

Efficient Network Embedding by Approximate Equitable Partitions

Giuseppe Squillace
IMT Lucca, Italy
giuseppe.squillace@imtlucca.it

Mirco Tribastone
IMT Lucca, Italy
mirco.tribastone@imtlucca.it

Max Tschaikowski
Aalborg University, Denmark
tschaikowski@cs.aau.dk

Andrea Vandin
Sant'Anna Pisa, Italy
DTU, Denmark
anvan@dtu.dk

Abstract—Structural network embedding is a crucial step in enabling effective downstream tasks for complex systems that aim to project a network into a lower-dimensional space while preserving similarities among nodes. We introduce a simple and efficient embedding technique based on approximate variants of equitable partitions. The approximation consists in introducing a user-tunable tolerance parameter relaxing the otherwise strict condition for exact equitable partitions that can be hardly found in real-world networks. We exploit a relationship between equitable partitions and equivalence relations for Markov chains and ordinary differential equations to develop a partition refinement algorithm for computing an approximate equitable partition in polynomial time. We compare our method against state-of-the-art embedding techniques on benchmark networks. We report comparable—when not superior—performance for visualization, classification, and regression tasks at a cost between one and three orders of magnitude smaller using a prototype implementation, enabling the embedding of large-scale networks that could not be efficiently handled by most of the competing techniques.

Index Terms—Equitable partitions, network embedding, backward equivalence, structural equivalence

I. INTRODUCTION

Network analysis is crucial in various fields, including, e.g., social networks [1], bioinformatics [2], and recommendation systems [3]. The increasing availability of data leads to complex networks that are difficult to interpret and challenging computationally. Network embedding tackles this problem by extracting properties of a network and encoding them in a lower-dimensional space [4]; one of the main objectives is to preserve the similarity of nodes in the embedded space [5].

Traditionally, the notion of similarity between nodes is based on network proximity. This approach is well suited for community detection but is not effective in capturing the *role* of a node in the network [6], i.e., patterns of interaction beyond simple local proximity. The notion of role, also known as position, is well studied in the social science literature (e.g., [7]). Roles partition nodes into classes according to some criteria such as structural [8], automorphic [9], and regular equivalence [10]. Roughly speaking, in structural equivalence,

two nodes have the same role if they interact with the same set of nodes. In automorphic equivalence, two nodes are equivalent if we can define an automorphism between them. Finally, two nodes are regular equivalent if they interact with the same variety of role classes (cf. Figure 1).

Structural equivalence is strict because it requires that related nodes interact with *the same* set of nodes; on the other hand, automorphic equivalence suffers from well-known difficulties in its computation. Here we propose a structural embedding method for undirected networks based on equitable partitions [12]. This concept is also known as exact regular coloration [13] and graph divisor [14]. The idea is to assign a role in the network to each node in such a way that nodes with the same role have the same number of neighbors for each role. It is a relaxation of automorphic equivalence [11], which can be computed efficiently in polynomial time in the size of the network [13]. It is stricter than regular equivalence because regular equivalence does not count the number of neighbors with the same role [8]. However, since it requires exact matching of neighbor counts, an equitable partition can still be too restrictive for real networks.

We develop a novel approach to compute a network embedding based on an approximate variant of equitable partitions that accepts differences in the number of neighbors by means of a parameter ε that, informally, determines the tolerance level for identifying similar nodes. We anticipate the results of our method on a simple network depicted in Figures 2 and 3. Figure 2 shows the equitable partition where node 2 is different from nodes 3 and 4 because it has a different degree. Our approximate variant, with $\varepsilon = 1$ accepts differences up to one node: the result is shown in Figure 3, where nodes 2, 3, and 4 (as well as their neighbors in the bottom layer) are identified to be approximately equivalent, thus capturing their

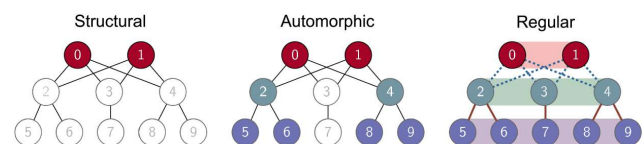


Fig. 1. Role equivalences (example from [11]). Nodes with the same color indicate the same role; white nodes have pairwise distinct roles.

This work was partially supported by Poul Due Jensen Foundation grant no. 883901, the Villum Investigator Grant S40S, the project SERICS (PE00000014), the Fsc regional Tuscan project AISLEA2 J54D23000780005, the project Tuscany Health Ecosystem (THE), CUP: B83C22003920001, the project SMaRT COOnSTRUCT (CUP J53C24001460006), within FAIR (CUP I53C22001380006). The last two projects are under the MUR National Recovery and Resilience Plan funded by the EU - NextGenerationEU.

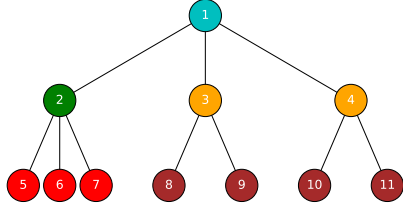


Fig. 2. Running example with equitable partitions.

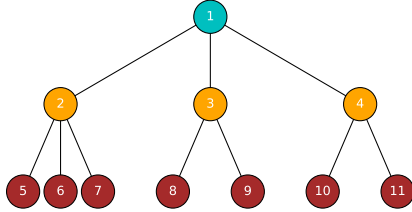


Fig. 3. Approximate variant of equitable partitions computed with ε -BE on the running example from Figure 2 with $\varepsilon = 1$.

intuitive structural similarity in the network.

Our starting point is the observation that, for undirected networks, the equitable partition definition corresponds to an equivalence relation known as backward equivalence (BE) [15]. BE extends exact lumpability [16], an aggregation method for Markov chains based on properties of its generator matrix, to arbitrary adjacency matrices [17] and nonlinear models [18]–[20]. Our network embedding method, which we call ε -BE, relaxes the equivalence criterion by means of the tolerance parameter ε .

A crucial advantage of BE, inherited from results available for Markov chains, is the availability of an efficient algorithm for computing it that runs in $O(m \log n)$ time [21], where m is the number of edges and n is the number of nodes. The algorithm is based on partition refinement [22]: given an initial partition of nodes, where each block represents a candidate role, it iteratively splits blocks until the BE criterion is met. In this paper, we develop an analogous algorithm for ε -BE, still based on [21]. Exploiting the partition refinement nature of the algorithm, we use it as an inner step of an iterative scheme that computes an ε -BE partition with increasing values of ε , keeping track of the positions discovered at previous iterations. As will be explained later in the paper, this avoids the degenerate cases where the algorithm may aggregate too little or too much when using a single-shot strategy with small or large enough, respectively, values of ε .

