

Simplifying Clustering with Graph Neural Networks

Filippo Maria Bianchi

UiT the Arctic University of Norway and NORCE Norwegian Research Centre
filippo.m.bianchi@uit.no

Abstract

The objective functions used in spectral clustering are generally composed of two terms: i) a term that minimizes the local quadratic variation of the cluster assignments on the graph and; ii) a term that balances the clustering partition and helps avoiding degenerate solutions. This paper shows that a graph neural network, equipped with suitable message passing layers, can generate good cluster assignments by optimizing only a balancing term. Results on attributed graph datasets show the effectiveness of the proposed approach in terms of clustering performance and computation time.

1 Introduction

Traditional clustering techniques partition the data directly in the input space by drawing regular boundaries to separate the clusters. This makes them unsuitable to handle complex data structures, such as images or time series, which lie in high-dimensional spaces where the relationships between samples are highly non-linear. Deep learning techniques allow to transform data samples into suitable representations, which can be partitioned into meaningful clusters [16]. Remarkably, end-to-end deep learning frameworks can directly map complex data directly into their cluster assignments [8].

Of particular interest for this work, are those data characterized by relationships, or interactions, among samples that are described by a graph. Graph Neural Networks (GNNs) are deep learning architectures specifically designed to process and make inference on such data [6]. Recently, GNNs have been adopted to cluster the nodes of an attributed graph based on their features and the graph topology. Inspired by spectral clustering algorithms, such GNNs optimize an unsupervised

loss composed of two terms: the first ensures that connected nodes are assigned to the same cluster; the second is a balancing term, which prevents degenerate solutions both by encouraging the samples to be assigned to only one cluster and the clusters to have similar size [1, 13].

In this work, I considerably simplify the clustering objective optimized by the previous GNN models by introducing a minimalist unsupervised loss, which consists only of a balancing term. The proposed loss is used to train a GNN composed of standard message passing layers that operate on a particular connectivity matrix. The empirical evaluation shows that the proposed model significantly reduces the computational complexity, while maintaining competitive clustering performance.

2 Background

Let a graph be represented by a tuple $G = \{\mathcal{V}, \mathcal{E}\}$, with node set \mathcal{V} and edge set \mathcal{E} . Let $|\mathcal{V}| = N$ and $|\mathcal{E}| = E$ be the number of nodes and edges, respectively. Each node i is associated with a feature vector $\mathbf{x}_i \in \mathbb{R}^F$. A graph is conveniently described by its adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and the node features matrix $\mathbf{X} \in \mathbb{R}^{N \times F}$.

2.1 Spectral Clustering

Graph clustering aims at partitioning the nodes in K subsets, so that the similarity between nodes in the same subset is maximized. The most famous graph clustering approach is spectral clustering, which relies on the k -way mincut objective to find a partition that minimizes the volume of edges crossing different clusters [14]. To avoid degenerate solutions, the objective function includes a balancing term that penalizes partitions where clusters have very unequal sizes. Specifically, the balanced

K -cut objective can be defined as a ratio of two set functions:

$$\min_{C_1, \dots, C_K} \sum_{k=1}^K \frac{\text{cut}(C_k, \bar{C}_k)}{\hat{B}(C_k)}, \quad (1)$$

where $\hat{B}(\cdot)$ is a set function that balances the size of the clusters in the partition. Depending on the choice of $\hat{B}(\cdot)$, one obtains different cuts, such as ratio cut and normalized cut [7, 14]. The numerator of (1) can be expressed in matrix form. To see that, first one rewrites $\text{cut}(C_k, \bar{C}_k)$ as

$$\sum_{i \in C_k, j \in \bar{C}_k} a_{ij}(1 - z_i z_j) \quad \text{s.t.} \quad z_i, z_j \in \{-1, 1\},$$

