

A Hierarchical Block Distance Model for Ultra Low-Dimensional Graph Representations

Nakis, Nikolaos; Çelikkanat, Abdulkadir; Lehmann, Sune; Mørup, Morten

Published in: IEEE Transactions on Knowledge and Data Engineering

Link to article, DOI: 10.1109/TKDE.2023.3304344

Publication date: 2024

Document Version Peer reviewed version

Link back to DTU Orbit

Citation (APA): Nakis, N., Çelikkanat, A., Lehmann, S., & Mørup, M. (2024). A Hierarchical Block Distance Model for Ultra Low-Dimensional Graph Representations. *IEEE Transactions on Knowledge and Data Engineering*, *36*(4), 1399 -1412. https://doi.org/10.1109/TKDE.2023.3304344

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- · You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A Hierarchical Block Distance Model for Ultra Low-Dimensional Graph Representations

Nikolaos Nakis, Abdulkadir Çelikkanat, Sune Lehmann, Morten Mørup

Abstract—Graph Representation Learning (GRL) has become central for characterizing structures of complex networks and performing tasks such as link prediction, node classification, network reconstruction, and community detection. Whereas numerous generative GRL models have been proposed, many approaches have prohibitive computational requirements hampering large-scale network analysis, fewer are able to explicitly account for structure emerging at multiple scales, and only a few explicitly respect important network properties such as homophily and transitivity. This paper proposes a novel scalable graph representation learning method named the Hierarchical Block Distance Model (HBDM). The HBDM imposes a multiscale block structure akin to stochastic block modeling (SBM) and accounts for homophily and transitivity by accurately approximating the latent distance model (LDM) throughout the inferred hierarchy. The HBDM naturally accommodates unipartite, directed, and bipartite networks whereas the hierarchy is designed to ensure linearithmic time and space complexity enabling the analysis of very large-scale networks. We evaluate the performance of the HBDM on massive networks consisting of millions of nodes. Importantly, we find that the proposed HBDM framework significantly outperforms recent scalable approaches in all considered downstream tasks. Surprisingly, we observe superior performance even imposing ultra-low two-dimensional embeddings facilitating accurate direct and hierarchical-aware network visualization and interpretation.

Index Terms—Latent Space Modeling, Complex Networks, Graph Representation Learning.

1 INTRODUCTION

N Etworks naturally arise in a plethora of scientific areas to model interactions between entities from physics to sociology and biology, with many instances such as collaboration, protein-protein interaction, and brain connectivity networks [1] to mention but a few. In recent years, Graph Representation Learning (GRL) approaches have attracted great interest due to their outstanding performance compared to classical techniques for arduous problems such as link prediction [2], node classification [3], [4], and community detection [5].

Numerous GRL methods have been proposed, see also [6] for a survey. The leading initial works are the random walk-based methods [4], [7]–[10], leveraging the Skip-Gram algorithm [11] to learn the node representations. Matrix factorization-based algorithms [6], [12] have also become prominent, extracting the embedding vectors by decomposing a designed feature matrix. Furthermore, neural network models [6], [13] have been proposed for graphstructured data, returning outstanding performance by combining node attributes and network structure when learning embeddings. Recent studies [14] aim to alleviate the computational burden of these algorithms through matrix sparsification tools [15], hierarchical representations [16], [17], or by fast hashing schemes [18].

Latent Space Models (LSMs) for the representation of graphs have been quite popular over the past years [19]– [25], especially for social networks analysis [26], [27] facilitating community extraction [28] and characterization of network polarization [29]. LSMs utilize the generalized linear model framework to obtain informative latent node embeddings while preserving network characteristics. The choice of latent effects in modeling the link probabilities between the nodes leads to different expressive capabilities characterizing network structure. In particular, in the Latent Distance Model (LDM) [30] nodes are placed closer in the latent space if they are similar or vice-versa. LDM obeys the triangle inequality and thus naturally represents transitivity [31], [32] ("a friend of a friend is a friend") and network homophily [33], [34] (a tendency where similar nodes are more likely to connect to each other than dissimilar ones). Homophily is a very well-known and well-studied effect appearing in social networks [31], [33], [34] and essentially describes the tendency for people to form connections with those that share similarities with themselves. Similarities can be drawn from meta-data (observed node attributes) and may refer to shared demographic properties, political opinions, etc. Homophily has been observed among a broad range of collaborations (see [32] for a complete overview). Homophily can also be accounted for based on the unobserved attributes as defined by the LDM as shown in [35]. Homophily explains prominent patterns as expressed in social networks in terms of transitivity, as well as, balance theory ("the enemy of my friend is an enemy") [36]. More specifically, in an LDM we can extend the meaning of similarity to some unobserved (latent) covariates, i.e., latent embeddings Z. The higher similarity between nodes translates here to a stronger relationship between the nodes and thereby a higher probability of observing connections. As a result, for two similar nodes $\{i, j\}$ the pairwise distance $|\mathbf{z}_i - \mathbf{z}_i|_2$ should be small which further implies that for a different node $\{k\}$ we obtain $|\mathbf{z}_i - \mathbf{z}_k|_2 \approx |\mathbf{z}_i - \mathbf{z}_k|_2$. The latter concludes that nodes $\{i, j\}$ are similar since they share similar relationships with the rest of the nodes.

1

The approach has been extended to bipartite networks in [37] by introducing mode-specific embedding vectors and community detection by endowing the LDM with a Gaus-

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

2

sian Mixture Model prior to promoting cluster structures in the latent space forming the latent position clustering model (LPCM) [35], [38]. However, the LDM is unable to account for possible stochastic equivalence as defined by the Stochastic Blockmodels [39], [40], i.e. ("groups of nodes defined by shared intra- and inter-group relationships") defining non positive-semi-definite latent representations. The LSMs were advanced to characterize such stochastic equivalence by imposing an Eigenmodel admitting negative eigenvalues [41]. These latent space methods are attractive due to their simplicity, as they define well-structured inference problems and are characterized by high explanatory power [25]. The time and space complexities are their main drawbacks as the likelihood function scales by the number of node pairs (i.e., quadratically in the number of nodes for a unipartite graph) typically addressed using subsampling strategies [42].

Many real-world networks are composed of structures emerging at multiple scales which can be expressed using hierarchical representations [43]. Several methods have thus been advanced to such hierarchical representations including stochastic block model approaches [44]-[49] as well as agglomerative [50]-[52] and recursive partitioning [53] procedures relying on various measures of similarity. Importantly, learning node representations characterizing structure at multiple scales of the network can facilitate network visualization and the understanding of the inner dynamics of networks. Hierarchical representation of bipartite networks is of special interest due to the fact that most unipartite hierarchical clustering algorithms do not generalize to the bipartite case beyond clustering each mode separately or transforming the bipartite network into a unipartite representation. In the work of [54], the authors used the spectral partitioning algorithm of [55] and then applied k-means on the spectral space to get initial biclusters which were followed by divisive bi-splits to create a dendrogram. In this case, the spectral embedding space was not constructed to reflect explicitly the clustering criterion. In addition to divisive procedures, agglomerative clustering has also been proposed for bipartite networks. In the work of [56] a multi-objective function was designed and combined with classical community construction algorithms. One limitation here is that the network should be transformed into a unipartite structure.

Despite the many advantages of hierarchical structures and block models, one major limitation remains to accurately account for homophily [41], which is a key characteristic of social networks. More specifically, block models have been extended to explicitly impose a community structure [57], [58] but notable this only provides withincluster homogeneity and thus homophily-like properties for the community relative to the other communities but not a hierarchy complying with such a structure. Whereas LPCM accounts for homophily it does not account for hierarchical structures and cluster structures are not strictly imposed beyond a prior to promoting the latent positions to form groups.

In this work, we propose a novel node representation learning approach, the Hierarchical Block Distance Model (HBDM)¹, as a reconciliation between hierarchical block

structures of different scales and network properties such as homophily and transitivity. In particular, we propose a framework combining embedding and hierarchical characterization for graph representation learning. Importantly, we design a hierarchical structure that respects a linearithmic total time and space complexity, in terms of the number of nodes (i.e., $\mathcal{O}(N \log N)$), and at the same time provides an accurate interpretable representation of structure at different scales. Using the HBDM, we embed moderate-sized and large-scale networks containing more than a million nodes and establish the performance of our model in terms of link prediction and node classification to existing prominent graph embedding approaches. We further highlight how the inferred hierarchical organization can facilitate accurate visualization of network structure even when using only D = 2 dimensional representations providing favorable performance in all the considered GRL tasks; link prediction, node classification, and network reconstruction. Additionally, we show how our proposed framework extends the hierarchical multi-resolution structure to bipartite networks and provides the characterization of communities at multiple scales.

2 THE HIERARCHICAL BLOCK DISTANCE MODEL

We first concentrate our study on undirected networks and later generalize our approach to bipartite graphs. We now provide the necessary definitions required throughout the paper. Let $\mathcal{G} = (V, E)$ be a graph where N := |V| is the number of nodes and $Y_{N \times N} = (y_{i,j}) \in \{0,1\}^{N \times N}$ be the adjacency matrix of the graph such that $y_{i,j} = 1$ if the pair $(i,j) \in E$ otherwise it is equal 0, for all $1 \leq i < j \leq N$. We denote the latent representations of nodes by $\mathbf{Z} = (z_{i,d}) \in \mathbb{R}^{N \times D}$ where each row vector, $\mathbf{z}_i \in \mathbb{R}^D$, indicates the corresponding embedding of node $i \in \mathcal{V}$ in a *D*-dimensional space.

