Spectral Clustering of Attributed Multi-relational Graphs

Ylli Sadikaj
Faculty of Computer
Science
University of Vienna
Vienna, Austria
ylli.sadikaj@univie.ac.at

Yllka Velaj
Faculty of Computer
Science
University of Vienna
Vienna, Austria
yllka.velaj@univie.ac.at

Sahar Behzadi
Faculty of Computer
Science
University of Vienna
Vienna, Austria
sahar.behzadi@univie.ac.at

Claudia Plant
Faculty of Computer
Science, ds:UniVie
University of Vienna
Vienna, Austria
claudia.plant@univie.ac.at

ABSTRACT

Graph clustering aims at discovering a natural grouping of the nodes such that similar nodes are assigned to a common cluster. Many different algorithms have been proposed in the literature: for simple graphs, for graphs with attributes associated to nodes, and for graphs where edges represent different types of relations among nodes. However, complex data in many domains can be represented as both attributed and multi-relational networks.

In this paper, we propose SpectralMix, a joint dimensionality reduction technique for multi-relational graphs with categorical node attributes. SpectralMix integrates all information available from the attributes, the different types of relations, and the graph structure to enable a sound interpretation of the clustering results. Moreover, it generalizes existing techniques: it reduces to spectral embedding and clustering when only applied to a single graph and to homogeneity analysis when applied to categorical data.

Experiments conducted on several real-world datasets enable us to detect dependencies between graph structure and categorical attributes, moreover, they exhibit the superiority of SpectralMix over existing methods.

CCS CONCEPTS

• **Theory of computation** → **Unsupervised learning and clustering.**

KEYWORDS

Graph embedding; Spectral clustering; Multi-relational graphs; Attributed graphs

ACM Reference Format:

Ylli Sadikaj, Yllka Velaj, Sahar Behzadi, and Claudia Plant. 2021. Spectral Clustering of Attributed Multi-relational Graphs. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, Singapore, 10 pages. <https://doi.org/10.1145/3447548.3467381>

1 INTRODUCTION

Complex data in many domains including biology, neuroscience, social and collaboration networks etc. can be represented as attributed multi-relational graphs. The nodes are often characterized by many

Figure 1: Example of a simple graph (left) and an attributed multi-relational graph (right). Different types of lines between nodes represent different relation types.

attributes and the links can be of different types, as depicted in Figure 1. In social networks, for instance, nodes represent users who are characterized by attributes like gender, age, hometown, hobbies etc., and edges describe the different types of links established by interaction on different platforms like Facebook and Twitter. In the field of neuroscience, the data obtained by neuroimaging techniques can be modelled using nodes for the anatomical brain regions, attributes for the region size and type of tissue, and edges for the structural and functional connectivity links, as measured by structural and functional Magnetic Resonance Imaging (MRI and fMRI) [15].

Despite the importance of attributed multi-relational graphs in modeling many real-world domains, they have received little attention, in particular in the task of clustering.

Clustering is one of the main research directions in data mining tasks. It aims at discovering a natural grouping of the data such that similar data objects are assigned to a common cluster. Clustering is very helpful in order to obtain an understanding of the major patterns in large datasets like attributed multi-relational networks. Existing approaches are mainly designed for simple graph clustering [13, 34, 35, 41, 42], for fuzzy clustering [18] that identifies overlapping clusters in which an object belongs with a membership probability, for clustering of networks with numerical attributed nodes [19].

Attributes and multiple types of relationships between nodes of a graph can yield complementary information. This information is valuable for detecting clusters of nodes with similar characteristics and can be used in different applications, e.g., targeting ads. Thus, we develop SpectralMix, a novel algorithm that provides a comprehensive approach to joint dimensionality reduction and spectral clustering of attributed multi-relational graphs with categorical node attributes. Given a multi-relational attributed graph G and the desired dimensionality d , SpectralMix defines a mapping $\Phi : G \rightarrow \mathbb{R}^d$ that minimizes the distance between nodes with similar attributes. The joint low-dimensional vector space generated by SpectralMix reduces noise, emphasizes the major patterns in

This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '21, August 14–18, 2021, Virtual Event, Singapore.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8332-5/21/08.

<https://doi.org/10.1145/3447548.3467381>

the data, and enables subsequent data mining. By application of the K -means algorithm we obtain a clustering inspired by spectral clustering methods. But we could also perform outlier detection or visualize the data when we select a 2D or 3D vector space. Moreover, other very important aspects of SpectralMix are that: there is only a single dimensionality parameter d that needs tuning, and SpectralMix generalises other existing techniques. When considering the setting of a single input graph without any node attributes as input data, our mapping is the same as Laplacian Eigenmaps [1]. When considering only categorical data without relational information, our mapping coincides with Homogeneity Analysis [5], a technique from statistics for categorical PCA [11].

To assess its performance, we evaluate SpectralMix based on clustering and data visualization tasks. The experiments conducted on several real-world networks enable us to detect dependencies between graph structure and the categorical attributes, moreover, they exhibit the superiority of SpectralMix over alternative baselines.

2 RELATED WORKS

Recently, the problem of clustering graphs with node attributes has received a lot of attention. The main goal in this field has been to improve the quality of the clustering by combining both structural and attribute properties.

While it may seem a natural extension, many methods [8, 33, 34] cannot be directly applied to attributed multi-relational graphs. They propose algorithms for detecting densely connected components [34, 35, 41, 42], or for identifying clusters of objects having similar connectivity [13, 19] but they either ignore the multiple types of edges [24, 26, 38] and/or deal with only one attribute type [20]. A technique, that has been used frequently to overcome these problems, is to project the attributed multi-relational graph to a single weighted graph where weights represent a combination of attribute and structural similarities, and then apply a clustering algorithm for weighted graphs with the aim of performing clustering and graph outlier detection [22] or identify strong community structures [25]. This projection, however, causes information loss and, as a consequence, it limits the clustering accuracy. Similarly, graph embeddings techniques have been designed mainly for single graphs such as FATNet [26] which integrates the global relations of the topology and attributes into robust representations.

Different works have presented results on spectral clustering on multi-layer networks [2, 6, 12], while very few, to the best of our knowledge, investigate spectral clustering methods on attributed graphs [27]. Whereas, there are methods not designed for graphs, which exploit the fusion of the cluster-separation information from all eigenvectors to achieve a better clustering [36] or perform clustering of numerical and categorical data based on homogeneity analysis [23].