Then,

$$\begin{aligned} \sum_{i,j} a_{ij}(1 - z_i z_j) &= \sum_{i,j} a_{ij} \left(\frac{z_i^2 + z_j^2}{2} - z_i z_j \right) \\ &= \frac{1}{2} \sum_i \left[\sum_j a_{ij} \right] z_i^2 + \frac{1}{2} \sum_j \left[\sum_i a_{ij} \right] z_j^2 \dots \\ &\quad - \sum_{i,j} a_{ij} z_i z_j \\ &= \frac{1}{2} \sum_i d_{ii} z_i^2 + \frac{1}{2} \sum_j d_{jj} z_j^2 - \mathbf{z}^T \mathbf{A} \mathbf{z} \\ &= \mathbf{z}^T \mathbf{D} \mathbf{z} - \mathbf{z}^T \mathbf{A} \mathbf{z} = \mathbf{z}^T \mathbf{L} \mathbf{z}, \end{aligned}$$

where \mathbf{L} is the graph Laplacian. The relaxation done in spectral clustering to handle the discrete optimization problem is:

$$\min_{\mathbf{z}_k \in \{-1, 1\}^N} \sum_{k=1}^K \frac{\mathbf{z}_k^T \mathbf{L} \mathbf{z}_k}{\hat{B}(C_k)} \rightarrow \min_{\mathbf{s}_k \in \mathbb{R}^N} \sum_{k=1}^K \frac{\mathbf{s}_k^T \mathbf{L} \mathbf{s}_k}{B(C_k)}, \quad (2)$$

where $B(C_k)$ is the continuous counterpart of $\hat{B}(C_k)$.

Remark 1: Besides the graph Laplacian, other operators matching the sparsity pattern of the adjacency matrix can be used in problem (2). One of such operators is the symmetrically normalized Laplacian, $\mathbf{L}_s = \mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$, which generally yields a different partition as the edges to be cut are weighted by the degree of their end nodes.

Remark 2: The term $\mathbf{s}^T \mathbf{L} \mathbf{s}$ in (2) measures the *local quadratic variation* (LQV) of \mathbf{s} on the graph,

which is the quadratic variation of \mathbf{s} across adjacent vertices. Laplacian smoothing minimizes LQV, by making similar the elements s_i and s_j if nodes i and j are connected. When using \mathbf{L}_s , the LQV is:

$$\mathbf{s}^T \mathbf{L}_s \mathbf{s} = \frac{1}{2} \sum_{(i,j) \in \mathcal{E}} a_{i,j} \left(\frac{s_i}{\sqrt{d_i}} - \frac{s_j}{\sqrt{d_j}} \right)^2. \quad (3)$$

2.2 Graph Neural Networks

The main building block of a GNN is the message passing (MP) layer that, first, combines the node features with those of the neighbors on the graph. Then, the aggregated features are mapped into a new representation by applying an affine transformation and a nonlinearity [3]. A basic MP layer is implemented as follows:

$$\mathbf{X}^{(l+1)} = \text{MP}(\mathbf{X}^{(l)}, \tilde{\mathbf{A}}) = \sigma(\tilde{\mathbf{A}} \mathbf{X}^{(l)} \Theta_l), \quad (4)$$

where $\mathbf{X}^{(l)}$ and $\mathbf{X}^{(l+1)}$ are, respectively, the input and output node features of the l -th MP layer, $\tilde{\mathbf{A}}$ is an operator matching the sparsity pattern of \mathbf{A} , σ is a nonlinear activation function, and Θ_l are trainable parameters.

2.3 Clustering with GNNs

MinCutPool [1] is a GNN layer that computes soft cluster assignments as:

$$\mathbf{S} = \text{softmax}(\text{MLP}(\bar{\mathbf{X}}, \Theta_{\text{MLP}})) \in \mathbb{R}^{N \times K}, \quad (5)$$

where K is the number of clusters, $\bar{\mathbf{X}}$ are node features generated by a stack of one or more MP layers, and $\text{MLP}(\cdot)$ denotes a multi-layer perceptron with trainable parameters Θ_{MLP} . The **softmax** function ensures that \mathbf{S} is a proper cluster assignment matrix, since $\mathbf{S} \mathbf{1} = \mathbf{1}$ and $0 \leq s_{i,j} \leq 1$.

To learn the cluster assignments, MinCutPool optimizes the following unsupervised loss:

$$\mathcal{L}_{mc} = \underbrace{-\frac{\text{Tr}(\mathbf{S}^T \tilde{\mathbf{A}} \mathbf{S})}{\text{Tr}(\mathbf{S}^T \tilde{\mathbf{D}} \mathbf{S})}}_{\mathcal{L}_q} + \underbrace{\left\| \frac{\mathbf{S}^T \mathbf{S}}{\|\mathbf{S}^T \mathbf{S}\|_F} - \frac{\mathbf{I}_K}{\sqrt{K}} \right\|_F}_{\mathcal{L}_b}, \quad (6)$$

where $\tilde{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$ and $\tilde{\mathbf{D}}$ is the degree matrix of $\tilde{\mathbf{A}}$. The first term, \mathcal{L}_q , minimizes the LQV, while \mathcal{L}_b is a balancing term that helps prevent degenerate solutions. Compared to problem (2), the LQV

and the balancing terms are summed rather than taking their ratio. This helps both to prevent numerical issues when \mathcal{L}_b gets too small and to keep \mathcal{L}_{mc} in a controlled range, which is desirable when the GNN must also minimize other losses.