The most natural choice for modeling homophily and transitivity can be found in the Latent Space Model (LSM) which defines an \mathbb{R}^D -dimensional latent space in which every node of the graph is characterized through the unobserved but informative node-specific variables $\{\mathbf{z}_i \in \mathbb{R}^D\}$. These variables are considered sufficient to describe and explain the underlying relationships between the nodes of the network. The probability of an edge occurring is considered conditionally independent given the unobserved latent positions. Consequently, the total probability distribution of the network can be written as:

$$P(Y|\boldsymbol{Z},\boldsymbol{\theta}) = \prod_{i < j}^{N} p(y_{i,j}|\mathbf{z}_i, \mathbf{z}_j, \boldsymbol{\theta}_{i,j}),$$
(1)

where θ denotes any potential additional parameters, such as covariate regressors. A popular and convenient parameterization of Equation (1) for binary data is through the logistic regression model [30], [35], [59], [60]. In contrast, we adopt the Poisson regression model similar to [60] under a generalized linear model framework for the LSM. The use of a Poisson likelihood for modeling binary relationships in a network does not decrease the predictive performance nor the ability of the model to detect the network structure, as shown in [61]. It also generalizes the analysis to integerweighted graphs. In addition, the exchange of the *logit* to a

¹For implementation details please visit: github.com/Nicknakis/HBDM.

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

3

log link function when transitioning from a Bernoulli to a Poisson model defines nice decoupling properties over the predictor variables in the likelihood [62], [63]. Utilizing the Poisson Latent Distance Model (LDM) of the LSM family framework, the rate of an occurring edge depends on a distance metric between the latent positions of the two nodes. In our formulation, we consider the LDM Poisson rate with node-specific biases or random-effects [35], [60] such that the expression for the Poisson rate becomes:

$$\lambda_{ij} = \exp\left(\gamma_i + \gamma_j - d(\mathbf{z}_i, \mathbf{z}_j)\right),\tag{2}$$

where $\gamma_i \in \mathbb{R}$ denotes the node-specific random-effects and $d_{ij}(\cdot, \cdot)$ denotes any distance metric obeying the triangle inequality $\{d_{ij} \leq d_{ik} + d_{kj}, \forall (i, j, k) \in V^3\}$. Considering variables $\{\mathbf{z}_i\}_{i \in V}$ as the latent characteristics, Equation (2) shows that similar nodes will be placed closer in the latent space, yielding a high probability of an occurring edge and thus modeling homophily and satisfies network transitivity and reciprocity through the triangle inequality whereas the node-specific bias can account for degree heterogeneity. The conventional LDM rate utilizing a global bias, γ^g , corresponds to the special case in which $\gamma_i = \gamma_j = 0.5\gamma^g$. As in [30], we presently adopt the Euclidean distance as the choice for the distance metric $d_{ij}(\cdot, \cdot)$.

2.1 Designing A Linearithmic Complexity

Our goal is to design a Hierarchical Block Model preserving homophily and transitivity properties with a total complexity allowing for the analysis of large-scale networks. Our HBDM, defines the rate of a link between each network dyad $\{i, j\} \in V \times V$ based on the Euclidean distance, as shown in Equation (2). Therefore, we can define a blockalike hierarchical structure by a divisive clustering procedure over the latent variables in the Euclidean space. The total optimization cost of such a model though should have a linearithmic upper bound complexity to make large-scale analysis feasible. Introducing a number of clusters K equal to the number of nodes N in the HBDM, leads to the same log-likelihood as of the standard LDM, defining a sum over each ordered pair of the network, as:

$$\log P(Y|\mathbf{\Lambda}) = \sum_{i < j} \left(y_{ij} \log(\lambda_{ij}) - \lambda_{ij} \right)$$
$$= \sum_{i < j: y_{ij} = 1} \log(\lambda_{ij}) - \sum_{i < j} \lambda_{ij}, \qquad (3)$$

For brevity, we presently ignore the linear scaling by dimensionality D of the above log-likelihood function. Notably, the link contribution $\sum_{y_{i,j}=1} \log(\lambda_{i,j})$ is responsible for positioning "similar" nodes closer in the latent space, expressing the desired homophily.

In addition, large networks are highly sparse [64] with the number of edges being proportional to the number of nodes in the network. As a result, the computation of the link contribution is relatively cheap, scaling linearithmic or sub-linearithmic (as shown in supplementary). Most importantly, the link term removes rotational ambiguity between the different blocks of the hierarchy (as discussed later). For these three reasons, no block structure is imposed on the calculation of the link contribution. The second term acts as the repelling force for dissimilar nodes and requires the computation of all node pairs scaling as $O(N^2)$ making the evaluation of the above likelihood infeasible for large networks. By enforcing a block structure, i.e., akin to stochastic block models [39], [40], [65], when grouping the nodes into K clusters we define the rate between block k and k' in terms of their distance between centroids. A simple block structure without a hierarchy would lead to the following non-link expression:

$$\sum_{ik}^{K} \sum_{i\in C_k} \sum_{j\in C_{k'}} \exp(\gamma_i + \gamma_j - ||\boldsymbol{\mu}_k - \boldsymbol{\mu}_{k'}||_2) \right), \quad (4)$$

where μ_k denotes the k'th cluster centroid of the set $C = \{C_1, \ldots, C_K\}$, and has absorbed the dependency over the variables $\boldsymbol{Z} \in \mathbb{R}^{N \times D}$. More specifically, the cluster centroids μ_k are implicit parameters defined as a function over the latent variables, as we will show later. Overall, considering the principle that connected and homophilic nodes will be placed closer in the latent space, this expression generalizes this principle by introducing a clustering procedure that obeys "cluster-homophily" and "clustertransitivity" over the latent space. More specifically, we can assume that closely related nodes will be positioned in the same cluster while related or interconnected clusters will also be positioned close in the latent space, providing an accurate block structure schema. As opposed to the LPCM where clustering structures are imposed through a prior, the above formulation strictly defines the clustering structure as shared overall proximity between blocks as defined by the distances between centroids of the formed groups.

2.1.1 A Hierarchical Representation

In order to obtain the desired hierarchical representation, we define hierarchical clustering via a divisive procedure. In detail, we organize the embedded clusters into a hierarchy using a tree structure, defining a cluster dendrogram. The root of the tree is a single cluster containing the total amount of latent variable embeddings Z. At every level of the tree, we perform partitioning until we obtain leaf nodes containing equal or less than the desired number of nodes, N_{leaf} . This number is chosen with respect to our linearithmic complexity upper bound and set as $N_{\text{leaf}} = \log N$, resulting in approximately K = N/log(N) total clusters. The treenodes belonging to a specific tree-level are considered the clusters for that specific tree height. Every novel partition of a non-leaf node is performed only on the set of points allocated to the parent tree-node (cluster). For every level of the tree, we consider the pairwise distances of datapoints belonging to different tree-nodes as the distance between the corresponding cluster centroids, as illustrated by Fig. 1 (ii). Based on these distances, we calculate the likelihood contribution of the blocks and continue with binary splits, down the tree, for the non-leaf tree-nodes. When all treenodes are considered as leaves, we calculate analytically the inner cluster pairwise distances for the corresponding likelihood contribution of analytical blocks, as shown in the



Fig. 1. Schematic representation of the distance matrix calculation for a hierarchical structure of the tree of height L = 3 and for the number of observations N = 64.

last part of Fig. 1 (ii). The latter analytical calculation comes at a linearithmic cost of $\mathcal{O}(KN_{\text{leaf}}^2) = \mathcal{O}(N \log N)$ while enforces the homophily and transitivity properties of the model since for the most similar nodes the HBDM behaves explicitly as the standard LDM.

We can thereby define a Hierarchical Block Distance Model with Random Effects (HBDM-RE) as:

$$\log P(Y|\mathbf{Z}, \boldsymbol{\gamma}) = \sum_{\substack{i < j \\ y_{i,j} = 1}} \left(\gamma_i + \gamma_j - ||\mathbf{z}_i - \mathbf{z}_j||_2 \right)$$
$$- \sum_{k=1}^{K_L} \left(\sum_{\substack{i < j \\ i,j \in C_k^{(L)}}} \exp\left(\gamma_i + \gamma_j - ||\mathbf{z}_i - \mathbf{z}_j||_2\right) \right)$$
$$- \sum_{l=1}^{L} \sum_{k=1}^{K_l} \sum_{k' > k}^{K_l} \left(\exp\left(-||\boldsymbol{\mu}_k^{(l)} - \boldsymbol{\mu}_{k'}^{(l)}||_2 \right) \times \left(\sum_{i \in C_k^{(l)}} \exp\left(\gamma_i\right) \right) \left(\sum_{j \in C_{k'}^{(l)}} \exp\left(\gamma_j\right) \right) \right), \tag{5}$$

where $l \in \{1, \ldots, L\}$ denotes the *l*'th dendrogram level, k_l is the index representing the cluster id for the different tree levels, and $\boldsymbol{\mu}_k^{(l)}$ the corresponding centroid. We also consider a Hierarchical Block Distance Model (HBDM) without the random effects setting $\gamma_i = 0.5\gamma^g$. For a multifurcating tree splitting in *K* clusters and having N/log(N) terminal nodes (clusters), the number of internal nodes are $\mathcal{O}(N/(K \log N))$ and each node needs to evaluate $\mathcal{O}(K^2)$ pairs providing an overall complexity of $\mathcal{O}(NK/\log N)$, thus $K \leq \log N^2$ to achieve $\mathcal{O}(N \log N)$ scaling [66].

