

---

# Robust Graph Embedding with Noisy Link Weights

---

Akifumi Okuno<sup>†,‡</sup>

okuno@sys.i.kyoto-u.ac.jp

Hidetoshi Shimodaira<sup>†,‡</sup>

shimo@i.kyoto-u.ac.jp

<sup>†</sup>Graduate School of Informatics, Kyoto University,

<sup>‡</sup>RIKEN Center for Artificial Intelligence Project (AIP)

## Abstract

We propose  $\beta$ -graph embedding for robustly learning feature vectors from data vectors and noisy link weights. A newly introduced *empirical moment  $\beta$ -score* reduces the influence of contamination and robustly measures the difference between the underlying correct expected weights of links and the specified generative model. The proposed method is computationally tractable; we employ a minibatch-based efficient stochastic algorithm and prove that this algorithm locally minimizes the empirical moment  $\beta$ -score. We conduct numerical experiments on synthetic and real-world datasets.

## 1 INTRODUCTION

In the past few decades, graph embedding (GE) that learns *feature vectors* of given graph nodes had high level of demand in a broad range of fields. Just an example among many, embedding social networks whose nodes and link weights represent users and their relationships, respectively, produces user feature vectors. Traditional multivariate statistical methods, such as clustering and classification, can then be applied to these feature vectors (Yan et al., 2007; Goyal and Ferrara, 2018), whereas these analysis methods, in general, cannot be applied directly to unprocessed graph nodes.

Classical GE typified by spectral GE (SGE; see Chung, 1997; Belkin and Niyogi, 2001) computes feature vectors so that their inner product similarities represent the link weights, and the locality-preserving projections (LPP; see He and Niyogi, 2004) extends SGE so that it additionally considers pre-obtained data vec-

tors of nodes. LPP computes feature vectors by linear transformation of data vectors; the local configuration of vectors is partially preserved through the transformation. Although SGE and LPP experimentally demonstrate reasonable performance, their computational complexity is high owing to eigendecomposition.

To reduce the high computational complexity, Tang et al. (2015) proposed a computationally efficient GE named large-scale information network embedding (LINE), which is based on the stochastic maximization of the likelihood over link weights. Although LINE, as well as SGE, does not consider pre-obtained data vectors of nodes, the graph embedding can be extended to utilize the data vectors by incorporating neural networks (Wang et al., 2016; Kipf and Welling, 2017; Dai et al., 2018). The graph embedding can be further extended to deal with multi-view setting (Okuno et al., 2018), where each vector is assigned to one of multiple types of vectors (e.g. image vectors, text vectors) with different dimensionalities.




Figure 1: **Proposed  $\beta$ -GE applied to noisy graph with four clusters. (a) corresponds to the existing likelihood-based GE. (c) reduces the influence of contamination in noisy link weights.** See Section 4.1 for detailed settings.

Although these GEs have been successful in many applications, their performance relies on the quality of observed link weights. However, these observed link weights, in practice, may contain noise. We especially call the link weights with noise as *noisy link weights*. As shown in Figure 1(a), the noise may degrade the GE’s performance; existing GE is hard to recover the underlying cluster structure from noisy link weights. To overcome this problem, GE should be robustified.

In order to capture the underlying link structure through the noisy link weights, noise-tolerant loss functions can be used. Although a variety of noise-tolerant loss functions have been considered in robust statistics, there have been only a few attempts at using such a noise-tolerant loss function in GE. To the best of the authors’ knowledge, only Pang and Yuan (2010) and Zhang et al. (2014) have incorporated noise-tolerant  $\ell_1$ - and  $\ell_{2,1}$ -norm-based loss functions into classical GE. Although these particular types of loss functions are experimentally demonstrated to alleviate the adverse effect of outlier data vectors, they are not theoretically guaranteed to be robust against noisy link weights.

