



# Unsupervised Learning of Joint Embeddings for Node Representation and Community Detection

Rayyan Ahmad Khan<sup>1,3</sup>(✉), Muhammad Umer Anwaar<sup>1,3</sup>, Omran Kaddah<sup>3</sup>, Zhiwei Han<sup>2</sup>, and Martin Kleinsteuber<sup>1,3</sup>

<sup>1</sup> Mercateo AG, Munich, Germany

<sup>2</sup> Fortiss GmbH, Munich, Germany

<sup>3</sup> Technical University of Munich, Munich, Germany  
rayyan.khan@tum.de

**Abstract.** In graph analysis community detection and node representation learning are two highly correlated tasks. In this work, we propose an efficient generative model called **J-ENC** for learning **Joint Embedding for Node representation and Community detection**. **J-ENC** learns a community-aware node representation, i.e., learning of the node embeddings are constrained in such a way that connected nodes are not only “closer” to each other but also share similar community assignments. This joint learning framework leverages community-aware node embeddings for better performance on these tasks: node classification, overlapping community detection and non-overlapping community detection. We demonstrate on several graph datasets that **J-ENC** effectively outperforms many competitive baselines on these tasks. Furthermore, we show that **J-ENC** not only has quite robust performance with varying hyperparameters but also is computationally efficient than its competitors.

## 1 Introduction

Graphs are flexible data structures that model complex relationships among entities, i.e. data points as nodes and the relations between nodes via edges. One important task in graph analysis is community detection, where the objective is to cluster nodes into multiple groups (communities). Each community is a set of densely connected nodes. The communities can be overlapping or non-overlapping, depending on whether they share some nodes or not. Several algorithmic [1, 5] and probabilistic approaches [9, 20, 34, 38] to community detection have been proposed. Another fundamental task in graph analysis is learning

---

R. A. Khan, M. U. Anwaar, O. Kaddah and Z. Han—Equal contribution.

---

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-86520-7\\_2](https://doi.org/10.1007/978-3-030-86520-7_2)) contains supplementary material, which is available to authorized users.

the node embeddings. These embeddings can then be used for downstream tasks like graph visualization [8, 28, 33, 34] and classification [3, 27].

In the literature, these tasks are usually treated separately. Although the standard graph embedding methods capture the basic connectivity, the learning of the node embeddings is independent of community detection. For instance, a simple approach can be to get the node embeddings via DeepWalk [23] and get community assignments for each node by using k-means or Gaussian mixture model. Looking from the other perspective, methods like Bigclam [36], that focus on finding the community structure in the dataset, perform poorly for node-representation tasks e.g. node classification. This motivates us to study the approaches that jointly learn community-aware node embeddings.

Recently several approaches, like CNRL [30], ComE [4], vGraph [26] etc., have been proposed to learn the node embeddings and detect communities simultaneously in a unified framework. Several studies have shown that community detection is improved by incorporating the node representation in the learning process [3, 18]. The intuition is that the global structure of graphs learned during community detection can provide useful context for node embeddings and vice versa.

The joint learning methods (CNRL, ComE and vGraph) learn two embeddings for each node. One node embedding is used for the node representation task. The second node embedding is the “context” embedding of the node which aids in community detection. As CNRL and ComE are based on Skip-Gram [22] and DeepWalk [23], they inherit “context” embedding from it for learning the neighbourhood information of the node. vGraph also requires two node embeddings for parameterizing two different distributions. In contrast, we propose learning a single community-aware node representation which is directly used for both tasks.

In this paper, we propose an efficient generative model called **J-ENC** for jointly learning both community detection and node representation. The underlying intuition behind **J-ENC** is that every node can be a member of one or more communities. However, the node embeddings should be learned in such a way that connected nodes are “closer” to each other than unconnected nodes. Moreover, connected nodes should have similar community assignments. Formally, we assume that for  $i$ -th node, the node embeddings  $\mathbf{z}_i$  are generated from a prior distribution  $p(\mathbf{z})$ . Given  $\mathbf{z}_i$ , the community assignments  $c_i$  are sampled from  $p(c_i|\mathbf{z}_i)$ , which is parameterized by node and community embeddings. In order to generate an edge  $(i, j)$ , we sample another node embedding  $\mathbf{z}_j$  from  $p(\mathbf{z})$  and respective community assignment  $c_j$  from  $p(c_j|\mathbf{z}_j)$ . Afterwards, the node embeddings and the respective community assignments of node pairs are fed to a decoder. The decoder ensures that embeddings of both the nodes and the communities of connected nodes share high similarity. This enables learning such node embeddings that are useful for both community detection and node representation tasks.

We validate the effectiveness of our approach on several real-world graph datasets. In Sect. 4, we show empirically that **J-ENC** is able to outperform the baseline methods including the direct competitors on all three tasks i.e. node classification, overlapping community detection and non-overlapping community detection. Furthermore, we compare the computational cost of training different

algorithms. **J-ENC** is up to 40x more time-efficient than its competitors. We also conduct hyperparameter sensitivity analysis which demonstrates the robustness of our approach. Our main contributions are summarized below:

- We propose an efficient generative model called **J-ENC** for joint community detection and node representation learning.
- We adopt a novel approach and argue that a single node embedding is sufficient for learning both the representation of the node itself and its context.
- Training **J-ENC** is extremely time-efficient in comparison to its competitors.

## 2 Related Work

### 2.1 Community Detection

Early community detection algorithms are inspired from clustering algorithms [35]. For instance, spectral clustering [29] is applied to the graph Laplacian matrix for extracting the communities. Similarly, several matrix factorization based methods have been proposed to tackle the community detection problem. For example, Bigclam [36] treats the problem as a non-negative matrix factorization (NMF) task. Another method CESNA [38] extends Bigclam by modelling the interaction between the network structure and the node attributes. Some generative models, like vGraph [26], Circles [20] etc., have also been proposed to detect communities in a graph.