Recently, many works in the area of graph embedding and clustering have been published and they can be categorized in the three groups: a) methods for single graphs with attributes, ANRL [39] and DGI [30]; b) methods for multi-relational graphs without attributes, DMGC [14] and CrossMNA [3]; and c) methods for attributed multi-relational graphs, HAN [32], MARINE [7], DMGI [21].

ANRL [39] is a deep neural network model designed for attributed networks representation learning which integrates network structural proximity and node attributes affinity into low-dimensional representation spaces. DGI [30] is an approach for learning unsupervised representations on graph-structured data and it relies on maximizing mutual information between patch representations and corresponding high-level summaries of graphs derived using established graph convolutional network architectures. DMGC [14] is a deep learning based multi-relational graph clustering method which aims at performing two procedures: graph clustering and cross-graph cluster association. CrossMNA [3] is an embedding method that integrates cross-network information to refine two embedding vectors for alignment tasks: inter-vector and intra-vector. MARINE [7] is an unsupervised embedding method for both simple and multi-relational networks that preserves the proximity and attributes information in addition to various types of relations. Furthermore, HAN [32] is a semi-supervised graph neural network framework that can be applied to graphs with various types of nodes and edges. It employs node-level attention and semantic-level attention simultaneously. DMGI [21] is an unsupervised network embedding method for attributed multi-relational graphs. It is a modified approach of DGI, and it jointly integrates the embeddings from multiple types of relations between nodes through a consensus regularization framework, and a universal discriminator.

To conclude, the previous methods differ from ours in the following ways: (i) they do not consider the existence of multiple edge types, (ii) they deal with only one type of attributes, (iii) their final node embedding is represented in higher dimensions, 100 or 200 dimensions, (iv) they can perform well only on one graph category: multigraphs with attributes or multigraphs without attributes.

SpectralMix outperforms the above methods because it exhibits good results when applied to graphs with one or many different edge types, to graphs with or without node attributes, and a combination of those. Moreover, the output from our algorithm can be represented in 2D or 3D space and has more meaningful cluster structure.

3 NOTATION AND PROBLEM DEFINITION

We consider an undirected node-attributed multi-relational network, (also referred to as multigraph), $G = (V, E, \mathcal{R})$ where $V = \{v_1, \dots, v_n\}$ represents the set of nodes. Each node v_i is characterized by a set $\mathcal{A} = \{A_1, \dots, A_c\}$ of c categorical attributes, i.e., $|\mathcal{A}| = c$, and c is the dimensionality of the attribute space. We denote categorical attributes by upper case letters. Let A_j for $j = 1, \dots, c$ be a categorical attribute having k_j distinct values, with $k_j = 2$ for binary attributes. For each categorical attribute, we allow k_j for $j = 1, \dots, c$ associated to it, to have a different value. Moreover, $v_i.A_j$ denotes the value or category of vertex v_i in attribute A_j . Thus, the total number of categories, denoted by C , of all c categorical attributes is $C = \sum_{j=1}^c k_j$.

We denote by $\mathcal{R}$ the set of different relation types and, therefore, $|\mathcal{R}|$ denotes the dimensionality of the network. The set of edges E consists of the tuples $e = (v_i, v_j, w_{ij}, r)$ where w_{ij} represents the weight of the edge of type r , and $r \in \mathcal{R}$ is the type of relation. We denote by E_r set of edges of type r , for $r = 1, \dots, |\mathcal{R}|$.

Our goal is to define a mapping $\Phi : G \rightarrow \mathbb{R}^d$, where d is the desired dimensionality, such that the distance between connected nodes with similar attributes is minimal. We define by $\mathcal{O}$ and $\mathcal{M}$ the matrices of size $n \times d$ and $C \times d$, respectively, where each entry o_{il} and m_{jl} for $i = 1, \dots, n$, $j = 1, \dots, C$ and $l = 1, \dots, d$, is a real value denoting the coordinate of the node or the categorical attribute according to the embedding.

Our algorithm, sketched in Algorithm 1, takes as input a multi-relational graph G , a categorical attribute matrix of size $n \times c$, and the desired dimensionality d of the embedding. It returns as output a d -dimensional feature space representation of data objects $\mathcal{O}$ and categories $\mathcal{M}$, and, if desired, also a clustering of the data objects. Hence, $\mathcal{O}$ and $\mathcal{M}$ are two matrices of size $n \times d$ and $C \times d$, respectively.

Our framework can also be applied to cluster multi-layer graphs where there are inter-layer edges between any pair of nodes. Given a multi-layer graph, we build a multi-relational graph with the same set of nodes and the same attributes associated to them. For the set of edges, we consider as many relations as there are layers, we also add relations representing the edges from nodes in one layer to the other layers, if those edges exist.

4 EMBEDDING MULTIDIMENSIONAL ATTRIBUTED GRAPHS

Our problem setting is characterized by mixed-type data of different modalities. In particular, we have a set of nodes which are characterized by multiple categorical attributes and linked by different relations. We integrate all the different modalities by a joint vector space representation. Our objective function for this embedding combines the ideas of spectral embedding of graphs with homogeneity analysis of categorical data. For n nodes we derive low-dimensional coordinates minimizing the following objective function:

$$\begin{aligned} \min_{\mathcal{O}, \mathcal{M}} \sum_{r=1}^{|\mathcal{R}|} \alpha_r \cdot \left(\sum_{e=(v_i, v_j, w_{ij}, r) \in E_r} w_{ij} \cdot \sum_{l=1}^d \|o_{il} - o_{jl}\|^2 \right) \\ + \sum_{i=1}^n \sum_{j=1}^{|\mathcal{A}|} \alpha_j \cdot \sum_{l=1}^d \|o_{il} - m_{jl}\|^2 \end{aligned}$$

subject to $\mathcal{O}^T \mathcal{O} = I_n$. (1)

As mentioned in the previous section, $\mathcal{O}$ is a $n \times d$ matrix of low-dimensional coordinates. Every row represents the coordinate vector for one data object (node). The first summand in Equation 1 represents the contribution of the relational part of the data, i.e., of the $|\mathcal{R}|$ -dimensional graph on the position of a data object. For every graph we consider all edges between v_i and vertices v_j and determine the coordinate matrix $\mathcal{O}$ such that the squared Euclidean distances are minimized. We denote by o_{il} the coordinate of the node v_i in the dimension l . As in spectral embedding techniques we place the vertices as close as possible to their neighbors, but now not considering a single graph but multiple different relations.

The second summand of the equation represents the categorical information. Categorical data establishes a bipartite graph between data objects and their categories. Every categorical variable establishes its own bipartite graph, where object v_i is connected to the

Figure 2: Example of an attributed multi-relational graph (left) and a bipartite graph for its categories for attribute 'Hometown'.

category or value of its attribute $v_i.A_j$. An example of a bipartite graph is shown in Figure 2.