2.1.2 Divisive partitioning using k-means with a Euclidean distance metric

Whereas the likelihood in Equation (5) can be directly minimized by assigning nodes to the clusters given by the tree structure, this evaluation for all *N* nodes scales prohibitively as $O(N^2/\log N)$. To reduce this scaling, we use a more efficient divisive partitioning procedure, minimizing the Euclidean norm $||\mu_{k_l} - \mu_{k'_l}||_2$. The divisive clustering procedure thus relies on the following Euclidean norm objective

$$J(\mathbf{r},\boldsymbol{\mu}) = \sum_{i=1}^{N} \sum_{k=1}^{K} r_{ik} ||\mathbf{z}_i - \boldsymbol{\mu}_k||_2,$$
(6)

where k denotes the cluster id, \mathbf{z}_i is the i'th data observation, r_{ik} the cluster responsibility/assignment, and $\boldsymbol{\mu}_k$ the cluster centroid.

This objective function is unfortunately not accounted for by existing k-means clustering algorithms relying on the squared Euclidean norm. We therefore presently derive an optimization procedure for k-means clustering with Euclidean norm utilizing the auxiliary function framework of [67] developed in the context of compressed sensing. We define an auxiliary function for (6) as:

4

$$J^{+}(\boldsymbol{\phi}, \mathbf{r}, \boldsymbol{\mu}) = \sum_{i=1}^{N} \sum_{k=1}^{K} r_{ik} \left(\frac{||\mathbf{z}_{i} - \boldsymbol{\mu}_{k}||_{2}^{2}}{2\phi_{ik}} + \frac{1}{2}\phi_{ik} \right), \quad (7)$$

where ϕ are the auxiliary variables. Thereby, minimizing Equation (7) with respect to ϕ_{nk} yields $\phi_{ik}^* = ||\mathbf{z}_i - \boldsymbol{\mu}_k||_2$ and by plugging ϕ_{ik}^* back to (6) we obtain $J^+(\phi^*, \mathbf{r}, \boldsymbol{\mu}) = J(\mathbf{r}, \boldsymbol{\mu})$ verifying that (7) is indeed a valid auxiliary function for (6). The algorithm proceeds by optimizing cluster centroids as $\boldsymbol{\mu}_k = \left(\sum_{i \in k} \frac{\mathbf{z}_i}{\phi_{ik}} / \sum_{i \in k} \frac{1}{\phi_{ik}}\right)$ and assigning points to centroids as $\arg \min_C = \sum_{k=1}^K \sum_{z \in C_k} \left(\frac{||z-\boldsymbol{\mu}_k||_2^2}{2\phi_k} + \frac{1}{2}\phi_k\right)$ upon which ϕ_k is updated. The overall complexity of this procedure is $\mathcal{O}(TKND)$ [68] where *T* is the number of iterations required to converge. As shown in [67], Equation (7) is a special case of a general algorithm for an $l_p(0$ norm minimization using an auxiliary function with thealgorithm converging faster the smaller*p*is. For a detailedstudy of the efficiency of the optimization procedure undersuch an auxiliary function, see [67].

A simple approach to construct the tree structure would be to use the above Euclidean k-means procedure to split the nodes into $K = N/\log(N)$ clusters and construct the tree according to agglomeration as in hierarchical clustering. Unfortunately, such a strategy is computationally prohibitive. For that, in the coarser level (first layer of the tree), we choose to split to the maximally allowed clusters of $K = \log N$ allowing scaling of $\mathcal{O}(N \log N)$. It would be tempting to continue splitting into $\log N$ clusters, however, for a balanced multifurcating tree with $N/\log N$ leaf clusters, it will result in a height scaling as $\mathcal{O}(\log N / \log \log N)$ and thus an overall complexity of $\mathcal{O}(N \log^2(N) / \log \log N)$ [66]. Whereas a balanced binary tree at all levels below the root results in a height scaling as $\mathcal{O}(\log N)$ providing an overall complexity when including the linear scaling by dimensionality *D* of $\mathcal{O}(DN \log N)$ (as each level of the tree defines $\mathcal{O}(DN)$ operations). Fig. 1 (i), illustrates the resulting tree² for a small problem of N = 64 nodes in which we first split into 4 ($\approx \log(64)$) clusters and subsequently create binary splits until each leaf cluster contains 4 ($\approx \log(64)$) nodes.

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

2.1.3 Expressing Homophily and Transitivity

A central component to preserve homophily and transitivity of the HBDM is not to approximate the link terms at the level of the block as in (hierarchical) SBMs but to calculate analytically the link contribution of the log-likelihood across the total hierarchy beyond the leaf/analytical blocks. In Fig., 2 (i) and (ii), two leaf clusters are illustrated and connected with a link. Assume that we only calculate the distance inside the blocks analytically and that both the link and nonlink contributions of pairs belonging to different clusters are approximated based on their centroids' distance. This essentially would allow for any rotation of each cluster for all clusters in the hierarchy since the inner-block distances (analytical), as well as the centroid distances, would not change by such rotations, yielding exactly the same likelihood (block-level rotational invariance). In that case, homophily would be violated as, e.g., the distance between nodes cand d would not necessarily be smaller than the distance with other inter-cluster pairs (ex: Fig. 2 (i)), showing that the rotation of the blocks substantially impacts the homophily properties of the HBDM. Calculating the link contributions between the different clusters analytically solves this ambiguity since the likelihood is penalized higher when nodes c and d are positioned in a non-rotational-aware way. The computational cost imposed by accounting for all the link terms analytically is that the model complexity depends on the number of edges of the network (a total block structure would strictly be linearithmic in complexity). Nevertheless,



(iii) Three latent block structures.

Fig. 2. The clusters within the dashed circles denote the leaf block structures. The red circles and blue rhombuses indicate the node embeddings and the centroids, respectively. Gray lines represent the links and the dashed lines the distance between the cluster centers.

²For visualization purposes only, we show equally sized clusters.

5

we show empirically in the supplementary that the number of network edges E scales linearly with $N \log N$ and thus this analytical term respects our complexity bound. In Fig. 2 (iii), we present clusters defining cases of block inter-connections of sparsely connected blocks ($\{C_1, C_3\}$, $\{C_1, C_2\}$) and densely connected blocks ($\{C_2, C_3\}$. Whereas the analytical inter-cluster links (blue lines) are responsible for fixing the block rotation the inter-cluster links also drive the cluster-level proximities of centroids ensuring cluster homophily and transitivity.

Pairwise distances in the HBDM stays invariant to rotation, reflection, and translation of the latent space due to its LDM inheritance [30] these isometries can be resolved via a Singular-Value-Decomposition procedure as provided in the supplementary. Whereas the analytical link term calculations provide rotational awareness to the HBDM clusters, we continue by investigating the conditions in which a continuous operation defining infinitesimal rotations (with respect to the cluster centroid) is admissible leaving the loss function of Equation (5) invariant to continuous rotations. In Lemma 2.1 (proof given in the supplementary material), we start our investigation of this problem by showing that blocks with a unique inter-cluster link connection reduce the clusters' degree of rotational freedom by one.

Lemma 2.1. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph and let \mathcal{C} be a cluster with its centroid located at $\mu \in \mathbb{R}^D$ having an edge $(i, j) \in \mathcal{E}$ for some $i \in \mathcal{C}$ and $j \in \mathcal{V} \setminus \mathcal{C}$ such that $\mathbf{z}_i \neq \mu$. If $\tilde{\mathbf{z}}_i = \mu + \mathbf{R}(\theta)(\mathbf{z}_i - \mu)$ such that $\mathbf{R}(\theta)$ is a rotation matrix acting on the embeddings of nodes in cluster \mathcal{C} , then the maximum degree of freedom of any infinitesimal λ_{ij} -invariant rotation is defined by $\theta \in \mathbb{R}^{D-2}$.

A direct consequence of Lemma 2.1 is that for a twodimensional embedding, there is no possible continuous rotation of a cluster having only one external edge. Since there is a path from one node to all others in a connected graph, every cluster must have at least one external link. For the general case of blocks having multiple inter-cluster edges, rotations preserving the total sum of pairwise distances among node embeddings are highly unlikely, as discussed in the supplementary. Consequently, we can for connected networks expect uniqueness of a (local) minima solutions with no continuous admissible rotations leaving the HBDM loss function of Equation (5) invariant.

TABLE 1

Complexity analysis of methods. N := |V| is the vertex set, |E|: edge set, \mathcal{W} : number of walks, \mathcal{L} : walk length, H: height of the hierarchical tree, D: node representation size, k: number of negative instances, q: order value, c: Chebyshev expansion order, γ : window size, α_1 and α_2 constants such as $\alpha_1, \alpha_2 \ll N$.