To obtain theoretically-guaranteed robust GE, we attempt to employ  $\beta$ -score, that robustly measures the difference between two non-negative functions.  $\beta$ -score is also called as density-power score (Kanamori and Fujisawa, 2014, 2015), and it is especially called as  $\beta$ -cross entropy (see, e.g. Futami et al., 2018) if the two non-negative functions are limited to probability density function (pdf) or probability mass function (pmf). It is well known that Kullback-Leibler divergence is expressed as the difference of two cross-entropies ( $\beta = 0$ ), and quite similarly  $\beta$ -divergence (Cichocki et al., 2010), which is also known as density-power divergence (Basu et al., 1998), is expressed as the difference of two  $\beta$ -cross entropies.

By specifying Poisson distribution for link weights, a probabilistic generative model of GE has been considered (Okuno et al., 2018), and  $\beta$ -score can be applied to this model. However, naively applying  $\beta$ -score to the probabilistic generative model of GE has two disadvantages:  $\beta$ -score in this setting is (i) computationally intractable, and (ii) sensitive to distributional misspecification of the generative model, because  $\beta$ -score measures the difference between two probabilistic models.

In order to avoid these two disadvantages, we introduce *moment  $\beta$ -score* that measures the difference in terms of the expected values instead of probability distributions of link weights. The moment  $\beta$ -score is estimated by *empirical moment  $\beta$ -score (EMBS)*, which is robust against noise in link weights, and is also carefully designed to be (i) computationally tractable, and (ii) free from the distributional specification. By incorporating EMBS into existing likelihood-based GE using the probabilistic generative model of link weights, we propose a theoretically proven robust GE named  *$\beta$ -graph embedding ( $\beta$ -GE)*. This new method naturally extends the likelihood-based GE because the EMBS reduces to the negative log-likelihood of link weights as  $\beta \downarrow 0$ .

Our contribution is summarized as follows.

- (1) We propose a novel EMBS, that is a robust score function against noise in link weights as shown in Theorem 3.1. EMBS is also computationally tractable and it is free from the distributional specification.
- (2) We propose  $\beta$ -GE by simply incorporating the newly proposed EMBS into existing likelihood-based GE (Section 3.3). As far as we know,  $\beta$ -GE is the first GE whose robustness is theoretically shown.
- (3) We propose an efficient minibatch-based stochastic algorithm, which can be proved to locally minimize the EMBS (Section 3.5).
- (4) We conduct experiments on both synthetic and real-world datasets (Section 4).

## 2 LIKELIHOOD-BASED GRAPH EMBEDDING

In this section, we review existing likelihood-based GE. Throughout this paper, we consider a setting similar to that presented in Okuno et al. (2018). Although only single-view graph embedding is considered below, its extension to multi-view setting is straightforward as explained in Section 3.7.

Our dataset consists of data vectors  $\{\mathbf{x}_i\}_{i=1}^n$  and link weights  $\{w_{ij}\}_{(i,j) \in \mathcal{I}_n}$ , where  $\mathcal{I}_n := \{(i, j) \in \mathbb{N}^2 \mid 1 \leq i < j \leq n\}$ . The  $p$ -dimensional data vector  $\mathbf{x}_i \in \mathcal{X}$  takes values in a compact set  $\mathcal{X} \subset \mathbb{R}^p$ . The link weight  $w_{ij} \in \{0, 1, 2, \dots\}$  is non-negative integer, and it represents the strength of association between  $\mathbf{x}_i$  and  $\mathbf{x}_j$ .

Similarly to existing methods (Perozzi et al., 2014; Tang et al., 2015; Grover and Leskovec, 2016; Kipf and Welling, 2016), which employ Bernoulli-based models for considering binary weights, Okuno et al. (2018) employs a Poisson-based model

$$w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j \stackrel{\text{indep.}}{\sim} \text{Po}(\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)) \quad (1)$$

by taking  $w_{ij}$  as random variables for all  $(i, j) \in \mathcal{I}_n$ . Note that the model (1) does not restrict weights  $\{w_{ij}\}$  to be binary. The conditional expected value  $E(w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j)$  is specified by a symmetric, continuous, and non-negative function  $\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)$ , which represents a similarity between two data vectors  $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$ , where  $\boldsymbol{\theta} \in \Theta$  is a parameter vector and  $\Theta \subset \mathbb{R}^q$  is a non-empty compact set. Although any symmetric, continuous, and non-negative function can be used as the expected weight function, in this paper, the function is specified as