We now also consider the c categorical graphs in determining the joint vector space embedding. Inspired by Homogeneity Analysis, our algorithm SpectralMix will discover for every category k_j of every categorical attribute A_j a position in $\mathbb{R}^d$. In Equation 1 we denote the position of the category of object v_i in dimension l by m_{il} . The matrix $\mathcal{M} \in C \times d$ contains these category coordinates as row vectors.

The coefficients α_r, α_j represent weighting factors for the relations types and the categorical attributes in the input data. It is possible to consider the unweighted objective function setting all to 1. However, the total number of edges in the multi-relational graph and the edges in the graphs generated by the attributes tend to be very different as the categorical data establishes bipartite graphs where every node is connected to only and exactly one category. As default weighting scheme we, therefore, suggest to set the weight factors such that every modality has the same weight in the low-dimensional representation.

In order to avoid trivial solutions, we require the matrix $\mathcal{O}$ to be column-orthonormal. Without this constraint, a minimum of the objective function would be achieved by mapping all vertices to one common location.

Our algorithm, reported in Algorithm 1, starts by initializing every element o_{il} for $i = 1, \dots, n$ and $l = 1, \dots, d$, in $\mathcal{O}$ with a real number selected at random. In this matrix, every row represents the data object associated to a vertex of the graph G , and every column is one of the coordinates in the embedding space.

Then, we compute the value of our objective function as described in Equation 1, and at each iteration we update the coordinates in $\mathcal{O}$ and $\mathcal{M}$ in order to obtain a smaller value for our objective. If this value does not decrease after the updates, then the algorithm converges and returns the final node embeddings for the data objects (nodes) and the attribute categories.

In particular, when we update the object coordinates we take into account the type of relations in the graph (lines 4 – 7) and the attributes (lines 8 – 10). For the graph contribution, we consider, for all the different relation types, the neighbors v_p of node v_i and the weight of these edges. We update each coordinate adding to it the product of the weighting factor α_r , the weight w_{ip} of the edge (v_i, v_p, w_{ip}, r) , the coordinate of the data object for node v_p , and divide by $\text{sum}_i(G)$ computed as the sum of the size of the neighborhood of node v_i in relation r . For the attribute contribution, we update each coordinate adding to it the product of the weighting

Algorithm 1: SpectralMix

Input: Graph G , attribute matrix, dimensionality d
Output: d -dimensional feature space representation of data objects O and categories $\mathcal{M}$, clusters $clu_1, \dots, clu_k$

```

1 Initialize matrices  $O$  and  $M$  randomly;
2 repeat
3     /* update object coordinates */
4     for  $i = 1, \dots, n$  do
5         for  $r = 1, \dots, |\mathcal{R}|$  do
6             for  $p = 1, \dots, |N_v|$  with  $N_v = \{v_p | (v_i, v_p) \in E_r\}$  do
7                 for  $l = 1, \dots, d$  do
8                      $o_{il} \leftarrow o_{il} + \alpha_r \cdot w_{ip} \cdot o_{pl} / \text{sum}_i(G)$ ;
9             for  $j = 1, \dots, C$  do
10                for  $l = 1, \dots, d$  do
11                     $o_{il} \leftarrow o_{il} + \alpha_i \cdot m_{jl} / \text{sum}_i(A)$ ;
12            Apply orthonormalization to ensure  $O^T O = I_n$ ;
13        /* update category coordinates */
14        for  $i = 1, \dots, n$  do
15            for  $l = 1, \dots, d$  do
16                for  $j = 1, \dots, C$  do
17                     $m_{jl} \leftarrow m_{jl} + o_{il} / \text{tot}_j$ ;
18 until convergence;
19 if clustering is desired then
20     perform k-Means on the rows of  $O$  with  $k = d - 1$ ;
21     return clusters  $clu_1, \dots, clu_k$ ;
22 else
23     return  $O, M$ ;

```

---

factor α_i for the bipartite graph generated by the category, the coordinate of the category for attribute A_j , and divide by $\text{sum}_i(A)$ computed as the sum of the weighting factor for category i . Then we apply the orthogonalization algorithm.

When we update the coordinates of the categories, we add to each of them the coordinate o_{il} of node v_i and divided by the total number of elements that belong to that category. We highlight that our algorithm has only one parameter d , desired dimensionality, to be tuned by human, while all the other parameters (weights) can be learned during optimization.

We provide now a proof of convergence of our algorithm.

THEOREM 4.1. *The SpectralMix algorithm converges in a finite number of iterations.*

PROOF. We say that our algorithm converges if the difference in the value of the objective function $\mathcal{F}$ between two consecutive iterations, t and $t+1$, of our algorithm is bounded by a small positive constant ϵ : $\mathcal{F}_{t+1} - \mathcal{F}_t < \epsilon$. The idea is to show that the two steps of the algorithm, i.e., the update of the object coordinates and the update of the attributes coordinates, decrease the objective function and that the difference $\mathcal{F}_{t+1} - \mathcal{F}_t$ tends to zero. Suppose that the algorithm proceeds from iteration t to iteration $t+1$, it suffices to show that :

$$\|o_{il} - o_{jl}\|_t^2 \geq \|o_{il} - o_{jl}\|_{t+1}^2 \text{ and } \|o_{il} - m_{il}\|_t^2 \geq \|o_{il} - m_{il}\|_{t+1}^2.$$

In lines 4 – 7 of Algorithm 1, the coordinates of a node are updated by taking into account the position of all its neighbors in the graph. Therefore, every node moves towards nodes that are closer in the space, reducing the distance $\|o_{il} - o_{jl}\|^2$. This observation also holds for the attributes as they are represented as nodes in bipartite graphs, in lines 12 – 15 the coordinates are updated so that these attributes moves towards the nodes to which they are associated with, therefore the distance $\|o_{il} - m_{il}\|^2$ decreases. This information is also considered while updating the coordinates of

the objects in lines 8 – 10. As the algorithm updates the coordinates of the nodes, the distance between two nodes v_i and v_j represented by the coordinates o_{il} and o_{jl} , decreases but their distance to other neighbors might increase. This does not affect the trend of the objective function. We can see that by distinguishing two cases: (a) Edge $e = (v_i, v_j, w_{ij}, r)$, for a given relation r , is the edge with highest weight w_{ij} for both v_i and v_j . In this case, when the distance between v_i and v_j decreases, due to the update in line 7, it produces a decrease in the value of $\mathcal{F}$. This difference compensates the distances that may increase from v_i and v_j to any of their neighbors v_k such that $w_{ij} > w_{ik}$ and $w_{ij} > w_{jk}$. (b) Edge $e = (v_i, v_j, w_{ij}, r)$, for a given relation r , is the edge with highest weight w_{ij} for v_i but not for v_j that is connected to v_k by an edge of weight $w_{jk} > w_{ij}$. In this case, the distance between v_i and v_j might not decrease, because as v_i moves towards v_j , the latter node moves towards v_k . Even in this case, the value of the objective function will decrease because the factor w_{jk} is higher than the other values used when updating the coordinates.