### 2.2 Node Representation Learning

Many successful algorithms which learn node representation in an unsupervised way are based on random walk objectives [10, 11, 23]. Some known issues with random-walk based methods (e.g. DeepWalk, node2vec etc.) are: (1) They sacrifice the structural information of the graph by putting over-emphasis on the proximity information [24] and (2) great dependence of the performance on hyperparameters (walk-length, number of hops etc.) [10, 23]. Some interesting GCN based approaches include graph autoencoders e.g. GAE and VGAE [17] and DGI [32].

### 2.3 Joint Community Detection and Node Representation Learning

In the literature, several attempts have been made to tackle both these tasks in a single framework. Most of these methods propose an alternate optimization process, i.e. learn node embeddings and improve community assignments with them and vice versa [4, 30]. Some approaches (CNRL [30], ComE [4]) are inspired from random walk, thus they inherit the issues discussed above. Others, like GEMSEC [25], are limited to the detection of non-overlapping communities. Some generative models like CommunityGAN [13] and vGraph [26] also jointly learn community assignments and node embeddings. CNRL, ComE and

vGraph require learning two embeddings for each node for simultaneously tackling the two tasks. Unlike them, **J-ENC** learns a single community-aware node representation which is directly used for both tasks.

It is pertinent to highlight that although both vGraph and **J-ENC** adopt a variational approach but the underlying models are quite different. vGraph assumes that each node can be represented as a mixture of multiple communities and is described by a multinomial distribution over communities, whereas **J-ENC** models the node embedding by a single distribution. For a given node, vGraph, first draws a community assignment and then a connected neighbor node is generated based on the assignment. Whereas, **J-ENC** draws the node embedding from prior distribution and then community assignment is conditioned on a single node only. In simple terms, vGraph also needs edge information in the generative process whereas **J-ENC** does not require it. **J-ENC** relies on the decoder to ensure that embeddings of the connected nodes and their communities share high similarity with each other.

### 3 Methodology

#### 3.1 Problem Formulation

Suppose an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with the adjacency matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$  and a matrix  $\mathbf{X} \in \mathbb{R}^{N \times F}$  of  $F$ -dimensional node features,  $N$  being the number of nodes. Given  $K$  as the number of communities, we aim to jointly learn the node embeddings and the community embeddings following a variational approach such that:

- One or more communities can be assigned to every node.
- The node embeddings can be used for both community detection and node classification.

#### 3.2 Variational Model

**Generative Model:** Let us denote the latent node embedding and community assignment for  $i$ -th node by the random variables  $\mathbf{z}_i \in \mathbb{R}^d$  and  $c_i$  respectively. The generative model is given by:

$$p(\mathbf{A}) = \int \sum_{\mathbf{c}} p(\mathbf{Z}, \mathbf{c}, \mathbf{A}) d\mathbf{Z}, \quad (1)$$

where  $\mathbf{c} = [c_1, c_2, \dots, c_N]$  and the matrix  $\mathbf{Z} = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N]$  stacks the node embeddings. The joint distribution in (1) is mathematically expressed as

$$p(\mathbf{Z}, \mathbf{c}, \mathbf{A}) = p(\mathbf{Z}) p_{\theta}(\mathbf{c}|\mathbf{Z}) p_{\theta}(\mathbf{A}|\mathbf{c}, \mathbf{Z}), \quad (2)$$

where  $\theta$  denotes the model parameters. Let us denote elements of  $\mathbf{A}$  by  $a_{ij}$ . Following existing approaches [14, 17], we consider  $\mathbf{z}_i$  to be *i.i.d* random variables.

Furthermore, assuming  $c_i|\mathbf{z}_i$  to be *i.i.d* random variables, the joint distributions in (2) can be factorized as

$$p(\mathbf{Z}) = \prod_{i=1}^N p(\mathbf{z}_i) \quad (3)$$

$$p_\theta(\mathbf{c}|\mathbf{Z}) = \prod_{i=1}^N p_\theta(c_i|\mathbf{z}_i) \quad (4)$$

$$p_\theta(\mathbf{A}|\mathbf{c}, \mathbf{Z}) = \prod_{i,j} p_\theta(a_{ij}|c_i, c_j, \mathbf{z}_i, \mathbf{z}_j), \quad (5)$$

where Eq. (5) assumes that the *edge decoder*  $p_\theta(a_{ij}|c_i, c_j, \mathbf{z}_i, \mathbf{z}_j)$  depends only on  $c_i, c_j, \mathbf{z}_i$  and  $\mathbf{z}_j$ .

**Inference Model:** We aim to learn the model parameters  $\theta$  such that  $\log(p_\theta(\mathbf{A}))$  is maximized. In order to ensure computational tractability, we introduce the approximate posterior

$$q_\phi(\mathbf{Z}, \mathbf{c}|\mathcal{I}) = \prod_i q_\phi(\mathbf{z}_i, c_i|\mathcal{I}) \quad (6)$$

$$= \prod_i q_\phi(\mathbf{z}_i|\mathcal{I})q_\phi(c_i|\mathbf{z}_i, \mathcal{I}), \quad (7)$$

where  $\mathcal{I} = (\mathbf{A}, \mathbf{X})$  if node features are available, otherwise  $\mathcal{I} = \mathbf{A}$ . We maximize the corresponding ELBO bound (for derivation, refer to the supplementary material), given by

$$\begin{aligned} \mathcal{L}_{\text{ELBO}} \approx & - \sum_{i=1}^N D_{KL}\left(q_\phi(\mathbf{z}_i|\mathcal{I}) \parallel p(\mathbf{z}_i)\right) \\ & - \sum_{i=1}^N \frac{1}{M} \sum_{m=1}^M D_{KL}\left(q_\phi(c_i|\mathbf{z}_i^{(m)}, \mathcal{I}) \parallel p_\theta(c_i|\mathbf{z}_i^{(m)})\right) \\ & + \sum_{(i,j) \in \mathcal{E}} \mathbb{E}_{(\mathbf{z}_i, \mathbf{z}_j, c_i, c_j) \sim q_\phi(\mathbf{z}_i, \mathbf{z}_j, c_i, c_j|\mathcal{I})} \left\{ \log\left(p_\theta(a_{ij}|c_i, c_j, \mathbf{z}_i, \mathbf{z}_j)\right) \right\}, \quad (8) \end{aligned}$$

where  $D_{KL}(\cdot|\cdot)$  represents the KL-divergence between two distributions. The distribution  $q_\phi(\mathbf{z}_i, \mathbf{z}_j, c_i, c_j|\mathcal{I})$  in the third term of Eq. (8) is factorized into two conditionally independent distributions i.e.

