



HM-LDM: A Hybrid-Membership Latent Distance Model

Nakis, Nikolaos; Çelikkanat, Abdulkadir; Mørup, Morten

Published in:

Proceedings of The Eleventh International Conference on Complex Networks and their Applications

Link to article, DOI:

[10.1007/978-3-031-21127-0_29](https://doi.org/10.1007/978-3-031-21127-0_29)

Publication date:

2023

Document Version

Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):

Nakis, N., Çelikkanat, A., & Mørup, M. (2023). HM-LDM: A Hybrid-Membership Latent Distance Model. In *Proceedings of The Eleventh International Conference on Complex Networks and their Applications* (Vol. 1077, pp. 350-363). Springer. https://doi.org/10.1007/978-3-031-21127-0_29

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.

HM-LDM: A Hybrid-Membership Latent Distance Model

Nikolaos Nakis, Abdulkadir Çelikkanat, and Morten Mørup

Section for Cognitive Systems,
Technical University of Denmark, Kongens Lyngby 2800, Denmark
nnak@dtu.dk, abce@dtu.dk, mmor@dtu.dk

Abstract. A central aim of modeling complex networks is to accurately embed networks in order to detect structures and predict link and node properties. The Latent Space Model (LSM) has become a prominent framework for embedding networks and includes the Latent Distance Model (LDM) and Eigenmodel (LEM) as the most widely used LSM specifications. For latent community detection, the embedding space in LDMs has been endowed with a clustering model whereas LEMs have been constrained to part-based non-negative matrix factorization (NMF) inspired representations promoting community discovery. We presently reconcile LSMs with latent community detection by constraining the LDM representation to the D -simplex forming the Hybrid-Membership Latent Distance Model (HM-LDM). We show that for sufficiently large simplex volumes this can be achieved without loss of expressive power whereas by extending the model to squared Euclidean distances, we recover the LEM formulation with constraints promoting part-based representations akin to NMF. Importantly, by systematically reducing the volume of the simplex, the model becomes unique and ultimately leads to hard assignments of nodes to simplex corners. We demonstrate experimentally how the proposed HM-LDM admits accurate node representations in regimes ensuring identifiability and valid community extraction. Importantly, HM-LDM naturally reconciles soft and hard community detection with network embeddings exploring a simple continuous optimization procedure on a volume constrained simplex that admits the systematic investigation of trade-offs between hard and mixed membership community detection.

Keywords: Latent Space Modeling, Community Detection, Non-negative Matrix Factorization, Graph Representation Learning.

1 Introduction

Networks naturally arise in the vast majority of scientific domains from physics to biology in order to model interactions among diverse entities with numerous instances such as collaboration, protein-protein, and brain connectivity networks [23]. Hence, graph analysis tools have become crucial to extract and analyze the underlying meaningful information from networks. In this direction, Graph

Representation Learning (GRL) [36] approaches have become a dominant way to carry out various downstream tasks such as node classification, link prediction, and community detection thanks to their superior performance compared to the classical techniques. GRL models mainly aim to map similar nodes in the network to close latent positions in a low dimension space by automatically learning corresponding node features [7].

The initial GRL works aimed to learn representations or features by simulating random walks over networks, taking inspiration from the Natural Language Processing field [4, 6, 25, 27, 31]. They mainly extract embeddings by optimizing the co-occurrence probability of node pairs within a certain distance through random walks. In recent years, we have witnessed a tremendous increase in the number of Graph Neural Networks (GNN) [7] methods with their usage in supervised tasks. They primarily rely on iterative message-passing operations of node attributes and hidden features around the surroundings of nodes for a given task. The matrix decomposition-based models [26, 27] are also a notable class of the GRL methods. They learn node embeddings by decomposing a designed target matrix based on first and higher-order proximity. However, few GRL methods rely on Non-negative Matrix Factorization (NMF), although it is a popular technique for unsupervised signal decomposition and approximation of multivariate non-negative data. NMF techniques have gathered lots of attention since they allow for structure retrieval through the latent factors of the imposed decomposition providing easy interpretable part-based representations [18].

Applications of NMF include network analysis allowing for efficient, unsupervised, and overlapping community detection, as well as GRL [2, 20, 33, 34]. Within the NMF formulation, various works have sought to define mixed-membership frameworks for analysis and community detection purposes. A Mixed-Membership Stochastic Block Model (MM-SBM) [1] has been linked to the symmetric-NMF decomposition with uniqueness guarantees [20]. Standard least-squares NMF optimization was exchanged to a Poisson likelihood optimization for obtaining the propensity of nodes belonging to different communities [2]. In addition, a GRL approach for overlapping communities was presented in [33] where NMF was utilized to discover Poisson distributed mixed-memberships. These works, design mixed-memberships vectors for part-based representations [18] projected in an NMF constructed space where node similarity, as well as, position and metric properties, can be abstract.

The Latent Space Models (LSMs) are also one of the most powerful ways to learn latent representations [22]. These methods employ generalized linear models for constructing latent node embeddings which express important network characteristics. More specifically, the LDM [11] employs the Euclidean norm for positioning similar nodes closer in the latent space, which comes as a direct consequence of the triangular inequality, naturally representing transitivity (*"a friend of a friend is a friend"*) and homophily (*a tendency where similar nodes are more likely to connect to each other than dissimilar ones*) properties. The LDM can be generalized through the Eigenmodel [10] that can account for stochastic equivalence (*"groups of nodes defined by shared intra- and inter-group*

relationships”) akin to the SBM [1] and the mixed membership SBM [1]. Furthermore, LDMs have been endowed with a clustering model imposing a Gaussian Mixture Model as prior forming the latent position clustering model [8, 29].