The resulting partition is employed to build the network embedding E . We are interested in producing a network embedding $E \in \mathbb{R}^{n \times d}$ where n is the number of the nodes, and d is the size of the embedding. This can be built by summing up the edges toward equivalent classes for each node. Table I shows the embeddings for the two partitions of Figures 2 and 3. The dimension of the embedding equals the number of different blocks in each partition. In particular, nodes with the same role exhibit similar vectors in the network embedding

Nodes	Embedding	Nodes	Embedding
{1}	[0,1,2,0,0]	{1}	[0,3,0]
{2}	[1,0,0,3,0]	{2}	[1,0,3]
{3,4}	[1,0,0,0,2]	{3,4}	[1,0,2]
{5,...,7}	[0,1,0,0,0]	{5,...,11}	[0,1,0]
{8,...,11}	[0,0,1,0,0]		

TABLE I
NODE REPRESENTATION IN THE EMBEDDING SPACE ASSOCIATED WITH THE EQUITABLE PARTITION OF FIGURE 2 (LEFT) AND FIGURE 3 (RIGHT).

and the ε -BE condition can be checked by computing the difference of the embedding vectors. For example, consider the nodes 2 and 3 of Table I (right). The difference of their embedding vectors is [0,0,1]. Each entry of this vector is $\leq \varepsilon$, and, for this reason, they present the same color in Figure 3.

To compare ε -BE against state of the art, we consider the best-performing methods to network embedding covering all three aforementioned categories presented in the recent review [4]. We compare the network embeddings across various machine-learning tasks such as classification, visualization, and regression using real-world networks with up to 10^5 nodes.

Overall, we show that ε -BE is able to provide competitive results, outperforming the other methods in many cases, with runtimes at least one order of magnitude faster. For larger networks, in particular, ε -BE provided node embeddings in a few minutes while most of the other methods timed out after 3 hours, thus promoting ε -BE as a valid alternative for the structural network embedding

Related work: Methods for structural network embeddings can be categorized into matrix factorization, random walk, and deep learning methods.

Matrix factorization methods capture structural properties by decomposing the adjacency matrix or other related matrices into lower-dimensional matrices. The most closely related approach to ours is RID ε Rs [23]. It is also based on a notion of ε -equitable refinement. However, the embedding method is substantially different. Indeed, instead of the iterative scheme proposed in this paper, RID ε Rs computes an embedding E_ε for each value of ε ranging from 0 to the average degree of the network d_{avg} . Since the overall dimension of the embedding can be large, a feature pruning technique is employed to reduce this dimension. Subsequently, RID ε Rs aggregates the pruned embedding into a single matrix $E = [E_0 | \dots | E_{d_{avg}}]$ and employs negative matrix factorization (NMF [24]) on E to generate the final embedding. RoIX [25] factorizes the network into a role membership matrix and a role-feature matrix. The matrix factorization is optimized by minimizing the reconstruction error. GLRD extends RoIX by adding further constraints in the previous optimization problem [26]. GraphWave is based on a physics interpretation [27]. It considers the application of the heat diffusion process on the network. Then it computes the graph diffusion kernel and gets the embedding using the characteristic function. xNetMF computes a similarity matrix computed between nodes, and

then it factorizes this matrix [28]. SEGK uses graph kernels to compute node structural similarities [29]. The standard one is the shortest path kernel, which uses the Nystrom method [30] for factorization. SPaE provides a structural embedding via the Laplacian eigenmaps method [31]. Approaches based on random walks explore the networks with paths of a given length. According to the taxonomy in [4], network embeddings are usually learned using methods from natural language processing. One popular technique is struc2vec [32]. It generates a multi-layer weighted context graph, and then it employs SkipGram [33] to learn the embedding. Role2vec assigns to each node a role and then builds a random walk where the generated sequence is replaced with role indicators [34]. NODE2BITS employs temporal random walks where the edges are sampled with non-decreasing timestamps [35]. The last category learns the embedding through deep learning techniques such as neural networks and graph neural networks. Among these, we mention DRNE [36]. It produces a network embedding based on the regular equivalence definition. The method is based on the training of a Long Short-Term Memory (LSTM) neural network. GAS employs GNNs exploiting the Weisfeiler-Lehman test [37]. Other works in this direction are GraLSP [38], which incorporates random walks in the GNN, and GCC [39], which uses a pre-trained Graph Neural Network model.

II. BACKGROUND

We consider an undirected graph with a set of nodes V denoted by a symmetric adjacency matrix $A = (a_{ij}) \in \mathbb{R}^{n \times n}$, where $a_{ij} = 1$ if there is an edge between node i and j . Node embedding corresponds to finding a matrix $E \in \mathbb{R}^{n \times d}$ where $d < n$. We are interested in an embedding that preserves the similarities between nodes with similar roles in the network.

We present the definition of BE. Originally developed for hypergraphs [15], here we provide a simplified version for graphs. Informally, a BE is a partition of the nodes such that nodes belonging to the same block present the same cumulative degrees towards equivalence classes.

Definition 1 (Backward equivalence, simplified to graphs from [15]). *For an adjacency matrix $A \in \mathbb{R}^{n \times n}$, a partition of the nodes \mathcal{H} is called a backward equivalence (BE) when*

$$\sum_{k \in B'} a_{ki} = \sum_{k \in B'} a_{kj}$$

for all blocks $B, B' \in \mathcal{H}$ and $i, j \in B$.

When simplified to graphs, BE is equivalent to the equitable partition definition presented in [23], as is apparent from comparing the previous and the forthcoming definitions.

Definition 2 (Equitable partition, adapted from [23]). *For a network with nodes V , and edges E , a partition \mathcal{H} of V is an equitable partition if and only if*

$$\text{deg}(v_i, B') = \text{deg}(v_j, B')$$

for all blocks $B, B' \in \mathcal{H}$, and $v_i, v_j \in B$, where

$$\text{deg}(v_l, B'') = |\{v_k | (v_l, v_k) \in E \text{ and } v_k \in B''\}|$$

The colors in Figure 2 identify the different blocks of a BE/equitable partition. We now relax the notion of BE to account for tolerances in node degrees. Here, we consider the tolerance ε as a measure to accept approximate and not only exact equitable partitions.

Definition 3 (ε -BE). *For an adjacency matrix $A \in \mathbb{R}^{n \times n}$, a partition of the nodes \mathcal{H} is called ε -BE partition when*

$$\left| \sum_{k \in B'} a_{ki} - \sum_{k \in B'} a_{kj} \right| \leq \varepsilon$$

for all blocks $B, B' \in \mathcal{H}$ and $i, j \in B$.

We can see that ε -BE corresponds to BE for $\varepsilon = 0$. However, by increasing ε , two nodes can be related even if they have similar degrees towards (approximately) equivalent blocks. The parameter ε tunes the level of similarity of the nodes in the same block. For instance, setting $\varepsilon = 1$ allows larger identifications of similar states in Figure 3 as opposed to the equitable partition in Figure 2.

This definition is inspired by a technique to reduce dynamical systems called ε -BDE [40]. The main difference between the two approaches lies in the use of the tolerance ε . In ε -BDE, the cumulative difference in the degree towards all equivalent blocks should be less than ε . Instead ε -BE is less strict as it requires the difference towards each block to be smaller than ε . With this change, it turns out that, for undirected networks, ε -BE corresponds to the notion of approximate equitable partition from [23].

We build a node embedding using ε -BE as follows.