$$q_\phi(\mathbf{z}_i, \mathbf{z}_j, c_i, c_j|\mathcal{I}) = q_\phi(\mathbf{z}_i, c_i|\mathcal{I})q_\phi(\mathbf{z}_j, c_j|\mathcal{I}). \quad (9)$$

### 3.3 Design Choices

In Eq. (3),  $p(\mathbf{z}_i)$  is chosen to be the standard gaussian distribution for all  $i$ . The corresponding approximate posterior  $q_\phi(\mathbf{z}_i|\mathcal{I})$  in Eq. (7), used as node embeddings encoder, is given by

$$q_\phi(\mathbf{z}_i|\mathcal{I}) = \mathcal{N}(\boldsymbol{\mu}_i(\mathcal{I}), \text{diag}(\boldsymbol{\sigma}^2_i(\mathcal{I}))). \quad (10)$$

The parameters of  $q_\phi(\mathbf{z}_i|\mathcal{I})$  can be learnt by any encoder network e.g. graph convolutional network [16], graph attention network [31], GraphSAGE [11] or even two matrices to learn  $\boldsymbol{\mu}_i(\mathcal{I})$  and  $\text{diag}(\boldsymbol{\sigma}^2_i(\mathcal{I}))$ . Samples are then generated using reparametrization trick [6].

For parameterizing  $p_\theta(c_i|\mathbf{z}_i)$  in Eq. (4), we introduce community embeddings  $\{\mathbf{g}_1, \dots, \mathbf{g}_K\}$ ;  $\mathbf{g}_k \in \mathbb{R}^d$ . The distribution  $p_\theta(c_i|\mathbf{z}_i)$  is then modelled as the softmax of dot products of  $\mathbf{z}_i$  with  $\mathbf{g}_k$ , i.e.

$$p_\theta(c_i = k|\mathbf{z}_i) = \frac{\exp(\langle \mathbf{z}_i, \mathbf{g}_k \rangle)}{\sum_{\ell=1}^K \exp(\langle \mathbf{z}_i, \mathbf{g}_\ell \rangle)}. \quad (11)$$

The corresponding approximate posterior  $q_\phi(c_i = k|\mathbf{z}_i, \mathcal{I})$  in Eq. (7) is affected by the node embedding  $\mathbf{z}_i$  as well as the neighborhood. To design this, our intuition is to consider the similarity of  $\mathbf{g}_k$  with the embedding  $\mathbf{z}_i$  as well as with the embeddings of the neighbors of the  $i$ -th node. The overall similarity with neighbors is mathematically formulated as the average of the dot products of their embeddings. Afterwards, a hyperparameter  $\alpha$  is introduced to control the bias between the effect of  $\mathbf{z}_i$  and the set  $\mathcal{N}_i$  of the neighbors of the  $i$ -th node. Finally, a softmax is applied, i.e.

$$q_\phi(c_i = k|\mathbf{z}_i, \mathcal{I}) = \text{softmax}\left(\alpha \langle \mathbf{z}_i, \mathbf{g}_k \rangle + (1 - \alpha) \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \langle \mathbf{z}_j, \mathbf{g}_k \rangle\right). \quad (12)$$

Hence, Eq. (12) ensures that graph structure information is employed to learn community assignments instead of relying on an extraneous node embedding as done in [4, 26]. Finally, the choice of edge decoder in Eq. (5) is motivated by the intuition that the nodes connected by edges have a high probability of belonging to the same community and vice versa. Therefore we model the edge decoder as:

$$p_\theta(a_{ij}|c_i = \ell, c_j = m, \mathbf{z}_i, \mathbf{z}_j) = \frac{\sigma(\langle \mathbf{z}_i, \mathbf{g}_m \rangle) + \sigma(\langle \mathbf{z}_j, \mathbf{g}_\ell \rangle)}{2}. \quad (13)$$

For better reconstructing the edges, Eq. (13) makes use of the community embeddings, node embeddings and community assignment information simultaneously. This helps in learning better node representations by leveraging the global information about the graph structure via community detection. On the other hand, this also forces the community assignment information to exploit the local graph structure via node embeddings and edge information.

### 3.4 Practical Aspects

The third term in Eq. (8) is estimated in practice using the samples generated by the approximate posterior. This term is equivalent to the negative of binary cross-entropy (BCE) loss between observed edges and reconstructed edges. Since community assignment follows a categorical distribution, we use Gumbel-softmax [12] for backpropagation of the gradients. As for the second term of Eq. (8), it is also enough to set  $M = 1$ , i.e. use only one sample per input node.

For inference, non-overlapping community assignment can be obtained for  $i$ -th node as

$$C_i = \arg \max_{k \in \{1, \dots, K\}} q_\phi(c_i = k | \mathbf{z}_i, \mathcal{I}). \quad (14)$$

To get overlapping community assignments for  $i$ -th node, we can threshold its weighted probability vector at  $\epsilon$ , a hyperparameter, as follows

$$C_i = \left\{ k \mid \frac{q_\phi(c_i = k | \mathbf{z}_i, \mathcal{I})}{\max_\ell q_\phi(c_i = \ell | \mathbf{z}_i, \mathcal{I})} \geq \epsilon \right\}, \quad \epsilon \in [0, 1]. \quad (15)$$

### 3.5 Complexity

Computation of dot products for all combinations of node and community embeddings takes  $\mathcal{O}(NKd)$  time. Solving Eq. (12) further requires calculation of mean of dot products over the neighborhood for every node, which takes  $\mathcal{O}(|\mathcal{E}|K)$  computations overall as we traverse every edge for every community. Finally, we need softmax over all communities for every node in Eq. (11) and Eq. (12) which takes  $\mathcal{O}(NK)$  time. Equation (13) takes  $\mathcal{O}(|\mathcal{E}|)$  time for all edges as we have already calculated the dot products. As a result, the overall complexity becomes  $\mathcal{O}(|\mathcal{E}|K + NKd)$ . This complexity is quite low compared to other algorithms designed to achieve similar goals [4, 39].