In this study, we propose a novel unsupervised representation learning method over graphs, namely, the Hybrid-Membership Latent Distance Model (HM-LDM), by bringing together the strengths of LDM and NMF. Specifically, the HM-LDM offers a reconciliation between part-based representations of networks and low-dimensional latent spaces satisfying similarity properties such as homophily and transitivity. The choice of these similarity properties is of high significance and one of the key characteristics behind GRL since they allow for easily interpretable discovery of network structure. Additionally, our proposed method permits us to capture the latent community structure of the networks using a simple continuous optimization procedure over the log-likelihood of the network. Notably, unlike most existing approaches imposing hard community memberships constraints, the assignment of community memberships in our proposed hybrid model can be controlled and altered through the simplex volume formed by the latent node representations. We extensively evaluate the performance of the proposed method in the ability to perform link prediction, as well as, community discovery over various networks of different types. We demonstrate that our model outperforms recent methods.

Source code: *Hybrid-Membership Latent Distance Model*.

2 Problem statement and proposed method

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph where \mathcal{V} shows the vertex set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ the edge set. We use $\mathbf{Y}_{N \times N} = (y_{i,j}) \in \{0, 1\}^{N \times N}$ to denote the adjacency matrix of the graph where $y_{i,j} = 1$ if the pair $(i, j) \in \mathcal{E}$ otherwise it is equal to 0 for all $1 \leq i < j \leq N := |\mathcal{V}|$. Our main goal is to learn a representation, $\mathbf{w}_i \in \mathbb{R}^D$, for each node $i \in \mathcal{V}$ in a lower dimensional space ($D \ll N$) such that similar nodes in the network should have close embeddings. More specifically, we concentrate on mapping the nodes into the unit D -simplex set, $\Delta^D \subset \mathbb{R}_+^{D+1}$. Therefore, the extracted node embeddings can convey information about latent community memberships. Many GRL approaches also do not provide identifiable or unique solution guarantees, so their interpretation highly depends on the initialization of the hyper-parameters. In this study, we will also address the identifiability problem and seek identifiable solutions which can only be achieved up to a permutation invariance, as reported in Def. 1.

Definition 1 (Identifiability). *An embedding matrix \mathbf{W} whose rows indicating the corresponding node representations is called an identifiable solution up to a permutation if it holds $\widetilde{\mathbf{W}} = \mathbf{W}\mathbf{P}$ for a permutation \mathbf{P} and a solution $\widetilde{\mathbf{W}} \neq \mathbf{W}$.*

We define a Poisson distribution over the adjacency matrix \mathbf{Y} of the network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ to be conditionally independent given the unobserved latent positions,

and write the log-likelihood function as follows:

$$\log P(\mathbf{Y}|\mathbf{A}) = \sum_{\substack{i < j \\ y_{ij}=1}} \log(\lambda_{ij}) - \sum_{i < j} \left(\lambda_{ij} + \log(y_{ij}!) \right). \quad (1)$$

where $\mathbf{A} = (\lambda_{ij})$ is the Poisson rate matrix which has absorbed the dependency over the model parameters. We here adopted a Poisson regression model similar to the work in [9]. In this study, we make use of a Poisson likelihood for modelling binary networks, as validated in [33].

We propose the Hybrid-Membership Latent Distance Model (HM-LDM) with a log-rate based on the ℓ^2 -norm as:

$$\log \lambda_{ij} = \left(\gamma_i + \gamma_j - \delta^p \cdot \|\mathbf{w}_i - \mathbf{w}_j\|_2^p \right), \quad (2)$$

where $\mathbf{w}_i \in [0, 1]^{D+1}$ and $\sum_{d=1}^{D+1} w_{id} = 1$, $\delta \in \mathbb{R}_+$ and $\gamma_i \in \mathbb{R}$ denotes the node-specific random-effects [9, 16] describing essentially the tendency of nodes to sending and receiving connections, accounting for degree heterogeneity. In addition, the norm degree $p \in \{1, 2\}$ controls the power of the ℓ^2 -norm and combined with the latent embeddings sum-to-one condition constrains the latent space to the D -simplex with size equal to δ . A remarkable property of Eq. (2), for $p = 2$, is that it resembles a positive Eigenmodel with random effects: $\tilde{\gamma}_i + \tilde{\gamma}_j + (\tilde{\mathbf{w}}_i \mathbf{A} \tilde{\mathbf{w}}_j^\top)$ where \mathbf{A} is a diagonal matrix having non-negative elements, i.e. $\tilde{\gamma}_i = \gamma_i - \delta^2 \cdot \|\mathbf{w}_i\|_2^2$, $\tilde{\gamma}_j = \gamma_j - \delta^2 \cdot \|\mathbf{w}_j\|_2^2$ and $\tilde{\mathbf{w}}_i \mathbf{A} \tilde{\mathbf{w}}_j^\top = 2\delta^2 \cdot \mathbf{w}_i \mathbf{w}_j^\top$ thus the squared Euclidean distance reconciles the conventional LDM and non-negativity constrained Eigenmodel. The squared Euclidean distance is not fully a metric but it still expresses homophily, leading to an interpretable latent space. Even though the triangle inequality is not exactly satisfied, it preserves the ordering of pairwise Euclidean distances, and it is highly preferred in applications since it is a strictly convex smooth function. By the well-known cosine formula, we have