$$\log \mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j) := \langle f_{\psi}(\mathbf{x}_i), f_{\psi}(\mathbf{x}_j) \rangle - \gamma, \quad (2)$$

where  $f_\psi : \mathbb{R}^p \rightarrow \mathbb{R}^K$  is arbitrary continuous map whose parameter vector is  $\psi$ ,  $\gamma \in \mathbb{R}$  is a scaling parameter regularizing the sparseness of  $\{w_{ij}\}$ , and  $\theta$  denotes  $(\psi, \gamma)$ . Typically, a NN or linear transformation is used as  $f_\psi$ , and the *inner product similarity* (IPS) between two NNs  $\langle f_\psi(\mathbf{x}_i), f_\psi(\mathbf{x}_j) \rangle$  is also known as siamese-style similarity (Bromley et al., 1994).

IPS is used for a wide variety of GEs (Wang et al., 2016; Kipf and Welling, 2016; Hamilton et al., 2017), and has been proven to approximate arbitrary positive-definite (PD) similarities (Okuno et al., 2018, Theorem 5.1). In addition, equation (2), which incorporates a new parameter  $\gamma \in \mathbb{R}$  into IPS, is specifically called *constantly-shifted IPS* (C-SIPS), and is proved to approximate conditionally PD similarities, a wider class of similarities than PD similarities (Okuno and Shimodaira, 2019, Theorem 4.2). Thus, C-SIPS has a sufficiently high representation capability.

Once the optimal  $\hat{\theta} = (\hat{\psi}, \hat{\gamma})$  is obtained, data vectors  $\{\mathbf{x}_i\}$  can be transformed to *feature vectors*  $\{\mathbf{y}_i\}$  as

$$\mathbb{R}^p \supset \mathcal{X} \ni \mathbf{x}_i \mapsto \mathbf{y}_i := f_{\hat{\psi}}(\mathbf{x}_i) \in \mathbb{R}^K \quad (3)$$

for all  $i = 1, 2, \dots, n$ . The IPS  $\langle \mathbf{y}_i, \mathbf{y}_j \rangle$  of obtained feature vectors  $\{\mathbf{y}_i\}$  indicates the similarity between data vectors  $\mathbf{x}_i, \mathbf{x}_j$ . This transformation usually reduces the dimensionality of data vectors from  $p$  to  $K (\leq p)$ . Thus, the obtained vectors are expected to reduce their redundancy while considering the information on link weights. Obtaining the vector representations  $\{\mathbf{y}_i\}$  by considering the link weights  $\{w_{ij}\}$  is formally called GE.

One simple way to obtain the optimal parameter  $\hat{\theta}$  is to minimize the negative log-likelihood

$$\begin{aligned} L_{0,n}(\theta) &:= -\log \mathbb{P}(\{w_{ij}\}_{(i,j) \in \mathcal{I}_n} \mid \{\mathbf{x}_i\}_{i=1}^n) \\ &= \sum_{(i,j) \in \mathcal{I}_n} \left\{ -w_{ij} \log \mu_\theta(\mathbf{x}_i, \mathbf{x}_j) + \mu_\theta(\mathbf{x}_i, \mathbf{x}_j) \right\}, \quad (4) \end{aligned}$$

which is based on the Poisson model (1). By minimizing this negative log-likelihood, we obtain the maximum likelihood estimator (MLE)

$$\hat{\theta}_0 = (\hat{\psi}_0, \hat{\gamma}_0) := \arg \min_{\theta \in \Theta} L_{0,n}(\theta),$$

which specifies the optimal NN  $f_{\hat{\psi}_0}$  and the optimal scaling factor  $\hat{\gamma}_0 \geq 0$ . The above-mentioned procedure using MLE to obtain feature vectors is called *likelihood-based GE* in this paper.

