, 126(2): 313–322.
- Shuman, D. I.; Faraji, M. J.; and Vandergheynst, P. 2016. A Multiscale Pyramid Transform for Graph Signals. *IEEE Transactions on Signal Processing*, 64(8): 2119–2134. Conference Name: IEEE Transactions on Signal Processing.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56): 1929–1958.
- Stanovic, S.; Gaüzère, B.; and Brun, L. 2022. Maximal Independent Vertex Set Applied to Graph Pooling. In Krzyzak, A.; Suen, C. Y.; Torsello, A.; and Nobile, N., eds., *Structural, Syntactic, and Statistical Pattern Recognition*, Lecture Notes in Computer Science, 11–21. Cham: Springer International Publishing. ISBN 978-3-031-23028-8.
- Tsitsulin, A.; Palowitch, J.; Perozzi, B.; and Müller, E. 2022. Graph Clustering with Graph Neural Networks. ArXiv:2006.16904 [cs, stat].
- Wang, Y.; Li, M.; Ma, Z.; Montufar, G.; Zhuang, X.; and Fan, Y. 2020. Haar Graph Pooling. *Proceedings of the International Conference on Machine Learning*, 1.
- Wei, V. 1981. A lower bound on the stability number of a simple graph. Technical report, Bell Laboratories Technical Memorandum.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.
- Yanardag, P.; and Vishwanathan, S. 2015. Deep Graph Kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, 1365–1374. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-3664-2.
- Ying, Z.; You, J.; Morris, C.; Ren, X.; Hamilton, W.; and Leskovec, J. 2018. Hierarchical Graph Representation Learning with Differentiable Pooling. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Yuan, H.; and Ji, S. 2019. StructPool: Structured Graph Pooling via Conditional Random Fields. In *International Conference on Learning Representations*.
- Zhang, L.; Wang, X.; Li, H.; Zhu, G.; Shen, P.; Li, P.; Lu, X.; Shah, S. A. A.; and Bennamoun, M. 2020. Structure-Feature based Graph Self-adaptive Pooling. In *Proceedings of The Web Conference 2020*, WWW '20, 3098–3104. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-7023-3.