$$\|\mathbf{w}_i - \mathbf{w}_j\|_2^2 = \|\mathbf{w}_i - \mathbf{w}_k\|_2^2 + \|\mathbf{w}_k - \mathbf{w}_j\|_2^2 + 2\|\mathbf{w}_i - \mathbf{w}_k\|_2 \|\mathbf{w}_k - \mathbf{w}_j\|_2 \cos(\theta),$$

where $\theta \in (-\pi/2, \pi/2)$ is the angle between the vectors $\mathbf{w}_i - \mathbf{w}_k$ and $\mathbf{w}_k - \mathbf{w}_j$. Notice that the third term also converges to 0 for similar nodes since we will have close representations. For the case where $\theta \in [-\pi, -\pi/2] \cup [\pi/2, \pi]$, it holds the triangle inequality: $\|\mathbf{w}_i - \mathbf{w}_j\|_2^2 \leq \|\mathbf{w}_i - \mathbf{w}_k\|_2^2 + \|\mathbf{w}_k - \mathbf{w}_j\|_2^2$.

The embedding vectors, $\{\mathbf{w}_i\}_{i=1}^N$ in Eq. (2), are constrained to non-negative values and to sum to one. Thereby, they reside on a simplex showing the participation of node $i \in \mathcal{V}$ over $D+1$ latent communities. Any LDM can be translated to the non-negative orthant without any loss in performance or in expressive capability. Non-negative embeddings do not affect the distance metric, as it is invariant to translation, as shown by Figure 1 (a). In addition, the D -dimensional non-negative orthant can be reconstructed by a large enough D -simplex. Based on these arguments, it is trivial to show that for large values of the δ parameter in Eq. (2), despite the sum-to-one constraint on the embeddings \mathbf{W} , we obtain an unconstrained LDM, as distances are unbounded when $\delta \rightarrow +\infty$. In this case,

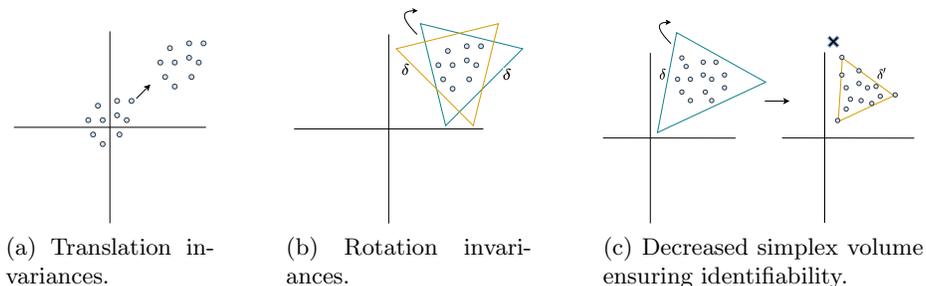


Fig. 1: A 2-dimensional latent space with the 2-simplex given as the green and yellow triangles, the blue points denote embedding positions of the LDM and δ is the simplex size.

the memberships defined by \mathbf{W} are not uniquely identifiable due to the distance invariance of rotation, as seen in Figure 1 (b). However, by shrinking the volume of the simplex (equivalent to decreasing δ), eventually the D -dimensional space of LDM will no longer be enclosed inside the D -simplex, forcing nodes to start populating the corners of this smaller simplex. We call a node *champion* if its latent representation is a standard binary unit vector.

Definition 2 (Community champion). *A node for a latent community is called champion if it belongs to the community (simplex corner) while forming a binary unit vector.*

The champion nodes are of great significance for identifiability because if every corner of the simplex is populated by at least one node (champion), then the solution of the model is identifiable (up to a permutation matrix) (Def. 1) as any random rotation does not leave the solution invariant anymore, as shown by Figure 1 (c). We observe then, that the scalar, δ , controls the type of memberships of the model and its expressive capabilities. Large enough values lead to the basic LDM but inherits its rotational invariance. Small values of δ lead to identifiable solutions and ultimately hard cluster assignments. Thereby, for very small values of δ , nodes are solely assigned to the simplex corners. Lastly, we can also find regimes of values for δ that offer identifiable solutions but also performance similar to LDM, defining a silver lining.

A different take on the identifiability of the model for $p = 2$, can also be given under the Non-negative Matrix Factorization (NMF) theory. This is easily shown by a re-parameterization of Eq. (2) by $\tilde{\gamma}_i + \tilde{\gamma}_j + 2\delta^2 \cdot (\mathbf{w}_i \mathbf{w}_j^\top)$ as described in Eq. (2). In this formulation, the product $\mathbf{W}\mathbf{W}^\top$ defines a symmetric NMF problem which is an identifiable and unique factorization (up to permutation invariance) when \mathbf{W} is full-rank and at least one node resides solely in each simplex corner, ensuring separability [12,20]. Under this NMF formulation, the product $\mathbf{w}_i \mathbf{w}_j^\top \in [0, 1]$ achieves its upper bound only if both nodes i and j reside in the same corner of the simplex. The parameter, δ , acts as a simple multiplicative factor in the first

Table 1: Network statistics; $|\mathcal{V}|$: # Nodes, $|\mathcal{E}|$: # Edges, $|\mathcal{K}|$: # Communities.