Method	Complexity
DEEPWALK	$\mathcal{O}\left(\gamma N \log\left(N\right) \mathcal{WLD}\right)$
NODE2VEC	$\mathcal{O}\left(\gamma N \mathcal{WLD}k\right)$
LINE	$\mathcal{O}\left(E Dk\right)$
NETMF	$\mathcal{O}(N^2D)$
NETSMF	$\mathcal{O}\left(E (\gamma+D)+ND^2+D^3\right)$
RANDNE	$\mathcal{O}\left(ND^2 + E Dq\right)$
LOUVAINNE	$\mathcal{O}\left(E \mathcal{H}+ND\right)$
ProNE	$\mathcal{O}\left(ND^2 + E c\right)$
VERSE	$\mathcal{O}\left(N(\mathcal{W}+kD)\right)$
HBDM	$\mathcal{O}\left(\alpha_2 N \log\left(N\right) D\right)$

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

2.1.4 Extension to Bipartite Networks

Our proposed frameworks, HBDM and HBDM-RE generalize to both directed and bipartite graphs. In the following, we provide the mathematical extension for the bipartite case (the directed network formulation of our proposed model can be considered a special case of the bipartite framework in which self-links are removed and thus omitted from the below log-likelihood). For a bipartite network with adjacency matrix $Y^{N_1 \times N_2}$ we can formulate the log-likelihood as:

$$\log P(Y|\mathbf{\Lambda}) = \sum_{\substack{i < j \\ y_{i,j}=1}} \left(\psi_i + \omega_j - ||\mathbf{w}_i - \mathbf{v}_j||_2 \right)$$
$$- \sum_{k_L=1}^{K_L} \left(\sum_{i,j \in C_{k_L}} \exp\left(\psi_i + \omega_j - ||\mathbf{w}_i - \mathbf{v}_j||_2\right) \right)$$
$$- \sum_{l=1}^{L} \sum_{k=1}^{K_l} \sum_{k'>k}^{K_l} \left(\exp\left(-||\boldsymbol{\mu}_k^{(l)} - \boldsymbol{\mu}_{k'}^{(l)}||_2 \right) \times \left(\sum_{i \in C_k^{(l)}} \exp\left(\psi_i\right) \right) \left(\sum_{j \in C_{k'}^{(l)}} \exp\left(\omega_j\right) \right) \right), \tag{8}$$

where $\{\boldsymbol{\mu}_{k}^{(l)}\}_{k=1}^{K_{L}}$ are the latent centroids which have absorbed the dependency of both sets of latent variables $\{\mathbf{w}_{i}, \mathbf{v}_{i}\}$ while we define the Poisson rate as:

$$\lambda_{ij} = \exp\left(\psi_i + \omega_j - d(\mathbf{w}_i, \mathbf{v}_j)\right),\tag{9}$$

where ψ_i and ω_j are the corresponding random effects and $\{\mathbf{w}_i, \mathbf{v}_j\}$ are the latent variables of the two disjoint sets of the vertex set of sizes N_1 and N_2 , respectively. In this setting, we use our divisive Euclidean distance hierarchical clustering procedure over the concatenation $\mathbf{z} = [\mathbf{w}; \mathbf{v}]$ of the two sets of latent variables. Therefore, we define an accurate hierarchical block structure for bipartite networks, with each block including nodes from both of the two disjoint modes. Here, a centroid is considered a leaf if the corresponding tree-cluster contains less than $\log(N_1)$ of the latent variables $\{\mathbf{w}_i\}_{i=1}^{N_1}$ or less than $\log(N_2)$ of $\{\mathbf{v}\}_{j=1}^{N_2}$.

2.1.5 Complexity Comparison

TABLE 1 provides a comparison between time complexities of several prominent GRL methods in terms of their Big O notation, similar to [69]. We observe that our proposed HBDM is positioned as one of the most competitive frameworks. In terms of space complexity, our model defines a linearithmic complexity contrary to the majority of the considered baselines which are usually characterized by a quadratic space complexity [69]. (For a more detailed discussion please visit the supplementary.)

3 EXPERIMENTS

We extensively evaluate the performance of our method compared to baseline graph representation learning approaches on networks of various sizes and structures. We have conducted all the experiments regarding the HBDM and HBDM-RE on a 32 GB Tesla V100 GPU machine with 5120 CUDA cores, and a 1380 MHz clock. For the HBDM and HBDM-RE models, we optimize the negative log-likelihood via the Adam [70] optimizer with learning rate $lr \in [0.01, 0.1]$. For both frameworks, we build the hierarchical structure by running the k-means procedure every t = 25 iterations. Experiments regarding the baselines have been conducted on an Intel Xeon Gold 6342 CPU with 24 cores, 2800 MHz clock, and 512 GB memory. The implementation for HBDM and HBDM-RE is GPU-focused using PyTorch 1.12.1, exploiting parallel computations (running the frameworks on a CPU machine leads to substantially higher runtimes). We argue, that runtime comparison in terms of real-time is a biased estimate between different models since it correlates highly with the programming language, parallelization schemes, etc. For that, we instead compare theoretical complexities in terms of their Big \mathcal{O} notation. In all TABLES, we denote with bold digits the bestperforming score while we underline the second-best.

Datasets: We have performed the experiments on ten undirected networks of various sizes and structures: a citation network (*Cora* [71]), social interaction graphs (*Facebook* [72], *YouTube* [73], [74], *Flickr* [74], *Flixster* [75]), product-label network (*Amazon* [73]) and collaboration networks (*Dblp* [76], *AstroPh* [77], *GrQc* [77], *HepTh* [77]). Each network is considered as unweighted for the consistency of the experiments. The detailed statistics of the networks are provided by TABLE 2. All of the considered networks have been widely adopted and extensively used as benchmarks in the GRL literature [78].

Baseline Methods: In our experiments, we have run various graph representation learning methods in order to evaluate the performance of our approach. The prominent GRL frameworks used in this study are: (i) DEEPWALK [7], (ii) NODE2VEC [4], (iii) LINE [8], (iv) NETMF [79]. In addition, we consider five scalable graph embedding approaches: (v) NETSMF [15], (vi) RANDNE [18], (vii) PRONE [14], (viii) LOUVAINNE [16], (ix) VERSE [69]. For more details see the supplementary material. In our analysis, we considered GRAPHSAGE [13] as a prominent member of the family of Graph Neural Networks (GNNs). Our study focuses on the setting where node meta-data are not available. In such a setting, GRAPHSAGE was characterized by a close-to-random performance and thus not presented.

3.1 Link Prediction

We report results for the area under the curve of the receiver operator characteristic (AUC). For the experimental setup, we follow the commonly applied strategy [4], [7] and remove half of the edges while keeping the residual network connected. This strategy is not feasible for largescale networks since checking if the residual network stays connected after each removed link results in a high runtime complexity. For that, we hide 30% of the edges for these networks and extract the giant component (after the link removal) which is treated as the residual network. Extensive details for the link prediction experiments, as well as, Precision-Recall AUC scores are given in the supplementary. Error bars across 5 re-runs for the following AUC scores were found to be on the scale of 10^{-3} and thus negligible.

Effectiveness and Efficiency of the Multi-Scale Approximation: In Fig. 3a, we provide an effectiveness analysis of the HBDM likelihood when contrasted with its full likelihood estimation evaluated on the moderate-sized network

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

TABLE 2 Statistics of undirected networks. N: number of nodes, |E|: number of edges.

	Cora	Dblp	AstroPh	GrQc	Facebook	HepTh	Amazon	YouTube	Flickr	Flixster
E	2,708	27,199	17,903	5,242	4,039	8,638	334,868	1,138,499	1,715,255	2,523,386
	5,278	66,832	197,031	14,496	88,234	24,827	925,876	2,990,443	15,555,042	7,918,801

of Facebook (results for more networks are provided in the supplementary material). We here observe that the HBDM likelihood essentially approximates the true full likelihood providing systematically slightly lower likelihood estimates which we attribute to the small structural differences between calculating the distances analytically versus in a hierarchical block manner. A close approximation to the true likelihood provides evidence for multi-scale structures that characterize networks, yielding a high effectiveness of the HBDM framework. In addition, in Fig. 3a we see fluctuations in the likelihood which is an immediate result of building the network hierarchy from scratch every 25'th iteration. Importantly, despite the fact that k-means is notoriously known to be an NP-hard problem [80], [81], we observe that rebuilding the hierarchy has a minimum effect on the value of the likelihood, highlighting the stability of the inferred hierarchy in the HBDM. Furthermore, TABLE 3 conveys information about the comparison between an HBDM (approx) framework where all link distances are approximated by the centroid distances and the proposed HBDM where link distances are calculated analytically. We witness for the Facebook network how the rotational awareness induced by explicitly accounting for all links in the likelihood (as explained in subsection 2.1.3) increases the predictive capability of the model and thus its efficiency (similar results were obtained for all networks).

Moderate-Sized Networks: Results for the moderatesized networks are given in TABLE 4. The symbol "-" indicates that the running time of the corresponding model takes more than 20 hours and "x" shows that the method is not able to run due to insufficient memory space. We observe that the HBDM and HBDM-RE perform significantly better or on par with the performance of the considered baseline approaches. In particular, the HBDM and HBDM-RE perform better than all the non-LDM baselines when D = 2. It highlights the superiority of LDMs in learning very low-dimensional network representations that accurately account for the network structure. We further observe that representing degree heterogeneity with random effects provides extended representational power as the HBDM-RE consistently outperforms the HBDM. Comparing our framework with the classic LDM-RE and LDM, we mostly see on-par results experimentally which we attribute to the hierarchical structure well-preserving properties of homophily and transitivity.

Large-Scale Networks: Results for the large-scale networks are given in TABLE 5. Again, we observe that HBDM-RE was on par with the most competitive baselines of NETSMF and VERSE while significantly outperforming the rest across networks. We also here find that the inclusion of random effects in the LDMs improves performance highlighting the importance of explicitly accounting for degree heterogeneity also for large networks. Notably, we

again detect very good performance for the HBDM-RE, but also for NETSMF and VERSE when utilizing the very low embedding dimension of D = 2.