## 4 Experiments

### 4.1 Synthetic Example

We start with a synthetic dataset, consisting of 3 communities with 5 points per community. This dataset is actually a random partition graph generated by the python package `networkx`. The encoder simply consists of two matrices that give  $\boldsymbol{\mu}_i(\mathcal{I})$  and  $\text{diag}(\boldsymbol{\sigma}_i^2(\mathcal{I}))$ . The results of the community assignments discovered by **J-ENC** are given in Fig. 1, where the node sizes are reciprocal to the confidence of **J-ENC** in the community assignments. We choose 3 communities for demonstration because the probabilistic community assignments in such case can be thought of as `rgb` values for coloring the nodes. It can be seen that **J-ENC** discovers the correct community structure. However, the two bigger nodes in the center can be assigned to more than one communities as **J-ENC** is not very confident in case of these nodes. This is evident from the colors that are a mix of red, green and blue. We now proceed to the experiments on real-world datasets.

## 4.2 Datasets

We have selected 18 different datasets ranging from 270 to 126,842 edges. For non-overlapping community detection and node classification, we use 5 the citation datasets [2, 40]. The remaining datasets [20, 37], used for overlapping community detection, are taken from SNAP repository [19]. Following [26], we take 5 biggest ground truth communities for youtube, amazon and dblp. Moreover, we also analyse the case of large number of communities. For this purpose, we prepare two subsets of amazon dataset by randomly selecting 500 and 1000 communities from 2000 smallest communities in the amazon dataset (Table 1).



**Fig. 1.** Visualization of community assignments discovered by **J-ENC** in the synthetic dataset of 15 points divided in three communities.

**Table 1.** Every dataset has  $|\mathcal{V}|$  nodes,  $|\mathcal{E}|$  edges,  $K$  communities and  $|F|$  features.  $|F| = \text{N/A}$  means that either the features were missing or not used.

Dataset	$ \mathcal{V} $	$ \mathcal{E} $	$K$	$ F $	Overlap
CiteSeer	3327	9104	6	3703	N
CiteSeer-full	4230	10674	6	602	N
Cora	2708	10556	7	1433	N
Cora-ML	2995	16316	7	2879	N
Cora-full	19793	126842	70	8710	N
fb0	333	2519	24	N/A	Y
fb107	1034	26749	9	N/A	Y
fb1684	786	14024	17	N/A	Y
fb1912	747	30025	46	N/A	Y
fb3437	534	4813	32	N/A	Y
fb348	224	3192	14	N/A	Y
fb414	150	1693	7	N/A	Y
fb698	61	270	13	N/A	Y
Youtube	5346	24121	5	N/A	Y
Amazon	794	2109	5	N/A	Y
Amazon500	1113	3496	500	N/A	Y
Amazon1000	1540	4488	1000	N/A	Y
Dblp	24493	89063	5	N/A	Y

## 4.3 Baselines

For overlapping community detection, we compare with the following competitive baselines: **MNMF** [34] learns community membership distribution by using joint non-negative matrix factorization with modularity based regularization. **BIGCLAM** [36] also formulates community detection as a non-negative matrix factorization (NMF) task. It simultaneously optimizes the model likelihood of observed links and learns the latent factors which represent community affiliations of nodes. **CESNA** [38] extends BIGCLAM by statistically modelling the interaction between the network structure and the node attributes. **Circles** [20] introduces a generative model for community detection in ego-networks by learning node similarity metrics for every community. **SVI** [9] formulates membership of nodes in multiple communities by a Bayesian model of networks.



**vGraph** [26] simultaneously learns node embeddings and community assignments by modelling the nodes as being generated from a mixture of communities. **vGraph+**, a variant further incorporates regularization to weigh local connectivity. **Come** [4] jointly learns community and node embeddings by using gaussian mixture model formulation. **CNRL** [30] enhances the random walk sequences (generated by DeepWalk, node2vec etc.) to jointly learn community and node embeddings. **CommunityGAN** (ComGAN) is a generative adversarial model for learning node embeddings such that the entries of the embedding vector of each node refer to the membership strength of the node to different communities. Lastly, we compare the results with the communities obtained by applying k-means to the learned embeddings of **DGI** [32].

For non-overlapping community detection and node classification, in addition to MNMF, DGI, CNRL, CommunityGAN, vGraph and ComE, we compare **J-ENC** with the following baselines: **DeepWalk** [23] makes use of SkipGram [22] and truncated random walks on network to learn node embeddings. **LINE** [27] learns node embeddings while attempting to preserve first and second order proximities of nodes. **Node2Vec** [10] learns the embeddings using biased random walk while aiming to preserve network neighborhoods of nodes. **Graph Autoencoder (GAE)** [17] extends the idea of autoencoders to graph datasets. We also include its variational counterpart i.e. **VGAE**. **GEMSEC** is a sequence sampling-based learning model which aims to jointly learn the node embeddings and clustering assignments.

#### 4.4 Settings

**For overlapping community detection**, we learn mean and log-variance matrices of 16-dimensional node embeddings. We set  $\alpha = 0.9$  and  $\epsilon = 0.3$  in all our experiments. Following [17], we first pre-train a variational graph autoencoder. We perform gradient descent with Adam optimizer [15] and learning rate = 0.01. Community assignments are obtained using Eq. (15). For the baselines, we employ the results reported by [26]. For evaluating the performance, we use *F1-score* and *Jaccard similarity*.