	<i>AstroPh</i> [19]	<i>GrQc</i> [19]	<i>Facebook</i> [19]	<i>HepTh</i> [19]	<i>Hamilton</i> [21]	<i>Amherst</i> [21]	<i>Rochester</i> [21]	<i>Mich</i> [21]
$ \mathcal{V} $	17,903	5,242	4,039	8,638	2,118	2,021	4,145	2,933
$ \mathcal{E} $	197,031	14,496	88,234	24,827	87,486	87,496	145,305	54,903
$ \mathcal{K} $	-	-	-	-	15	15	19	13

term of the objective function of HM-LDM, given in Eq. (1), while in the second term acts as a power of the exponential function. For small values of δ , the model is biased towards hard latent community assignments of nodes since similar nodes achieve high rates only when they belong to the same latent community (simplex corner). On the other hand, nodes heading towards the simplex corners for large values of δ lead to an exponential change in the second term of the log-likelihood function given in Eq. (1). Thus, a possible hard allocation of dissimilar nodes to the same community penalizes the likelihood severely. For this reason, high order of δ benefits mixed-membership allocations.

3 Experimental evaluation

We proceed by evaluating the efficiency and performance of the proposed method. In our set-up, we make use of networks with unknown community structures, as well as, with ground-truth communities. We employ the former networks to validate the ability of our framework to discover identifiable latent structures and predict missing links. The latter networks are used to verify that the HM-LDM discovers communities successfully. We consider multiple social and scientific collaboration networks as shown by Table 1. We treat all networks as unweighted and undirected.

For the training of HM-LDM we optimize the log-likelihood function of Eq. (1) via the Adam optimizer [15] with learning rate $lr \in [0.01, 0.1]$. The node-specific random effects vector $\gamma \in \mathbb{R}^N$ is randomly initialized and then tuned alone by optimizing a Poisson log-likelihood with a rate as $\log \lambda_{ij} = \gamma_i + \gamma_j$. Next, the latent embeddings matrix \mathbf{W} is initialized based on the eigenvalues obtained by the spectral decomposition of the normalized Laplacian matrix of the network [13, 24]. In all experiments, we compare against unsupervised methods, and we do not include GNNs since they perform poorly in unsupervised tasks due to the over-smoothing effect [35].

Link prediction: For the link prediction experiments, we follow the well-established strategy [6, 25] and remove 50% of the network edges while keeping the residual network connected. The removed edges combined with a sample of the same number of node pairs (which are not the edges of the original network) construct the negative instances for the testing set. We utilize the residual network to learn the node embeddings.

We consider four networks with unknown community structures and asses performance across different dimensions. In Table 2, we compare the results of our method with other prominent GRL and NMF approaches in terms of the

Table 2: Area Under Curve (AUC-ROC) scores for varying representation sizes.

Dimension (D)	<i>AstroPh</i>			<i>GrQc</i>			<i>Facebook</i>			<i>HepTh</i>		
	8	16	32	8	16	32	8	16	32	8	16	32
DEEPWALK [25]	.945	.950	.952	.919	.916	.929	.986	.986	.984	.874	.867	.873
NODE2VEC [6]	.950	<u>.962</u>	<u>.957</u>	.897	.913	.930	<u>.988</u>	<u>.988</u>	<u>.987</u>	.881	.882	.881
LINE [31]	.909	.938	.947	.920	.925	.919	.981	.987	.983	.873	.886	.882
NETMF [27]	.813	.823	.839	.860	.866	.877	.935	.963	.971	.792	.806	.821
NETSMF [26]	.891	.901	.919	.837	.858	.886	.975	.981	.985	.809	.822	.836
LOUVAINNE [3]	.813	.811	.819	.868	.875	.873	.958	.961	.963	.874	.867	.873
PRONE [37]	.907	.929	.947	.885	.911	.921	.971	.982	.987	.827	.846	.859
NNSD [30]	.861	.882	.891	.792	.808	.828	.908	.927	.935	.756	.779	.796
MNMF [32]	.893	.925	.943	.911	.928	.937	.965	.978	.982	.857	.880	.891
BIGCLAM [34]	.500	.723	.810	.752	.769	.780	.744	.722	.647	.776	.700	.748
SYMMNMF [17]	.767	.779	.800	.729	.772	.835	.933	.942	.951	.696	.727	.766
HM-LDM($p = 1$)	<u>.956</u>	.952	.952	.944	.948	.951	.982	.979	.974	.916	.921	.924
HM-LDM($p = 2$)	.972	.973	.963	<u>.940</u>	<u>.942</u>	<u>.946</u>	.992	.993	.993	<u>.908</u>	<u>.910</u>	<u>.911</u>

Area Under Curve-Receiver Operating Characteristic (AUC-ROC). All baselines have been tuned and feature vectors for dyads are constructed based on binary operators (average, Hadamard, weighted-L1, weighted-L2) [6]. For these constructed feature vectors we further train a logistic regression model with L_2 regularization to make predictions. In particular, for the baselines we choose the hyperparameter settings for each model, as well as, the binary operator for which the logistic regression predictions return the maximum AUC-ROC score.