### 3 $\beta$ -GRAPH EMBEDDING

Although the existing likelihood-based GE explained in Section 2 had some success, MLEs in general are

susceptible to contamination in data, because log-likelihood can be strongly influenced by noise. In Section 3.1, we review  $\beta$ -score and related divergence which have been developed in robust statistics for robustifying the log-likelihood. However, naively applying the  $\beta$ -score to the model (1) has two disadvantages as explained in Section 3.2. To overcome the drawbacks, we introduce *moment  $\beta$ -score* that robustly measures the difference between the underlying correct expected weights of links and the specified generative model, and propose its empirical estimation called *empirical moment  $\beta$ -score* (EMBS) in Section 3.3. Then, we propose robust  $\beta$ -GE equipped with EMBS. In Section 3.4, we theoretically prove that EMBS is robust against noise in link weights. In Section 3.5, we introduce a minibatch-based efficient stochastic algorithm, and prove that the algorithm locally minimizes EMBS. In Section 3.6, we discuss the selection of  $\beta$ . Finally, we extend  $\beta$ -GE to a multi-view setting in Section 3.7.

#### 3.1 $\beta$ -Score for Non-negative Functions

Our idea for robustifying the likelihood-based GE is to employ  $\beta$ -score, that is originally called as density-power score (Kanamori and Fujisawa, 2014, 2015).  $\beta$ -score robustly measures the difference between two non-negative functions as follows.

We consider a random variable  $\mathbf{z}$  taking a value in some set  $\mathcal{Z}$ , and consider a set of non-negative functions  $\mathcal{P}(\mathcal{Z}) := \{f : \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}\}$ . For non-negative functions  $g, f, \nu \in \mathcal{P}(\mathcal{Z})$ , we define  $\beta$ -score as

$$u_\beta^{\mathbf{z}}(g, f; \nu) := \begin{cases} -\sum_{\mathbf{z} \in \mathcal{Z}} g(\mathbf{z}) \nu(\mathbf{z}) \frac{f(\mathbf{z})^\beta - 1}{\beta} + \sum_{\mathbf{z} \in \mathcal{Z}} \nu(\mathbf{z}) \frac{f(\mathbf{z})^{1+\beta}}{1+\beta} & (\mathbf{z} \text{ is discrete}) \\ -\int_{\mathcal{Z}} g(\mathbf{z}) \nu(\mathbf{z}) \frac{f(\mathbf{z})^\beta - 1}{\beta} d\mathbf{z} + \int_{\mathcal{Z}} \nu(\mathbf{z}) \frac{f(\mathbf{z})^{1+\beta}}{1+\beta} d\mathbf{z} & (\mathbf{z} \text{ is continuous}) \end{cases}$$

where  $\beta > 0$  is a user-specified tuning parameter. For any fixed  $g \in \mathcal{P}(\mathcal{Z})$ , the minimizer  $f \in \mathcal{P}(\mathcal{Z})$  of  $u_\beta^{\mathbf{z}}(g, f; \nu)$  is known to be  $f = g$ , in the sense that  $f(\mathbf{z}) = g(\mathbf{z})$  over the support of  $\nu$ . If  $\nu(\mathbf{z}) = 1$  for all  $\mathbf{z} \in \mathcal{Z}$ , we abbreviate above  $\beta$ -score by  $u_\beta^{\mathbf{z}}(g, f)$ .

Let us consider the special case where the functions  $f, g$  are restricted to be pdf or pmf, and they are denoted by  $q, p$ , respectively. Then,  $\beta$ -score  $u_\beta^{\mathbf{z}}(q, p)$  is called  $\beta$ -cross entropy (see, e.g. Futami et al., 2018), and  $D_\beta(q, p) := u_\beta^{\mathbf{z}}(q, p) - u_\beta^{\mathbf{z}}(q, q)$  is known as  $\beta$ -divergence (Cichocki et al., 2010) or density-power divergence (Basu et al., 1998), which belongs to the Bregman-divergence family (Bregman, 1967). In particular,  $D_\beta(q, p)$  reduces to Kullback-Leibler divergence as  $\beta \downarrow 0$ .