This argument holds for any dimension l , any relation in $\mathcal{R}$, and for the categorical attributes represented as nodes in a bipartite graph. Hence, the value of the objective function will decrease at each iteration and when the decrement is smaller than ϵ , the algorithm terminates. $\square$

Complexity analysis. Lines 4 – 7 in Algorithm 1 compute the contribution provided by relational information, its time complexity is $O(|\mathcal{R}| \cdot m \cdot d)$, where m is the number of edges. Lines 8 – 10 compute the contribution given by the categorical data, whose time complexity is $O(C \cdot d)$. These operations are repeated for each node in the graph, thus, the time complexity of lines 4 – 10 for updating the object coordinates is $O(n(|\mathcal{R}| \cdot m \cdot d + C \cdot d))$. In line 11 we apply the Gram-Schmidt orthonormalization algorithm with complexity $O(n \cdot d^2)$. In lines 12 – 15, the algorithm updates the coordinates of the categorical data objects, it requires $O(n \cdot C \cdot d)$. Finally, in lines 17 – 19 we use K-Means algorithm with complexity $O(t \cdot d^2 \cdot n)$ setting K to $d - 1$. We denote by t the number of iterations of K-Means. The total time complexity of Algorithm 1 is $O(I(n \cdot |\mathcal{R}| \cdot m \cdot d + n \cdot d^2 + n \cdot C \cdot d) + t \cdot d^2 \cdot n)$ where I is the number of iterations the algorithm needs for converging.

5 EXPERIMENTS

In this section, we evaluate the performance of SpectralMix against a variety of state-of-the-art methods. Our focus is on node embedding and clustering tasks. Due to space limitations, we only can show a selection of the results.

Datasets. To compare SpectralMix with other baselines we use four datasets, which are multigraphs with node attributes, and two datasets, which are multigraphs without node attributes. To make fair comparisons with HAN [32], and DMGI [21], which are the most relevant baseline methods for attributed multigraphs, we evaluate our proposed algorithm on the same datasets used in [32] and in [21], i.e., ACM, and IMDB. Similarly, we use the datasets, DBLP and Flickr, that are used for evaluating the performance of DMGC [14], which is the most relevant method for multi-relational graphs without node attributes.

ACM [21] is a multigraph with 2 types of relations where nodes represent papers and edges are inferred via Authors collaborations

(PAP) and Subjects of the paper (PSP). The task is to cluster the data into three clusters (Database, Wireless Communication, and Data Mining).

IMDB [21] is a multigraph with 2 types of relations where nodes represent movies and edges are inferred via movie-actors (MAM) and movie-directors (MDM) relations. The task is to cluster the data into three clusters (Action, Comedy, Drama) according to movie genre.

DBLP [14] is a multigraph of 4 relations, that can be considered as 4 graphs: a) collaboration graph – consists of edges between authors who have published a paper together, b) citation graph – consists of edges between authors, an edge represents a citation between authors, c) co-citation graph – consists of edges between papers, an edge represents that papers have cited the same source, and d) authorship graph – consists of edges of authors and papers, an edge represents an author who has published a paper. The task is to cluster the author and paper nodes into three cluster (AI, computer graphics, and computer networks) according to research areas.

Flickr [14] is a dataset that does not have features and consists of two relations: a) friendship – represents an edge between users who are friends on social network, b) tag-similarity – represents an edge between users who have tag similarities among them. The task for this dataset is to cluster the user nodes into seven cluster, according to social groups.

Brain Networks (BN) can be defined as weighted graphs, where nodes are the brain regions, edges are connections between regions, and weights denote distances between regions. In our experiments, we use a publicly available dataset project [4] that contains data of 52 Typically Developed (TD) children and 49 children suffering from Autism Spectrum Disorder (ASD). After preprocessing [28], we obtain, for each subject, a graph with 116 nodes. We also use a second dataset [31] that provides structural and functional brain networks from 27 schizophrenic (SCZ) patients and 27 matched healthy adults (HL). The brain was divided into 83 approximately equally sized regions. For each node we are given the following attributes: average apparent diffusion coefficient, fiber density, generalized fractional anisotropy, length of fibers and number of fibers. We perform two types of experiments. For the first dataset, in order to obtain a multi-relational graph, we aggregate the graphs of each subject in the groups ASD and TD. These multi-relational graphs are defined over the same set of nodes as the original graphs, while the weight and the type of relation of each edge are those of each subject in the original groups. For the second dataset, we select at random a subject in each group, i.e., SCZ and HL, and we build a multi-relational graph where we have two edge types: the functional and structural connectivity measured with MRI and fMRI. We describe this results in details in Section 5.

Baseline methods. We compare SpectralMix with the state-of-the-art methods described in Section 2, i.e., ANRL [39], DGI [30], DMGC [14], CrossMNA [3], HAN [32], MARINE [7], and DMGI [21]. We report their main characteristics in Table 1.

For fair comparisons with methods *not* designed for multigraphs, ANRL and DGI, we obtained the final node embedding matrix $\mathbf{Z}$ of a multi-relational graph by computing the average of the node embeddings obtained from each single graph, similarly as in [21]. For the other two methods in the second category that ignore the

Table 1: Properties of the compared methods

	Multigraphs	Attribute	Unsupervised
ANRL	X	✓	✓
DGI	X	✓	✓
DMGC	✓	X	✓
CrossMNA	✓	X	✓
HAN	✓	✓	X
MARINE	✓	✓	✓
DMGI	✓	✓	✓
SpectralMix	✓	✓	✓

node attributes, i.e., DMGC and CrossMNA, we concatenated the raw attribute matrix $\mathbf{X}$ to the learned node embeddings $\mathbf{Z}$ i.e., $\mathbf{Z} \leftarrow [\mathbf{Z}; \mathbf{X}]$, similarly as in [21].

