

STERLING: Synergistic Representation Learning on Bipartite Graphs

Baoyu Jing¹, Yuchen Yan¹, Kaize Ding², Chanyoung Park³,
Yada Zhu⁴, Huan Liu⁵, Hanghang Tong¹

¹University of Illinois at Urbana-Champaign

²Northwestern University

³Korea Advanced Institute of Science & Technology

⁴MIT-IBM Watson AI Lab, IBM Research

⁵Arizona State University

baoyuj2@illinois.edu, yuchen5@illinois.edu, kaize.ding@northwestern.edu, cy.park@kaist.ac.kr,
yzhu@us.ibm.com, huanliu@asu.edu, htong@illinois.edu

Abstract

A fundamental challenge of bipartite graph representation learning is how to extract informative node embeddings. Self-Supervised Learning (SSL) is a promising paradigm to address this challenge. Most recent bipartite graph SSL methods are based on contrastive learning which learns embeddings by discriminating positive and negative node pairs. Contrastive learning usually requires a large number of negative node pairs, which could lead to computational burden and semantic errors. In this paper, we introduce a novel synergistic representation learning model (STERLING) to learn node embeddings without negative node pairs. STERLING preserves the unique local and global synergies in bipartite graphs. The local synergies are captured by maximizing the similarity of the inter-type and intra-type positive node pairs, and the global synergies are captured by maximizing the mutual information of co-clusters. Theoretical analysis demonstrates that STERLING could improve the connectivity between different node types in the embedding space. Extensive empirical evaluation on various benchmark datasets and tasks demonstrates the effectiveness of STERLING for extracting node embeddings.

Introduction

The bipartite graph is a powerful representation formalism to model interactions between two types of nodes, which has been used in various real-world applications. In recommender systems (Wang et al. 2021a; Wei and He 2022; Zhou et al. 2021), users, items and their interactions (e.g., buy) naturally formulate a bipartite graph; in drug discovery (Pavlopoulos et al. 2018), chemical interactions between drugs and proteins also formulate a bipartite graph; in information retrieval (He et al. 2016), clickthrough between queries and web pages can be modeled by a bipartite graph.

A fundamental challenge for bipartite graphs is how to extract informative node embeddings that can be easily used for downstream tasks (e.g., link prediction). In recent years, Self-Supervised Learning (SSL) has become a prevailing paradigm to learn embeddings without human-annotated labels (Wu et al. 2021b; Zheng et al. 2022). Despite its great performance on downstream tasks (e.g., node classification), most of the methods are designed for homogeneous graphs

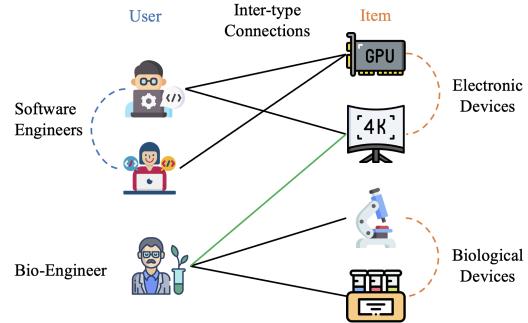


Figure 1: Example of a bipartite graph and its unique local and global properties. Locally, the dashed curves are *implicit* intra-type connections. Globally, the green line shows the inter-connection between the co-clusters, i.e., bio-engineer and electronic devices.

(You et al. 2020; Zhu et al. 2021; Feng et al. 2022a; Zhou et al. 2022; Zheng et al. 2021; Wang et al. 2023; Ding et al. 2023) and heterogeneous graphs (Park et al. 2020; Jing, Park, and Tong 2021; Wang et al. 2021b; Fu et al. 2020; Yan et al. 2022). These methods are usually sub-optimal to bipartite graphs (Gao et al. 2018), and therefore, several methods have been specifically proposed for bipartite graphs. BiNE (Gao et al. 2018) and BiANE (Huang et al. 2020) learn embeddings by maximizing the similarity of neighbors sampled by random walks; NeuMF (He et al. 2017) and NGCF (Wang et al. 2019) train neural networks by reconstructing the edges; BiGI (Cao et al. 2021), SimGCL (Yu et al. 2022) and COIN (Jing et al. 2022b) further improve the quality of embeddings via contrastive learning.

The aforementioned methods are mainly based on contrastive learning, which learns node embeddings by discriminating positive node pairs (e.g., local neighbors) and negative node pairs (e.g., unconnected nodes). The success of contrastive learning heavily relies on the careful treatment of large-scale negative node pairs (e.g., adaptive selection of negative samples) (Grill et al. 2020). There still lacks a principled mechanism to efficaciously construct desirable negative node pairs, which may cause computational burden (Thakoor et al. 2021; Zheng, Zhu, and He 2023) and

semantic errors (Li, Jing, and Tong 2022). Recently, BGRL (Thakoor et al. 2021) and AFGRL (Lee, Lee, and Park 2022) propose to learn node embeddings without negative pairs for homogeneous graphs via bootstrapping. However, these non-contrastive methods cannot be effectively applied to bipartite graphs due to their incapability of capturing unique properties of bipartite graphs.

As shown in Fig. 1, bipartite graphs have their unique local and global properties. Locally, besides the explicit inter-type connections (solid lines), the implicit intra-type synergies are also important (the dashed curves). In Fig. 1, it is very likely that a user (the girl software engineer) will buy the item (the monitor) that was bought by a similar user (the boy software engineer). Globally, the two node types are inter-connected, and thus their cluster-level semantics are inherently synergistic. For example, users and items in Fig. 1 can be clustered based on their professions (software engineers and bio-engineer) and usage (electronic devices and biological devices). The relationship among these co-clusters is not a simple one-to-one correspondence, e.g., software engineers \sim electronic devices and bio-engineer \sim biological devices. In fact, they are inherently interconnected: the bio-engineer cluster is also connected to the electronic device cluster. Neither independently nor equally treating these co-clusters could capture such a synergy.

In this paper, we introduce a novel synergistic representation learning model (STERLING) for bipartite graphs. Compared with bipartite contrastive learning methods, STERLING is a non-contrastive SSL method that does not require negative node pairs. Compared with SSL methods on general graphs, STERLING captures both local and global properties of bipartite graphs. For the local synergies, we maximize the similarity of the positive node embedding pairs. When creating positive node pairs, not only do we consider the inter-type synergies (e.g., connected users and items), but also the intra-type synergies (e.g., similar users). For the global synergies, we introduce a simple end-to-end deep co-clustering model to produce co-clusters of the two node types, and we capture the global cluster synergies by maximizing the mutual information of the co-clusters. We further present the theoretical analysis and empirical evaluation of STERLING. In theoretical analysis, we prove that maximizing the mutual information of co-clusters increases the mutual information of the two node types in the embedding space. This theorem indicates that maximizing the mutual information of co-clusters could improve the connectivity of two node types in the embedding space. In empirical evaluation, we extensively evaluate STERLING on various real-world datasets and tasks to demonstrate its effectiveness.

The major contributions are summarized as follows:

- A novel SSL model (STERLING) is proposed for bipartite graphs, which preserves local and global synergies of bipartite graphs and does not require negative node pairs.
- Theoretical analysis shows that STERLING improves the connectivity of the two node types in embedding space.
- Extensive evaluation is conducted to demonstrate the effectiveness of the proposed STERLING.