Similarly to MinCutPool, DMoN [13] optimizes a loss composed of an LQV and a balancing term:

$$\mathcal{L}_{dm} = -\underbrace{\frac{\text{Tr}(\mathbf{S}^T \tilde{\mathbf{A}} \mathbf{S})}{2E}}_{\mathcal{L}_m} + \underbrace{\frac{\sqrt{K}}{N} \left\| \sum_i \mathbf{S}_i^T \right\|_F}_{\mathcal{L}_r} - 1, \quad (7)$$

where $\tilde{\mathbf{A}} = \mathbf{A} - \mathbf{d}^T \mathbf{d}$ and \mathbf{d} is the degree vector of \mathbf{A} . The term \mathcal{L}_m pushes strongly connected components to the same cluster, while \mathcal{L}_r is a regularization term that penalizes the degenerate solutions.

3 Proposed approach

The cluster assignments \mathbf{S} , computed as in (5), can be optimized by minimizing:

$$\mathcal{L} = -\text{Tr}(\sqrt{\mathbf{S}^T \mathbf{S}}) \quad (8)$$

The proposed loss simplifies \mathcal{L}_{mc} and \mathcal{L}_{dm} considerably as it consists only of a balancing term. Such a simplification offers the following advantages:

- The computational complexity is reduced, as less operations are needed to compute \mathcal{L} .
- Fewer competing terms in the loss can ease the training and speed-up the convergence.
- There are no ratios in \mathcal{L} , which could cause numerical instability during training.

Despite its simplicity, the proposed loss can still yield an optimal clustering assignment. The key insights that motivated its design are presented in the following.

3.1 Removal of the LQV term

The absence of the LQV term in the loss is compensated by the presence of the MP layers that generate the features $\tilde{\mathbf{X}}$ used to compute the cluster

assignments \mathbf{S} in (5). In particular, consider the following MP layer:

$$\mathbf{X}^{(l+1)} = \sigma \left(\left[\mathbf{I} - \delta (\mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}) \right] \mathbf{X}^{(l)} \Theta_l \right) \quad (9)$$

where $\tilde{\mathbf{A}} = \mathbf{I} - \delta (\mathbf{I} - \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}})$ is an operator matching the sparsity pattern of the graph and δ is an hyperparameter. When $\delta = 1$, Eq. 9 reduces to $\mathbf{X}^{(l+1)} = \sigma(\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \mathbf{X}^{(l)} \Theta_l)$, which is very similar to the update equation used in popular MP layers [2]. When $\delta = 0$, the node features are not aggregated with those of the neighbors and the MP layer becomes equivalent to a dense layer of an MLP.

Theorem 1. *The MP layer in (9) minimizes the LQV of the node features \mathbf{X} .*

Proof. Let

$$\|\mathbf{x}\|_{LQV} = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N a_{i,j} \left(\frac{x_i}{\sqrt{d_i}} - \frac{x_j}{\sqrt{d_j}} \right)^2$$

To minimize the LQV, we first compute the derivative with respect to the i -th component:

$$\begin{aligned} (\partial \|\mathbf{x}\|_{LQV})_i &= \frac{1}{\sqrt{d_i}} \sum_j a_{i,j} \left(\frac{x_i}{\sqrt{d_i}} - \frac{x_j}{\sqrt{d_j}} \right) \\ &= \frac{x_i}{d_i} \sum_j a_{i,j} - \sum_j \frac{a_{i,j}}{\sqrt{d_i} \sqrt{d_j}} x_j \\ &= (\mathbf{I} \mathbf{x})_i - (\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \mathbf{x})_i \end{aligned}$$

The whole gradient is given by:

$$\nabla(\|\mathbf{x}\|_{LQV}) = (\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}) \mathbf{x}$$