$\beta$ -score has also been used for unnormalized models (Kanamori and Fujisawa, 2015) defined as  $cp(\mathbf{z})$  where  $c \geq 0$  is a scaling parameter and  $p(\mathbf{z})$  is pdf or pmf. For contaminated  $q = c_0q_0 + (1 - c_0)r$  with outlier distribution  $r$  and contamination rate  $1 - c_0$ , the  $\beta$ -score  $\min_{c \in (0,1]} u_\beta^z(q, cp)$  is equivalent to the “ $\gamma$ -divergence” between  $q$  and  $p$ , which robustly measures the difference between  $q_0$  and  $p$  (Jones et al., 2001; Fujisawa and Eguchi, 2008).

Although  $\beta$ -score can measure the difference between arbitrary non-negative functions, it has only been applied to probability distributions and unnormalized models as seen above. We call  $\beta$ -score considering only these cases as *probability  $\beta$ -score*. This is different from our usage of  $\beta$ -score introduced in Section 3.3.

### 3.2 Two Disadvantages of Naively Applying $\beta$ -Score to GE

Probability  $\beta$ -score can be employed for GE. Considering a conditional pdf  $q$  of  $w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j$  and the probabilistic generative model  $p_\theta$  defined in (1), we may minimize

$$E_{\mathcal{X}^2}(u_\beta^{(w_{12} \mid \mathbf{x}_1, \mathbf{x}_2)}(q, p_\theta)), \quad (5)$$

whose empirical estimation is

$$\begin{aligned} & \frac{1}{|\mathcal{I}_n|} \sum_{(i,j) \in \mathcal{I}_n} u_\beta^{(w \mid \mathbf{x}_i, \mathbf{x}_j)}(\hat{q}_{ij}, p_\theta) \\ &= -\frac{1}{|\mathcal{I}_n|} \sum_{(i,j) \in \mathcal{I}_n} \frac{p_\theta(w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j)^\beta - 1}{\beta} \\ & \quad + \frac{1}{|\mathcal{I}_n|} \sum_{(i,j) \in \mathcal{I}_n} \sum_{w \in \mathbb{N}_0} \frac{p_\theta(w \mid \mathbf{x}_i, \mathbf{x}_j)^{1+\beta}}{1 + \beta}, \quad (6) \end{aligned}$$

where  $\hat{q}_{ij}(w \mid \mathbf{x}_i, \mathbf{x}_j)$  takes value 1 if  $w = w_{ij}$  and 0 otherwise. (6) is a slight generalization of Ghosh et al. (2013) that applies probability  $\beta$ -score to linear regression, and it asymptotically converges in probability to (5) as  $n \rightarrow \infty$ . See Lemma A.2 in Supplement A.1 for the convergence.

However, estimating the probabilistic generative model  $p_\theta$  by minimizing (6) has the following two disadvantages:

- (i) The last term in (6) is computationally intractable because of the infinite summation  $\sum_{w \in \mathbb{N}_0}$ .
- (ii)  $w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j \sim q$  is required to follow Poisson distribution (1) for correctly estimating the probabilistic model, but the model is only approximation to the underlying distribution in reality.

Regarding (i), many of existing studies such as Ghosh et al. (2013) only consider normal linear regression, so that the corresponding term can be analytically calculated. As for non-normal setting, the infinite-summation in (6) similarly appears in eq. (2.4) of Kawashima and Fujisawa (2018), that applies  $\gamma$ -divergence to sparse Poisson regression, and they compute the term by the finite-sum approximation instead.

Regarding (ii), although other probabilistic model can be used as  $p_\theta$ , its estimation is still sensitive to the distributional misspecification as long as the  $\beta$ -score is naively applied to the user-specified probabilistic generative model  $p_\theta$ .