Related Work

Graph Embedding. Graphs are ubiquitous in real-world applications, e.g., social network (Goyal and Ferrara 2018; Zheng et al. 2023; Yan, Zhang, and Tong 2021; Zeng et al. 2023a), finance (Zhou et al. 2020; Jing, Tong, and Zhu 2021) and natural language processing (Wu et al. 2021a; Jing et al. 2021; Yan et al. 2021). A fundamental challenge of graph representation learning is to extract informative node embeddings (Wu et al. 2021b; Zeng et al. 2023c). Early methods DeepWalk (Perozzi, Al-Rfou, and Skiena 2014), node2vec (Grover and Leskovec 2016), LINE (Tang et al. 2015) use the random walk to sample node pairs and maximizes their similarities. Contrastive learning methods, e.g., DGI (Velickovic et al. 2019) GraphCL (You et al. 2020), GCA (Zhu et al. 2021) and ARIEL (Feng et al. 2022b), learn embeddings by discriminating positive and negative node pairs. Recently, some non-contrastive methods are proposed. BGRL (Thakoor et al. 2021) and AFGRL (Lee, Lee, and Park 2022) learn embeddings via bootstrapping, and GraphMAE (Hou et al. 2022) learns embeddings by reconstructing the full graph from the masked graph. All of these methods are designed for homogeneous graphs, yet many real-world graphs are heterogeneous (Hu et al. 2020; Yan et al. 2023a; Zeng et al. 2023b). Metapath2vec (Dong, Chawla, and Swami 2017) extends node2vec to heterogeneous graphs. DMGI (Park et al. 2020) and HDMI (Jing, Park, and Tong 2021) extends DGI to heterogeneous graphs. HeCo (Wang et al. 2021b) and X-GOAL (Jing et al. 2022a) introduces co-contrastive learning and prototypical contrastive learning for heterogeneous graphs. Although they achieved impressive performance on downstream tasks (e.g., classification), they are not tailored for bipartite graphs and usually have sub-optimal performance compared with bipartite graph methods on tasks such as recommendation and link prediction.

Bipartite Graph Embedding. Bipartite graphs have been widely used to model interactions between two disjoint node sets (He et al. 2020; Mao et al. 2021; Yan et al. 2023b; Zhang et al. 2023). Early methods such as BiNE (Gao et al. 2018) and BiANE (Huang et al. 2020) learn node embeddings based on biased random walks. IGE (Zhang et al. 2017) learns node embeddings based on the direct connection between nodes and edge attributes. PinSage (Ying et al. 2018) combines graph convolutional network with random walks. Collaborative filtering methods are also popular for bipartite graphs (Wei et al. 2020). NeuMF (He et al. 2017) is the first neural collaborative filtering method. NGCF (Wang et al. 2019) incorporates high-order collaborative signals to improve the quality of embeddings. DirectAU (Wang et al. 2022) learns node embeddings from the perspective of alignment and uniformity. In recent years, contrastive learning has been applied to bipartite graph. BiGI (Cao et al. 2021) extends DGI to from homogeneous graphs to bipartite graphs. SimGCL (Yu et al. 2022) and COIN (Jing et al. 2022b) use InfoNCE (Oord, Li, and Vinyals 2018) as the loss. AdaGCL (Jiang, Huang, and Huang 2023) uses contrastive loss as an auxiliary signal for the recommendation task. Different from these contrastive methods, STERLING does not require any negative node pairs, and further explores the synergies among co-clusters.

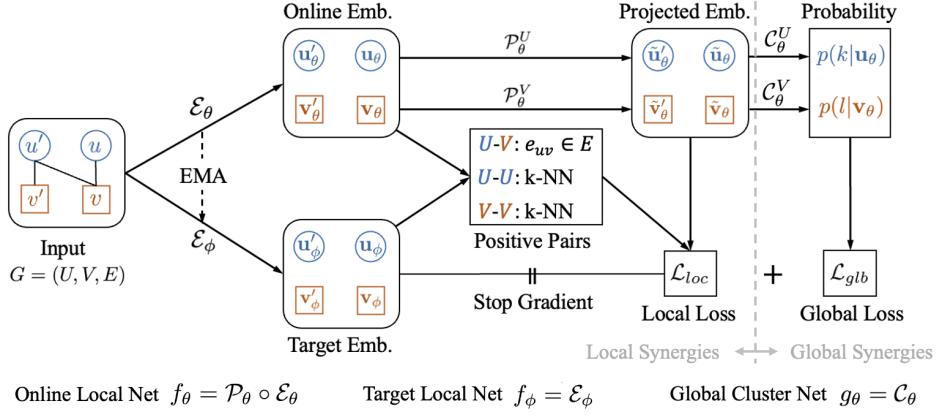


Figure 2: Overview of STERLING. \mathcal{E} , \mathcal{P} and \mathcal{C} are the encoder, projector and cluster network. θ and ϕ are parameters of the online and target networks. θ is updated by optimizing objectives, ϕ is updated via Exponential Moving Average (EMA) of θ . For details, please refer to the methodology section.

Co-Clustering. Co-clustering aims to partition rows and columns of a co-occurrence matrix into co-clusters simultaneously. In practice, co-clustering algorithms usually have impressive improvements over traditional one-way clustering algorithms (Xu et al. 2019). CCInfo (Dhillon, Mallela, and Modha 2003) co-clusters matrices via a mutual information based objective. SCC (Dhillon 2001) is based on spectral analysis. BCC (Shan and Banerjee 2008), LDCC (Shafiei and Milios 2006) and MPCCE (Wang et al. 2011) are Bayesian approaches. CCMOD (Ailem, Role, and Nadif 2015) obtains co-clusters by maximizing graph modularity. SCMK (Kang, Peng, and Cheng 2017) is a kernel based method. DeepCC (Xu et al. 2019) is a deep learning based co-clustering method, which uses an auto-encoder to extract embeddings, a deep Gaussian mixture model to obtain co-cluster assignments, and a fixed input prior distribution to regularize co-clusters. Different from DeepCC, STERLING uses non-contrastive SSL to extract embeddings which can be used for various downstream tasks, and its loss function for co-clustering is simple and the joint distribution of different nodes is learned from data.

Preliminary

Self-Supervised Learning on Bipartite Graphs. Given a bipartite graph $G = (U, V, E)$, where U, V are disjoint node sets, and $E \subseteq U \times V$ is the edge set, the task is to extract informative node embeddings $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{|U| \times d}$ from G . We use u, v and \mathbf{u}, \mathbf{v} to denote elements for U, V and \mathbf{U}, \mathbf{V} .

Co-Clustering. Given a bipartite graph G , co-clustering maps U, V into $N_K \ll |U|, N_L \ll |V|$ clusters via function c_θ , which produces the probabilities of cluster assignments $p(k|u), p(l|v)$ for nodes u, v . Here k, l are cluster indices. We use K, L to denote random variables of co-clusters.