The following update minimizes the LQV of the node features with gradient descent:

$$\begin{aligned} \mathbf{x}^{(l+1)} &= \mathbf{x}^{(l)} - \delta \nabla(\|\mathbf{x}\|_{LQV})^{(l)} \\ &= \mathbf{x}^{(l)} - \delta (\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}) \mathbf{x}^{(l)} \end{aligned}$$

where δ indicates the gradient step. \square

Clearly, there is a difference in minimizing the LQV of \mathbf{X} rather than the LQV of \mathbf{S} directly, as done by the terms \mathcal{L}_q in (6) and \mathcal{L}_m in (7). Nevertheless, being (5) a smooth function, if two node features \mathbf{x}_i and \mathbf{x}_j are similar they will likely be mapped into similar cluster assignments \mathbf{s}_i and \mathbf{s}_j .

3.2 Optimality of the proposed balancing term

The purpose of the balancing term is to encourage a partition where:

- the nodes are assigned with high confidence to only one cluster;
- the number of elements in each cluster is approximately the same.

Typical degenerate solutions are those that violate one these two requirements. In particular, the first degenerate solution is when samples are uniformly assigned to all cluster, i.e., $\mathbf{s}_i = [1/K, 1/K, \dots, 1/K]$, $i = 1, \dots, N$. The second degenerate solution occurs when all samples are assigned to the same cluster, e.g., $\mathbf{s}_i = [1, 0, \dots, 0]$, $i = 1, \dots, N$.

It is straightforward to see that the first degenerate solution is avoided when \mathcal{L} in (8) is minimized. The trace is the sum of the elements along the diagonal of $\sqrt{\mathbf{S}^T \mathbf{S}}$. Since $0 \leq s_{i,j} \leq 1$, the trace is maximized when each cluster assignment \mathbf{s}_i has one entry equal to 1, i.e., when $s_{i,j} \in \{0, 1\}$. On the other hand, if a sample i is assigned to more than one cluster, due to the `softmax` normalization \mathbf{s}_i will have two or more non-zero entries with values less than 1. As a consequence, at least one value in the off-diagonal of $\sqrt{\mathbf{S}^T \mathbf{S}}$ would be non-zero and the trace would be smaller.

The second degenerate solution is avoided when the clusters assume equal size, which is ensured by the following theorem.

Theorem 2. *The optimum of the problem*

$$\max_{\mathbf{S} \mathbf{1} = \mathbf{1}, s_{i,j} \in \{0,1\}} \text{Tr} \left(\sqrt{\mathbf{S}^T \mathbf{S}} \right) \quad (10)$$

is a balanced partition with clusters of size $\frac{N}{K}$.

Proof. Let $\mathbf{S}^T \mathbf{S} = \mathbf{C} \in \mathbb{R}^{K \times K}$. The conditions $\mathbf{S} \mathbf{1} = \mathbf{1}$ and $s_{i,j} \in \{0,1\}$ imply the constraint $\sum_{i=1}^K c_{i,i} = N$, being $c_{i,i}$ the volume of samples assigned to cluster i . The solution of the constrained optimization problem can be found using Lagrange multipliers. Specifically, the Lagrangian is

$$L(\mathbf{C}, \lambda) = \sum_{i=1}^K \sqrt{c_{i,i}} - \lambda \left(\sum_{i=1}^K c_{i,i} - N \right).$$

To be an optimum, $\frac{\partial L(\mathbf{C}, \lambda)}{\partial c_{i,i}} = 0$ must hold for each i . Hence, by taking the derivative one obtains

$$\frac{1}{2\sqrt{c_{i,i}}} - \lambda = 0 \rightarrow c_{i,i} = \frac{1}{4\lambda^2} \quad (11)$$

Solving for λ can be done by considering the constraint $\sum_{i=1}^K c_{i,i} = N$, which gives

$$\sum_{i=1}^K \frac{1}{4\lambda^2} = N \rightarrow \lambda = \sqrt{\frac{K}{4N}}. \quad (12)$$

Finally, substituting (12) in (11) gives $c_{i,i} = \frac{N}{K}$. \square

3.3 Numerical computation

The matrix $\mathbf{S}^T \mathbf{S}$ is symmetric and is either positive definite or positive semi-definite. The latter case occurs when at least one cluster is completely empty, i.e., $s_{i,k} = 0$ for each vertex i . Either way, the matrix $\sqrt{\mathbf{S}^T \mathbf{S}}$ is unique and real. The matrix square root is computed by first reducing the matrix to quasi-triangular form with the real Schur decomposition. The square root of the quasi-triangular matrix is then computed directly.