In contrast, for our models we adopt an unbiased evaluation, and we choose the first of the considered δ values which keeps the solution identifiable (at least one champion per community), as δ decreases. We note though, the existence of identifiable regimes with higher predictive power. Furthermore, predictions and AUC-ROC scores for HM-LDM, can be obtained directly (without the use of a logistic regression model) and are based on the learned Poisson rates λ_{ij} of the test set pairs $\{i, j\}$. The true dimensions for HM-LDM are $D + 1$ but reported as D since they express the true number of model parameters, for a fair comparison with the baselines. For our method, we show the mean performance over five independent runs (error bars were found to be in the scale of 10^{-3} and thus not presented).

Comparing now the results with the non-NMF models, we observe that our HM-LDM (either $p = 1$ or $p = 2$) outperforms the baselines and in most cases significantly, returning favorable results. For the NMF models, we see mostly a big performance gap with the HM-LDM, showcasing the existence of regimes for δ where we can successfully achieve identifiable community memberships while also exhibiting the link prediction power of the LDM. (AUC Precision-Recall scores are similar to the AUC-ROC scores and thus not presented)

Performance and simplex sizes: In Figure 2, we provide the link prediction performance as a function of δ^2 in terms of the AUC-ROC scores across various latent dimensions, networks and for both $p = 1$ and $p = 2$. We here

observe that small δ values provide the minimum scores. This phenomenon is anticipated due to the fact that homophily properties are not sufficiently met (except within clusters) due to the very small simplex volume that these low δ values define. Rethinking HM-LDM with $p = 2$ as a positive Eigenmodel, we can also notice how the positivity constraint on the Λ diagonal matrix does not allow for stochastic equivalence properties which would essentially boost performance even on low simplex volumes. As we increase the values of δ , we naturally reach the performance of an unconstrained LDM. Comparing now, the squared and simple ℓ^2 -norm metric we observe that the former converges to performance saturation more rapidly.

Type and quality of latent memberships: In order to understand how the size of the simplex affects the membership types of HM-LDM, we provide in Figure 3 the total network percentage of community champions as a function of δ^2 across various latent dimensions. As expected, for very small values of δ almost 100% of nodes are assigned solely to a unique simplex corner, yielding hard cluster assignments. As we increase δ , we observe that more and more nodes are assigned with mixed-memberships; on the other hand, the number of champions goes to zero across all dimensions for large values of δ . Contrasting again, the different powers p of the HM-LDM formulation, we notice that the decrease in community champions is steeper for $p = 2$. This also explains why the squared ℓ^2 choice leads to faster convergence in the AUC-ROC, as the model converges faster to the classic LDM. Overall, it is evident that the $p = 2$ HM-LDM needs smaller simplex volumes to be identifiable. We continue with assessing unique latent structures of HM-LDM. For that purpose, in Figure 4 we provide the reorganized adjacency matrices with respect to the community allocations of HM-LDM (for mixed-memberships we assign a node based on the maximum membership). We witness how HM-LDM successfully discovers latent commu-

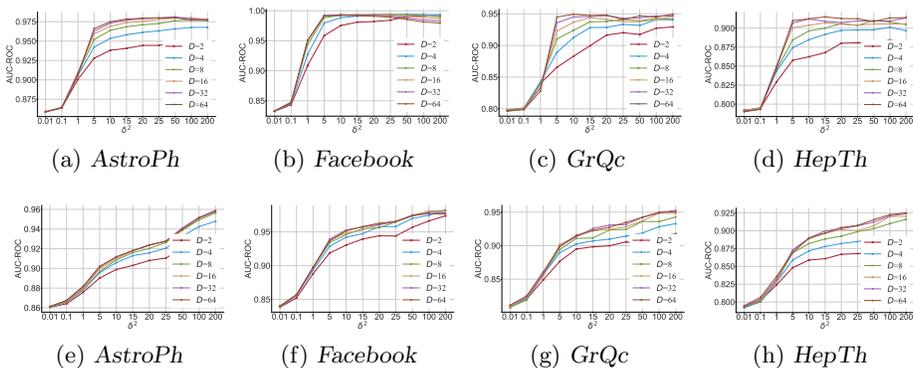


Fig. 2: AUC-ROC scores as a function of δ^2 across dimensions for HM-LDM. Top row: $p = 2$. Bottom row $p = 1$.

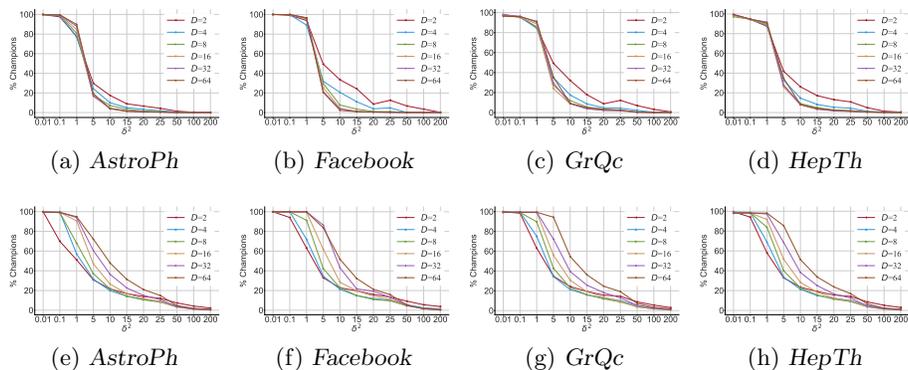


Fig. 3: Total community champions (%) in terms of δ^2 across dimensions for HM-LDM. Top row: $p = 2$. Bottom row $p = 1$.