**For non-overlapping community detection**, since the default implementations of most the baselines use 128 dimensional embeddings, for we use  $d = 128$  for fair comparison. Equation (14) is used for community assignments. For vGraph, we use the code provided by the authors. We employ *normalized mutual information (NMI)* and *adjusted random index (ARI)* as evaluation metrics.

**For node classification**, we follow the training split used in various previous works [16, 32, 40], i.e. 20 nodes per class for training. We train logistic regression using LIBLINEAR [7] solver as our classifier and report the evaluation results on rest of the nodes. For the algorithms that do not use node features, we train the classifier by appending the raw node features with the learnt embeddings. For evaluation, we use *F1-macro* and *F1-micro* scores.

All the reported results are the average over five runs. Further implementation details can be found in: [https://github.com/RayyanRiaz/gnn\\_comm\\_det](https://github.com/RayyanRiaz/gnn_comm_det).

## 4.5 Discussion of Results

Tables 2 and 3 summarize the results of the performance comparison for the overlapping community detection tasks.

**Table 2.** F1 scores (%) for overlapping communities. Best and second best values are bold and underlined respectively.

Dataset	MNMF	Bigclam	CESNA	Circles	SVI	vGraph	vGraph+	ComE	CNRL	ComGan	DGI	J-ENC
fb0	14.4	29.5	<b>28.1</b>	28.6	28.1	24.4	26.1	31.1	11.5	<b>35.0</b>	27.4	<u>34.7</u>
fb107	12.6	39.3	37.3	24.7	26.9	28.2	31.8	39.7	20.2	<u>47.5</u>	35.8	<b>59.7</b>
fb1684	12.2	50.4	51.2	28.9	35.9	42.3	43.8	<u>52.9</u>	38.5	47.6	42.8	<b>56.4</b>
fb1912	14.9	34.9	34.7	26.2	28.0	25.8	<u>37.5</u>	28.7	8.0	35.6	32.6	<b>45.8</b>
fb3437	13.7	19.9	20.1	10.1	15.4	20.9	22.7	21.3	3.9	<u>39.3</u>	19.7	<b>50.2</b>
fb348	20.0	49.6	53.8	51.8	46.1	55.4	53.1	46.2	34.1	<u>55.8</u>	54.7	<b>58.2</b>
fb414	22.1	58.9	60.1	48.4	38.9	64.7	<u>66.9</u>	55.3	25.3	43.9	56.9	<b>69.6</b>
fb698	26.6	54.2	58.7	35.2	40.3	54.0	<u>59.5</u>	45.8	16.4	58.2	52.2	<b>64.0</b>
Youtube	59.9	43.7	38.4	36.0	41.4	50.7	52.2	<u>65.5</u>	51.4	43.6	47.8	<b>67.3</b>
Amazon	38.2	46.4	46.8	53.3	47.3	53.3	53.2	50.1	<u>53.5</u>	51.4	44.7	<b>58.1</b>
Amazon500	30.1	52.2	57.3	46.2	41.9	<u>61.2</u>	60.4	59.8	38.4	59.3	33.8	<b>67.6</b>
Amazon1000	19.3	28.6	30.8	25.9	21.6	<u>54.3</u>	47.3	50.3	27.1	52.7	37.7	<b>60.5</b>
Dblp	21.8	23.6	35.9	36.2	33.7	39.3	39.9	<u>47.1</u>	46.8	34.9	44.0	<b>53.9</b>

**Table 3.** Jaccard scores (%) for overlapping communities. Best and second best values are bold and underlined respectively.

Dataset	MNMF	Bigclam	CESNA	Circles	SVI	vGraph	vGraph+	ComE	CNRL	ComGan	DGI	J-ENC
fb0	08.0	18.5	17.3	18.6	17.6	14.6	15.9	19.5	06.8	<u>24.1</u>	16.8	<b>24.7</b>
fb107	06.9	27.5	27.0	15.5	17.2	18.3	21.7	28.7	11.9	<u>38.5</u>	25.3	<b>46.8</b>
fb1684	06.6	38.0	38.7	18.7	24.7	29.2	32.7	<u>40.3</u>	25.8	37.9	38.8	<b>42.5</b>
fb1912	08.4	24.1	23.9	16.7	20.1	18.6	<u>28.0</u>	18.5	04.6	13.5	22.5	<b>37.3</b>
fb3437	07.7	11.5	11.7	05.5	09.0	12.0	13.3	12.5	02.0	<u>33.4</u>	11.6	<b>36.2</b>
fb348	11.3	35.9	40.0	39.3	33.6	<u>41.0</u>	40.5	34.4	21.7	23.2	<u>41.8</u>	<b>43.5</b>
fb414	12.8	47.1	47.3	34.2	29.3	51.8	<u>55.9</u>	42.2	15.4	53.6	46.4	<b>58.4</b>
fb698	16.0	41.9	45.9	22.6	30.0	43.6	<u>47.7</u>	33.8	09.6	46.9	42.1	<b>50.4</b>
Youtube	46.7	29.3	24.2	22.1	28.7	34.3	34.8	<u>52.5</u>	35.5	44.0	32.7	<b>53.3</b>
Amazon	25.2	35.1	35.0	36.7	36.4	36.9	36.9	34.6	<u>38.7</u>	38.0	29.1	<b>41.9</b>
Amazon500	20.8	51.2	53.8	47.2	45.0	59.1	<u>59.6</u>	58.4	41.1	57.3	23.3	<b>64.9</b>
Amazon1000	20.3	26.8	28.9	24.9	23.6	<u>54.3</u>	49.7	52.0	26.9	54.1	23.2	<b>57.1</b>
Dblp	20.9	13.8	22.3	23.3	20.9	25.0	25.1	27.9	<u>32.8</u>	25.0	29.2	<b>37.3</b>

**Table 4.** Non-overlapping community detection results. Best and second best values are bold and underlined respectively.