The computational cost of the proposed loss is dominated by the matrix multiplication $\mathbf{S}^T \mathbf{S}$ and by the Schur decomposition. The first costs $\mathcal{O}(NK^2)$, while the second costs $\mathcal{O}(K^3)$. Since $N > K$, the total cost is $\mathcal{O}(NK^2)$. On the other hand, the cost in MinCutPool and DMoN is dominated by the numerators of the LQV terms in (6) and (7), whose complexity is $\mathcal{O}(N^2K + NK^2)$. When using sparse operations, the latter cost is reduced to $\mathcal{O}(EK + NK^2)$, which is still greater than the complexity of the proposed loss.

4 Experimental evaluation

Dataset	N	E	F	K
Cora	2,708	10,556	1,433	7
Citeseer	3,327	9,104	3,703	6
Pubmed	19,717	88,648	500	3
DBLP	17,716	105,734	1,639	4

Table 1: Details of the citation datasets: number of nodes (N), number of edges (E), dimension of the feature vectors (F), and number of classes/clusters (K).

The clustering performance are tested on four citation datasets, whose details are reported in Tab. 4. In each dataset, the number of clusters K is set equal to the number of classes to which the nodes belong to. The distribution of the node classes for each dataset, depicted in Fig. 1, shows that the classes are rather unbalanced, which makes the clustering problem non-trivial.




Figure 1: Distribution of the node classes.

The proposed architecture, referred to as *Just Balance GNN* (JBGNN), consists of a stack of MP layers, which outputs $\bar{\mathbf{X}}$, followed by an MLP, which outputs \mathbf{S} . The MP layers are implemented by a GCN [9] operating on the connectivity matrix $\bar{\mathbf{A}}$ defined in (9). The whole JBNN is optimized only by minimizing the loss in (8). The following hyperparameters are used for each dataset: δ : 0.85; MP activation: ReLU; MP channels: 64; MP layers: 10; MLP activation: ReLU; MLP channels: 16; MLP hidden layers: 1; learning rate: 5e-5; epochs: 2,000. The software implementation of JBNN is available online ¹.

The performance of JBGNN are compared against three classes of methods. The first are algorithms that generate node embeddings based only on the adjacency matrix. The node embeddings are then clustered with k -means. Representatives of this category are Spectral Clustering (SC), DeepWalk [11], Node2vec [5], and NetMF [12]. The second class of methods are neural networks that generate node embeddings by accounting both for the adjacency matrix and for the node features. Also in this case, the learned embeddings are clustered with k -means in a post-processing step. The chosen representatives for this category are the Graph AutoEncoder (GAE) and Variational Graph AutoEncoder (VGAE) [10]. Finally, the last class of methods are end-to-end GNN models that directly generate soft cluster assignments \mathbf{S} by accounting both for the graph connectivity and the node features. The hard cluster assignments are computed as $\bar{\mathbf{s}} = \text{argmax}(\mathbf{S})$. DiffPool [15], DMoN [13],

MinCutPool [1], and the proposed JBGNN belong to this class. To make the comparison fair between DiffPool, DMoN, MinCutPool, and JBGNN, the GNN architectures are configured to have the same capacity (number of layers and trainable parameters) and are trained for the same number of epochs. The only difference is in the loss and in how the cluster assignments \mathbf{S} are computed (Diffpool uses an MP layer rather than the MLP in (5) to compute \mathbf{S}).

The clustering performance is evaluated in terms of the normalized mutual information (NMI) between true class labels \mathbf{y} and cluster assignments $\bar{\mathbf{s}}$ and the cluster accuracy (ACC), defined as

$$\text{ACC} = \frac{1}{N} \sum_{i=1}^N \theta(y_i, h(\bar{s}_i)), \quad (13)$$

where $h(\cdot)$ maps the hard-cluster assignment \bar{s}_i to the best matching class label according to Kuhn-Munkres algorithm, and θ is the Heaviside step function, which is 1 when y_i and $h(\bar{s}_i)$ are equal, and zero otherwise. The results are reported in Tab. 2. Methods with stochastic components are trained and evaluated 10 different times.