Methodology

Overview of STERLING

An overview of STERLING is shown in Fig. 2. For local synergies, our idea is to obtain the node embeddings via

an encoder \mathcal{E}_θ and maximize the similarity of positive node embedding pairs without minimizing the similarity of negative pairs (e.g., randomly sampled node pairs). The positive pairs are selected based on both inter-type and intra-type synergies. The inter-type $U-V$ positive pairs are the connected $u-v$ pairs ($e_{uv} \in E$). The intra-type $U-U$ (or $V-V$) positive pairs are selected based on the k-NN $u-u$ (or $v-v$) pairs. The pair $(\mathbf{u}_\theta, \mathbf{v}_\theta)$ in Fig. 2 is an exemplar positive pair as they are connected in G : $e_{uv} \in E$. However, directly maximizing the similarity of $(\mathbf{u}_\theta, \mathbf{v}_\theta)$ without minimizing the similarity of negative pairs may result in mode collapse (e.g., mapping all nodes to the same embedding) (Grill et al. 2020). To address this issue, following (Grill et al. 2020), we use a target encoder \mathcal{E}_ϕ to obtain the bootstrapped target embeddings $(\mathbf{u}_\phi, \mathbf{v}_\phi)$, and also project the original online embeddings $(\mathbf{u}_\theta, \mathbf{v}_\theta)$ into $(\tilde{\mathbf{u}}_\theta, \tilde{\mathbf{v}}_\theta)$ via a projector $\mathcal{P}_\theta = (\mathcal{P}_\theta^U, \mathcal{P}_\theta^V)$, which could potentially add noise to the original embeddings $(\mathbf{u}_\theta, \mathbf{v}_\theta)$. Then we maximize the similarity of $(\tilde{\mathbf{u}}_\theta, \mathbf{v}_\phi)$ and $(\tilde{\mathbf{v}}_\theta, \mathbf{u}_\phi)$. We denote $f_\theta = \mathcal{P}_\theta \circ \mathcal{E}_\theta$ as the online local network and $f_\phi = \mathcal{E}_\phi$ as the target local network. f_θ and f_ϕ are updated alternatively. f_θ is trained by maximizing the similarity of positive node pairs \mathcal{L}_{loc} and f_ϕ is updated via the Exponential Moving Average (EMA) of θ : $\theta \leftarrow \tau\phi + (1-\tau)\theta, \tau \in [0, 1]$.

For global synergies, we use a co-clustering network $g_\theta = \mathcal{C}_\theta = (\mathcal{C}_\theta^U, \mathcal{C}_\theta^V)$ to obtain the co-cluster probabilities $p(k|\mathbf{u}_\theta), p(l|\mathbf{v}_\theta)$. Since neural networks \mathcal{P} are treated as an injective function in practice (Vincent et al. 2008), and thus $p(k|\tilde{\mathbf{u}}_\theta) = p(k|\mathbf{u}_\theta), p(l|\tilde{\mathbf{v}}_\theta) = p(k|\mathbf{v}_\theta)$. The global objective \mathcal{L}_{glb} is to maximize the mutual information of the co-clusters K and L . The global objective \mathcal{L}_{glb} is also used to update f_θ . Note that co-clustering and mutual information calculation does not require negative node pairs.

After training, we use the online embeddings $\mathbf{u}_\theta, \mathbf{v}_\theta$ and co-cluster probabilities $p(k|\mathbf{u}_\theta), p(l|\mathbf{v}_\theta)$ for downstream tasks. $c_\theta = g_\theta \circ f_\theta$ is the final co-clustering function.

In the following content, we introduce local objective and global objective in detail. Then we introduce the overall objective and provide theoretically analysis.

Local Objective

There are two kinds of local synergies among nodes in bipartite graphs: the explicit inter-type (U - V) and the implicit intra-type synergies (U - U , V - V). The local objective \mathcal{L}_{loc} captures both inter and intra-type synergies by maximizing the similarities of the inter and intra-type positive node pairs.

Inter-Type Synergies. Node embeddings \mathbf{u}_θ and \mathbf{v}_θ should be similar if u, v are connected $e_{uv} \in E$. Rather than directly maximizing the similarity of $(\mathbf{u}_\theta, \mathbf{v}_\theta)$, we maximize the similarity of the projected online embedding $\tilde{\mathbf{u}}_\theta$ and the target embedding \mathbf{v}_ϕ , as well as $\tilde{\mathbf{v}}_\theta$ and \mathbf{u}_ϕ :

$$\mathcal{L}_{uv} = -\left(\frac{\tilde{\mathbf{u}}_\theta^T \mathbf{v}_\phi}{\|\tilde{\mathbf{u}}_\theta\| \cdot \|\mathbf{v}_\phi\|} + \frac{\tilde{\mathbf{v}}_\theta^T \mathbf{u}_\phi}{\|\tilde{\mathbf{v}}_\theta\| \cdot \|\mathbf{u}_\phi\|}\right) \quad (1)$$

Note that when updating θ via the above objective function (as well as the following objective functions), ϕ is fixed. ϕ is updated via EMA after θ is updated. This practice could effectively avoid the mode collapse problem (Grill et al. 2020).

Intra-Type Synergies. If $u, u' \in U$ are highly correlated, then their embeddings should have a high similarity score¹. We determine the similarity of u, u' from the perspectives of both graph structure and hidden semantics.

For the structure information, we use the Adamic-Adar (AA) index (Adamic and Adar 2003), which is the inverse log frequency (or degree) of the shared neighbors of u, u' :

$$s_{aa}(u, u') = \sum_{v \in V_{uu'}} \frac{1}{\log(d_v)} \quad (2)$$

where $V_{uu'} \subset V$ is the set of v that connects with both u and u' ; d_v is the degree of v . Essentially, the AA index calculates the second-order structure proximity and it has two characteristics: (1) the more common neighbors u and u' share, the higher the AA index will be; (2) the lower the degree d_v of the shared neighbor v , the higher the importance of v .

For the semantic information, given a node u , we determine its semantic similarity with another node u' by:

$$s_{emb}(u, u') = \frac{\mathbf{u}_\theta^T \mathbf{u}'_\phi}{\|\mathbf{u}_\theta\| \cdot \|\mathbf{u}'_\phi\|} \quad (3)$$

The final similarity score between u and u' is thus:

$$s(u, u') = s_{aa}(u, u') \cdot s_{emb}(u, u') \quad (4)$$

Given $u \in U$, we select its k-Nearest-Neighbors (k-NN), i.e., top-K similar nodes, from U to construct positive pairs. Similar to \mathcal{L}_{uv} in Equation (1), the loss for each u is:

$$\mathcal{L}_u = \frac{-1}{N_{knn}} \sum_{u' \in k\text{-NN}(u)} \left(\frac{\tilde{\mathbf{u}}_\theta^T \mathbf{u}'_\phi}{\|\tilde{\mathbf{u}}_\theta\| \cdot \|\mathbf{u}'_\phi\|} + \frac{\tilde{\mathbf{u}}_\theta'^T \mathbf{u}_\phi}{\|\tilde{\mathbf{u}}_\theta'\| \cdot \|\mathbf{u}_\phi\|} \right) \quad (5)$$

where N_{knn} is the number of the selected neighbors. Note that for each $v \in V$, we use the same strategy as $u \in U$ described above to obtain \mathcal{L}_v :

$$\mathcal{L}_v = \frac{-1}{N_{knn}} \sum_{v' \in k\text{-NN}(v)} \left(\frac{\tilde{\mathbf{v}}_\theta^T \mathbf{v}'_\phi}{\|\tilde{\mathbf{v}}_\theta\| \cdot \|\mathbf{v}'_\phi\|} + \frac{\tilde{\mathbf{v}}_\theta'^T \mathbf{v}_\phi}{\|\tilde{\mathbf{v}}_\theta'\| \cdot \|\mathbf{v}_\phi\|} \right) \quad (6)$$

¹For clarity, we only use U - U synergies to describe our method. V - V synergies are captured in the same way as U - U .