Alg.	NMI(%)					ARI(%)				
	CiteSeer	CiteSeer-full	Cora	Cora-ML	Cora-full	CiteSeer	CiteSeer-full	Cora	Cora-ML	Cora-full
MNMF	14.1	09.4	19.7	37.8	42.0	02.6	00.4	02.9	24.1	06.1
DeepWalk	08.8	15.4	39.7	43.2	48.5	09.5	16.4	31.2	33.9	22.5
LINE	08.7	13.0	32.8	42.3	40.3	03.3	03.7	14.9	32.7	11.7
Node2Vec	14.9	22.3	39.7	39.6	48.1	08.1	10.5	25.8	27.9	18.8
GAE	17.4	55.1	39.7	<u>48.3</u>	48.3	14.1	50.6	29.3	41.8	18.3
VGAE	16.3	48.4	40.8	<u>48.3</u>	47.0	10.1	40.6	34.7	<u>42.5</u>	17.9
DGI	<u>37.8</u>	<u>56.7</u>	<u>50.1</u>	46.2	39.9	<b>38.1</b>	<u>50.8</u>	<u>44.7</u>	42.1	12.1
GEMSEC	11.8	11.1	27.4	18.1	10.0	00.6	01.0	04.8	01.0	00.2
CNRL	13.6	23.3	39.4	42.9	47.7	12.8	20.2	31.9	32.5	<u>22.9</u>
ComGAN	03.2	16.2	05.7	11.5	15.0	01.2	04.9	03.2	06.7	00.6
vGraph	09.0	07.6	26.4	29.8	41.7	05.1	04.2	12.7	21.6	14.9
ComE	18.8	32.8	39.6	47.6	<u>51.2</u>	13.8	20.9	34.2	37.2	19.7
<b>J-ENC</b>	<b>38.5</b>	<b>59.0</b>	<b>52.7</b>	<b>56.3</b>	<b>55.2</b>	<u>35.2</u>	<b>60.3</b>	<b>45.1</b>	<b>49.8</b>	<b>28.8</b>

**Table 5.** Node classification results. Best and second best values are bold and underlined respectively.

Alg.	F1-macro (%)					F1-micro (%)				
	CiteSeer	CiteSeer-full	Cora	Cora-ML	Cora-full	CiteSeer	CiteSeer-full	Cora	Cora-ML	Cora-full
MNMF	57.4	68.6	60.9	64.2	30.4	60.8	68.1	62.7	64.2	32.9
DeepWalk	49.0	56.6	69.7	75.8	41.7	52.0	57.3	70.2	75.6	<u>48.3</u>
LINE	55.0	60.2	68.0	75.3	39.4	57.7	60.0	68.3	74.6	42.1
Node2Vec	55.2	61.0	71.3	78.4	<u>42.3</u>	57.8	61.5	71.4	78.6	48.1
GAE	57.9	<u>79.9</u>	71.2	76.5	36.6	61.6	<u>79.6</u>	73.5	77.6	41.8
VGAE	59.1	74.4	70.4	75.2	32.4	62.2	74.4	72.0	76.4	37.7
DGI	<u>62.6</u>	<b>82.1</b>	71.1	72.6	16.5	<u>67.9</u>	<b>81.8</b>	73.3	75.4	21.1
GEMSEC	37.5	53.3	60.3	70.6	35.8	39.4	53.5	59.4	72.5	38.9
CNRL	50.0	58.0	70.4	77.8	41.3	53.2	57.9	70.4	78.4	45.9
ComGAN	55.9	65.7	56.6	62.5	27.7	59.1	64.9	58.5	62.8	29.4
vGraph	30.8	28.5	44.7	59.8	33.4	32.1	28.5	44.6	62.3	37.6
ComE	59.6	69.9	<u>71.6</u>	<u>78.5</u>	42.2	63.1	70.2	<u>74.2</u>	<u>79.5</u>	47.8
<b>J-ENC</b>	<b>64.8</b>	76.8	<b>73.1</b>	<b>80.2</b>	<b>43.1</b>	<b>68.2</b>	77.0	<b>75.6</b>	<b>82.0</b>	<b>49.6</b>

First, we note that our proposed method **J-ENC** outperforms the competitors on all datasets in terms of Jaccard score as well as F1-score, with the dataset (*fb0*) being the only exception where **J-ENC** is the second best. These results demonstrate the capability of **J-ENC** to learn multiple community assignments quite well and hence reinforces our intuition behind the design of Eq. (12).

Second, we observe that there is no consistent performing algorithm among the competitive methods. That is, excluding **J-ENC**, the best performance is achieved by vGraph/vGraph+ on 5, ComGAN on 4 and ComE on 3 out of 13 datasets in terms of F1-score. A similar trend can be seen in Jaccard Similarity. It is worth noting that all the methods, which achieve the second-best perfor-

mance, are solving the task of community detection and node representation learning jointly.

Third, we observe that vGraph+ results are generally better than vGraph. This is because vGraph+ incorporates a regularization term in the loss function which is based on Jaccard coefficients of connected nodes as edge weights. However, it should be noted that this preprocessing step is computationally expensive for densely connected graphs.

Table 4 shows the results on non-overlapping community detection. First, we observe that MNMF, DeepWalk, LINE and Node2Vec provide a good baseline for the task. However, these methods are not able to achieve comparable performance on any dataset relative to the frameworks that treat the two tasks jointly. Second, **J-ENC** consistently outperforms all the competitors in NMI and ARI metrics, except for *CiteSeer* where it achieves second best ARI. Third, we observe that GCN based models i.e. GAE, VGAE and DGI show competitive performance. That is, they achieve second best performance in all the datasets except *CiteSeer*. In particular, DGI achieves second best NMI results in 3 out of 5 datasets and 2 out of 5 datasets in terms of ARI. Nonetheless, DGI results are not very competitive in Table 2 and Table 3, showing that while DGI can be a good choice for learning node embeddings for attributed graphs with non-overlapping communities, it is not the best option for non-attributed graphs or overlapping communities.