The computational complexity of the GNN models for clustering is measured in terms of training times. Tab. 3 reports the seconds used by each model to perform one step of gradient descent and, in red, the percentage increments from the times used by JBGNN. Times are measured on an Nvidia RTX A6000. Finally, Fig. 2 reports the evolution of the losses and the NMI during training. LQV losses are in orange and balance losses are in blue.

5 Discussion and conclusions

The clustering performance of JBGNN is competitive with respect to the other end-to-end GNN clustering methods: it achieves a comparable ACC and NMI on Cora and Pubmed, and it is the best overall performing method on DBLP. Diffpool is the worst-performing method in this category, as it does not optimize a proper clustering objective. On Citeseer, MinCutPool outperforms every other method, including JBGNN, by a large margin. As discussed in Section 3.1, JBGNN minimizes the LQV of $\bar{\mathbf{X}}$ while MinCutPool directly minimizes the LQV of \mathbf{S} , which could make a significant difference in certain datasets.

¹<https://github.com/FilippoMB/Simplifying-Clustering-with-Graph-Neural-Networks>

	Cora		Citeseer		Pubmed		DBLP	
	ACC	NMI	ACC	NMI	ACC	NMI	ACC	NMI
SC	0.298	0.028	0.217	0.014	0.589	0.182	0.458	0.023
DeepWalk	0.229 (0.020)	0.064 (0.024)	0.193 (0.002)	0.004 (0.001)	0.361 (0.001)	0.001 (0.000)	0.266 (0.001)	0.001 (0.000)
Node2vec	0.229 (0.025)	0.060 (0.029)	0.194 (0.003)	0.004 (0.001)	0.362 (0.001)	0.001 (0.000)	0.272 (0.001)	0.001 (0.000)
NetMF	0.389	0.251	0.277	0.127	0.448	0.058	0.455	0.037
GAE	0.464 (0.062)	0.327 (0.051)	0.381 (0.038)	0.162 (0.029)	0.588 (0.071)	0.235 (0.044)	0.416 (0.035)	0.111 (0.028)
VGAE	0.572 (0.054)	0.437 (0.028)	0.360 (0.038)	0.156 (0.034)	0.610 (0.060)	0.245 (0.043)	0.507 (0.047)	0.212 (0.021)
DiffPool	0.472 (0.010)	0.306 (0.005)	0.336 (0.007)	0.180 (0.008)	0.418 (0.002)	0.084 (0.001)	0.370 (0.042)	0.045 (0.043)
DMoN	0.488 (0.063)	0.357 (0.042)	0.364 (0.043)	0.196 (0.029)	0.559 (0.042)	0.192 (0.048)	0.590 (0.039)	0.334 (0.026)
MinCut	0.534 (0.041)	0.406 (0.029)	0.497 (0.049)	0.295 (0.029)	0.572 (0.034)	0.208 (0.014)	0.538 (0.033)	0.297 (0.024)
JBGNN	0.457 (0.025)	0.351 (0.128)	0.334 (0.019)	0.140 (0.024)	0.564 (0.023)	0.223 (0.013)	0.607 (0.008)	0.359 (0.008)

Table 2: Clustering performance metrics for each dataset. For method with stochastic components, the mean and the standard deviation (in brackets) obtained from 10 independent runs is reported.

	Cora	Citeseer	Pubmed	DBLP
Diffpool	0.009	0.030	0.234	0.229
	+200%	+42%	+1376%	+477%
DMoN	0.006	0.028	0.030	0.061
	+100%	+33%	+76%	+27%
MinCut	0.006	0.029	0.030	0.061
	+100%	+38%	+76%	+27%
JBGNN	0.003	0.021	0.017	0.048

Table 3: Training times (seconds/step). In red, the increments in training times with respect to JBGNN.

About the performance of the other methods, SC, DeepWalk, Node2vec, and NetMF always achieve worse ACC and NMI. This is expected, as they build embeddings based only on the graph connectivity. On the other hand, GAE and VGAE also account for the node features and achieve superior performance. Compared to the end-to-end GNN clustering methods, GAE and VGAE follow a significantly different approach: they optimize a link reconstruction loss to learn node embeddings, which are then clustered in post-processing with k -means. Despite the combination of GAE and VGAE embeddings with k -means yields very good performance, it has the disadvantage of being a two-step procedure. In addition, the k -means is computationally expensive and is not designed for out-of-samples predictions. On the other hand, the GNN-based approaches directly output soft cluster assignments, which are efficiently evaluated at inference time and can also be used to implement differentiable operations, such as graph pooling [4].