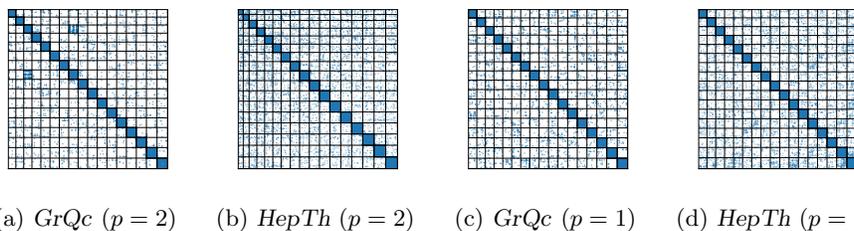


Fig. 4: Ordered adjacency matrices based on the memberships of a $D = 16$ dimensional HM-LDM with δ values ensuring identifiability.

nities, facilitating part-based network representations while choosing appropriate δ regimes ensure identifiability.

Experiments using real ground-truth communities: In order to assess the ability of HM-LDM to discover informative communities, we make use of four networks providing ground-truth community labels. For the NMF-based methods, including ours, we test the ability of the algorithms to detect valid structures by comparing the inferred memberships with the ground-truth community labels while we set the latent dimensions to be equal to the total number of communities. For the GRL approaches which do not define memberships, we extract latent embeddings and use *k-means* (average over 20 runs for robustness) to obtain memberships. We report the Normalized Mutual Information (NMI) score, as well as, the Adjusted Rand Index (ARI), both measures have been validated for community quality assessment in [5]. Again, all the baselines have been tuned individually for each network in terms of their hyperparameters. In contrast, for our HM-LDM, we do not perform any tuning and we just set $\delta = 1$ for all networks since this choice provides in general informative and mostly hard

Table 3: Normalized Mutual Information (NMI) and Adjusted Rand Index (ARI) scores for networks with ground-truth communities.

Metric	<i>Amherst</i>		<i>Rochester</i>		<i>Mich</i>		<i>Hamilton</i>	
	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI
DEEPWALK [25]	.498	.347	.348	.205	.207	.157	.447	.303
NODE2VEC [6]	.535	.375	.364	.223	.217	.161	.481	.348
LINE [31]	.549	.452	.365	.217	.249	.192	.499	.411
NETMF [27]	.491	.330	.377	.243	.237	.136	.456	.297
NETSMF [26]	<u>.562</u>	.408	<u>.381</u>	.228	<u>.242</u>	.169	.494	.391
LOUVAINNE [3]	<u>.562</u>	.395	.347	.204	.175	.114	.475	.334
PRONE [37]	.536	.443	.356	.312	.229	<u>.200</u>	.478	.396
NNSD [30]	.295	.243	.168	.116	.064	.035	.335	.285
MNMF [32]	.542	.362	.324	.171	.188	.102	.466	.287
BIGCLAM [34]	.091	.066	.028	.022	.024	.015	.053	.041
SYMMNMF [17]	.596	.397	.308	.175	.207	.088	.437	.341
HM-LDM($p = 1$)	<u>.562</u>	<u>.502</u>	.400	.392	.228	.205	.527	<u>.485</u>
HM-LDM($p = 2$)	.539	.506	<u>.384</u>	<u>.373</u>	.217	.183	<u>.507</u>	.504

cluster assignments. For our method and the classic LDMs, we report scores averaged over five independent runs in each of which we run the algorithm five times extracting the model with the lowest training loss to remove the effect of local-minimas. We summarize our findings in Table 3, where we witness mostly favorable or on-par performance of HM-LDM with all of the competitive baselines for the NMI metric. For the ARI metric, we observe that our framework outperforms significantly the baselines in all of the considered networks.

Comparison with the LDM: We further investigate the performance of HM-LDM against the LDM, including random effects for a fair comparison and for both normal and squared ℓ^2 -norm LDM-RE and LDM-RE- $(\ell^2)^2$, respectively. Towards that aim, in Table 4 and Table 5 we provide the performance scores for the link prediction and clustering tasks of each model. We here witness that constraining the latent space in identifiable simplex volumes leads to a minor decrease in the predictive power, in terms of the AUC-ROC. For the community detection task, we see favorable NMI scores while the HM-LDM leads to considerably higher ARI scores. Comparing the classical LDM with HM-LDM for $\delta^2 = 10^3$ provides on-par link-prediction performance but the clustering scores drop significantly. This is expected as for large simplex volumes the HM-LDM approximates almost exactly the LDM with the cost of identifiability.

Extension to bipartite networks: Finally, we showcase the extension of our HM-LDM framework to the analysis of bipartite networks. This is straightforward by introducing a different set of latent variables for the two disjoint sets of nodes, as defined by the bipartite structure. In particular, HM-LDM for $p = 2$, simply extends the symmetric NMF formulation, obtained for the undirected networks, to the non-symmetric NMF specification. In Figure 5, we provide the re-ordered adjacency matrix with respect to the community allocations defined by the learned embeddings of HM-LDM for a *Drug-Gene* [19]

Table 4: HM-LDM and LDM-RE comparison for the link prediction task.