Local Objective Function. During training, given a connected pair (u, v) , its local objective is:

$$\mathcal{L}_{loc} = \lambda_{uv} \mathcal{L}_{uv} + \lambda_u \mathcal{L}_u + \lambda_v \mathcal{L}_v \quad (7)$$

where \mathcal{L}_{uv} , \mathcal{L}_u and \mathcal{L}_v are obtained from Eq. (1)(5)(6).

Global Objective

The two types of nodes U, V are inherently correlated, and thus so as their clusters K, L , as illustrated by the green link in Fig. 1. Jointly co-clustering U and V usually produce better results than traditional one-side clustering (Xu et al. 2019). In this paper, we introduce a simple end-to-end co-clustering algorithm to capture the global cluster synergy by maximizing the mutual information of the co-clusters $I(K; L)$. According to the definition of mutual information, to calculate $I(K; L)$, we need to obtain the joint distribution $p(k, l)$ and marginal distributions $p(k)$ and $p(l)$. We decompose $p(k, l)$ by two other easy-to-obtain distributions $p(u, v)$ and $p(k, l|u, v)$ via $p(k, l) = \sum_{u, v} p(k, l|u, v)p(u, v)$. In the following content, we first introduce the joint distribution $p(u, v)$, then introduce the conditional probability $p(k, l|u, v)$, and finally introduce the global objective \mathcal{L}_{glob} .

Joint Distribution $p(u, v)$. The joint distribution $p(u, v)$ characterizes the connectivity of u, v . Instead of simply deriving $p(u, v)$ from the edges E and treating $p(u, v)$ as a fixed prior in (Xu et al. 2019), STERLING learns $p(u, v)$ to encode both structure and semantic information.

For the structure information, we build an n -hop metapath (Dong, Chawla, and Swami 2017) between u and v to find potential links between them. The 1-hop U - V metapath is the original U - V graph, and the 2-hop U - V metapath is the U - V - U - V graph. We denote \mathbf{A}_{meta} as the adjacency matrix of the n -hop u - v metapath.

For the semantic information, we construct \mathbf{A}_{emb} from the extracted node embeddings:

$$\mathbf{A}_{emb} = \frac{1}{2} [\delta(\mathbf{U}_\theta \mathbf{V}_\phi^T) + \delta(\mathbf{U}_\phi \mathbf{V}_\theta^T)] \quad (8)$$

where $\mathbf{U}_\theta, \mathbf{V}_\theta$ are online embeddings and $\mathbf{U}_\phi, \mathbf{V}_\phi$ are target embeddings, δ is an activation function. We further filter out noisy connections by resetting the small values as 0:

$$\mathbf{A}_{emb} = \max(\mathbf{A}_{emb}, \mu + \alpha\sigma) \quad (9)$$

where μ and σ are the mean and standard deviation of \mathbf{A}_{emb} , and α is a tunable threshold.

Finally, the joint distribution $p(U, V)$ is given by:

$$p(U, V) = \frac{1}{Z} \mathbf{A}_{meta} \odot \mathbf{A}_{emb} \quad (10)$$

where Z is a normalization factor, \odot is Hadamard product.

Conditional Distribution $p(k, l|u, v)$. As for $p(k, l|u, v)$, we obtain it via neural networks. We first extract the online node embeddings \mathbf{U}_θ and \mathbf{V}_θ from the input bipartite graph $G = (U, V, E)$ via \mathcal{E}_θ . Then we apply the function $\mathcal{C}_\theta \circ \mathcal{P}_\theta = (\mathcal{C}_\theta^U \circ \mathcal{P}_\theta^U, \mathcal{C}_\theta^V \circ \mathcal{P}_\theta^V)$ over \mathbf{u}_θ and \mathbf{v}_θ to obtain the co-cluster probabilities $p(k|\mathbf{u}_\theta)$ and $p(l|\mathbf{v}_\theta)$. Since neural networks are usually treated as deterministic and injective functions in practice (Vincent et al. 2008), we have

$p(k, l|u, v) = p(k, l|\mathbf{u}_\theta, \mathbf{v}_\theta)$. Furthermore, since $\mathcal{C}_\theta^U \circ \mathcal{P}_\theta^U$ and $\mathcal{C}_\theta^V \circ \mathcal{P}_\theta^V$ have separate sets of parameters, it is natural to have $p(k, l|\mathbf{u}_\theta, \mathbf{v}_\theta) = p(k|\mathbf{u}_\theta)p(l|\mathbf{v}_\theta)$. Hence, we have $p(k, l|u, v) = p(k|\mathbf{u}_\theta)p(l|\mathbf{v}_\theta)$.

Global Objective Function. Combining the conditional distribution $p(k, l|u, v) = p(k|\mathbf{u}_\theta)p(l|\mathbf{v}_\theta)$ with the joint distribution $p(u, v)$ obtained from Equation (10), we have:

$$p(k, l) = \sum_{u, v} p(k|\mathbf{u}_\theta)p(l|\mathbf{v}_\theta)p(u, v) \quad (11)$$

Then we could easily obtain the marginal distributions $p(k) = \sum_l p(k, l)$, $p(l) = \sum_k p(k, l)$. Since $p(k, l)$, $p(k)$ and $p(l)$ are calculated by neural networks, we can directly maximize $I(K; L)$. The global loss \mathcal{L}_{glb} is given by:

$$\mathcal{L}_{glb} = -I(K; L) = -\sum_{k=1}^{N_K} \sum_{l=1}^{N_L} p(k, l) \log \frac{p(k, l)}{p(k)p(l)} \quad (12)$$

Overall Objective

The overall objective function of STERLING is:

$$\mathcal{L} = \mathcal{L}_{loc} + \mathcal{L}_{glb} \quad (13)$$

Theoretical Analysis

Theorem 1 shows that $I(\mathbf{U}_\theta; \mathbf{V}_\theta)$ is lower bounded by $I(K; L)$, indicating that maximizing $I(K; L)$ (or minimizing \mathcal{L}_{glb}) could improve the connectivity of \mathbf{U}_θ and \mathbf{V}_θ in the embedding space. This theorem is corroborated by visualization results in Fig. 4a-4b in the Experiment section. Please refer to Appendix for the proof.

Theorem 1 (Information Bound). *The mutual information $I(\mathbf{U}_\theta; \mathbf{V}_\theta)$ of embeddings \mathbf{U}_θ and \mathbf{V}_θ is lower-bounded by the mutual information of co-clusters $I(K; L)$:*

$$I(K; L) \leq I(\mathbf{U}_\theta; \mathbf{V}_\theta) \quad (14)$$

Experiments

Experimental Setup

Data. Table 1 shows the summary of datasets. ML-100K and Wiki are processed by (Cao et al. 2021), where Wiki has two splits (50%/40%) for training. IMDB, Cornell and Citeceer are document-keyword bipartite graphs (Xu et al. 2019).