Definition 4 (ε -BE embedding). *Let us consider an adjacency matrix $A \in \mathbb{R}^{n \times n}$ and an ε -BE partition $\mathcal{H} = \{B_1, \dots, B_d\}$ with $d < n$. We define the network embedding matrix $E \in \mathbb{R}^{n \times d}$ as follows:*

$$E_{ik} = \sum_{j \in B_k} a_{ij} \quad (1)$$

for all $i = \{1, \dots, n\}$ and $k = \{1, \dots, d\}$.

The matrix E collects the encoding of each node in the low dimensional space \mathbb{R}^d . The value E_{ik} corresponds to the number of edges connecting node i to nodes in the equivalence class B_k . Notice that, for $\varepsilon = 0$, BE-equivalent nodes have the same vector in the embedded space E .

III. COMPUTATION OF ITERATIVE ε -BE

A. Partition-refinement Algorithm

We present an efficient algorithm to find an ε -BE partition based on the Markov chain lumping algorithm from [21]. We report the pseudocode in Algorithm 1. The algorithm starts with an initial partition and refines it, splitting the blocks, until it gets a partition where the ε -BE condition is satisfied. In general, the initial partition can be used to avoid the aggregation of undesired nodes; indeed, nodes in different blocks of the initial partition cannot be aggregated. Throughout this paper, initial partitions always consist of a single block containing all nodes. We will see how the notion of initial

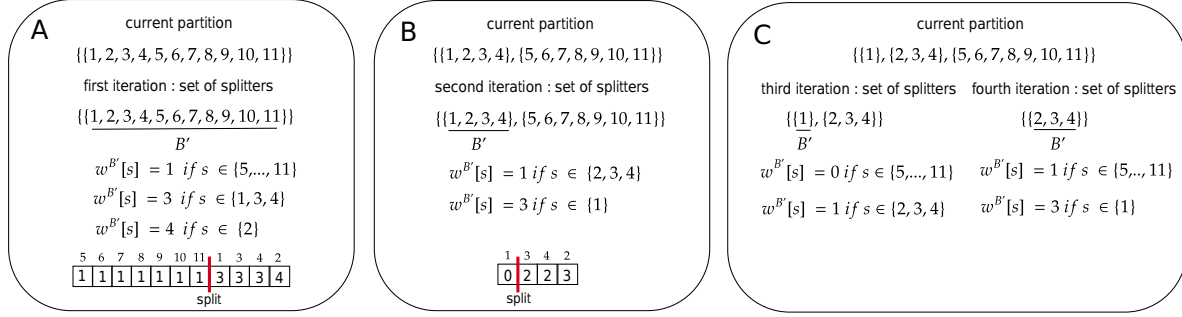


Fig. 4. Execution of Algorithm 1. We indicate the current splitter by underlining it. The weight $w^{B'}$ denotes the number of edges connected to elements of the splitter B' . The arrays at the bottom show the splitting phase of a certain block B with respect to the current splitter. The red line indicates where the block is split because it does not respect the tolerance ε .

partition is instrumental to obtain an iterative version of the algorithm in Section III-B.

The algorithm iteratively scans a list of *splitters*, i.e., partition blocks. The list is initialized with the initial partition. It removes each splitter from the list and checks whether all blocks satisfy the condition for the splitter. The blocks not satisfying the condition are refined such that the ε -BE condition is satisfied for the splitter; the newly identified blocks are added to the list of splitters. The computation terminates when the list of splitters is empty.

For an illustrative explanation, Figure 4 reports the computation of ε -BE for the running example from Figure 2 with $\varepsilon = 1$. It is divided into three phases (panels A, B, and C in Figure 4). In panel A, we show the set of splitters U as the blocks in the initial partition (line 1). The procedure removes a block from U and initializes S_T , the set of nodes connected to any node in the current splitter (lines 3-5). In Figure 4 we indicate the current splitter with B' . Then, the algorithm computes the $w^{B'}[s]$ as the number of edges connecting the node s to a node in the current splitter B' (lines 6-14). In panel A, $w^{B'}[s]$ is computed for the first block. The nodes $\{5, \dots, 11\}$ are connected to one node belonging to B' , $\{1, 3, 4\}$ to three nodes, and $\{2\}$ to four nodes.

The blocks containing at least a node connected to the current splitter are retrieved and stored in B_T (lines 15-20). For each block B in B_T , the algorithm sorts $w^{B'}[s]$ with $s \in B$ and splits them into blocks (lines 21-30). Here, differently than [21], we allow a tolerance ε in the aggregation of the nodes.

Figure 4 shows the splitting process in the bottom part of panel A. The nodes s belonging to B are sorted by $w^{B'}[s]$. If the difference between the maximum and minimum nodes is greater than ε , it splits the block. A red line indicates the split. For instance, when node $\{1\}$ is considered, the difference is greater than $\varepsilon = 1$ and it splits the block.

In line 23, the algorithm finds a possible majority candidate (pmc) in the array $w^{B'}[s]$ with $s \in B$. That is, a value x that is present in more than half of the positions of the array; in the case where no such value exists, the pmc can be any value present in the array. The pmc can be computed in linear

Algorithm 1 ε -BE

Require: Adjacency matrix A , initial partition \mathcal{H}_{in} , tolerance $\varepsilon \geq 0$.

- 1: $U = \mathcal{H}_{in}$; $B_T = \{\}$;
- 2: **while** $U \neq \emptyset$ **do**
- 3: $w^{B'}[s] = 0$ for each node s
- 4: $B' =$ remove a block from U .
- 5: $S_T = \{\}$
- 6: **for** $s' \in B'$ **do**
- 7: **for** s connected to s' **do**
- 8: **if** $w^{B'}[s] == 0$ **then**
- 9: $S_T = S_T \cup \{s\}$ $w^{B'}[s] = A[s][s']$
- 10: **else**
- 11: $w^{B'}[s] = w^{B'}[s] + A[s][s']$
- 12: **end if**
- 13: **end for**
- 14: **end for**
- 15: **for** $s \in S_T$ **do**
- 16: $B =$ the block that contains s
- 17: **if** B not in B_T **then**
- 18: $B_T = B_T \cup \{B\}$
- 19: **end if**
- 20: **end for**
- 21: **while** $B_T \neq \emptyset$ **do**
- 22: $B =$ remove a block from B_T
- 23: find possible majority candidate of $w^{B'}[s]$ for $s \in B$
- 24: sort B without nodes with $w^{B'}[s] = pmc$
- 25: add nodes with $w^{B'}[s] = pmc$ to B
- 26: partition B in B_1, B_2, \dots, B_l respect to tolerance ε
- 27: **if** $l > 1$ **then**
- 28: add B_1, B_2, \dots, B_l to U
- 29: **end if**
- 30: **end while**
- 31: **for** $s \in S_T$ **do**
- 32: $w^{B'}[s] = 0$
- 33: **end for**
- 34: **end while**

Algorithm 2 Iterative ε -BE.