Bipartite Networks: We validate the performance of our proposed framework for bipartite structures by reporting the AUC score. We perform our experiments on three bipartite networks: (1) Drug-Gene [82] ($N_1 = 5,017$, $N_2 = 2,324, |E| = 15,138$), (2) GitHub [83] ($N_1 = 56,519$, $N_2 = 120, 867, |E| = 440, 237$), and (3) Gottron-Reuters [84] $(N_1 = 21,557, N_2 = 38,677, |E| = 1,464,182)$ following the same experimental setting as in the undirected case of the moderate-sized networks (network details are given in the supplementary). We provide the results in TABLE 7 where we witness how the random-effects formulation of HBDM-RE and VERSE outperform all the baselines and in most cases significantly. Another interesting observation is that the random-effects model has a notably higher performance than the corresponding global bias model, as the three studied networks have a high degree of heterogeneity.

3.2 Hyperparemeter sensitivity

We here study the effect of hyperparameters introduced by the HBDM frameworks. Contrary to many GRL approaches, our models only define three hyperparameters which include the embedding dimensionality D, the number of training iterations, and the learning rate lr for the optimizer. In Fig. 3b, we view the predictive performance as a function of latent dimension D and here, in general, attain modest improvements in the predictive performance when increasing the embedding dimensions from D = 2 to D = 8 with no further improvements increasing to D = 128, highlighting the efficiency in which HBDM and HBDM-RE utilize very low-dimensional representations. Fig. 3c, shows the effect that the learning rate has on performance. We here witness that very small choices $lr \approx 0.001$ define a very slowly increasing performance. Medium magnitude choices of $lr \in [0.005, 0.01]$ define faster convergence while the optimum choices defining very rapid performance saturation exist in the $lr \in [0.05, 0.1]$ regime. In Fig. 3d, we investigate the convergence of the best performing HBDM-RE D = 8 for the large networks, and we witness that the

TABLE 3

AUC-ROC scores for varying dimension sizes on the *Facebook* network for a model approximating the link terms (top two rows) and for the proposed model which calculates analytically the link terms (bottom two rows).

Dimension (D)	2	3	8	32	64	128
HBDM (approx)	.656	.797	.946	.943	.940	.945
HBDM-RE (approx)	.802	.838	.909	.932	.940	.942
HBDM	.980	.986	.986	.987	.986	.985
HBDM-Re	.986	.990	.988	.989	.989	.989

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015



Fig. 3. (a) NLL comparison between HBDM and LDM for *Facebook* with D = 2. (b) AUC-ROC performance over various networks for varying embedding sizes. (c) Performance sensitivity over different learning rate choices for the optimizer in terms of AUC-ROC scores for the *Facebook* network. (d) AUC scores of HBDM-RE in terms of iterations sensitivity for large-scale networks.



Fig. 4. Micro-F1 scores for the classification task considering different low-dimensions and training set ratios for the DBLP network.

TABLE 4 AUC scores for representation sizes of 2 and 8 over moderate-sized networks.

-	AstroPh		GrQc		Facebook		HepTh		Cora		DBLP	
Dimension (D)	2	8	2	8	2	8	2	8	2	8	2	8
DEEPWALK	.831	.945	.845	.919	.958	.986	.773	.874	.684	.782	.803	.939
NODE2VEC	.825	.950	.809	.884	.914	.988	.780	.881	.640	.776	.803	.945
LINE	.632	.910	.688	.920	.751	.980	.659	.874	.634	.521	.625	.503
NETMF	.800	.814	.830	.860	.872	.935	.757	.792	.629	.739	.838	.858
NETSMF	.828	.891	.756	.805	.907	.976	.705	.810	.605	.737	.766	.857
RANDNE	.524	.554	.534	.560	.614	.657	.519	.509	.508	.556	.508	.517
LOUVAINNE	.798	.813	.861	.868	.957	.958	.774	.874	.767	.747	.900	.904
ProNE	.768	.907	.818	.883	.900	.971	.678	.823	.675	.764	.813	.924
VERSE	.899	.974	.885	.941	.970	.992	.844	.910	.749	.760	.910	.955
LDM	.925	х	.915	.943	.989	.991	.855	.919	.780	.786	.918	x
LDM-RE	.943	x	.925	.944	.990	.992	.869	.917	.770	.787	.926	х
HBDM	.920	.960	.917	.944	.980	.986	.853	.915	.786	.792	.919	.956
HBDM-RE	939	964	926	953	986	988	871	924	774	795	930	963

model rapidly converges. After a couple of thousand iterations (very-scalable regime), we already obtain competitive performance for link prediction, which then gently increases until convergence. Our hyperparameter sensitivity analysis focuses on the predictive performance in the downstream task of link prediction. Since our method defines a likelihood over the network, the link predictive performance here shows how well the proposed framework characterizes generalizable patterns of network structure and we therefore focus the analysis on this aspect rather than node classification. If the network structure complies with the node classes, we can expect the node classification performance to follow the same behavior as the link prediction task. Potential disagreement in classification scores against the

TABLE 5 AUC for varying representation sizes over the large-scale networks.

8





Fig. 5. YouTube network—HBDM-RE runtime complexity in miliseconds (ms) as a function of increasing sample sizes of network nodes (in terms of $N \log N$) until the sample set becomes the node set of the total graph. The y-axis showcases the average runtime over 100 iterations of the forward pass while the shaded areas provide standard deviations, as a measure of uncertainty. Runtimes are presented across D = 2, 3, 8 dimensions while we also show the runtime when the inferred hierarchy of the HBDM-RE is created from scratch versus when it is kept static.

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015



Fig. 6. 2-D true embedding space versus 128D t-SNE constructed space. For TES, we provide AUC-ROC for the network reconstruction (NR) task.

TABLE 6 Micro-F₁ scores varying dimension sizes for two moderate and large-scale networks.

	Cora				DBLP			Amazon			YouTube		
Dimension (D)	2	3	128	2	3	128	2	3	8	2	3	8	
DEEPWALK	.502	.712	.838	.519	.605	.822	.231	.596	.929	.293	.351	.413	
NODE2VEC	.419	.658	.835	.448	.540	.815	.096	.305	.895	-	-	-	
LINE	.197	.191	.794	.328	.294	.771	.005	.003	.003	.185	.134	.177	
NETMF	.389	.653	.835	.654	.707	.742	х	x	х	х	х	х	
NETSMF	.554	.705	.842	.622	.732	.829	.387	.649	.845	.317	.361	.397	
RANDNE	.271	.337	.731	.406	.473	.718	.223	.411	.787	.211	.226	.277	
LOUVAINNE	.804	.811	.801	.780	.812	.825	.970	.971	.974	.362	.360	.359	
ProNE	.450	.611	.830	.574	.634	.825	.420	.750	.933	.218	.274	.379	
VERSE	.471	.719	.828	.518	.565	.757	.078	.416	.949	.243	.305	.393	
LDM	.810	.802	.774	х	х	х	х	х	x	х	х	х	
LDM-RE	.802	.803	.796	х	х	х	х	х	х	х	х	х	
HBDM	.789	.807	.816	.812	.814	.772	.970	.971	.931	.320	.366	.414	
HBDM-RE	.805	.813	.818	.805	.822	.808	.956	.955	.931	.326	.367	.414	

link prediction scores implies that the node labels do not follow the network structure, and such discrepancy would be network specific rather than a limitation of the method to be investigated.

3.3 Node classification

We assess the performance of the proposed framework in the uni/multi-label classification task and provide the Micro-F₁ scores in TABLE 6 (Macro-F₁ scores are reported in the supplementary). Scores are defined as the mean value over 10 random shuffles defining the training and test sets. Standard deviations as error bars were found in the scale of 10^{-3} and thus not presented. For the experimental setup, we randomly pick 50% of nodes as the training set and use the rest as the testing set. For an accurate comparison across different methods, we used two simple classifiers, a linear (logistic/multinomial regression classifier) and a non-linear (linearithmic k-nearest neighbors (*kNN*) classifier), and report the highest scores. We found that all methods benefit

from using *kNN*. The number of neighbors was set to k = 10(similar results were obtained with higher choices for k as well). Lastly, we report the average Micro- F_1 scores across 10 repeated trials. Results for the uni-labeled Cora and DBLP networks are reported in the two leftmost columns of TABLE 6. We observe that HBDM-RE and HBDM significantly outperform the baselines for the regimes of D = 2, 3with only LOUVAINNE being competitive. Results for largescale and multi-labeled networks Amazon and YouTube are provided by the two rightmost columns in TABLE 6. Again, the proposed framework outperforms the baselines for the low-dimensional regime with LOUVAINNE being on par. Comparing our framework with the classic LDM-RE and LDM, we again see an on-par performance which we attribute to the HBDM well preserving the intrinsic properties of homophily and transitivity. We further investigate the effect that the amount of training data has on classification performance. In Fig. 4 we provide the performance across multiple training size ratios and consider ultra-low dimensional embeddings of D = 2, 3, 8 for the *DBLP* network. We here observe that for the cases of D = 2, 3, our frameworks significantly outperform all the baselines with only LOU-VAINNE being competitive. Increasing the dimensionality to D = 8, the baseline models are defined with enough capacity to be competitive while HBDM and HBDM-RE return favorable results.

3.4 Across tasks comparison:

We have considered multiple downstream tasks in each of which various baselines were found to be competitive against our HBDM frameworks. In general, the HBDM is characterized by the most consistent performance across tasks, especially for low dimensions. NETSMF was most competitive in the large-scale networks but underperformed

9



Fig. 7. Amazon network dendrogram, embedding space and ordered adjacency matrices for the learned D = 2 embeddings of HBDM-RE and various levels (L) of the hierarchy.

TABLE 7 AUC scores for varying representation sizes over three bipartite networks.