We underline that SpectralMix can operate on multi-relational graphs with or without attributes and it obtains a joint embedding for all the relation types of the graphs. SpectralMix is applicable to multi-relational graphs but it is not limited to them because the learned node embeddings that we obtain for single graphs are the same as the final embeddings of Laplacian Eigenmaps.

Evaluation Metrics. SpectralMix is an unsupervised method and does not need any labeled data for training. We focus on the task of node clustering, which is a classical task for unsupervised methods. We compute the most commonly used metrics to assess the quality of the clustering results [32]: Normalized Mutual Information (NMI) [9], and Adjusted Rand Index (ARI) [9], where a value equal to 1 indicates a perfect clustering.

Experimental setup. The final node embeddings obtained from SpectralMix are suitable for spectral clustering, i.e. subsequent application of K-Means and for visualization. Therefore, we apply the K-Means on the final node embeddings to perform node clustering by setting K to the number of clusters of the dataset, and we use K-Means++ for initialization. Since the performance of K-Means is affected by initial centroids, we repeat the process for 100 times and report the average results. In the comparison with the alternative baselines, we keep the same setup and parameters that are made publicly available by the authors.

Performance evaluation. The experimental results of six public datasets are shown in Table 2, the best results are marked in bold. It is easy to notice that SpectralMix outperforms all the alternative methods. Moreover, the performance of SpectralMix improves with the volume of data. As we can see from the results, we obtain the highest NMI and ARI values for the ACM dataset which is the graph with the largest number of edges and has 1,870 attributes. Thus, SpectralMix achieves state-of-the-art performance on ACM dataset, with 5% performance lift in NMI, and 12% in ARI, compared to the results from the competitors.

SpectralMix outperforms other baselines on other real-world datasets too. For the IMDB dataset the improvement of SpectralMix in terms of NMI is 11% and ARI is 13%, and for multigraphs without attributes, DBLP and Flickr, the percentages of improvement in terms of NMI are 21% and 94%, respectively, compared to the best performer from competitors for each dataset.

Considering the results in Table 2, we observe that the flaws of all baseline methods is that they perform well only for one category

Table 2: Performance on node clustering on real-world datasets (NA: Not Applicable).

	ACM		IMDB		DBLP		Flickr		BrainNetwork-TD		BrainNetwork-ASD	
	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI
ANRL	0.46	0.41	0.11	0.11	0.12	0.1	0.05	0.05	NA	NA	NA	NA
DGI	0.63	0.61	0.15	0.14	0.08	0.05	0.09	0.09	0.13	0.04	0.13	0.04
CrossMNA	0.32	0.31	0.15	0.15	0.27	0.22	0.23	0.19	NA	NA	NA	NA
DMGC	0.42	0.37	0.04	0.04	0.39	0.32	0.33	0.29	NA	NA	NA	NA
HAN	0.64	0.62	0.15	0.14	0.04	0.08	0.1	0.12	0.18	0.07	0.16	0.05
MARINE	0.46	0.43	0.03	0.02	0.18	0.11	0.22	0.14	NA	NA	NA	NA
DMGI	0.68	0.65	0.19	0.15	0.09	0.04	0.1	0.07	0.12	0.29	0.13	0.02
SpectralMix	0.71	0.77	0.21	0.17	0.47	0.34	0.64	0.51	0.7	0.62	0.6	0.51

of graphs. The strength of SpectralMix is that it is applicable to different graph types: multi-relational graphs with and without attributes. Therefore, it outperforms the most recent and well known methods in both categories, respectively. Our advantage lies in the importance of considering separately the weights that each graph structure and each category of the node attributes have in our embedding. Also, SpectralMix performs well when applied to single graphs.

In particular, we observe that the DMGI [21] performance drops when it is applied to multi-relational graphs without attributes, and this is due to the crucial role that node attributes play for DMGI. Similarly, the performance of all baselines which are designed for attributed graphs, such as ANRL [39], HAN [32], and MARINE [7], is poor on multi-relational graphs without attributes. Albeit, MARINE integrates the node attributes, we observe that very different node categories can impact its performance, which is the case for its performance drop on the IMDB dataset. Whereas, MARINE performance improves when applied to multi-relational graphs without attributes, and it occurs as MARINE is more focused on preserving link structures. As expected HAN [32] performs well on attributed multi-relational graphs and there is a drop in performance for the other two datasets, multi-relational graphs without attributes. The lack of node properties has affected the performance drop, since the node level mechanism of HAN uses information from both graph structure and node attributes.

Experimental results on brain networks, are presented for SpectralMix and only three competitors, DMGI, HAN, and DGI. For other methods, experiments could not be conducted as those methods are not able to process the data due to their limitations, such as not being able to process decimal or negative values for edge weights or node attributes. In case of brain networks, the gap in performance is even higher between SpectralMix and other methods, and one of the factors is that baseline methods have to train and validate their models in order to have a satisfiable model. Since for the brain networks the number of nodes and edges is very limited, it has an impact on the performance of these baselines methods.

Visualization. To further demonstrate the effectiveness of SpectralMix over comparison methods, we show in Figure 4 the visualization of all methods on ACM dataset. For the sake of fair comparison we apply visualization tool t-SNE [29] on the node embeddings obtained from all methods. We set the same parameter values for visualization for all embeddings. The nodes are colored according

to the research area: Database, Wireless Communication, and Data Mining. From Figure 4, we can observe that ANRL, CrossMNA, DMGC, and MARINE do not perform well as the nodes belonging to different research areas are mixed with each other. Also, we can note that SpectralMix, DMGI, HAN, and DGI are able to better distinguish different types of nodes. It is clear that SpectralMix is able to separate the nodes in different research areas with more distinct boundaries and higher number of nodes clustered correctly. Additionally, it is noticeable on the SpectralMix node embedding in Figure 4h that there is a distinct group of nodes, on the right side, which represent the noise data or *outliers*. These nodes are only connected with each other in PAP relation type of ACM dataset and do not share the same properties or attributes with other nodes. Thus, we can see that SpectralMix is robust against outliers.

It is worth mentioning that the outlier nodes represent papers in both the types of relations in the graph of ACM, i.e., PAP and PSP. These papers have very few authors in common with other papers, i.e., they are connected by very few edges to other nodes in the PAP graph. Whereas, they are written by the same small group of authors, i.e., the edges between the outliers form a complete graph.

We emphasize the importance of node attributes in the final node embeddings in Figure 3. We visualize the final node embeddings without node attributes in Figure 3a and with node attributes in Figure 3b. In Figure 3a the cluster structure is not clear and nodes from two different clusters are closer in the lower dimensional space. Whereas, if we integrate the available information from node attributes we can improve our final node embeddings, Figure 3b

Figure 3: Visualization of SpectralMix on ACM dataset.