Evaluation. For recommendation to a given user u , the score of an item v is determined by the similarity of $\mathbf{u}_\theta, \mathbf{v}_\theta$, and then items with top-K scores are selected. We use F1, Normalized Discounted Cumulative Gain (NDCG), Mean Average Precision (MAP), and Mean Reciprocal Rank (MRR) as metrics. For link prediction, given the learned embeddings $\mathbf{U}_\theta, \mathbf{V}_\theta$, and edges E , we train a logistic regression classifier, and then evaluate it on the test data. We use Area Under ROC Curves (AUC) as the metric. For co-clustering, we assign the cluster with the highest probability as the cluster assignment for the given node, and use Normalized Mutual Information (NMI) and accuracy (ACC) as the metrics.

Baselines. Three groups of baselines are used: (1) **Bipartite Graph**: random-walk methods BiNE (Gao et al. 2018), PinSage (Ying et al. 2018); matrix completion methods GCMC (Berg, Kipf, and Welling 2017), IGMC (Zhang and

Dataset	Task	$ U $	$ V $	$ E $	# Class
ML-100K	Recommendation	943	1,682	100,000	-
Wikipedia	Link Prediction	15,000	3,214	64,095	-
IMDB	Co-Clustering	617	1878	20,156	17
Cornell	Co-Clustering	195	1,703	18,496	5
CiteSeer	Co-Clustering	3,312	3,703	105,165	6

Table 1: Summary of the datasets.

Chen 2019); collaborative filtering methods NeuMF (He et al. 2017), NGCF (Wang et al. 2019) DirectAU (Wang et al. 2022); contrastive learning methods BiGI (Cao et al. 2021), SimGCL (Yu et al. 2022), COIN (Jing et al. 2022b);

(2) **Graph**: traditional methods DeepWalk (Perozzi, Al-Rfou, and Skiena 2014), LINE (Tang et al. 2015), Node2vec (Grover and Leskovec 2016), VGAE (Kipf and Welling 2016); contrastive methods GraphCL (You et al. 2020), non-contrastive methods AFGRL (Lee, Lee, and Park 2022), GraphMAE (Hou et al. 2022). For heterogeneous graph methods, we compare with the random-walk method Metapath2vec (Dong, Chawla, and Swami 2017), and contrastive methods DMGI (Park et al. 2020), HDMI (Jing, Park, and Tong 2021), HeCo (Wang et al. 2021b). (3) **Co-Clustering**: traditional methods CCIInfo (Dhillon, Mallela, and Modha 2003), SCC (Dhillon 2001), CCMOD (Ailem, Role, and Nadif 2015) and SCMK (Kang, Peng, and Cheng 2017); the SOTA deep learning method DeepCC (Xu et al. 2019).

Implementation. The encoder \mathcal{E} is a simple L -layer message passing model $\mathbf{u}^{(l+1)} = \text{AGG}(\mathbf{u}^{(l)}, \{\mathbf{v}^{(l)} : e_{uv} \in E\})$, where AGG is an aggregation function. $\mathbf{v}^{(l+1)}$ is obtained in a similar way. The projector \mathcal{P} is either a Multi-Layer Perceptron (MLP) or identity mapping. The cluster network \mathcal{C} is a MLP, and its final activation function is softmax. We set $N_K = N_L$ for co-clusters. We perform grid search over several hyper-parameters such as N_{knn} , N_K , α , the number of layers, and embedding size d . We set δ as absolute activation. Please refer to Appendix for more details.

Overall Performance

Recommendation. The results on ML-100K are shown in the left part of Table 2. The upper/middle/lower groups of baselines are homogeneous/heterogeneous/bipartite methods respectively. Comparing the best-performing baselines of the three groups, we observe that (1) the contrastive homogeneous and heterogeneous methods (GraphCL and HDMI) are competitive to each other, and non-contrastive method AFGRL performs better than contrastive methods; (2) the bipartite graph method SimGCL is significantly better than GraphCL/AFGRL/HDMI. This observation demonstrates that homogeneous and heterogeneous graph methods perform sub-optimally on bipartite graphs as they do not capture the synergies of bipartite graphs. STERLING further outperforms SimGCL over all metrics, indicating the superiority of the non-contrastive SSL approach over the contrastive approaches for bipartite graphs.

Link Prediction. The results are shown in the right part of Table 2, where the upper/middle/lower parts are homogeneous/heterogeneous/bipartite methods. We could observe:

Method	ML-100K										Wiki(50%) Wiki(40%)	
	F1@10	NDCG@3	NDCG@5	NDCG@10	MAP@3	MAP@5	MAP@10	MRR@3	MRR@5	MRR@10	AUC	AUC
DeepWalk	14.20	7.17	9.32	13.13	2.72	3.54	4.92	43.86	46.83	48.75	87.19	81.60
LINE	13.71	6.52	8.57	12.37	2.45	3.26	4.67	44.16	44.37	46.30	66.69	64.28
Node2vec	14.13	7.69	9.91	13.41	3.07	3.90	5.19	44.80	48.02	49.78	89.37	88.41
VGAE	11.38	6.43	8.18	10.93	2.35	2.95	3.94	39.39	42.32	43.68	87.81	86.32
GraphCL	19.46	10.13	13.24	18.17	4.17	5.65	8.04	58.04	60.67	61.97	94.40	93.67
GraphMAE	21.28	11.35	14.67	20.11	4.66	6.22	9.01	62.96	65.12	66.12	95.12	94.48
AFGRL	22.04	11.70	15.16	20.85	4.75	6.44	9.34	65.18	66.95	67.95	94.80	94.22
Metapath2vec	14.11	7.88	9.87	13.35	2.85	3.71	5.08	45.49	48.74	49.83	87.20	86.75
DMGI	19.58	10.16	13.13	18.31	3.98	5.33	7.82	59.33	61.37	62.71	93.02	92.01
HDMI	20.51	11.07	14.32	19.42	4.59	6.18	8.67	62.25	64.38	65.44	94.18	93.57
HeCo	19.65	11.18	14.15	18.98	4.73	6.26	8.74	58.68	60.02	61.23	94.39	93.72
PinSage	21.68	10.95	14.51	20.27	4.52	6.18	9.13	62.56	64.77	65.76	94.27	92.79
BiNE	14.83	7.69	9.96	13.79	2.87	3.80	5.24	48.14	50.94	52.51	94.33	93.15
GC-MC	20.65	10.88	13.87	19.21	4.41	5.84	8.43	60.60	62.21	63.53	91.90	91.40
IGMC	18.81	9.21	12.20	17.27	3.50	4.82	7.18	56.89	59.13	60.46	92.85	91.90
NeuMF	17.03	8.87	11.38	15.89	3.46	4.54	6.45	54.42	56.39	57.79	92.62	91.47
NGCF	21.64	11.03	14.49	20.29	4.49	6.15	9.11	62.56	64.62	65.55	94.26	93.06
DirectAU	21.04	11.13	14.27	19.65	4.76	6.21	8.79	59.99	62.53	63.80	94.62	93.98
BiGI	23.36	12.50	15.92	22.14	5.41	7.15	10.50	66.01	67.70	68.78	94.91	94.08
COIN	24.78	13.48	17.37	23.62	5.71	7.82	11.34	70.58	72.14	72.76	95.30	94.53
SimGCL	25.19	13.51	17.62	24.08	5.73	7.94	11.62	71.31	73.02	73.77	95.22	94.62
STERLING	25.54	14.01	18.23	24.37	6.06	8.40	11.93	71.99	73.55	74.27	95.48	95.04

Table 2: Performance (%) of top-K recommendation on ML-100K (left), and link prediction on Wikipedia (right).