	D	rug-Gei	ne		GitHub)	Gottron-Reuters		
Dimension (D)	2	3	8	2	3	8	2	3	8
DEEPWALK	.673	.843	.878	.762	.853	.902	.673	.769	.905
NODE2VEC	.758	.814	.793	.724	.823	.876	.694	.766	.830
LINE	.798	.836	.867	.805	.766	.902	.715	.696	.850
NETMF	.576	.598	.742	.711	.711	.708	.747	.747	.730
NETSMF	.839	.838	.796	.846	.847	.857	.874	.934	.941
RANDNE	.536	.551	.613	.615	.651	.707	.769	.808	.872
LOUVAINNE	.760	.767	.779	.694	.702	.735	.654	.648	.679
ProNE	.667	.765	.831	.676	.771	.840	.606	.725	.909
VERSE	.910	.913	.922	.943	.952	.959	.962	.966	.967
HBDM	.798	.836	.889	.849	.869	.905	.941	.949	.950
HBDM-RE	.872	.891	.914	.932	.934	.937	.964	.967	.973

in the moderate-sized networks and node classification. VERSE was the most competitive baseline across tasks but massively underperformed in node classification for low dimensions. Furthermore, LOUVAINNE had high performance in node classification but underperformed in link prediction. Models such as NETSMF and VERSE can express structural (stochastic) equivalence while our HBDM explicitly expresses homophily and transitivity. This can explain the occasional higher performance of these baselines in the link prediction task.

3.5 Runtime complexity:

We assess the runtime complexity of the HBDM-RE framework in terms of increasing network sizes. In Fig. 5, we consider the *YouTube* network and show the runtime complexity in milliseconds (ms) as a function of increasing sample sizes of network nodes (in terms of $N \log N$) until the whole network is recovered. Runtimes are presented as

the average across 100 iterations of the forward pass while the shaded areas provide the standard deviation. Runtimes are presented across D = 2, 3, 8 dimensions while we also compare the runtimes when the HBDM-RE builds the hierarchical structure via the k-means procedure and when the hierarchical structure is kept fixed. We here observe, that the runtime increases almost linearly as we increase the number of nertwork nodes. Comparing the runtime when creating the hierarchy via the k-means procedure from scratch versus keeping the dendrogram static from past iterations shows a significant decrease in runtime for the latter (in the experiments we create the hierarchy every 25'th iteration). Thus, the main bottleneck of the HBDM-RE approach is the computations required for the proposed Euclidean k-means procedure. Despite being deemed outside of the scope of this paper, such a bottleneck can be addressed by exploring existing procedures scaling conventional squared Euclidean k-means by avoiding unnecessary distance calculations [85] or by the use of binary space partitioning trees [86]. Such improved scaling would even admit utilization of nonbinary splits beyond the root node improving the accuracy of the hierarchical approximation.

10

3.6 Network visualization

The graph representation learning literature mainly focuses on embeddings with dimensionality greater than D = 2 and 3. As a direct consequence, network visualizations rely on dimensionality reduction frameworks, typically using the t-distributed Stochastic Neighbor Embedding (t-SNE) [87]. In order to verify the quality of the t-SNE constructed Space (t-SNES), we provide the labeled-colored True Embedding Space (TES) in Fig. 6 for D = 2, as well as for D = 128 mapped to D = 2 via the use of t-SNE for *Cora* and *DBLP*. We see that the HBDM-RE frameworks

11

provide highly informative embeddings with no need for dimensionality reduction, unlike the rest of the baselines. This is also verified from the optimal performance in network reconstruction, HBDM-RE can successfully express the network structure using just D = 2. In Fig. 7 we provide the hierarchical block structure constructed by the HBDM-RE for the Amazon network. For visualization, we used the average within-cluster Euclidean distance to the centroid $\left(\Delta\{A,B\} = \frac{1}{N_A + N_B} \sum_{i \in C_A, C_B} \|\boldsymbol{z}_i - \boldsymbol{\mu}_A \|_2\right),$ as a linkage function to form a post-processing agglomerative clustering, for ordering the initial $\log N$ centroids. In Fig. 7 (i), we provide the dendrogram which denotes the agglomeration result in the top-level with red lines. The dendrogram continues with the hierarchical splits of our HBDM-RE where each color indicates the initial $\log N$ blocks. The y-axis of the dendrogram represents the binary logarithm of the Sum of Euclidean Distances, *Log2-SED* = $\log_2\left(\sum_{i \in C_k^{(l)}} \|\boldsymbol{z}_i - \boldsymbol{\mu}_k^{(l)}\|_2\right)$. Moreover, Fig. 7 (ii) conveys the corresponding latent space, colored based on the coarse $\log N$ split, revealing directly interpretable and accurate network representations. In Fig. 7 (iii), (iv), (v) and (vi) we showcase the organized adjacency matrices, based on the 2dimensional HBDM-RE learned hierarchy for various levels L of the tree. We here, observe the representation power of the extracted hierarchy from just a 2-dimensional HBDM-RE defining communities and their sub-communities at finer and finer details.

For the bipartite case, we show how HBDM-RE can enhance our understanding of the bipartite structure at multiple scales and levels. Similar to the undirected case, Fig. 8 (i), indicates the dendrogram of the imposed hierarchy, enriched with agglomeration for a coarse level block ordering and proximity for the GitHub network. In addition, Fig. 8 (ii), provides the corresponding latent space, colored based on the coarse $\log N$ split. Notably, no dimensionalityreduction is necessary to define accurate network representation in the latent space of the two disjoint populations and visually access and express node similarity. In Fig. 8 (iii), (iv) and (v), we exhibit how the multi-scale structure evolves through different levels of the hierarchy defined by HBDM-RE, showcasing how a joint bi-clustering for complex network embeddings naturally can be obtained, with no need for post-processing steps. Our HBDM, can thus accurately characterize bipartite networks and successfully uncover their hierarchical block structure efficiently.

4 DISCUSSION

We developed the HBDM, a scalable reconciliation of latent distance models and their ability to account for homophily and transitivity with hierarchical representations of network structures. We demonstrated how the proposed HBDM provides favorable network representations by: (1) Operating with a Euclidean distance metric providing an intuitive human perception of node similarity. (2) Naturally representing multiscale hierarchical structure based on its block structure and carefully designed clustering procedure optimized in terms of Euclidean distances. (3) Directly and consistently operating in D = 2, 3 with high performance. (4) Performing well on all considered downstream tasks

highlighting its ability to account for the underlying network structure. Importantly, the inferred hierarchical structure admits community discovery at multiple scales as highlighted by the inferred dendrograms and ordered adjacency matrices, and naturally extends to the characterization of communities of bipartite networks.

Our finding of ultra-low dimensional accurate characterizations of network structures supports the findings in [88] in which a logistic PCA model was found to enable exact low-dimensional recovery of multiple real-world networks. Whereas the work of [88] focuses on exact network reconstruction we find that generalizable patterns can be well extracted in ultra-low dimensional representations with performance saturating after just D = 8 dimensions for all networks considered. Whereas [88] found that their lowdimensional space did not perform well in classification tasks we observed strong node classification performance by the low-dimensional representations provided by HBDM. Importantly, for node classification, we observed better performance using KNN as opposed to simple linear classification based on logistic/multinomial regression typically used for node classification. This highlights that whereas most GRL works use linear classifiers there is no guarantee that the embedding space will be linearly separable and performance should therefore be compared to non-linear classifiers as they may provide more favorable performance as observed in this study.

Recent pioneering works [89], [90] have drawn significant attention of the research community by questioning the conventional embedding space preference. It is well known that many real-world networks show power-law degree distribution, or they can consist of latent hierarchical inner structures. Therefore, Euclidean space might not always be appropriate to represent such complex network architectures. It might also require higher-dimensional spaces to show comparable performance in the GRL tasks. The works of [89], [90] demonstrated that hyperbolic spaces, such as the Poincare disk model, can provide substantial benefits over the Euclidean space. The presented HBDM model naturally extends to other distance measures and future studies should explore how the HBDM can be extended to hierarchical representations beyond Euclidean geometry.

Covariate information plays an important role in the outstanding performance of GRL methods and especially GNNs. In the current LSM literature, side information is accounted for by extra regressors in the logit/log link functions expressing the likelihood of a dyad being connected. Using the Mahalanobis distance imposing a blockdiagonal covariance matrix (see supplementary), the proposed HBDM can naturally incorporate covariate information directly to the latent space and notably construct multiscale structures via the enriched and concatenated embedding of the latent variables and the covariate information. Our analysis presently did not explore side information and this is also why we did not include comparisons to prominent GNN-based approaches as these procedures do not provide favorable performance when only learning from the graph structure itself. As such, we observed (not shown) poor performance of GraphSage [91] when only having access to the graph structure in the present setup. The HBDM operates on static networks and thus is not naturally

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

12



Fig. 8. GitHub network dendrogram, embedding space and ordered adjacency matrices for the learned D = 2 embeddings of HBDM-RE and various levels (L) of the hierarchy.

an inductive model. Nevertheless, potential new emerging nodes can be projected into the inferred latent space by fixing the embeddings of nodes present in the training set while optimizing the new nodes for their locations in the learned latent space. We leave a comparison of such a strategy against naturally inductive models such as GNNs for future work.

Our discoveries highlight the existence and importance of hierarchical multi-scale structures in complex networks. The across hierarchy re-ordered adjacency matrices given by HBDM, manifest sub-communities inside of what already appears as a strongly connected community. This points to how delicate the task of defining communities is and the importance of accounting for communities at multiple scales, as enabled by the HBDM. Importantly, these results generalize for bipartite networks where multi-scale geometric representations, joint hierarchical structures, and community discovery are arduous tasks.