### 3.3 Proposed $\beta$ -Graph Embedding

In order to avoid the two disadvantages of probability  $\beta$ -score explained in Section 3.2, here we introduce *moment  $\beta$ -score* that applies  $\beta$ -score to the expected value of  $\mathbf{z}$  instead of  $p(\mathbf{z})$ . For graph embedding, we consider the conditional expectation of link weights  $g(\mathbf{x}_1, \mathbf{x}_2) := E(w_{12} \mid \mathbf{x}_1, \mathbf{x}_2)$ , and apply  $\beta$ -score to it as

$$u_\beta^{(\mathbf{x}_1, \mathbf{x}_2)}(g, \mu_\theta; \nu) = E_{\mathcal{X}^2} \left( -g(\mathbf{x}_1, \mathbf{x}_2) \frac{\mu_\theta(\mathbf{x}_1, \mathbf{x}_2)^\beta - 1}{\beta} + \frac{\mu_\theta(\mathbf{x}_1, \mathbf{x}_2)^{1+\beta}}{1 + \beta} \right),$$

where  $\nu$  is pdf of  $(\mathbf{x}_1, \mathbf{x}_2) \in \mathcal{X}^2$ .  $\beta > 0$  is a user-specified tuning parameter, and  $\mu_\theta$  is specified as CSIPS (2).

Moment  $\beta$ -score can be empirically estimated by *empirical moment  $\beta$ -score (EMBS)*, that is

$$L_{\beta,n}(\theta) := \sum_{(i,j) \in \mathcal{I}_n} \left\{ -w_{ij} \frac{\mu_\theta(\mathbf{x}_i, \mathbf{x}_j)^\beta - 1}{\beta} + \frac{\mu_\theta(\mathbf{x}_i, \mathbf{x}_j)^{1+\beta}}{1 + \beta} \right\} \quad (7)$$

in the sense that  $L_{\beta,n}(\theta) \xrightarrow{P} u_\beta^{(\mathbf{x}_1, \mathbf{x}_2)}(g, \mu_\theta; \nu)$  as  $n \rightarrow \infty$ , under some assumptions. See Lemma A.3 in Supplement A.1 for the convergence.

Noisy link weights can be modeled by defining the conditional expectation  $g(\mathbf{x}_i, \mathbf{x}_j) = E(w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j)$  as the sum of underlying weight  $\mu_*(\mathbf{x}_i, \mathbf{x}_j)$  and noise  $\eta_*(\mathbf{x}_i, \mathbf{x}_j)$ . As moment  $\beta$ -score robustly measures the difference between  $\mu_*$  and  $\mu_\theta$  even if  $g$  is contaminated by  $\eta_*$ , EMBS is robust against noise in link weights. See Section 3.4 for details.

In addition to the robustness of EMBS against noisy link weights, EMBS simultaneously avoids the two disadvantages of probability  $\beta$ -score explained in Section 3.2. Regarding (i), i.e. the computational intractability, the numerical optimization of  $L_{\beta,n}(\theta)$  is not difficult since infinite summation is not involved.



Regarding (ii), i.e. the lack of robustness against distributional misspecification, EMBS is free from this problem because it does not require distributional specification.

We replace the negative log-likelihood  $L_{0,n}(\boldsymbol{\theta})$  defined in (4) with EMBS  $L_{\beta,n}(\boldsymbol{\theta})$  defined in (7). The remaining procedure is all the same: we define the  $\beta$ -estimator as a minimizer of EMBS

$$\hat{\boldsymbol{\theta}}_{\beta} = (\hat{\boldsymbol{\psi}}_{\beta}, \hat{\gamma}_{\beta}) := \arg \min_{\boldsymbol{\theta} \in \Theta} L_{\beta,n}(\boldsymbol{\theta}), \quad (8)$$

and the  $\beta$ -estimator defines feature vectors  $\mathbf{y}_{\beta,i} := f_{\hat{\boldsymbol{\psi}}_{\beta}}(\mathbf{x}_i)$ ,  $i = 1, 2, \dots, n$  by substituting  $\hat{\boldsymbol{\psi}}_{\beta}$  into (3). The whole procedure described above, that obtains feature vectors using the  $\beta$ -estimator (8), is called  *$\beta$ -graph embedding ( $\beta$ -GE)*. Since EMBS reduces to the negative log-likelihood as  $\beta \downarrow 0$ ,  $\beta$ -GE naturally extends the likelihood-based GE, and its robustness can be formally proven as shown in the following section.