The results for node classification are presented in Table 5. **J-ENC** achieves best F1-micro and F1-macro scores on 4 out of 5 datasets. We also observe that GCN based models i.e. GAE, VGAE and DGI show competitive performance, following the trend in results of Table 4. Furthermore, we note that the node classification results of CommunityGan (ComGAN) are quite poor. We think a potential reason behind it is that the node embeddings are constrained to have same dimensions as the number of communities. Hence, different components of the learned node embeddings simply represent the membership strengths of nodes for different communities. The linear classifiers may find it difficult to separate such vectors.

## 4.6 Hyperparameter Sensitivity


We study the dependence of **J-ENC** on  $\epsilon$  and  $\alpha$  by evaluating on four datasets of different sizes: *fb698* ( $N = 61$ ), *fb1912* ( $N = 747$ ), *amazon1000* ( $N=1540$ ) and *youtube* ( $N = 5346$ ).

**Effect of  $\epsilon$ :** We sweep for  $\epsilon = \{0.1, 0.2, \dots, 0.9\}$ . For demonstrating effect of  $\alpha$ , we fix  $\epsilon = 0.3$  and sweep for  $\alpha = \{0.1, 0.2, \dots, 0.9\}$ . The average results of five runs for  $\epsilon$  and  $\alpha$  are given in Fig. 2a and Fig. 2b respectively. Overall **J-ENC** is quite robust to the change in the values of  $\epsilon$  and  $\alpha$ . In case of  $\epsilon$ , we see a general trend of decrease in performance when the threshold  $\epsilon$  is set quite high e.g.  $\epsilon > 0.7$ . This is because the datasets contain overlapping communities and a very high  $\epsilon$  will cause the algorithm to give only the most probable community assignment instead of potentially providing multiple communities per node. However, for a large part of sweep space, the results are almost consistent.

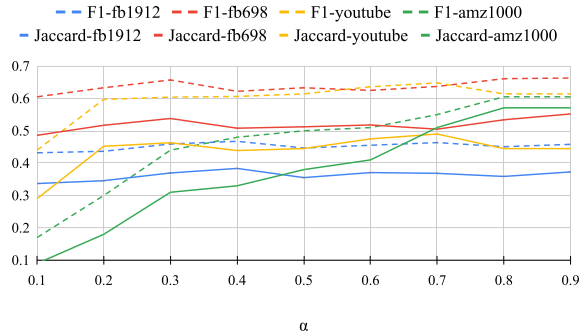
**Effect of  $\alpha$ :** When  $\epsilon$  is fixed and  $\alpha$  is changed, the results are mostly consistent except when  $\alpha$  is set to a low value. Equation (12) shows that in such a case the node itself is almost neglected and **J-ENC** tends to assign communities based upon neighborhood only, which may cause a decrease in the performance. This effect is most visible in *amazon1000* dataset because it has only 1.54 points on average per community. This implies a decent chance for neighbours of a point of being in different communities. Thus, sole dependence on the neighbors will most likely result in poor results.

#### 4.7 Training Time

Now we compare the training times of different algorithms in Fig. 3. As some of the baselines are more resource intensive than others, we select aws instance type `g4dn.4xlarge` for fair comparison of training times. For vGraph, we train for 1000 iterations and for **J-ENC** for 1500 iterations. For all other algorithms




(a) Effect of  $\epsilon$ . Overall a slight decrease in scores can be observed after  $\epsilon = 0.7$  mark.



(b) Effect of  $\alpha$ . The scores generally tend to decrease for small values of  $\alpha$ .

**Fig. 2.** Effect of hyperparameters on the performance. F1 and Jaccard scores are in solid and dashed lines respectively.



**Fig. 3.** Comparison of running times of different algorithms. We can see that **J-ENC** outperforms the direct competitors. The time on y-axis is in log scale.

we use the default parameters as used in Sect. 4.4. We observe that the methods that simply output the node embeddings take relatively less time compared to the algorithms that jointly learn node representations and community assignments e.g. **J-ENC**, vGraph and CNRL. Among these algorithms **J-ENC** is the most time efficient. It consistently trains in less time compared to its direct competitors. For instance, it is about 12 times faster than ComE for *CiteSeer-full* and about 40 times faster compared to vGraph for *Cora-full* dataset. This provides evidence for lower computational complexity of **J-ENC** in Sect. 3.5.

#### 4.8 Visualization

Our experiments demonstrate that a single community-aware node embedding is sufficient to aid in both the node representation and community assignment tasks. This is also qualitatively demonstrated by graph visualizations of node embeddings (obtained via t-SNE [21]) and inferred communities for two datasets, fb107 and fb3437, presented in Fig. 4.



**Fig. 4.** Graph visualization with community assignments (better viewed in color)

## 5 Conclusion

We propose a scalable generative method **J-ENC** to simultaneously perform community detection and node representation learning. Our novel approach learns a single community-aware node embedding for both the representation of the node and its context. **J-ENC** is scalable due to its low complexity, i.e.  $\mathcal{O}(|\mathcal{E}|K + NKd)$ . The experiments on several graph datasets show that **J-ENC** consistently outperforms all the competitive baselines on node classification, overlapping community detection and non-overlapping community detection tasks. Moreover, training the **J-ENC** is highly time-efficient than its competitors.