The HBDM uses the LDM and thus is good at characterizing transitivity and homophily at a node and cluster level, whereas the random effects enable accounting for degree heterogeneity. Notably, the HBDM suffers from the limitations of the LDM and is thus unable to model stochastic equivalence. Future work should therefore investigate hierarchical structures imposed on more flexible GRL procedures enabling stochastic equivalence and contrast the performance when accounting for stochastic equivalence to the existing hierarchical methods based on the SBM [44], [45], [47]–[49], [52].

In conclusion, we proposed the Hierarchical Block Distance Model (HBDM), a scalable reconciliation of network embeddings using the latent distance model (LDM) and hierarchical characterizations of structure at multiple scales via a novel clustering framework. Notably, the model mimics the behavior of the LDM where the use of homophily and transitivity is most important while scaling in complexity by $\mathcal{O}(DN \log N)$. We analyzed thirteen networks

from moderate sizes to large-scale with the HBDM having favorable performance when compared to existing scalable embedding procedures. In particular, we observed that the HBDM well predicts links and node classes providing accurate network visualizations and characterization of structure at multiple scales. Our results demonstrate that favorable performance can be achieved using ultra-low (i.e. D = 2) embedding dimensions and a scalable hierarchical representation that accounts for homophily and transitivity.

ACKNOWLEDGMENTS

We would like to express sincere appreciation and thank the reviewers for their constructive feedback and their insightful comments. We would also like to thank Louis Boucherie, Lasse Mohr Mikkelsen, and Giorgio Giannone for the valuable and fruitful discussions. The authors gratefully acknowledge the Independent Research Fund Denmark for supporting this work [grant number: 0136-00315B].

REFERENCES

- [1] M. E. J. Newman, "The structure and function of complex networks," SIAM Review, vol. 45, no. 2, pp. 167-256, 2003.
- D. Liben-Nowell and J. Kleinberg, "The link prediction problem for social networks," in CIKM, 2003, p. 556-559.
- [3] L. Getoor and B. Taskar, Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning). The MIT Press, 2007.
- A. Grover and J. Leskovec, "Node2Vec: Scalable feature learning [4] for networks," in *KDD*, 2016, pp. 855–864. S. Fortunato, "Community detection in graphs," *Physics Reports*,
- [5] vol. 486, no. 3, pp. 75-174, 2010.
- D. Zhang, J. Yin, X. Zhu, and C. Zhang, "Network representation [6] learning: A survey," IEEE Transactions on Big Data, vol. 6, no. 1, pp. 3–28, 2020.
- B. Perozzi, R. Al-Rfou, and S. Skiena, "DeepWalk: Online learning [7] of social representations," CoRR, vol. abs/1403.6652, 2014.
- J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "LINE: [8] Large-scale information network embedding," in WWW, 2015, pp. 1067-1077.
- A. Çelikkanat and F. D. Malliaros, "Exponential family graph [9] embeddings," in AAAI, 2020, pp. 3357-3364.

- [10] D. Nguyen and F. D. Malliaros, "BiasedWalk: Biased sampling for representation learning on graphs," in Big Data, 2018, pp. 4045-4053.
- [11] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in NIPS, 2013, pp. 3111–3119. [12] S. Cao, W. Lu, and Q. Xu, "GraRep: Learning graph representa-
- tions with global structural information," in CIKM, 2015, pp. 891-900
- [13] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in NIPS, 2017.
- [14] J. Zhang, Y. Dong, Y. Wang, J. Tang, and M. Ding, "Prone: Fast and scalable network representation learning," in IJCAI, 7 2019, pp. 4278-4284.
- [15] J. Qiu, Y. Dong, H. Ma, J. Li, C. Wang, K. Wang, and J. Tang, "NetSMF: Large-scale network embedding as sparse matrix factorization," in WWW, 2019, pp. 1509-1520.
- [16] A. K. Bhowmick, K. Meneni, M. Danisch, J.-L. Guillaume, and B. Mitra, "LouvainNE: Hierarchical louvain method for high quality and scalable network embedding," in WSDM, 2020, pp. 43-51.
- [17] H. Chen, B. Perozzi, Y. Hu, and S. Skiena, "HARP: hierarchical representation learning for networks," in AAAI, 2018, pp. 2127-2134.
- [18] Z. Zhang, P. Cui, H. Li, X. Wang, and W. Zhu, "Billion-scale network embedding with iterative random projection," in ICDM, 2018, pp. 787-796.
- [19] Matias, Catherine and Robin, Stéphane, "Modeling heterogeneity in random graphs through latent space models: a selective review*," ESAIM: Proc., vol. 47, pp. 55-74, 2014. [Online]. Available: https://doi.org/10.1051/proc/201447004 [20] D. K. Sewell and Y. Chen, "Latent space models for
- dynamic networks," Journal of the American Statistical Association, vol. 110, no. 512, pp. 1646–1657, 2015. [Online]. Available: https://doi.org/10.1080/01621459.2014.988214
- [21] M. Salter-Townshend and T. H. McCormick, "Latent space models for multiview network data," *The Annals of Applied Statistics*, vol. 11, no. 3, pp. 1217 – 1244, 2017. [Online]. Available: https://doi.org/10.1214/16-AOAS955
- [22] L. Zhu, D. Guo, J. Yin, G. V. Steeg, and A. Galstyan, "Scalable temporal latent space inference for link prediction in dynamic social networks," IEEE Transactions on Knowledge and Data Engineering,
- vol. 28, no. 10, pp. 2765–2777, 2016.
 [23] A. Çelikkanat, N. Nakis, and M. Mørup, "Piecewise-velocity model for learning continuous-time dynamic
 "2022 [Online] Available: https://doi.org/10.1016/j. continuous-time dynamic node representations," 2022. [Online]. Available: https: //arxiv.org/abs/2212.12345
- [24] P. Sarkar and A. W. Moore, "Dynamic social network analysis using latent space models," in NIPS, 2006, pp. 1145-1152.
- [25] A. L. Smith, D. M. Asta, and C. A. Calder, "The Geometry of Continuous Latent Space Models for Network Data," Statistical *Science*, vol. 34, no. 3, pp. 428 – 453, 2019. [26] J. Sosa and L. Buitrago, "A review of latent space models for social
- networks," CoRR, vol. abs/2012.02307, 2020.
- [27] B. Kim, K. Lee, L. Xue, and X. Niu, "A review of dynamic network models with latent variables," 2017.
- [28] N. Nakis, A. Çelikkanat, and M. Mørup, "Hm-ldm: A hybridmembership latent distance model," 2022. [Online]. Available: https://arxiv.org/abs/2206.03463
- [29] N. Nakis, A. Çelikkanat, L. Boucherie, C. Djurhuus, F. Burmester, D. M. Holmelund, M. Frolcová, and M. Mørup, "Characterizing polarization in social networks using the signed relational latent distance model," 2023. [Online]. Available: https://arxiv.org/ abs/2301.09507
- [30] P. D. Hoff, A. E. Raftery, and M. S. Handcock, "Latent space approaches to social network analysis," JASA, vol. 97, no. 460, pp. 1090-1098, 2002.
- [31] H. Louch, "Personal network integration: transitivity and homophily in strong-tie relations," Social Networks, vol. 22, no. 1, pp. 45-64, 2000. [Online]. Available: https://www.sciencedirect. com/science/article/pii/S0378873300000150
- [32] C. Zhang, Y. Bu, Y. Ding, and J. Xu, "Understanding scientific collaboration: Homophily, transitivity, and preferential attachment," J. Assoc. Inf. Sci. Technol., vol. 69, no. 1, p. 72-86, jan 2018. [Online]. Available: https://doi.org/10.1002/asi.23916 [33] P. BLOCK and T. GRUND, "Multidimensional homophily in
- friendship networks," Network Science, vol. 2, no. 2, p. 189-212, 2014.