Figure 4: Visualization on ACM dataset.

yields the importance to integrate the node attributes in the final joint embedding, and this property of SpectralMix allows us to achieve a better performance on the clustering task.

Run time evaluation. For 4 datasets analysed in Table 2, we show the run time evaluation for all methods in Figure 5. We do not include the brain networks on the run time evaluation as we do not have results on these datasets for all the alternative baselines. For the alternative baselines, we report the run time needed to converge. It is noticeable that SpectralMix is faster than other methods on three datasets, i.e., IMDB, Flickr, and DBLP. In the ACM dataset,

SpectralMix is slower than other methods and this occurs due to higher dimensionality of the output. Although, we recall that our algorithm is able to achieve better performance as showed in Table 2. Whereas, for other baseline methods, i.e. MARINE and CrossMNA, the main reason for their run time is the high number of edges in ACM dataset.

Our method, SpectralMix, outperforms also the alternative baselines on datasets without node attributes, Flickr, and DBLP. The strength of SpectralMix is highlighted by these experiments. It outperforms baseline methods on different types of graphs: multi-relational graphs with node attributes and multi-relational graphs without node attributes.

Convergence. In Section 4 we state that the value of our objective function decreases and when the difference in the value $\mathcal{F}$ between two consecutive iterations is smaller than ϵ , our algorithm converges. We denote the number of these iterations by I and we show experimentally that, in practice, SpectralMix converges with few iterations. For the two largest datasets, ACM and IMDB, I is equal to 20 and 60 respectively, for $\epsilon = 0.0001$. In Figure 7 we can see that the value $\mathcal{F}$ of the objective function drops and remains almost constant even if we let the algorithm run for 5000 iterations. Similarly, the performance of SpectralMix measured by the values of NMI when I ranges in $[1, 5000]$ remain constant after reaching convergence.

Parameter Analysis. As we have mentioned in Section 3, there is only one parameter d , desired dimensionality, that can be tuned for SpectralMix¹. To show the impact of dimensionality on SpectralMix performance, we use variations of synthetic datasets, differing on number of relations between nodes. Figure 6 shows the

Figure 5: Run time evaluation.

¹The values selected for d are reported in the Appendix.

NMI values and run time of SpectralMix on synthetic dataset, with different dimensional embeddings as outputs. On these variations of synthetic datasets the NMI value is 1 for all different dimensional embeddings with d in the range $[2, 64]$, and is not affected by different numbers of relations between nodes starting from 2 up to 10. We note that the running time increases slightly by having more relations between nodes and higher embeddings, but this is a common feature for methods that use all the available information for computations as the complexity increases too.

Figure 8 shows the impact of dimensionality on clustering performance and number of iterations SpectralMix needs to converge for two real-world datasets. As the value of d increases the drop on performance is noticeable on IMDB dataset. For the ACM dataset we see that the performance changes over different dimensions and the best performance is achieved when we set dimensionality $d = 9$. The main reasons are that: ACM is more dense and, PAP and PSP, have different number of edges (29, 281 and 2, 210, 761 respectively). Thus, this causes the runtime for ACM dataset to be higher than for other datasets, as shown in Figure 5.

Ablation study. To measure the impact of each component of SpectralMix, we conduct ablation studies on the largest dataset, ACM, and we report the results in Table 3. We observe the following facts: 1) As expected, the node attributes help to improve the node clustering performance, therefore, node attributes have significant role for representation learning of nodes. 2) If we set the same weight factors for all attribute categories then the impact is very high, as we see from the drop on the performance on clustering. 3) If we set the same weight factor for all different relation types, then the performance is affected but not in the same scale as it is for attribute categories. 4) The second summand, attribute contribution, in Eq. 1 indeed plays a significant role in optimization. Thus, the ablation study demonstrates that by assigning the different weights to relation types and categories of node attributes, SpectralMix can learn a meaningful node embedding.

Case Study: Brain Network. In Figure 9 we plot the results obtained with SpectralMix. We plot the clusters obtained for an healthy subject $\mathcal{G}_{HL}$ (top) and those obtained for a subject $\mathcal{G}_{SCZ}$

Figure 6: Performance analysis on synthetic data.

Figure 7: Objective function convergence on two real-world datasets: IMDB (left) and ACM (right).

Figure 8: Parameter analysis on real-world datasets.

affected by schizophrenia (bottom) in an axial and sagittal view of the brain. We run experiments considering different number of clusters in the range $[2, 6]$, and we present the result with the highest NCut [10] value, therefore we plot the results for three clusters. We first notice that in $\mathcal{G}_{HL}$ there is a large cluster, i.e., clu_1 , while in $\mathcal{G}_{SCZ}$ many regions in the temporal and frontal lobe belong to a different cluster, i.e., clu_3 . These findings are consistent with works in neuroscience [16, 40] affirming that some specific cognitive deficits present in schizophrenia can be attributed to parietal and frontal lobe dysfunction. Moreover other neuroimaging studies [17] have provided evidence that schizophrenia is associated with reductions in the volume of tissue in the cerebrum and cerebellum, and a disruption of the neural connectivity between them. These results, are coherent with the clustering found by SpectralMix where nodes in the cerebellar area belong to the same cluster in $\mathcal{G}_{SCZ}$ as they are characterized by reduced cerebral volume while in the healthy subject, the same nodes are assigned to 3 different clusters.

6 FUTURE WORK

There are different open problems worth investigating. For the interpretation of the algorithm, it is interesting to quantify how

Table 3: Result for ablation studies of SpectralMix.

ACM Dataset		NMI	ARI
SpectralMix		0.71	0.77
1) No attributes		0.51	0.42
2) Attributes Contribution: $\alpha_1 = \dots = \alpha_{ \mathcal{A} }$		0.01	0
3) Relations Contribution: $\alpha_1 = \dots = \alpha_{ \mathcal{R} }$		0.55	0.61
4) Objective Function:	$\alpha_{\mathcal{R}} > \alpha_{\mathcal{A}}$	0.65	0.71
	$\alpha_{\mathcal{R}} < \alpha_{\mathcal{A}}$	0.69	0.75

Figure 9: Clustering for the graph of a subject in the HL group (top) and for the graph of a subject in the SCZ group (bottom).

much of the link information of each modality is preserved in each Eigenvector. Furthermore, we are interested in answering the question: Are there any dense sub-graphs of a certain link type which are associated with a certain subspace of the attributes? Moreover, we can extend SpectralMix to perform other data mining tasks such as link prediction, node classification, and to infer labels of unlabeled nodes in sparsely labeled multi-graphs, extending the semi-supervised learning method in [37].