Require: Adjacency matrix A , initial partition of nodes \mathcal{H}_{in} , an initial tolerance ε_0 , step size $\delta \geq 0$, maximum tolerance $\Delta \geq 0$.

```
1:  $\varepsilon = \varepsilon_0$ 
2: while  $\varepsilon \leq \Delta$  do
3:    $\mathcal{H}_\varepsilon = \varepsilon$ -BE( $A, \mathcal{H}_{in}, \varepsilon$ )
4:    $\mathcal{H}_{in} = \text{joinSingletons}(\mathcal{H}_\varepsilon)$ 
5:    $\varepsilon = \varepsilon + \delta$ 
6: end while
7: return  $\mathcal{H}_\varepsilon$ 
```

time and is useful to reduce the number of nodes to be sorted in line 24. Here the algorithm sorts the nodes in B except the ones with $w^{B'}[s] = pmc$. For a detailed discussion of the algorithm, see [21].

After the splitting process, in panel B of Figure 4, the current partition is composed of two blocks. The first block $\{1, \dots, 4\}$, where the difference is 0, and the second block $\{5, \dots, 11\}$, where the difference is 1. On line 27, we evaluate whether the number of new blocks exceeds 1. If this condition holds true, the new blocks are incorporated into U as additional splitters. When l equals 1, B is not split into any new blocks, so we avoid adding itself to U .

The algorithm considers a new splitter B' and tries to split the first block. It considers the sorted $w^{B'}$ and it finds that node 1 can be separated. At this point, the set of splitters and the partition will be updated. In phase C, we consider the remaining splitters. No actual refinements are performed, hence the procedure terminates. Therefore, the output is a partition with three blocks as pictured in Figure 3.

B. Iterative ε -BE

One of the crucial aspects of Algorithm 1 is the choice of the tolerance ε : increasing this value leads to partition with few, large blocks, but it may collapse the network too aggressively. For instance, running it with $\varepsilon = 3$ on the running example in Figure 2 results in the trivial partition composed of a single block. For an effective network embedding, we propose an iterative strategy. The pseudocode is shown in Algorithm 2.

It requires the following parameters: an initial tolerance, denoted by ε_0 , a step size δ , which determines the increment of ε at each step, and Δ , the maximum value of ε to be considered. The main loop iterates in lines 2-6 until it reaches the maximum threshold Δ . In line 3, Algorithm 1 is called to find the partition up to a certain threshold ε . The so-computed partition is stored in \mathcal{H}_ε (line 3). In line 4, all the singleton blocks in the current partition \mathcal{H}_ε are joined into one block for the next iteration. The intuition is to attempt node aggregation for smaller values of ε first. If that fails, i.e., nodes are eventually outputted as singletons blocks, the merging of such nodes is used to attempt aggregation for the larger ε values in the next iterations. Finally, in line 5, the algorithm increments the value of ε by δ .

Figure 5 provides an illustrative example. It shows a network with two cliques connected by edges; the one on the left side is a complete clique (nodes 1, 2, 4, 3), while the one on the right has some missing edges (nodes 7, 8, 9, 10). Intuitively, one expects to group the elements of the complete clique with $\varepsilon = 0$, while for the incomplete one, we expect to find an association of the cliques' nodes considering a tolerance greater than 0. We employ the Algorithm 2 in the example in Figure 5. We set up $\varepsilon_0 = 0$, $\delta = 1$ and $\Delta = 1$. this corresponds to two iterations, with ε equal to 0 and 1, shown in panels A and B of Figure 5, respectively.

Panel A shows the initial partition \mathcal{H}_{in} and the coloration for $\varepsilon = 0$. As expected, the nodes in the left clique present the same role. In panel B the initial partition \mathcal{H}_{in} presents the merged singleton blocks $\{3\}$, $\{7\}$, $\{8\}$, $\{9\}$, $\{10\}$ of \mathcal{H}_ε at the previous iteration of panel A.

In the subsequent iteration, the algorithm computes the partition with $\varepsilon = 1$, considering as new initial partition the one computed in the previous iteration. The results are depicted in panel B, where, due to the tolerance parameter ε , nodes within the incomplete clique are denoted with the same color. This example demonstrates how the algorithm can collect similar groups of nodes while gradually increasing the tolerance ε .

To show the advantages brought by the iterative scheme, panel C shows the output of a *one-shot* application of Algorithm 1 directly with $\varepsilon = 1$. This comparison shows that the iterative algorithm is able to better assign the roles to nodes in the two cliques, while the one-shot application identifies nodes 5, 6, 8, and 9, in contrast with their intuitively different roles in the network.

C. Complexity

The original partition-refinement algorithm from [21] runs in $O(m \log(n))$ time, where n is the number of nodes and m is the number of edges. The logarithmic term is achieved by discarding the largest block in each iteration. Unfortunately, for an approximate reduction, as presented here, discarding the largest block does not assure correctness (i.e., the partition does not satisfy the ε -BE condition). Consequently, the complexity of the Algorithm 1 is $O(mn)$. The splitting of the blocks according to the tolerance ε can be implemented in linear time once the nodes are sorted for their weights w and does affect the overall complexity.

Algorithm 1 is employed as an inner step in Algorithm 2. An important point to consider is how many times lines 2-6 are repeated. In this case, the maximum threshold Δ and the step size δ have to be considered. The inner loop will be repeated $\lceil \Delta/\delta \rceil$ times. Since every element of the adjacency matrix is either 0 or 1, the minimum δ one can consider is 1, so that complexity is $O(mn\Delta)$. Despite this, the parameter Δ is always orders of magnitude smaller than n to avoid trivial results where all the nodes belong to the same block.

IV. NUMERICAL EXPERIMENTS

In this section, we present the experimental results achieved by ε -BE compared with state-of-the-art methods for network

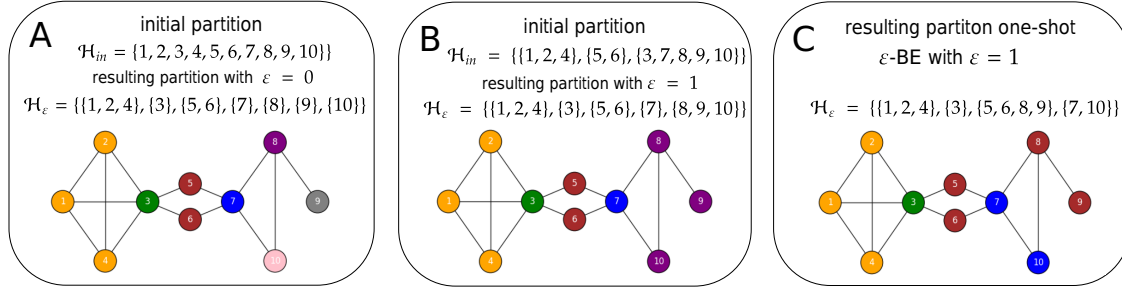


Fig. 5. Execution of the iterative ε -BE on a network composed of a complete clique on the left, and another on the right with missing edges. We show the first and the second iteration with ε equal to 0 and 1 in panels A and B. For each panel, we indicate the initial partition and the resulting one after the application of ε -BE. The resulting partition of a phase will be employed as the initial partition in the following phase after merging the singleton blocks. The initial partition \mathcal{H}_{in} in panel B corresponds to \mathcal{H}_ε in panel A where the singleton blocks are merged into a single block. We depict the network with the different colors assigned by the resulting partition. Panel C represents the result of the application of Algorithm 1 without the iterative scheme.