- [34] Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications, ser. Structural Analysis in the Social Sciences. Cambridge University Press, 2012.
- [35] P. N. Krivitsky, M. S. Handcock, A. E. Raftery, and P. D. Hoff, "Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models," Social Networks, vol. 31, no. 3, pp. 204 – 213, 2009.
- [36] D. Cartwright and F. Harary, "Structural balance: a generalization
- of heider's theory." *Psychological review*, vol. 63 5, pp. 277–93, 1956. [37] N. Friel, R. Rastelli, J. Wyse, and A. E. Raftery, "Interlocking directorates in irish companies using a latent space model for bipartite networks," PNAS, vol. 113, no. 24, pp. 6629–6634, 2016.
- M. Handcock, A. Raftery, and J. Tantrum, "Model-based clustering [38] for social networks," J R Stat Soc Ser A Stat Soc, vol. 170, pp. 301 -354, 03 2007.
- [39] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: First steps," Social networks, vol. 5, no. 2, pp. 109-137, 1983.
- [40] K. Nowicki and T. A. B. Snijders, "Estimation and prediction for stochastic blockstructures," JASA, vol. 96, no. 455, pp. 1077-1087, 2001.
- [41] P. D. Hoff, "Modeling homophily and stochastic equivalence in symmetric relational data," in NIPS, 2007, p. 657-664.
- [42] A. Raftery, X. Niu, P. Hoff, and K. Y. Yeung, "Fast inference for the latent space network model using a case-control approximate likelihood," J Comput Graph Stat, vol. 21, 10 2012.
- [43] E. Ravasz and A.-L. Barabási, "Hierarchical organization in complex networks," Phys. Rev. E, vol. 67, p. 026112, 2003.
- [44] A. Clauset, C. Moore, and M. E. Newman, "Hierarchical structure and the prediction of missing links in networks," Nature, vol. 453, no. 7191, pp. 98-101, 2008.
- [45] D. M. Roy, C. Kemp, V. Mansinghka, and J. Tenenbaum, "Learning annotated hierarchies from relational data," in NIPS, vol. 19, 2007.
- [46] D. M. Roy and Y. Teh, "The mondrian process," in NIPS, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds., vol. 21. Curran Associates, Inc., 2008.
- T. Herlau, M. Mørup, M. N. Schmidt, and L. K. Hansen, "Detecting [47] hierarchical structure in networks," in CIP. IEEE, 2012, pp. 1-6.
- [48] T. Herlau, M. Mørup, and M. Schmidt, "Modeling temporal evolution and multiscale structure in networks," in ICML, 2013, pp. 960-968.
- [49] T. P. Peixoto, "Hierarchical block structures and high-resolution model selection in large networks," Physical Review X, vol. 4, no. 1, 2014.
- [50] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," J. Stat. Mech. Theory Exp., vol. 2008, no. 10, p. P10008, 2008.
- [51] Y.-Y. Ahn, J. P. Bagrow, and S. Lehmann, "Link communities reveal multiscale complexity in networks," nature, vol. 466, no. 7307, pp. 761-764, 2010.
- [52] C. Blundell and Y. W. Teh, "Bayesian hierarchical community discovery," in NIPS, vol. 26, 2013.
- [53] T. Li, L. Lei, S. Bhattacharyya, K. Van den Berge, P. Sarkar, P. J. Bickel, and E. Levina, "Hierarchical community detection by recursive partitioning," JASA, pp. 1-18, 2020.
- [54] M. Wieling and J. Nerbonne, "Hierarchical spectral partitioning of bipartite graphs to cluster dialects and identify distinguishing features," in TextGraphs-5, 2010, pp. 33 - 41.
- [55] I. S. Dhillon, "Co-clustering documents and words using bipartite spectral graph partitioning," in KDD, 2001, p. 269-274.
- [56] Q. Cai and J. Liu, "Hierarchical clustering of bipartite networks based on multiobjective optimization," IEEE Trans. Netw. Sci. Eng., vol. 7, no. 1, pp. 421–434, 2020.
- [57] M. Rosvall and C. T. Bergstrom, "An information-theoretic framework for resolving community structure in complex networks," *PNAS*, vol. 104, no. 18, pp. 7327–7331, 2007.
- [58] M. Mørup and M. N. Schmidt, "Bayesian community detection," Neural computation, vol. 24, no. 9, pp. 2434-2456, 2012.
- [59] M. Handcock, A. Raftery, and J. Tantrum, "Model-based clustering for social networks," J. R. Stat. Soc., vol. 170, pp. 301 354, 2007.
- [60] P. D. Hoff, "Bilinear mixed-effects models for dyadic data," JASA, vol. 100, no. 469, pp. 286-295, 2005.
- [61] D. K. Wind and M. Mørup, "Link prediction in weighted networks," in 2012 IEEE International Workshop on Machine Learning for Signal Processing, 2012, pp. 1–6.
- [62] B. Karrer and M. E. Newman, "Stochastic blockmodels and com-munity structure in networks," *Physical review E*, vol. 83, no. 1, p. 016107, 2011.

JOURNAL OF LATEX CLASS FILES, VOL. 14, NO. 8, AUGUST 2015

- [63] T. Herlau, M. N. Schmidt, and M. Mørup, "Infinite-degreecorrected stochastic block model," Physical review E, vol. 90, no. 3, p. 032819, 2014.
- [64] A.-L. Barabási and M. Pósfai, Network science. Cambridge University Press, 2016.
- [65] H. C. White, S. A. Boorman, and R. L. Breiger, "Social structure from multiple networks. i. blockmodels of roles and positions,' American journal of sociology, vol. 81, no. 4, 1976.
- [66] S. S. Epp, Discrete Mathematics with Applications, 4th ed. USA: Brooks/Cole, 2010.
- [67] H. Tsutsu and Y. Morikawa, "An lp norm minimization using auxiliary function for compressed sensing," in Proc. Int. Multiconf. Comp. Sci. Inf. Technol., 2012.
- [68] J. A. Hartigan and M. A. Wong, "Algorithm AS 136: A K-Means clustering algorithm," Applied Statistics, vol. 28, no. 1, pp. 100-108, 1979
- [69] A. Tsitsulin, D. Mottin, P. Karras, and E. Müller, "VERSE," in Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18. ACM Press, 2018. [Online]. Available: https://doi.org/10.1145%2F3178876.3186120
- [70] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017.
- [71] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad, "Collective classification in network data," AI magazine, 2008.
- [72] J. Leskovec and J. J. Mcauley, "Learning to discover social circles
- [72] J. Leskovec and J. J. Malary, "Learning to calculate the second s tems, vol. 42, no. 1, pp. 181–213, Jan 2015.
- [74] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and analysis of online social networks," in IMC, 2007.
- [75] R. Zafarani and H. Liu, "Social computing data repository at ASU," 2009.
- [76] B. Perozzi, V. Kulkarni, H. Chen, and S. Skiena, "Don't walk, skip! online learning of multi-scale network embeddings," in ASONAM, 2017, pp. 258–265.
- [77] J. Leskovec, J. Kleinberg, and C. Faloutsos, "Graph evolution: Densification and shrinking diameters," ACM Trans. Knowl. Discov. Data, vol. 1, no. 1, 2007.
- [78] J. Leskovec and A. Krevl, "SNAP Datasets: Stanford large network dataset collection," Jun. 2014.
- [79] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang, "Network embedding as matrix factorization: Unifying DeepWalk, LINE, PTE, and Node2Vec," in WSDM, 2018.
- [80] S. Dasgupta, "The hardness of k-means clustering," 2008.
- [81] M. Mahajan, P. Nimbhorkar, and K. Varadarajan, "The planar kmeans problem is np-hard," Theoretical Computer Science, vol. 442, pp. 13-21, 2012.
- [82] D. S. Wishart, Y. D. Feunang, A. C. Guo, E. J. Lo, A. Marcu, J. R. Grant, T. Sajed, D. Johnson, C. Li, Z. Sayeeda, N. Assempour, I. Iynkkaran, Y. Liu, A. Maciejewski, N. Gale, A. Wilson, L. Chin, R. Cummings, D. Le, A. Pon, C. Knox, and M. Wilson, "DrugBank 5.0: a major update to the DrugBank database for 2018," Nucleic Acids Research, vol. 46, no. D1, pp. D1074–D1082, 11 2017.
- [83] S. Chacon, "2009 github challenge," 2009. [Online]. Available: https://github.blog/2009-07-29-the-2009-github-contest/
- [84] D. D. Lewis, Y. Yang, T. G. Rose, and F. Li, "Rcv1: A new benchmark collection for text categorization research," J. Mach. Learn. Res., vol. 5, pp. 361-397, 2004.
- [85] C. Elkan, "Using the triangle inequality to accelerate k-means," in Proceedings of the 20th international conference on Machine Learning (*ICML-03*), 2003, pp. 147–153.
- [86] D. Pettinger and G. Di Fatta, "Space partitioning for scalable kmeans," in 2010 Ninth International Conference on Machine Learning and Applications. IEEE, 2010, pp. 319-324.
- [87] L. van der Maaten and G. Hinton, "Visualizing data using t-sne," Journal of Machine Learning Research, vol. 9, no. 86, pp. 2579–2605, 2008.
- [88] S. Chanpuriya, C. Musco, K. Sotiropoulos, and C. E. Tsourakakis, "Node embeddings and exact low-rank representations of complex networks," CoRR, vol. abs/2006.05592, 2020.
- [89] M. D. Ben Chamberlain and J. Clough, "Neural embeddings of graphs in hyperbolic space," in MLG Workshop, 2017.
- [90] M. Nickel and D. Kiela, "Poincaré embeddings for learning hierarchical representations," in NIPS, vol. 30, 2017.

[91] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," Advances in neural information processing systems, vol. 30, 2017.



Nikolaos Nakis is currently a Ph.D. student at the Section for Cognitive Systems of the Technical University of Denmark. He received his BS in physics from the National and Kapodistrian University of Athens and his MS degree in mathematical modeling and computation from the Technical University of Denmark. His research mainly focuses on machine learning applied to complex systems and graph representation learning.



Abdulkadir Çelikkanat is currently a postdoctoral researcher at the Section for Cognitive Systems of the Technical University of Denmark. He completed his Ph.D. at the Centre for Visual Computing of CentraleSupélec, Paris-Saclay University, and he was also a member of the OPIS team at Inria Saclay. Before his Ph.D. studies, he received his Bachelor degree in Mathematics and Master's degree in Computer Engineering from Bogaziçi University. His research mainly focuses on the analysis of graph-

structured data. In particular, he is interested in graph representation learning and its applications for social and biological networks.



Sune Lehmann Sune's work focuses on a quantitative understanding of social systems based on massive data sets. A physicist by training, his research draws on approaches from the physics of complex systems, machine learning, and statistical analysis. He works on large-scale behavioral data and while Sune's primary focus is on modeling complex networks, his research has made substantial contributions on topics such as human mobility, sleep, academic performance, complex contagion, epidemic spreading, and be-

havior on Twitter. He is the author of multiple high-impact papers and his research has won various prizes.



Morten Mørup received the MS and PhD degrees in applied mathematics from the Technical University of Denmark, where he is currently professor at the Section for Cognitive Systems at DTU Compute. He has been an associate editor of the IEEE Transactions on Signal Processing and his research interests include machine learning, neuroimaging, and complex network modeling.