REFERENCES

- [1] M. Belkin and P. Niyogi. 2001. Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic* (Vancouver, British Columbia, Canada) (NIPS'01). MIT Press, Cambridge, MA, USA, 585–591.
- [2] P. Chen and A. O. Hero. 2017. Multilayer Spectral Graph Clustering via Convex Layer Aggregation: Theory and Algorithms. *IEEE Tran. on Signal and Info. Proc. over Networks* 3, 3 (2017), 553–567.
- [3] X. Chu, X. Fan, D. Yao, Z. Zhu, J. Huang, and J. Bi. 2019. Cross-Network Embedding for Multi-Network Alignment. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 273–284. <https://doi.org/10.1145/3308558.3313499>
- [4] C. Craddock, Y. Benhajali, C. Chu, F. Chouinard, A. Evans, A. Jakab, B. S. Khundrakpam, J. D. Lewis, Q. Li, M. Milham, C. Yan, and P. Bellec. 2013. The neuro bureau preprocessing initiative: Open sharing of preprocessed neuroimaging data and derivatives. *Front. Neurosci.* (2013).
- [5] J. de Leeuw and P. Mair. 2009. Gif Methods for Optimal Scaling in R: The Package homals. *Journal of Statistical Software, Articles* 31, 4 (2009), 1–21.
- [6] X. Dong, P. Frossard, P. Vanderghenst, and N. Nefedov. 2012. Clustering With Multi-Layer Graphs: A Spectral Perspective. *IEEE Transactions on Signal Processing* 60, 11 (2012), 5820–5831.
- [7] M.-H. Feng, C.-C. Hsu, C.-T. Li, M.-Y. Yeh, and S.-D. Lin. 2019. MARINE: Multi-Relational Network Embeddings with Relational Proximity and Node Attributes. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 470–479. <https://doi.org/10.1145/3308558.3313715>
- [8] S. Günnemann, I. Färber, S. Raubach, and T. Seidl. 2013. Spectral Subspace Clustering for Graphs with Feature Vectors. In *ICDM*.
- [9] L. Hubert and P. Arabie. 1985. Comparing partitions. *Journal of Classification* 2 (12 1985), 193–218. Issue 1. <https://doi.org/10.1007/BF01908075>
- [10] J. Shi and J. Malik. 2000. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 8 (2000), 888–905.
- [11] I. T. Jolliffe. 1986. *Principal Components in Regression Analysis*. Springer New York, New York, NY, 129–155. https://doi.org/10.1007/978-1-4757-1904-8_8
- [12] A. Karaaslanli and S. Aviyente. 2019. Strength Adjusted Multilayer Spectral Clustering. In *2019 IEEE MLSP*. 1–6.
- [13] L. Akoglu, H. Tong, B. Meeder, and C. Faloutsos. 2012. *PICS: Parameter-free Identification of Cohesive Subgroups in Large Attributed Graphs*.
- [14] J. Luo, D. and Ni, S. Wang, Y. Bian, X. Yu, and X. Zhang. 2020. Deep Multi-Graph Clustering via Attentive Cross-Graph Association. In *Proceedings of the Thirteenth ACM International Conference on Web Search and Data Mining*. ACM.
- [15] J. Meier, P. Tewarie, A. Hillebrand, L. Douw, B. W. van Dijk, S. M. Stufflebeam, and P. Van Mieghem. 2016. A Mapping Between Structural and Functional Brain Networks. *Brain connectivity* 6 4 (2016), 298–311.
- [16] A. Mubarik and H. Tohid. 2016. Frontal lobe alterations in schizophrenia: a review. *Trends in psychiatry and psychotherapy* 38 4 (2016), 198–206.
- [17] G. Okugawa. 2013. *Cerebellum and Schizophrenia – The Cerebellum Volume Reduction Theory of Schizophrenia*. Springer Netherlands, 1907–1921.
- [18] A. Papadopoulos, G. Pallis, and M. D. Dikaikos. 2016. Weighted clustering of attributed multi-graphs. *Computing* 99 (12 2016).
- [19] A. Papadopoulos, G. Pallis, and M. D. Dikaikos. 2013. Identifying Clusters with Attribute Homogeneity and Similar Connectivity in Information Networks. In *IEEE/WIC/ACM*.
- [20] A. Papadopoulos, D. Rafailidis, G. Pallis, and M. D. Dikaikos. 2015. Clustering Attributed Multi-graphs with Information Ranking. In *DEXA*.
- [21] C. Park, D. Kim, J. Han, and H. Yu. 2020. Unsupervised Attributed Multiplex Network Embedding. (2020).
- [22] B. Perozzi, L. Akoglu, P. I. Sánchez, and E. Müller. 2014. Focused Clustering and Outlier Detection in Large Attributed Graphs. In *Proceedings of the KDD '14*. 1346–1355. <https://doi.org/10.1145/2623330.2623682>
- [23] C. Plant. 2012. Dependency Clustering across Measurement Scales (KDD '12). New York, NY, USA, 361–369. <https://doi.org/10.1145/2339530.2339589>
- [24] M. Shiga, I. Takigawa, and H. Mamitsuka. 2007. A Spectral Clustering Approach to Optimally Combining Numerical vectors with a Modular Network. In *SIGKDD*.
- [25] K. Steinhauser and N. V. Chawla. 2008. Community Detection in a Large Real-World Social Network. In *Social Computing, Behavioral Modeling, and Prediction*. Springer US, Boston, MA, 168–175.
- [26] X. Sun, Y. Yu, Y. Liang, J. Dong, C. Plant, and C. Böhm. 2021. Fusing attributed and topological global-relations for network embedding. *Inf. Sci.* 558 (2021), 76–90. <https://doi.org/10.1016/j.ins.2021.01.012>
- [27] F. Tang, J. Su, and Y. Wang. 2019. Spectral clustering-based community detection using graph distance and node attributes. *Computational Statistics* 35 (07 2019).
- [28] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot. 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 15, 1 (2002), 273 – 289.
- [29] L. J. P. van der Maaten and G. E. Hinton. 2008. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research* 9 (2008), 2579–2605.
- [30] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm. 2019. Deep Graph Infomax. In *International Conference on Learning Representations*.
- [31] J. Vohryzek, Y. Aleman-Gomez, A. Griffa, J. Raoul, M. Cleusix, P. S. Baumann, P. Conus, K. D. Cuenod, and P. Hagmann. 2020. *Structural and functional connectomes from 27 schizophrenic patients and 27 matched healthy adults*.
- [32] Wang X., Ji H., Shi C., Wang B., Cui P., Yu P., and Ye Y. 2019. Heterogeneous Graph Attention Network. *WWW* (2019).
- [33] Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng. 2012. A Model-based Approach to Attributed Graph Clustering. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*.
- [34] Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng. 2014. GBAGC: A General Bayesian Framework for Attributed Graph Clustering. *ACM Trans. Knowl. Discov. Data* 9, 1 (2014), 5:1–5:43.
- [35] J. Yang, J. McAuley, and J. Leskovec. 2013. Community Detection in Networks with Node Attributes. In *2013 IEEE ICDM*. 1151–1156.
- [36] W. Ye, S. Goebel, C. Plant, and C. Böhm. 2016. FUSE: Full Spectral Clustering. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016*. 1985–1994.
- [37] W. Ye, L. Zhou, D. Mautz, C. Plant, and C. Böhm. 2017. Learning from Labeled and Unlabeled Vertices in Networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Halifax, NS, Canada) (KDD '17). New York, NY, USA, 1265–1274.
- [38] W. Ye, L. Zhou, X. Sun, C. Plant, and C. Böhm. 2017. Attributed Graph Clustering with Unimodal Normalized Cut. In *ECML/PKDD*.
- [39] Z. Zhang, H. Yang, J. Bu, S. Zhou, P. Yu, J. Zhang, M. Ester, and C. Wang. 2018. ANRL: Attributed Network Representation Learning via Deep Neural Networks. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. 3155–3161. <https://doi.org/10.24963/ijcai.2018/438>
- [40] S.-Y. Zhou, M. Suzuki, T. Takahashi, H. Hagino, Y. Kawasaki, M. Matsui, H. Seto, and M. Kurachi. 2007. Parietal lobe volume deficits in schizophrenia spectrum disorders. *Schizophrenia Research* 89, 1 (2007), 35 – 48.
- [41] Y. Zhou, H. Cheng, and J. X. Yu. 2009. Graph Clustering Based on Structural/Attribute Similarities. *Proc. VLDB Endow.* 2, 1 (Aug. 2009), 718–729.
- [42] Y. Zhou, H. Cheng, and J. X. Yu. 2010. Clustering Large Attributed Graphs: An Efficient Incremental Approach. In *2010 IEEE ICDM*.