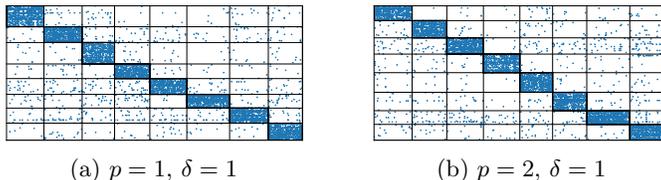
Dimension (D)	<i>AstroPh</i>			<i>GrQc</i>			<i>Facebook</i>			<i>HepTh</i>		
	8	16	32	8	16	32	8	16	32	8	16	32
LDM-RE	.973	.974	.979	.949	.952	.954	.993	.994	.992	.920	.923	.923
HM-LDM($p = 1, \delta^2 = \text{identifiable}$)	.956	.952	.952	.944	.948	.951	.982	.979	.974	.916	.921	.924
HM-LDM($p = 1, \delta^2 = 10^3$)	.967	.967	.965	.956	.955	.951	.985	.986	.987	.932	.931	.926
LDM-RE- $(\ell^2)^2$.979	.978	.976	.944	.944	.945	.990	.990	.991	.913	.912	.909
HM-LDM($p = 2, \delta^2 = \text{identifiable}$)	.972	.973	.963	.940	.942	.946	.992	.993	.993	.908	.910	.911
HM-LDM($p = 2, \delta^2 = 10^3$)	.984	.983	.980	.948	.946	.946	.991	.991	.992	.920	.918	.913

Table 5: HM-LDM and LDM-RE comparison for the clustering task.

Metric	<i>Amherst</i>		<i>Rochester</i>		<i>Mich</i>		<i>Hamilton</i>	
	NMI	ARI	NMI	ARI	NMI	ARI	NMI	ARI
LDM-RE	.548	.366	.391	.212	.230	.132	.491	.320
HM-LDM($p = 1, \delta^2 = \text{identifiable}$)	.562	.502	.400	.392	.228	.205	.527	.485
HM-LDM($p = 1, \delta^2 = 10^3$)	.439	.386	.308	.303	.176	.133	.405	.377
LDM-RE- $(\ell^2)^2$.546	.370	.393	.211	.231	.137	.497	.327
HM-LDM($p = 2, \delta^2 = \text{identifiable}$)	.539	.506	.384	.373	.217	.183	.507	.504
HM-LDM($p = 2, \delta^2 = 10^3$)	.240	.133	.206	.119	.116	.056	.232	.209

network ($|\mathcal{V}| = 7,341$, $|\mathcal{E}| = 15,138$) where we observe a clear block structure. Importantly, the HM-LDM offers identifiable joint embedding representations, mixed memberships, and community discovery for bipartite networks, tasks considered to be non-trivial and arduous.

Complexity analysis: The HM-LDM framework requires the computation of the node pairwise distance matrix and consequently scales prohibitively as $\mathcal{O}(N^2)$ in time and space. Fortunately, there are various ways of scaling HM-LDM for the analysis of large-scale networks. One way is through unbiased estimators of the log-likelihood given by Eq. (1). This is possible through random sampling a set of network nodes S (per iteration) and taking a gradient step based on the log-likelihood of the block defined by the sampled node-set, returning an $\mathcal{O}(S^2)$ space and time complexity. Another option is through the case-control approach [28] scaling on the number of network edges as $\mathcal{O}(E)$.

Fig. 5: *Drug-Gene* ordered adjacency matrices based on HM-LDM with $D = 8$.

Lastly, the Hierarchical Block Distance Model (HBDM) [22] is an attractive option where gradient steps over the model parameters are based on a hierarchical approximation of the likelihood of the whole network. The HBDM model scales linearly as $\mathcal{O}(N \log N)$ both in space and time while also offering hierarchical characterizations of structures at multiple scales.

4 Conclusion and future work

In this paper, we have proposed the HM-LDM that reconciles network embedding and latent community detection. The approach utilizes both the normal and squared Euclidean distance model where the latter integrated the non-negativity constrained Eigenmodel with the Latent Distance Model. We demonstrated that the model could be constrained to the simplex without losing expressive power. The reduced simplex provides unique representations, ultimately resulting in hard clustering of nodes to communities when the simplex is sufficiently shrunk. Notably, the proposed HM-LDM combines network homophily and transitivity properties with latent community detection enabling explicit control of soft and hard assignment through the volume of the induced simplex. We observed favorable link prediction performance in regimes in which the HM-LDM provides unique representations while enabling the ordering of the adjacency matrix in terms of prominent latent communities. Finally, we showed the ability of the model to extract correct community structures across multiple networks and showcased how the analysis extends to bipartite networks. Future work should compare the performance of HM-LDM against classical non-embedding methods such as the Degree Corrected Stochastic Block Model (DC-SBM) [14] or the Mixed Membership Stochastic Block Model (MM-SBM) [1]. Such a comparison is of particular interest since DC-SBM accounts for degree heterogeneity while MM-SBM for soft assignments, two important properties of HM-LDM.