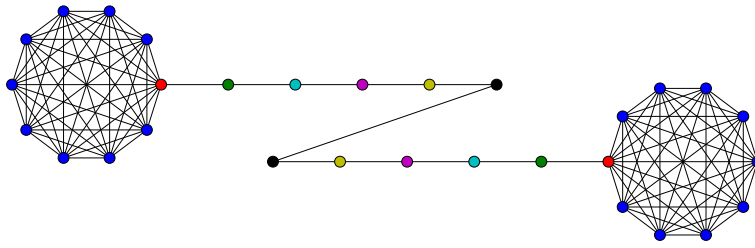


Fig. 6. Barbell role assignment.

embedding. All the experiments and datasets with the implementation of our method are accessible on the GitHub page.¹ We considered the best methods reported in the recent review [4]. Specifically, we chose Graphwave [27], SEGK [29] and RID ε Rs [23] as representatives of matrix factorization techniques, struc2vec [32] as a random-walks-based approach, and DRNE [36] and GAS [37] as representatives of deep learning methodologies. The implementation details for each method were taken from GitHub repositories mentioned in the respective publications or in [4].² Similarly to [4], [27], [29], the comparison was conducted on synthetic and real networks considering three machine-learning tasks: visualization, classification, and regression. All experiments were run on a machine equipped with an Intel(R) Xeon(R) CPU E7-4830 with 500 GB RAM. We set up a timeout of 3 hours for all the computations reported in this section.

A. Set-up

Most of the methods in network embedding allow the user to select the desired embedding dimension d for the network. However, choosing the optimal dimension is challenging [41]. In this study, we followed [29] and [27] in adopting the default embedding dimensions recommended for each method. Specifically, we set d to the following values: 20 (SEGK), 100 (Graphwave), 64 (DRNE), and 128 (struct2vec and GAS).

¹GitHub link: <https://github.com/EmbNet01/EmbNetworks.git>

²GitHub repositories accessed on 01/06/2023.

Regarding RID ε Rs, the embedding dimension was determined through the NMF method, optimizing the minimum description length criterion. In addition to the size of the embedding, each method requires specifying other parameters. We set up all these parameters with their default values.

ε -BE yields a discrete embedding based on the partition computed by Algorithm 2. The reduction with the maximum number of blocks corresponds to a partition where every node is a singleton, resulting in an embedding equivalent to the adjacency matrix itself. In this scenario, similar nodes are defined as those that are connected to the exact same set of nodes. Conversely, considering a sufficiently large ε , the partition with the minimum number of blocks groups all nodes into a single block, resulting in an embedding dimension d equal to 1. Here, each entry represents the sum of the elements in the respective columns, and similar nodes are defined as those with comparable degrees. These represent two extreme reductions of the embedding space. Given that the primary objective of network embedding is to offer an encoding where d is significantly smaller than N , we set the parameters to achieve a reduction of at least 70% of the network dimension N .

B. Visualization

Network visualization is a fundamental task, particularly in scenarios where the roles of network nodes are predefined. The primary goal is to project these nodes into a two-dimensional

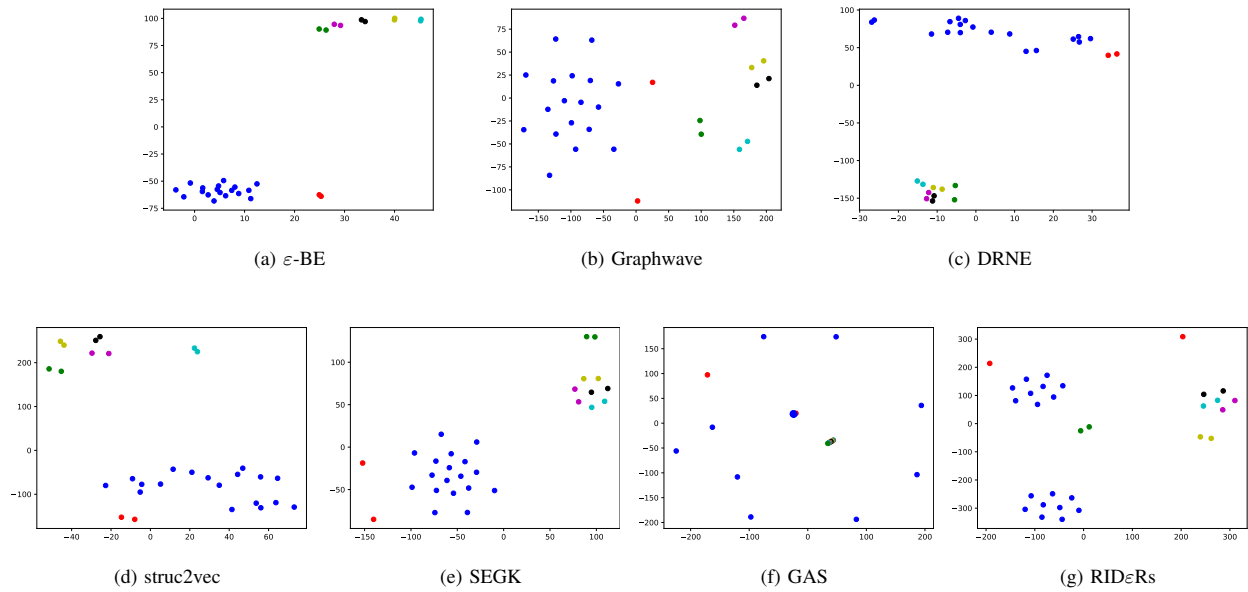


Fig. 7. Barbell visualization with t-SNE.