### 3.4 Robustness against Noise in Link Weights

Thus far, we have described likelihood-based GE using the model (1) and proposed  $\beta$ -GE. However, the model (1) does not explain how noisy link weights are actually produced. For that reason, in this section, we first formally define the probabilistic model that considers noisy link weights. Then, we explain why moment  $\beta$ -score, that is used in  $\beta$ -GE, is robust against noise in link weights. The following explanation for moment  $\beta$ -score is an adaptation of that for probability  $\gamma$ -score given in Lemma 3.1 of Fujisawa and Eguchi (2008).

We consider the generative model of  $w_{ij} \geq 0$  up to the first and second moments. For describing noisy link weights, *expected noise*  $\eta_*(\mathbf{x}_1, \mathbf{x}_2)$  is added to the underlying *correct expected weight*  $\mu_*(\mathbf{x}_1, \mathbf{x}_2)$ , where  $\mu_*$  and  $\eta_*$  are non-negative functions over  $\mathcal{X}^2$ .

$$\begin{aligned} E(w_{ij} \mid \mathbf{x}_i, \mathbf{x}_j) &= \mu_*(\mathbf{x}_i, \mathbf{x}_j) + \eta_*(\mathbf{x}_i, \mathbf{x}_j) \\ &=: g(\mathbf{x}_i, \mathbf{x}_j), \end{aligned} \quad (9)$$

$$E(w_{ij}^2 \mid \mathbf{x}_i, \mathbf{x}_j) < \infty, \quad (10)$$

$$\mathbf{x}_i \stackrel{\text{i.i.d.}}{\sim} Q \quad (11)$$

for all  $(i, j) \in \mathcal{I}_n$ . The support of the density function  $Q(\mathbf{x})$  is denoted as  $\text{supp}Q \subset \mathcal{X}$ . The amount of noise is measured by

$$\alpha := E_{\mathcal{X}^2}(\eta_*(\mathbf{x}_1, \mathbf{x}_2)), \quad (12)$$

where  $E_{\mathcal{X}^2}$  indicates expectation with respect to the joint density  $\nu(\mathbf{x}_1, \mathbf{x}_2) := Q(\mathbf{x}_1)Q(\mathbf{x}_2)$ . We implicitly assume that  $\mu_*(\mathbf{x}_1, \mathbf{x}_2)$  is sufficiently small when  $\eta_*(\mathbf{x}_1, \mathbf{x}_2)$  is large. More specifically, supposing

$$E_{\mathcal{X}^2}(\eta_*(\mathbf{x}_1, \mathbf{x}_2)\mu_*(\mathbf{x}_1, \mathbf{x}_2)^{\beta_0}) < \varepsilon_*, \quad (13)$$

we assume that  $\varepsilon_*$  is sufficiently small for an appropriately large  $\beta_0 > 0$ . This corresponds to the assumption (\*) for  $\nu_f$  in Fujisawa and Eguchi (2008). We also assume that the model  $\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)$  is correctly specified:

$$\exists \boldsymbol{\theta}_* \in \Theta \text{ such that } \mu_{\boldsymbol{\theta}_*}(\mathbf{x}_1, \mathbf{x}_2) = \mu_*(\mathbf{x}_1, \mathbf{x}_2) \quad (14)$$

for all  $(\mathbf{x}_1, \mathbf{x}_2) \in (\text{supp}Q)^2$ .

Following the above settings,  $\beta$ -GE robustly infers the correct expected weight  $\mu_*$  from noisy link weights  $\{w_{ij}\}$  whose conditional expectation is contaminated by noise as (9).