Metric	Dataset	SCC	CCMod	CCInfo	SCMK	DeepCC	STERLING
NMI	IMDB	25.5	21.6	18.7	18.4	26.8	33.4
	Cornell	28.8	18.9	20.6	25.7	35.4	37.5
	Citeseer	15.2	16.9	17.7	21.1	29.8	31.6
ACC	IMDB	25.2	24.7	23.0	18.4	23.3	34.7
	Cornell	58.9	55.5	56.6	49.6	68.7	73.4
	Citeseer	37.4	44.7	43.0	50.2	59.3	63.7

Table 3: NMI (upper) & ACC (lower) for co-clustering.

(1) STERLING has the best overall performance; (2) homogeneous/heterogeneous approaches are sub-optimal on bipartite graphs; (3) the non-contrastive method (STERLING) is better than contrastive method (COIN/SimGCL).

Co-Clustering. The results in Tables 3 show that among all the baseline methods, DeepCC achieves the highest NMI and ACC scores on all datasets. DeepCC obtains node embeddings by reconstructing the input matrix, and trains cluster networks based on the fixed prior distribution derived from the input matrix. STERLING has further improvements over DeepCC, indicating that (1) the non-contrastive SSL is better than the re-construction based representation learning, and (2) the learned joint distribution $p(U; V)$ helps improve the performance over the fixed prior distribution.

Ablation Study

Components in \mathcal{L} . In the upper part of Table 4, we study the impact of each component in the loss function. (1) The global synergies \mathcal{L}_{glb} are important for all datasets. Note that for the Cornell dataset (co-clustering task), removing \mathcal{L}_{glb} means we use a dummy un-trained \mathcal{C}_θ , which is a random deterministic function mapping similar node embeddings to

Method	ML-100K (F1@10)	Wiki(40%) (AUC)	Cornell (NMI)
STERLING	25.54	95.04	37.51
w/o \mathcal{L}_{glb}	25.12	94.37	26.02
w/o \mathcal{L}_v	21.67	94.89	37.06
w/o \mathcal{L}_u	25.38	94.74	35.74
w/o \mathcal{L}_{uv}	0.20	94.76	30.93
w/o \mathbf{A}_{meta}	25.41	95.01	26.61
w/o \mathbf{A}_{emb}	25.15	94.42	37.43
w/o noise filter(Eq.(9))	25.17	94.70	36.82
abs \rightarrow ReLU(Eq.(8))	25.20	94.84	36.53
w/o s_{aa}	20.32	94.23	32.75
w/o s_{emb}	24.84	94.84	37.10
AA \rightarrow Co-HITS	24.94	94.90	35.87

Table 4: Ablation study.

similar cluster distributions. An NMI score of 26.02 rather than 0 means that the local loss \mathcal{L}_{loc} can discover the clusters of node embeddings to a certain degree. (2) The intra-types synergies \mathcal{L}_u and \mathcal{L}_v have different impacts for different datasets. For ML-100K, \mathcal{L}_v is more important, and for Wiki and Cornell, \mathcal{L}_u has a higher impact. (3) The inter-type synergies \mathcal{L}_{uv} are indispensable for all datasets. Surprisingly, for ML-100K (recommendation task) the model can barely recommend correct items to given users. These results imply that for simpler tasks, which only require class/cluster-level predictions, such as link prediction (0/1 classification) and co-clustering, the implicit intra-type synergies \mathcal{L}_u , \mathcal{L}_v and global synergies could provide a large amount of the information needed. For harder tasks requiring precise element-level predictions, e.g., recommendation (ranking items for a given users), explicit inter-type information \mathcal{L}_{uv} is required.

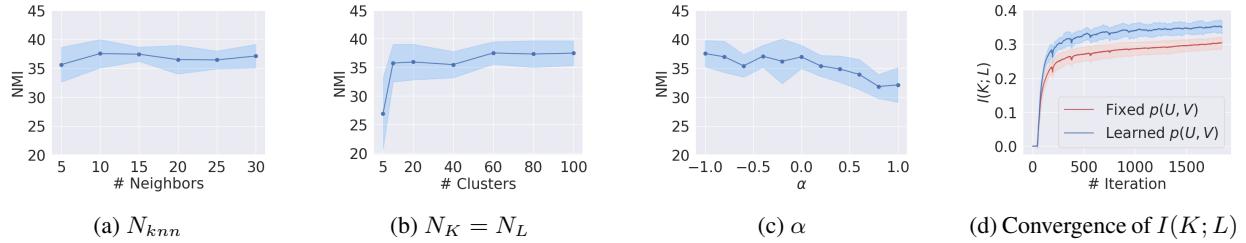
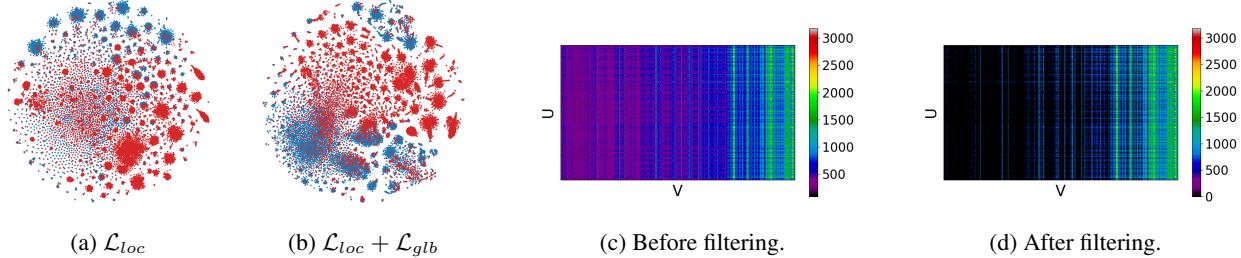
Figure 3: (a-c) Sensitivity analysis and (d) convergence of $I(K; L)$ on the Cornell dataset.

Figure 4: (a-b) T-SNE visualization on Wiki (40%). (c-d) Visualization of noise filter on ML-100K.

Method	ML-100K (F1@10)	Wiki(40%) (AUC)	Cornell (NMI)
STERLING	25.54	95.04	37.51
$\mathbf{A}_{emb} = \mathbf{U}_\theta \mathbf{V}_\phi^T$	25.37	94.48	36.92
$\mathbf{A}_{emb} = \mathbf{U}_\phi \mathbf{V}_\theta^T$	25.46	94.91	35.01
$\mathbf{A}_{emb} = \mathbf{U}_\theta \mathbf{V}_\theta^T$	25.49	94.71	36.57
$\mathbf{A}_{emb} = \mathbf{U}_\phi \mathbf{V}_\phi^T$	25.46	94.36	36.09
$s_{emb} = \frac{1}{2}(\mathbf{u}_\theta^T \mathbf{v}_\phi + \mathbf{u}_\phi^T \mathbf{v}_\theta)$	25.40	94.97	35.84
$s_{emb} = \mathbf{u}_\phi^T \mathbf{v}_\theta$	25.46	94.93	36.82
$s_{emb} = \mathbf{u}_\theta^T \mathbf{v}_\theta$	25.34	95.03	37.34
$s_{emb} = \mathbf{u}_\phi^T \mathbf{v}_\phi$	25.43	94.86	35.13

Table 5: Different Variants of \mathbf{A}_{emb} and s_{emb} .