## References

1. Ahn, Y.Y., Bagrow, J.P., Lehmann, S.: Link communities reveal multiscale complexity in networks. *Nature* **466**(7307), 761–764 (2010)
2. Bojchevski, A., Günnemann, S.: Deep Gaussian embedding of graphs: unsupervised inductive learning via ranking. arXiv preprint [arXiv:1707.03815](https://arxiv.org/abs/1707.03815) (2017)
3. Cao, S., Lu, W., Xu, Q.: GraRep: learning graph representations with global structural information. In: Proceedings of the 24th ACM International on Conference on Information and Knowledge Management (CIKM 2015), pp. 891–900. Association for Computing Machinery, New York (2015). <https://doi.org/10.1145/2806416.2806512>
4. Cavallari, S., Zheng, V.W., Cai, H., Chang, K.C.C., Cambria, E.: Learning community embedding with community detection and node embedding on graphs. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, pp. 377–386 (2017)
5. Derényi, I., Palla, G., Vicsek, T.: Clique percolation in random networks. *Phys. Rev. Lett.* **94**(16), 160202 (2005)
6. Doersch, C.: Tutorial on variational autoencoders. arXiv preprint [arXiv:1606.05908](https://arxiv.org/abs/1606.05908) (2016)
7. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: LIBLINEAR: a library for large linear classification. *J. Mach. Learn. Res.* **9**(Aug), 1871–1874 (2008)
8. Gao, S., Denoyer, L., Gallinari, P.: Temporal link prediction by integrating content and structure information. In: Proceedings of the 20th ACM International Conference on Information and Knowledge Management, pp. 1169–1174 (2011)
9. Gopalan, P.K., Blei, D.M.: Efficient discovery of overlapping communities in massive networks. *Proc. Natl. Acad. Sci.* **110**(36), 14534–14539 (2013)
10. Grover, A., Leskovec, J.: node2vec: scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 855–864 (2016)
11. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: Advances in Neural Information Processing Systems, pp. 1024–1034 (2017)
12. Jang, E., Gu, S., Poole, B.: Categorical reparameterization with Gumbel-Softmax. arXiv preprint [arXiv:1611.01144](https://arxiv.org/abs/1611.01144) (2016)
13. Jia, Y., Zhang, Q., Zhang, W., Wang, X.: CommunityGAN: community detection with generative adversarial nets. In: The World Wide Web Conference, pp. 784–794 (2019)

14. Khan, R.A., Anwaar, M.U., Kleinsteuber, M.: Epitomic variational graph autoencoder (2020)
15. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
16. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
17. Kipf, T.N., Welling, M.: Variational graph auto-encoders. arXiv preprint [arXiv:1611.07308](https://arxiv.org/abs/1611.07308) (2016)
18. Kozdoba, M., Mannor, S.: Community detection via measure space embedding. In: Proceedings of the 28th International Conference on Neural Information Processing Systems (NIPS 2015), vol. 2, pp. 2890–2898. MIT Press, Cambridge (2015)
19. Leskovec, J., Krevl, A.: SNAP datasets: stanford large network dataset collection, June 2014. <http://snap.stanford.edu/data>
20. Leskovec, J., McAuley, J.J.: Learning to discover social circles in ego networks. In: Advances in Neural Information Processing Systems, pp. 539–547 (2012)
21. Maaten, L.V.D., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(Nov), 2579–2605 (2008)
22. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint [arXiv:1301.3781](https://arxiv.org/abs/1301.3781) (2013)
23. Perozzi, B., Al-Rfou, R., Skiena, S.: DeepWalk: online learning of social representations. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 701–710 (2014)
24. Ribeiro, L.F., Saverese, P.H., Figueiredo, D.R.: struc2vec: learning node representations from structural identity. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 385–394 (2017)
25. Rozemberczki, B., Davies, R., Sarkar, R., Sutton, C.: GEMSEC: graph embedding with self clustering. In: Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, pp. 65–72 (2019)
26. Sun, F.Y., Qu, M., Hoffmann, J., Huang, C.W., Tang, J.: vGraph: a generative model for joint community detection and node representation learning. In: Advances in Neural Information Processing Systems, pp. 514–524 (2019)
27. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: LINE: large-scale information network embedding. In: Proceedings of the 24th International Conference on World Wide Web, pp. 1067–1077 (2015)
28. Tang, J., Aggarwal, C., Liu, H.: Node classification in signed social networks. In: Proceedings of the 2016 SIAM International Conference on Data Mining, pp. 54–62. SIAM (2016)
29. Tang, L., Liu, H.: Leveraging social media networks for classification. *Data Min. Knowl. Discov.* **23**(3), 447–478 (2011). <https://doi.org/10.1007/s10618-010-0210-x>
30. Tu, C., et al.: A unified framework for community detection and network representation learning. *IEEE Trans. Knowl. Data Eng.* **31**(6), 1051–1065 (2018)
31. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. arXiv preprint [arXiv:1710.10903](https://arxiv.org/abs/1710.10903) (2017)
32. Velickovic, P., Fedus, W., Hamilton, W.L., Liò, P., Bengio, Y., Hjelm, R.D.: Deep graph infomax (2019)
33. Wang, D., Cui, P., Zhu, W.: Structural deep network embedding. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1225–1234 (2016)



34. Wang, X., Cui, P., Wang, J., Pei, J., Zhu, W., Yang, S.: Community preserving network embedding. In: AAAI, vol. 17, pp. 203–209 (2017)
35. Xie, J., Kelley, S., Szymanski, B.K.: Overlapping community detection in networks: the state-of-the-art and comparative study. *ACM Comput. Surv. (CSUR)* **45**(4), 1–35 (2013)
36. Yang, J., Leskovec, J.: Overlapping community detection at scale: a nonnegative matrix factorization approach. In: Proceedings of the sixth ACM International Conference on Web Search and Data Mining, pp. 587–596 (2013)
37. Yang, J., Leskovec, J.: Defining and evaluating network communities based on ground-truth. *Knowl. Inf. Syst.* **42**(1), 181–213 (2013). <https://doi.org/10.1007/s10115-013-0693-z>
38. Yang, J., McAuley, J., Leskovec, J.: Community detection in networks with node attributes. In: 2013 IEEE 13th International Conference on Data Mining, pp. 1151–1156. IEEE (2013)
39. Yang, L., Cao, X., He, D., Wang, C., Wang, X., Zhang, W.: Modularity based community detection with deep learning. *IJCAI* **16**, 2252–2258 (2016)
40. Yang, Z., Cohen, W., Salakhudinov, R.: Revisiting semi-supervised learning with graph embeddings. In: International Conference on Machine Learning, pp. 40–48. PMLR (2016)