In terms of computational complexity, JBGNN

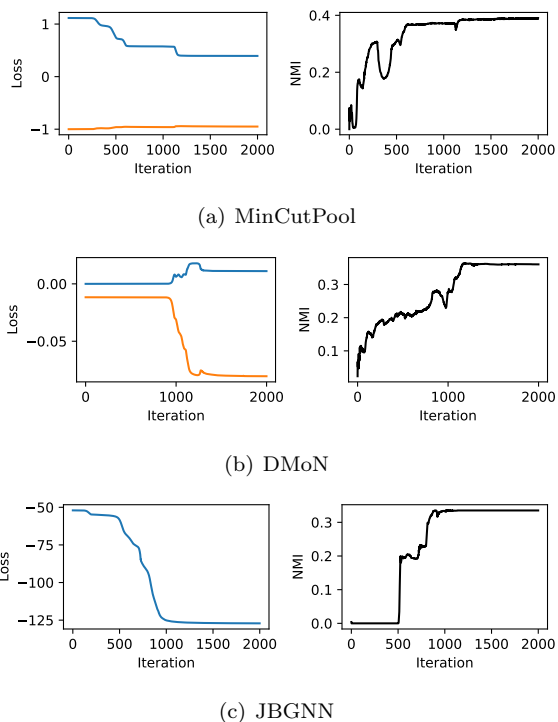


Figure 2: Evolution of the losses and NMI when training MinCutPool, DMoN and JBGNN on Cora.

outperforms all other GNN clustering methods thanks to its simplicity and efficient formulation. In particular, JBGNN achieves a speed improvement ranging from 27% to 1,376% compared to the other methods across the four datasets, making it particularly suitable for large-scale applications. Referring to Fig. 2, MinCutPool and DMoN start

to converge between 1000 and 1500 epochs, while JBGNN generally converges earlier. This, together with the faster updates, makes training the JBGNN significantly faster.

A natural extension of this work is to test JBGNN in other tasks, such as to implement graph pooling [4] in a deep GNN architecture for graph classification. As a concluding remark, it should be noted that the proposed clustering objective relies on the assumption that all clusters have equal size. This is reasonable, since clustering is an unsupervised task and the actual sizes are usually unknown. However, if information about cluster size is available, a clustering objective that puts a bias towards an unbalanced partition could be used instead.

Acknowledgments I gratefully acknowledge the support of Nvidia Corporation with the donation of the two RTX A6000 GPUs used in this work.

References

- [1] F. M. Bianchi, D. Grattarola, and C. Alippi. Spectral clustering with graph neural networks for graph pooling. In *International Conference on Machine Learning*, pages 874–883. PMLR, 2020.
- [2] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems*, 29, 2016.
- [3] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*. PMLR, 2017.
- [4] D. Grattarola, D. Zambon, F. M. Bianchi, and C. Alippi. Understanding pooling in graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [5] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
- [6] W. L. Hamilton. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- [7] M. Hein and S. Setzer. Beyond spectral clustering-tight relaxations of balanced graph cuts. In *NIPS*. Citeseer, 2011.
- [8] M. Kampffmeyer, S. Løkse, F. M. Bianchi, L. Livi, A.-B. Salberg, and R. Jenssen. Deep divergence-based approach to clustering. *Neural Networks*, 113:91–101, 2019.
- [9] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations*, 2016.
- [10] T. N. Kipf and M. Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- [11] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2014.
- [12] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proceedings of the 11th ACM international conference on web search and data mining*, 2018.
- [13] A. Tsitsulin, J. Palowitch, B. Perozzi, and E. Müller. Graph clustering with graph neural networks. *arXiv preprint arXiv:2006.16904*, 2020.
- [14] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [15] Z. Ying, J. You, C. Morris, X. Ren, W. Hamilton, and J. Leskovec. Hierarchical graph representation learning with differentiable pooling. *Advances in neural information processing systems*, 31, 2018.
- [16] S. Zhou, H. Xu, Z. Zheng, J. Chen, J. Bu, J. Wu, X. Wang, W. Zhu, M. Ester, et al. A comprehensive survey on deep clustering: Taxonomy, challenges, and future directions. *arXiv preprint arXiv:2206.07579*, 2022.