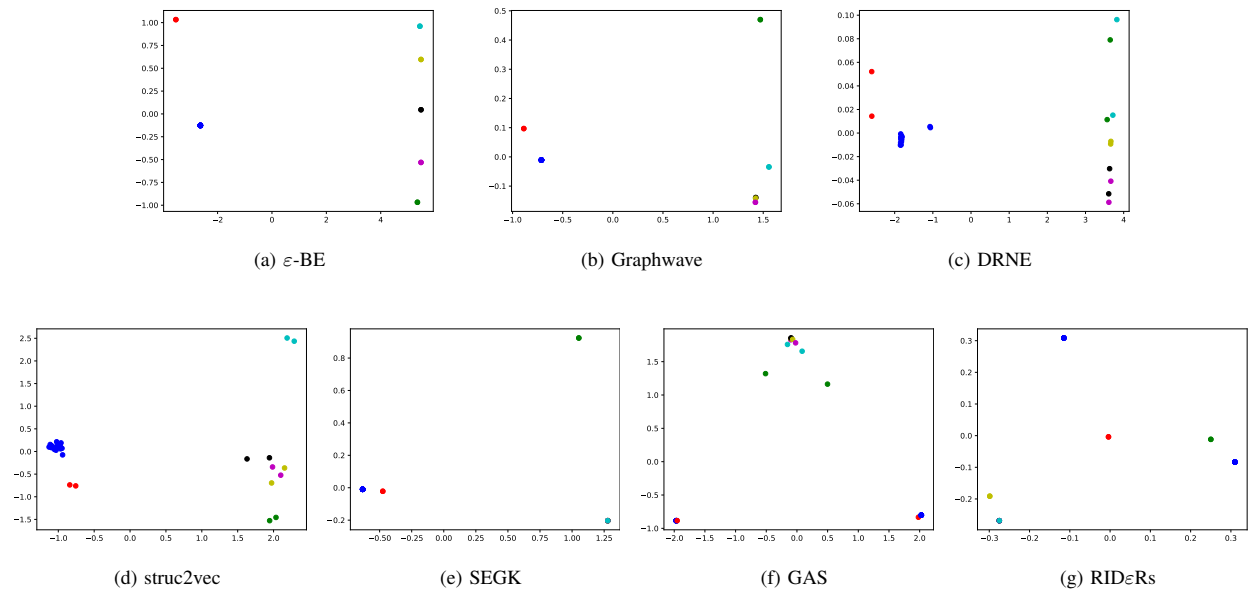


Fig. 8. Barbell visualization with PCA.

space for visualization purposes. An effective embedding technique should position nodes with similar roles in close proximity within this reduced-dimensional space.

ε -BE, RID ε Rs, and Graphwave are not directly suitable for visualization. For this reason, we further reduce the embedding space of dimension d into two dimensions using PCA and t-SNE [42]. The former is a deterministic approach based on matrix factorization. The latter is a probabilistic method based on a variation of Stochastic Neighbor Embedding [43].

We consider the well-known barbell graph depicted in Figure 6, a standard benchmark in the literature for evaluating visualization tasks [27], [29], [32]. It consists of two cliques connected by a chain of nodes, where each node is assigned a distinct role represented by different colors. The nodes within each clique share the same role, except for the red nodes, which are additionally connected to the chain. Conversely, the roles within the chain are specular. To demonstrate the efficacy of ε -BE, we compute it with $\varepsilon = 0$.

Figure 7 shows the visualizations by t-SNE reduction achieved by all the different methods. We observe notable differences in the performance of different network embedding approaches. GAS and RID ε Rs exhibit poorer performance than the other methods, especially GAS, where the majority of nodes are clustered closely together in the embedding space, making it challenging to discern distinct groups. In contrast, SEGK and Graphwave have superior performance by effectively distinguishing nodes within cliques. However, they tend to separate the red nodes into two separate regions within the embedding space. On the other hand, struc2vec and DRNE perform well, providing a good spatial organization. Interestingly, our approach yields results similar to struc2vec with a cleaner separation within the 2-dimensional space.

Considering the inherent probabilistic nature of t-SNE, we also considered a deterministic approach using PCA, shown in Figure 8. In this case, nodes with the same encoding in the embedded space are mapped in the same point in the 2-dimensional space. As a consequence, all nodes with the same role are mapped onto the same point for ε -BE. Graphwave, RID ε Rs, and SEGK show similar behavior, although they differentiate chain nodes less effectively. The other methods present a more sparse allocation without nodes that overlap perfectly. This is due to the fact that they rely on random walks and deep learning techniques. Consequently, it is unlikely to yield embedding vectors with identical components across these approaches. As a result, the application of PCA to visualize the embeddings generated by these methods yields less uniform results compared to other techniques.

Overall, for this case study ε -BE approach emerges as the optimal strategy for accurate visualization using PCA.

C. Classification

The objective of classification is to train a model that can accurately assign appropriate roles to new nodes. We proceed by computing the embedding space for each network, followed by a 5-fold cross-validation to assess the accuracy of role

Network	N	d	δ	Δ	ε_0
Brazil airport	131	38	1	3	0
Europe airport	399	96	1	4	0
USA airport	1190	325	1	2	0
actor	7779	201	2	8	2
Film	27312	177	6	30	6

TABLE II
PARAMETER SETTINGS FOR THE ITERATIVE ε -BE ALGORITHM.

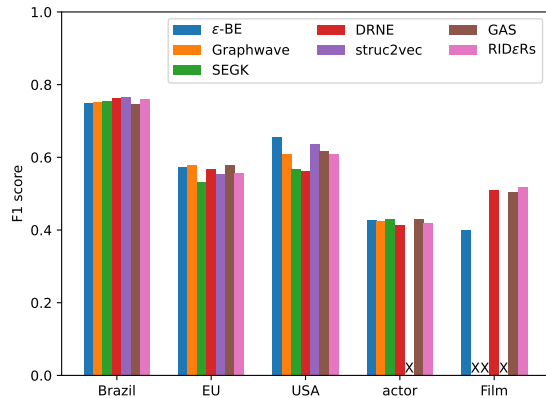


Fig. 9. Airport networks classification results. The bar indicated with 'X' refers to methods that timed out (3h).

prediction. To evaluate classification performance, we compute the average F1 score.

We consider different real-world networks. We started with the air traffic networks from Brazil, the United States, and Europe, presented in [32]; additionally, we use an actor co-occurrence network and the Film network from [4]. Air traffic networks are widely recognized within the network community and serve as standardized benchmarks for comparative evaluations [32].

All such networks come with ground truth labels that specify the role of each node. For each method, we computed the corresponding network embedding. We show in Table II, together with the original network size N and the embedding size d , the parameters utilized for computing the network embeddings with our approach. We averaged the results over 50 5-fold cross-validations and reported them in Figure 9.

In all the networks, ε -BE is able to achieve comparable results in the classification task. This is in line with recent findings claiming that there is currently no method that clearly outperforms the state of the art in this task [4]. The Film network proved challenging; however, three other methods (Graphwave, SEGK, and struc2vec) timed out in this case.

D. Regression

In regression, the objective is to predict a continuous or numerical value based on input features. Following [36], we

Method	Brazil		Europe		United States		Actor		Film	
	EIG	BET	EIG	BET	EIG	BET	EIG	BET	EIG	BET
ε -BE	8.47E-05	2.79E-02	1.23E-05	7.56E-03	1.28E-05	4.12E-02	9.93E-03	1.73E-02	2.30E-03	8.10E-03
Graphwave	5.30E-04	4.08E-02	8.51E-04	1.17E-02	2.95E-03	4.48E-02	2.36E-02	5.87E-02	TO	TO
SEGK	4.90E-03	6.17E-02	3.38E-03	1.41E-02	1.05E-02	5.12E-02	2.29E-02	6.06E-02	TO	TO
DRNE	3.41E-03	4.81E-02	5.97E-04	9.89E-03	3.51E-03	4.86E-02	5.58E-03	5.56E-02	1.02E-02	1.30E-02
struc2vec	3.68E-03	4.57E-02	2.47E-03	7.19E-03	5.66E-03	4.81E-02	TO	TO	TO	TO
GAS	7.88E-02	4.19E-01	1.56E-02	2.69E-02	1.62E-02	6.21E-02	1.87E-02	6.18E-01	1.04E-02	1.25E-02
RID ε Rs	6.01E-03	5.03E-02	5.09E-03	1.50E-02	1.26E-02	5.35E-02	1.67E-02	6.01E-02	8.21E-3	1.10E-02

TABLE III
MEAN SQUARED ERRORS FOR THE NETWORK REGRESSION TASK. EIG AND BET STAND FOR EIGENVECTOR AND BETWEENNESS CENTRALITY. WE HIGHLIGHT THE BEST AND SECOND-BEST RESULTS IN GREEN AND BLUE, RESPECTIVELY. TO: TIMEOUT (3 H)

define the target score for regression by incorporating two well-recognized centrality measures: eigenvector and betweenness centrality. The idea is to use the encoding of each node in the embedding space as a predictor for the centrality measures in the original network.