Components in \mathcal{L}_{glb} . The results in the middle part of Table 4 show that: (1) \mathbf{A}_{meta} has more impact for Cornell and \mathbf{A}_{emb} is more influential for ML-100K and Wiki; (2) noise filtering is useful; (3) surprisingly, the absolute activation performs better than ReLU. We believe this is because the datasets only record strength of connections but do not distinguish the sign (positive/negative) of connections.

Components in \mathcal{L}_{loc} . The results in the lower part of Table 4 show that (1) both s_{aa} , s_{emb} are crucial, and s_{aa} plays a more important role; (2) the AA index is better than another popular index Co-HITS (Deng, Lyu, and King 2009) since the AA index will raise the weights of the low-degree neighbors, and lower the weights of the high-degree neighbors.

Other Results

Sensitivity. The sensitivity analysis on the Cornell dataset is shown in Fig. 3a-3c. The optimal numbers of k-NN, co-clusters and noise threshold are: $N_{knn} \approx 10$, $N_K = N_L \geq 60$, and $\alpha \in [-1, 0]$. The optimal N_K is larger than the real number of classes (5), implying over-clustering is beneficial.

Convergence of MI. Fig. 3d shows the $I(K; L)$ of each training iteration on the Cornell dataset. $I(K; L)$ will converge, using either the fixed or learned $p(U, V)$. However, the learned $p(U, V)$ results in higher $I(K; L)$ in the end.

Visualization of Emb. T-SNE (Van der Maaten and Hinton 2008) visualization for embeddings of Wiki(40%) test data is shown in Fig. 4a-4b. Each embedding is the concatenation of a given $(\mathbf{u}_\theta, \mathbf{v}_\theta)$ pair, which is the input of the logistic regression classifier. These embeddings are labeled with 1/0, indicating whether a pair is true/false. The two colors in Fig. 4a-4b correspond to the two classes. Fig. 4a-4b show that \mathcal{L}_{glb} further helps discover the underlying connectivity of $(\mathbf{u}_\theta, \mathbf{v}_\theta)$ and better separate embeddings of the two classes than \mathcal{L}_{loc} alone, which corroborates Theorem 1.

Visualization of Noise Filtering. We visualize the effect of noise filtering (Eq.(9)) on ML-100K in Fig. 4c-4d. The weak connections can be effectively filtered out since the purple dots in Fig. 4c are removed (becomes black) after filtering in Fig. 4d. The matrix densities (the ratio of non-zero elements) in Fig. 4c and 4d are 100% and 35.66% respectively.

Different Variants. We show the results of different variants of \mathbf{A}_{emb} and s_{emb} (normalization is dropped for clarity) in Table 5, which indicate that the variants used in STERLING has the best overall performance.

Conclusion

In this paper, we introduce a novel non-contrastive SSL method STERLING for bipartite graphs, which preserves both local inter/intra-type synergies and global co-cluster synergies. Theoretical analysis indicates that STERLING could improve the connectivity of the two node types in the embedding space. Empirical evaluation shows that the node embeddings extracted by STERLING have SOTA performance on various downstream tasks.

Acknowledgments

This work is partly supported by NSF (#2229461), DARPA (HR001121C0165), NIFA (2020-67021-32799), MIT-IBM Watson AI Lab, and IBM-Illinois Discovery Accelerator Institute. The content of the information in this document does not necessarily reflect the position or the policy of the Government or Amazon, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. Dr. Chanyoung Park is supported by the IITP grant funded by the Korea government (MSIT) (RS-2023-00216011).

References

Adamic, L. A.; and Adar, E. 2003. Friends and neighbors on the web. *Social networks*.

Ailem, M.; Role, F.; and Nadif, M. 2015. Co-clustering document-term matrices by direct maximization of graph modularity. In *CIKM*.

Berg, R. v. d.; Kipf, T. N.; and Welling, M. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*.

Cao, J.; Lin, X.; Guo, S.; Liu, L.; Liu, T.; and Wang, B. 2021. Bipartite graph embedding via mutual information maximization. In *WSDM*.

Deng, H.; Lyu, M.; and King, I. 2009. A generalized co-hits algorithm and its application to bipartite graphs. In *KDD*.

Dhillon, I. S. 2001. Co-clustering documents and words using bipartite spectral graph partitioning. In *KDD*.

Dhillon, I. S.; Mallela, S.; and Modha, D. S. 2003. Information-theoretic co-clustering. In *KDD*.

Ding, K.; Wang, Y.; Yang, Y.; and Liu, H. 2023. Eliciting structural and semantic global knowledge in unsupervised graph contrastive learning. In *AAAI*.

Dong, Y.; Chawla, N. V.; and Swami, A. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *KDD*.

Feng, S.; Jing, B.; Zhu, Y.; and Tong, H. 2022a. Adversarial graph contrastive learning with information regularization. In *TheWebConf*.

Feng, S.; Jing, B.; Zhu, Y.; and Tong, H. 2022b. ARIEL: Adversarial Graph Contrastive Learning. *arXiv preprint arXiv:2208.06956*.

Fu, D.; Xu, Z.; Li, B.; Tong, H.; and He, J. 2020. A view-adversarial framework for multi-view network embedding. In *CIKM*.

Gao, M.; Chen, L.; He, X.; and Zhou, A. 2018. Bine: Bipartite network embedding. In *SIGIR*.

Goyal, P.; and Ferrara, E. 2018. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*.

Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent-a new approach to self-supervised learning. *NeurIPS*.

Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *KDD*.

He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*.

He, X.; Gao, M.; Kan, M.-Y.; and Wang, D. 2016. Birank: Towards ranking on bipartite graphs. *TKDE*.

He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *TheWebConf*.

Hou, Z.; Liu, X.; Cen, Y.; Dong, Y.; Yang, H.; Wang, C.; and Tang, J. 2022. Graphmae: Self-supervised masked graph autoencoders. In *KDD*.

Hu, Z.; Dong, Y.; Wang, K.; and Sun, Y. 2020. Heterogeneous graph transformer. In *TheWebConference*.

Huang, W.; Li, Y.; Fang, Y.; Fan, J.; and Yang, H. 2020. Biane: Bipartite attributed network embedding. In *SIGIR*.

Jiang, Y.; Huang, C.; and Huang, L. 2023. Adaptive graph contrastive learning for recommendation. In *KDD*.

Jing, B.; Feng, S.; Xiang, Y.; Chen, X.; Chen, Y.; and Tong, H. 2022a. X-GOAL: Multiplex Heterogeneous Graph Prototypical Contrastive Learning. In *CIKM*.

Jing, B.; Park, C.; and Tong, H. 2021. Hdmi: High-order deep multiplex infomax. In *TheWebConf*.

Jing, B.; Tong, H.; and Zhu, Y. 2021. Network of tensor time series. In *TheWebConf*.

Jing, B.; Yan, Y.; Zhu, Y.; and Tong, H. 2022b. COIN: Co-Cluster Infomax for Bipartite Graphs. *NeurIPS GLFrontiers*.

Jing, B.; You, Z.; Yang, T.; Fan, W.; and Tong, H. 2021. Multiplex Graph Neural Network for Extractive Text Summarization. In *EMNLP*.

Kang, Z.; Peng, C.; and Cheng, Q. 2017. Twin learning for similarity and clustering: A unified kernel approach. In *AAAI*.

Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*.

Lee, N.; Lee, J.; and Park, C. 2022. Augmentation-free self-supervised learning on graphs. In *AAAI*.