APPENDIX: IMPLEMENTATION AND REPRODUCIBILITY

The source code of SpectralMix is publicly available ².

Our source code is implemented in Java and Python. We use Java to implement the main part of the SpectralMix algorithm and we have used Python to run K-Means on node embeddings obtained from SpectralMix and other methods too. To run the code call: `java -jar -SpectralMix.jar` with the following argument *source file name*, which is the name of the Matlab file where all the information about adjacency matrix (matrices) and node attributes are stored. Example: `java -jar SpectralMix.jar flickr`.

To change parameters about desired dimensionality or number of iterations for each dataset, please use the class *Sociopatterns.java* (package *datautil*).

Also, to conduct experiments on synthetic datasets and to generate synthetic datasets please use the class *DataGeneratorMain.java*.

The datasets are stored in the folder named *data*. Embeddings for each dataset are stored in the folder named *embeddings*.

First we run the Java source code implementation, then we get the final embeddings, and we need to run Python script for K-Means on the final embeddings to get the cluster labels for each node, for this purpose we run the Python script stored on K-Means python notebook.

All experimental data are open source. Table 4 summarizes their characteristics. A brief description for the construction of the first four datasets (*ACM*, *IMDB*, *DBLP*, *Flickr*) is provided below.

- *ACM* [21]: The heterogeneous graph is constructed by using information from 3,025 papers (P), 5,835 authors (A) and 56 subjects (S). There are 1870 attributes for each paper.
- *IMDB* [21]: The constructed heterogeneous graph includes information from 3,550 movies (M), 4,441 actors (A), and 1,726 directors (D). There are 2000 attributes for each movie.
- *DBLP* [14]: The collaboration graph has 2,401 author nodes and 8,703 edges. The citation graph has 6,000 paper nodes and 10,003 edges. The co-citation graph has 6,000 paper nodes and 141,996 edges. The authorship graph has 64,096 edges.
- *Flickr* [14]: Each relation type involves 10,364 users as nodes. The friendship relation type has 401,302 edges. The tag-similarity relation type has 125,547 edges, where each edge represents the tag similarities between two users. Two nodes in these two relations (graphs) are linked if they refer to the same user.

Table 4: Statistics of datasets.

datasets	#relations	#nodes	#edges	#atts	#clusters
ACM	2	8,916	2,240,042	1,870	3
IMDB	2	9,717	80,216	2,000	3
DBLP	4	14,401	224,798	0	3
Flickr	2	20,728	537,213	0	7
BN-TD	51	116	66,744	6	3
BN-ASD	49	116	62,907	6	3

²<https://gitlab.cs.univie.ac.at/yllis19cs/spectralmixpublic>

Experimental Setup: Parameter Settings

We summarize details about the parameter settings of SpectralMix. The only tuning parameter is d and here we provide the desired dimensionality for each dataset:

- *ACM*: Dimensionality = 9.
- *IMDB*: Dimensionality = 2.
- *Flickr*: Dimensionality = 11.
- *DBLP*: Dimensionality = 2.
- *BrainNet-Healthy*: Dimensionality = 2.
- *BrainNet-Autism*: Dimensionality = 2.

Moreover, we set the number of iterations for convergence to 100 for all datasets. In the experiments, we notice that the algorithm converges with less than 100 iterations.

More details on how to run experiments are available at the provided link on the paper.

For the comparison methods, we use the source code and the parameter settings provided by authors of each method and can be found in the following links:

- *ANRL* [39]: python implementation is publicly available at <https://github.com/cszhangzhen/ANRL>
- *DGI* [30]: python implementation is publicly available at <https://github.com/PetarV-/DGI>
- *CrossMNA* [3]: python implementation is publicly available at <https://github.com/ChuXiaokai/CrossMNA>
- *DMGC* [14]: python implementation is publicly available at <https://github.com/flyingdoog/DMGC>
- *HAN* [32]: python implementation is publicly available at <https://github.com/Jhy1993/HAN>
- *MARINE* [7]: python implementation is publicly available at <https://github.com/ntumslab/MARINE>
- *DMGI* [21]: python implementation is publicly available at <https://github.com/pcy1302/DMGI>