The goal is to compare methods with respect to their ability to preserve the centrality measures in the lower dimensional space. This is significant because the centrality measures have established correlations with the fundamental notions of role and position within a network [36].

We train a linear regression allocating 20% of the nodes as the designated test set and employing the remainder for training. We use the trained linear regression model to predict the centrality score for each node in the test set. In line with [36], we employ the Mean Square Error (MSE) metric to quantify the discrepancy between the predicted and the true centrality. The former corresponds to the centrality predicted for each node in the test set. The latter corresponds to the centrality of these nodes in the real networks where their encoding corresponds to the rows of the adjacency matrix A . To normalize for the different scales inherent to centrality measures, the resultant value is divided by the average value.

The analysis results are shown in Table III. ε -BE consistently exhibits optimal performance in eigenvector and betweenness centrality across nearly all network instances. This outcome is aligned with the intrinsic properties of the BE reduction methodology, which effectively preserves eigenvector centrality, Katz centrality, and PageRank centrality, as reported in the work [44]. Although RID ε Rs should inherit this property too, since it is based on approximate equitable partitions, it demonstrated good performance (second-best) for eigenvalue centrality only in the Film network, while it was outperformed by the competition in the other networks. Finally, we observe that ε -BE is highly competitive since it is the method with always the best or second-best performance. Struc2vec, Graphwave, and DRNE display good performances but not comparable with our approach. Additionally, some of them timed out when analyzing the larger networks.

E. Scalability

We empirically analyzed the scalability of ε -BE by comparing its running times against the other methods on networks

Method	Networks				
	Brazil	USA	BlogCatalog	Brightkite	GitHub
	$n = 131$	1190	10312	58689	177316
ε -BE	2.94E-1	6.49E-1	6.24E+0	7.68E+1	4.81E+2
Graphwave	2.52E+0	1.03E+1	1.47E+3	TO	TO
SEGK	1.92E+1	4.58E+3	TO	TO	TO
DRNE	1.07E+1	2.50E+1	1.96E+2	1.26E+3	TO
struc2vec	1.83E+1	9.93E+3	TO	TO	TO
GAS	2.30E+0	7.95E+0	4.90E+2	4.98E+2	2.65E+3
RID ε Rs	3.51E+1	2.73E+2	3.96E+2	7.38E+3	TO

TABLE IV
NETWORK EMBEDDING RUNNING TIMES (IN S). TO: TIMEOUT (3 H).

of increasing size. For this, we also considered three larger networks, BlogCatalog [45], Brightkite [46], and GitHub [47].

Results are reported in Table IV, showing that ε -BE consistently demonstrated superior runtime performance. Specifically, SEGK and struc2vec timed out with the BlogCatalog network, while Graphwave timed out with the Brightkite network. On the GitHub network, DRNE encountered a runtime error, which stopped the embedding process prematurely. RID ε Rs timed out only with the GitHub network, but its running times are at least two orders of magnitude larger than ε -BE (while it performed significantly better than ε -BE only in the classification task of the Film network). The GitHub network could be effectively handled only by ε -BE and GAS. Although GAS exhibited competitive running times (which are, however, one order of magnitude slower), its performance was inferior, except for the classification of the Film network.

V. CONCLUSION

Structural network embeddings are useful for tasks that exploit properties of similarities of the roles of the nodes in a network. Our network embedding is based on ε -BE, which partitions the nodes of a network into approximately equitable blocks using an efficient partition refinement algorithm embedded in an iterative framework that controls the impact of the tolerance parameter ε . The experimental evaluation on visualization, classification, and regression tasks demonstrate that ε -BE consistently achieves comparable or superior performance compared to existing methods at shorter running times compared to other state-of-the-art methods.

While ε -BE currently addresses undirected binary networks, it appears to be extensible to directed and weighted networks by suitably accommodating the notion of equitable partition for both in- and out-neighbors. Additionally, we plan to enhance ε -BE by considering different initial partitions based on the problem context. While this study adopts an initial partition where all nodes fall into a single block to ensure a fair comparison with related works, incorporating problem-specific initial partitions may lead to improved performance.