Acknowledgements

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

References

1. Airoldi, E.M., Blei, D.M., Fienberg, S.E., Xing, E.P.: Mixed membership stochastic blockmodels. *J Mach Learn Res* 9(65), 1981–2014 (2008)
2. Ball, B., Karrer, B., Newman, M.E.J.: An efficient and principled method for detecting communities in networks. *CoRR* abs/1104.3590 (2011)

3. Bhowmick, A.K., Meneni, K., Danisch, M., Guillaume, J.L., Mitra, B.: LouvainNE: Hierarchical louvain method for high quality and scalable network embedding. In: WSDM. pp. 43–51 (2020)
4. Çelikkanat, A., Malliaros, F.D.: Exponential family graph embeddings. In: AAAI. pp. 3357–3364 (2020)
5. Chakraborty, T., Dalmia, A., Mukherjee, A., Ganguly, N.: Metrics for community analysis: A survey (2016)
6. Grover, A., Leskovec, J.: Node2Vec: Scalable feature learning for networks. In: KDD. pp. 855–864 (2016)
7. Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.* 40(3), 52–74 (2017)
8. Handcock, M.S., Raftery, A.E., Tantrum, J.M.: Model-based clustering for social networks. *J R Stat Soc Ser A Stat Soc.* 170(2), 301–354 (2007)
9. Hoff, P.D.: Bilinear mixed-effects models for dyadic data. *JASA* 100(469), 286–295 (2005)
10. Hoff, P.D.: Modeling homophily and stochastic equivalence in symmetric relational data (2007)
11. Hoff, P.D., Raftery, A.E., Handcock, M.S.: Latent space approaches to social network analysis. *JASA* 97(460), 1090–1098 (2002)
12. Huang, K., Sidiropoulos, N.D., Swami, A.: Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition. *IEEE Trans. Signal Process* 62(1), 211–224 (2014)
13. Jianbo Shi, Malik, J.: Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8), 888–905 (2000)
14. Karrer, B., Newman, M.E.: Stochastic blockmodels and community structure in networks. *Physical review E* 83(1), 016107 (2011)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)
16. Krivitsky, P.N., Handcock, M.S., Raftery, A.E., Hoff, P.D.: Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social Networks* 31(3), 204 – 213 (2009)
17. Kuang, D., Ding, C., Park, H.: Symmetric nonnegative matrix factorization for graph clustering. In: SDM (2012)
18. Lee, D.D., Seung, H.S.: Learning the parts of objects by nonnegative matrix factorization. *Nature* 401, 788–791 (1999)
19. Leskovec, J., Krevl, A.: SNAP Datasets: Stanford large network dataset collection (Jun 2014)
20. Mao, X., Sarkar, P., Chakrabarti, D.: On mixed memberships and symmetric non-negative matrix factorizations. In: ICML. vol. 70 (2017)
21. Mucha, P., Porter, M.: Social structure of facebook networks. *Physica A: Statistical Mechanics and its Applications* 391, 4165–4180 (08 2012)
22. Nakis, N., Çelikkanat, A., Jørgensen, S.L., Mørup, M.: A hierarchical block distance model for ultra low-dimensional graph representations (2022)
23. Newman, M.E.J.: The structure and function of complex networks. *SIAM Review* 45(2), 167–256 (2003)
24. Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: Analysis and an algorithm. In: Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic. p. 849–856. NIPS’01, MIT Press, Cambridge, MA, USA (2001)
25. Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: KDD. p. 701–710 (2014)

26. Qiu, J., Dong, Y., Ma, H., Li, J., Wang, C., Wang, K., Tang, J.: NetSMF: Large-scale network embedding as sparse matrix factorization. In: WWW (2019)
27. Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., Tang, J.: Network embedding as matrix factorization: Unifying DeepWalk, LINE, PTE, and Node2Vec. In: WSDM. pp. 459–467 (2018)
28. Raftery, A.E., Niu, X., Hoff, P.D., Yeung, K.Y.: Fast inference for the latent space network model using a case-control approximate likelihood. *Journal of Computational and Graphical Statistics* 21(4), 901–919 (2012)
29. Ryan, C., Wyse, J., Friel, N.: Bayesian model selection for the latent position cluster model for social networks. *Network Science* 5(1), 70–91 (2017)
30. Sun, B.J., Shen, H., Gao, J., Ouyang, W., Cheng, X.: A non-negative symmetric encoder-decoder approach for community detection. In: CIKM (2017)
31. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: LINE: Large-scale information network embedding. In: WWW. pp. 1067–1077 (2015)
32. Wang, X., Cui, P., Wang, J., Pei, J., Zhu, W., Yang, S.: Community preserving network embedding. In: AAAI (2017)
33. Wind, D.K., Mørup, M.: Link prediction in weighted networks. In: 2012 IEEE Int. Workshop MLSP. pp. 1–6 (2012)
34. Yang, J., Leskovec, J.: Overlapping community detection at scale: A nonnegative matrix factorization approach. In: WSDM (2013)
35. Yang, L., Gu, J., Wang, C., Cao, X., Zhai, L., Jin, D., Guo, Y.: Toward unsupervised graph neural network: Interactive clustering and embedding via optimal transport. In: ICDM (2020)
36. Zhang, D., Yin, J., Zhu, X., Zhang, C.: Network representation learning: A survey. *IEEE Trans. Big Data* 6(1) (2020)
37. Zhang, J., Dong, Y., Wang, Y., Tang, J., Ding, M.: Prone: Fast and scalable network representation learning. In: IJCAI (2019)