Li, B.; Jing, B.; and Tong, H. 2022. Graph Communal Contrastive Learning. In *TheWebConf*.

Mao, K.; Zhu, J.; Xiao, X.; Lu, B.; Wang, Z.; and He, X. 2021. UltraGCN: ultra simplification of graph convolutional networks for recommendation. In *CIKM*.

Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.

Park, C.; Kim, D.; Han, J.; and Yu, H. 2020. Unsupervised attributed multiplex network embedding. In *AAAI*.

Pavlopoulos, G. A.; Kontou, P. I.; Pavlopoulou, A.; Bouyioukos, C.; Markou, E.; and Bagos, P. G. 2018. Bipartite graphs in systems biology and medicine: a survey of methods and applications. *GigaScience*, 7.

Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *KDD*.

Shafiei, M. M.; and Milios, E. E. 2006. Latent Dirichlet co-clustering. In *ICDM*.

Shan, H.; and Banerjee, A. 2008. Bayesian co-clustering. In *ICDE*.

Tang, J.; Qu, M.; Wang, M.; Zhang, M.; Yan, J.; and Mei, Q. 2015. Line: Large-scale information network embedding. In *TheWebConf*.

Thakoor, S.; Tallec, C.; Azar, M. G.; Azabou, M.; Dyer, E. L.; Munos, R.; Veličković, P.; and Valko, M. 2021. Large-scale representation learning on graphs via bootstrapping. *arXiv preprint arXiv:2102.06514*.

Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *JMLR*.

Velickovic, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. *ICLR*.

Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*.

Wang, C.; Yu, Y.; Ma, W.; Zhang, M.; Chen, C.; Liu, Y.; and Ma, S. 2022. Towards representation alignment and uniformity in collaborative filtering. In *KDD*.

Wang, H.; Jing, B.; Ding, K.; Zhu, Y.; and Zhou, D. 2023. Characterizing Long-Tail Categories on Graphs. *arXiv preprint arXiv:2305.09938*.

Wang, P.; Laskey, K.; Domeniconi, C.; and Jordan, M. 2011. Nonparametric bayesian co-clustering ensembles. In *SDM*.

Wang, S.; Hu, L.; Wang, Y.; He, X.; Sheng, Q. Z.; Orgun, M. A.; Cao, L.; Ricci, F.; and Yu, P. S. 2021a. Graph learning based recommender systems: A review. *arXiv preprint arXiv:2105.06339*.

Wang, X.; He, X.; Wang, M.; Feng, F.; and Chua, T.-S. 2019. Neural graph collaborative filtering. In *SIGIR*.

Wang, X.; Liu, N.; Han, H.; and Shi, C. 2021b. Self-supervised heterogeneous graph neural network with co-contrastive learning. In *KDD*.

Wei, T.; and He, J. 2022. Comprehensive fair meta-learned recommender system. In *KDD*.

Wei, T.; Wu, Z.; Li, R.; Hu, Z.; Feng, F.; He, X.; Sun, Y.; and Wang, W. 2020. Fast adaptation for cold-start collaborative filtering with meta-learning. In *ICDM*.

Wu, L.; Chen, Y.; Ji, H.; and Liu, B. 2021a. Deep learning on graphs for natural language processing. In *SIGIR*.

Wu, L.; Lin, H.; Gao, Z.; Tan, C.; and Li, S. Z. 2021b. Self-supervised on graphs: Contrastive, generative, or predictive. *arXiv e-prints*, arXiv-2105.

Xu, D.; Cheng, W.; Zong, B.; Ni, J.; Song, D.; Yu, W.; Chen, Y.; Chen, H.; and Zhang, X. 2019. Deep co-clustering. In *SDM*.

Yan, Y.; Chen, Y.; Chen, H.; Xu, M.; Das, M.; Yang, H.; and Tong, H. 2023a. From Trainable Negative Depth to Edge Heterophily in Graphs. In *NeurIPS*.

Yan, Y.; Jing, B.; Liu, L.; Wang, R.; Li, J.; Abdelzaher, T.; and Tong, H. 2023b. Reconciling Competing Sampling Strategies of Network Embedding. In *NrurIPS*.

Yan, Y.; Liu, L.; Ban, Y.; Jing, B.; and Tong, H. 2021. Dynamic knowledge graph alignment. In *AAAI*.

Yan, Y.; Zhang, S.; and Tong, H. 2021. Bright: A bridging algorithm for network alignment. In *TheWebConf*.

Yan, Y.; Zhou, Q.; Li, J.; Abdelzaher, T.; and Tong, H. 2022. Dissecting Cross-Layer Dependency Inference on Multi-Layered Inter-Dependent Networks. In *CIKM*.

Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph convolutional neural networks for web-scale recommender systems. In *KDD*.

You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *NeurIPS*.

Yu, J.; Yin, H.; Xia, X.; Chen, T.; Cui, L.; and Nguyen, Q. V. H. 2022. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *SIGIR*.

Zeng, Z.; Du, B.; Zhang, S.; Xia, Y.; Liu, Z.; and Tong, H. 2023a. Hierarchical Multi-Marginal Optimal Transport for Network Alignment. *arXiv preprint arXiv:2310.04470*.

Zeng, Z.; Zhang, S.; Xia, Y.; and Tong, H. 2023b. PARROT: Position-Aware Regularized Optimal Transport for Network Alignment. In *TheWebConf*.

Zeng, Z.; Zhu, R.; Xia, Y.; Zeng, H.; and Tong, H. 2023c. Generative graph dictionary learning. In *ICML*.

Zhang, M.; and Chen, Y. 2019. Inductive matrix completion based on graph neural networks. *arXiv preprint arXiv:1904.12058*.

Zhang, Y.; Xiong, Y.; Kong, X.; and Zhu, Y. 2017. Learning node embeddings in interaction graphs. In *CIKM*.

Zhang, Z.; Liu, J.; Zhao, K.; Yang, S.; Zheng, X.; and Wang, Y. 2023. Contrastive learning for signed bipartite graphs. In *SIGIR*.

Zheng, L.; Fu, D.; Maciejewski, R.; and He, J. 2021. Deeper-GXX: deepening arbitrary GNNs. *arXiv preprint arXiv:2110.13798*.

Zheng, L.; Xiong, J.; Zhu, Y.; and He, J. 2022. Contrastive Learning with Complex Heterogeneity. In *KDD*.

Zheng, L.; Zhou, D.; Tong, H.; Xu, J.; Zhu, Y.; and He, J. 2023. FairGen: Towards Fair Graph Generation. *arXiv preprint arXiv:2303.17743*.

Zheng, L.; Zhu, Y.; and He, J. 2023. Fairness-aware Multi-view Clustering. In *SDM*.

Zhou, D.; Zhang, S.; Yildirim, M. Y.; Alcorn, S.; Tong, H.; Davulcu, H.; and He, J. 2021. High-order structure exploration on massive graphs: A local graph clustering perspective. *TKDD*.

Zhou, D.; Zheng, L.; Fu, D.; Han, J.; and He, J. 2022. MentorGNN: Deriving Curriculum for Pre-Training GNNs. In *CIKM*.

Zhou, D.; Zheng, L.; Han, J.; and He, J. 2020. A Data-Driven Graph Generative Model for Temporal Interaction Networks. In *KDD*.

Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2021. Graph contrastive learning with adaptive augmentation. In *TheWebConf*.