REFERENCES

- [1] L. Freeman, "The development of social network analysis," *A Study in the Sociology of Science*, vol. 1, no. 687, pp. 159–167, 2004.
- [2] G. A. Pavlopoulos, M. Secrier, C. N. Moschopoulos, T. G. Soldatos, S. Kossida, J. Aerts, R. Schneider, and P. G. Bagos, "Using graph theory to analyze biological networks," *BioData mining*, vol. 4, pp. 1–27, 2011.
- [3] L. Lü, M. Medo, C. H. Yeung, Y.-C. Zhang, Z.-K. Zhang, and T. Zhou, "Recommender systems," *Physics reports*, vol. 519, no. 1, pp. 1–49, 2012.
- [4] P. Jiao, X. Guo, T. Pan, W. Zhang, Y. Pei, and L. Pan, "A survey on role-oriented network embedding," *IEEE Transactions on Big Data*, vol. 8, no. 4, pp. 933–952, 2022.
- [5] P. Cui, X. Wang, J. Pei, and W. Zhu, "A survey on network embedding," *IEEE TKDE*, vol. 31, no. 5, pp. 833–852, 2018.
- [6] J. Jin, M. Heimann, D. Jin, and D. Koutra, "Toward understanding and evaluating structural node embeddings," *ACM TKDD*, vol. 16, no. 3, pp. 1–32, 2021.
- [7] R. K. Merton, *Social theory and social structure*. Simon and Schuster, 1968.
- [8] D. R. White and K. P. Reitz, "Graph and semigroup homomorphisms on networks of relations," *Social Networks*, vol. 5, no. 2, pp. 193–234, 1983.
- [9] S. P. Borgatti and M. G. Everett, "Notions of position in social network analysis," *Sociological Methodology*, vol. 22, pp. 1–35, 1992.
- [10] F. Lorrain and H. C. White, "Structural equivalence of individuals in social networks," *The Journal of mathematical sociology*, vol. 1, no. 1, pp. 49–80, 1971.
- [11] R. Jin, V. E. Lee, and H. Hong, "Axiomatic ranking of network role similarity," in *Proceedings of 17th ACM SIGKDD KDD*, 2011, pp. 922–930.
- [12] C. D. Godsil, "Compact graphs and equitable partitions," *Linear Algebra and its Applications*, vol. 255, no. 1-3, pp. 259–266, 1997.
- [13] M. G. Everett and S. P. Borgatti, "Exact colorations of graphs and digraphs," *Social networks*, vol. 18, no. 4, pp. 319–331, 1996.
- [14] D. M. Cvetković, M. Doob, and H. Sachs, *Spectra of graphs: theory and applications*. A. Press, Ed., 1980.
- [15] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin, "Maximal aggregation of polynomial dynamical systems," *PNAS*, vol. 114, no. 38, pp. 10029 – 10034, 2017.
- [16] P. Buchholz, "Exact and ordinary lumpability in finite markov chains," *Journal of applied probability*, vol. 31, no. 1, pp. 59–75, 1994.
- [17] L. Cardelli, M. Tribastone, M. Tschaikowski, and A. Vandin, "Symbolic computation of differential equivalences," in *Proceedings of the 43rd ACM SIGPLAN-SIGACT POPL*, 2016.
- [18] S. Tognazzi, M. Tribastone, M. Tschaikowski, and A. Vandin, "EGAC: a genetic algorithm to compare chemical reaction networks," in *GECCO*, P. A. N. Bosman, Ed. ACM, 2017, pp. 833–840.
- [19] N. Gast, L. Bortolussi, and M. Tribastone, "Size expansions of mean field approximation: Transient and steady-state analysis," *ACM SIGMETRICS Performance Evaluation Review*, vol. 46, no. 3, pp. 25–26, 2019.
- [20] L. Cardelli, I. C. Pérez-Verona, M. Tribastone, M. Tschaikowski, A. Vandin, and T. Waizmann, "Exact maximal reduction of stochastic reaction networks by species lumping," *Bioinform.*, vol. 37, no. 15, pp. 2175–2182, 2021.
- [21] A. Valmari and G. Franceschinis, "Simple $O(m \log n)$ time markov chain lumping," in *Proceedings of 16th TACAS*, vol. 6015, 2010, pp. 38–52.
- [22] R. Paige and R. E. Tarjan, "Three partition refinement algorithms," *SIAM Journal on computing*, vol. 16, no. 6, pp. 973–989, 1987.
- [23] P. V. Gupte, B. Ravindran, and S. Parthasarathy, "Role discovery in graphs using global features: Algorithms, applications and a novel evaluation strategy," in *2017 IEEE 33rd ICDE*, 2017, pp. 771–782.
- [24] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [25] K. Henderson, B. Gallagher, T. Eliassi-Rad, H. Tong, S. Basu, L. Akoglu, D. Koutra, C. Faloutsos, and L. Li, "Rolx: structural role extraction & mining in large graphs," in *Proceedings of the 18th ACM SIGKDD KDD*, 2012, pp. 1231–1239.
- [26] S. Gilpin, T. Eliassi-Rad, and I. Davidson, "Guided learning for role discovery (GLRD) framework, algorithms, and applications," in *Proceedings of the 19th ACM SIGKDD KDD*, 2013, pp. 113–121.
- [27] C. Donnat, M. Zitnik, D. Hallac, and J. Leskovec, "Learning structural node embeddings via diffusion wavelets," in *Proceedings of 24th ACM SIGKDD KDD*, 2018, pp. 1320–1329.
- [28] M. Heimann, H. Shen, T. Safavi, and D. Koutra, "Regal: Representation learning-based graph alignment," in *Proceedings of the 27th ACM CIKM*, 2018, pp. 117–126.
- [29] G. Nikolentzos and M. Vazirgiannis, "Learning structural node representations using graph kernels," *IEEE TKDE*, vol. 33, no. 5, pp. 2045–2056, 2019.
- [30] P. Drineas and M. W. Mahoney, "Approximating a gram matrix for improved kernel-based learning," in *18th Annual Conference on Learning Theory, COLT*, 2005, pp. 323–337.
- [31] Y. Pei, G. Fletcher, and M. Pechenizkiy, "Joint role and community detection in networks via l_2, l_1 norm regularized nonnegative matrix trifactORIZATION," in *Proceedings of the 2019 IEEE/ACM ASONAM*, 2019, pp. 168–175.
- [32] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo, "struc2vec: Learning node representations from structural identity," in *Proceedings of 23rd ACM SIGKDD KDD*, 2017, pp. 385–394.
- [33] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [34] N. K. Ahmed, R. A. Rossi, J. B. Lee, T. L. Willke, R. Zhou, X. Kong, and H. Eldardiry, "Role-based graph embeddings," *IEEE TKDE*, vol. 34, no. 5, pp. 2401–2415, 2020.
- [35] D. Jin, M. Heimann, R. A. Rossi, and D. Koutra, "Node2bits: Compact time- and attribute-aware node representations for user stitching," in *Joint European conference on Machine Learning and Knowledge Discovery in Databases*, 2019, pp. 483–506.
- [36] K. Tu, P. Cui, X. Wang, P. S. Yu, and W. Zhu, "Deep recursive network embedding with regular equivalence," in *Proceedings of 24th ACM SIGKDD KDD*, 2018, pp. 2357–2366.
- [37] X. Guo, W. Zhang, W. Wang, Y. Yu, Y. Wang, and P. Jiao, "Role-oriented graph role oriented-encoder guided by structural information," in *Database Systems for Advanced Applications: 25th International Conference, DASFAA 2020*, 2020, pp. 466–481.
- [38] Y. Jin, G. Song, and C. Shi, "GraLSP: Graph neural networks with local structural patterns," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 4361–4368.
- [39] J. Qiu, Q. Chen, Y. Dong, J. Zhang, H. Yang, M. Ding, K. Wang, and J. Tang, "Gcc: Graph contrastive coding for graph neural network pre-training," in *Proceedings of 26th ACM SIGKDD KDD*, 2020, pp. 1150–1160.
- [40] L. Cardelli, G. Squillace, M. Tribastone, M. Tschaikowski, and A. Vandin, "Formal lumping of polynomial differential equations through approximate equivalences," *Journal of Logical and Algebraic Methods in Programming*, vol. 134, p. 100876, 2023.
- [41] W. Gu, A. Tandon, Y.-Y. Ahn, and F. Radicchi, "Principled approach to the selection of the embedding dimension of networks," *Nature Communications*, vol. 12, no. 1, p. 3772, 2021.
- [42] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [43] G. E. Hinton and S. Roweis, "Stochastic neighbor embedding," *Advances in neural information processing systems*, vol. 15, 2002.
- [44] S. Tognazzi, M. Tribastone, M. Tschaikowski, and A. Vandin, "Differential equivalence yields network centrality," in *ISOLA*, 2018, pp. 186–201.
- [45] L. Tang and H. Liu, "Relational learning via latent social dimensions," in *Proceedings of 15th ACM SIGKDD KDD*, 2009, pp. 817–826.
- [46] E. Cho, S. A. Myers, and J. Leskovec, "Friendship and mobility: user movement in location-based social networks," in *Proceedings of 17th ACM SIGKDD KDD*, 2011, pp. 1082–1090.
- [47] S. Chacon, "The 2009 github contest." <https://github.com/blog/466-the-2009-github-contest>.