For identifying the robustness of  $\beta$ -GE, we consider a restricted parameter set

$$\begin{aligned} \Theta_{\varepsilon} &:= \{\boldsymbol{\theta} \in \Theta \mid \\ &E_{\mathcal{X}^2}(\eta_*(\mathbf{x}_1, \mathbf{x}_2)\mu_{\boldsymbol{\theta}}(\mathbf{x}_1, \mathbf{x}_2)^{\beta_0}) < \varepsilon\} \end{aligned} \quad (15)$$

that satisfies  $\boldsymbol{\theta}_* \in \Theta_{\varepsilon_*} \subseteq \Theta_{\varepsilon}$  for  $\varepsilon \geq \varepsilon_*$  if (13) and (14) hold.

**Theorem 3.1.** Suppose that  $\varepsilon \geq \varepsilon_*$ ,  $\beta \in (0, \beta_0]$ , and (9)–(15) hold. Then, there exists a function  $M(\boldsymbol{\theta}) \geq 0$  such that, for all  $\boldsymbol{\theta} \in \Theta_{\varepsilon}$ ,

$$\begin{aligned} L_{\beta,n}(\boldsymbol{\theta}) &= u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(\mu_*, \mu_{\boldsymbol{\theta}}; \nu) + \alpha\beta^{-1} \\ &\quad - M(\boldsymbol{\theta})\varepsilon^{\beta/\beta_0} + O_p(1/\sqrt{n}), \end{aligned} \quad (16)$$

$$M(\boldsymbol{\theta}) \leq \alpha^{1-\beta/\beta_0}\beta^{-1}. \quad (17)$$

**Proof:** The law of large numbers shown in Lemma A.3 of Supplement indicates  $L_{\beta,n}(\boldsymbol{\theta}) = u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(g, \mu_{\boldsymbol{\theta}}; \nu) + O_p(1/\sqrt{n})$ , and a simple calculation leads to (16) as

$$\begin{aligned} u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(g, \mu_{\boldsymbol{\theta}}; \nu) &= u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(\mu_* + \eta_*, \mu_{\boldsymbol{\theta}}; \nu) \\ &= u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(\mu_*, \mu_{\boldsymbol{\theta}}; \nu) - E_{\mathcal{X}^2}\left(\eta_*(\mathbf{x}_1, \mathbf{x}_2)\frac{\mu_{\boldsymbol{\theta}}(\mathbf{x}_1, \mathbf{x}_2)^{\beta} - 1}{\beta}\right) \\ &= u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(\mu_*, \mu_{\boldsymbol{\theta}}; \nu) + \underbrace{\beta^{-1}E_{\mathcal{X}^2}(\eta_*(\mathbf{x}_1, \mathbf{x}_2))}_{=\alpha\beta^{-1}} \\ &\quad - \underbrace{\beta^{-1}E_{\mathcal{X}^2}(\eta_*(\mathbf{x}_1, \mathbf{x}_2)\mu_{\boldsymbol{\theta}}(\mathbf{x}_1, \mathbf{x}_2)^{\beta})\varepsilon^{-\beta/\beta_0}}_{=:M(\boldsymbol{\theta})}\varepsilon^{\beta/\beta_0}. \end{aligned}$$

Then (17) follows from the inequality in (15) and Lyapunov's inequality as shown in Lemma A.4 of Supplement.  $\square$

Theorem 3.1 asserts that, EMBS  $L_{\beta,n}(\boldsymbol{\theta})$  approximates  $u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(\mu_*, \mu_{\boldsymbol{\theta}}; \nu)$  up to the constant  $\alpha\beta^{-1}$  as long as  $\varepsilon > 0$  is specified to be sufficiently small, even if the correct expected weight  $\mu_*$  is contaminated by expected noise  $\eta_*$ . As  $\alpha\beta^{-1}$  is constant with respect to  $\boldsymbol{\theta}$ , and  $u_{\beta}^{(\mathbf{x}_1, \mathbf{x}_2)}(\mu_*, \mu_{\boldsymbol{\theta}})$  is minimized at  $\boldsymbol{\theta} = \boldsymbol{\theta}_*$ , minimization of EMBS leads to robust estimation of the correct expected weight  $\mu_*$ .

### 3.5 Minibatch-Based Stochastic Optimization Algorithm

To minimize the EMBS (7), we consider gradient descent-based approaches. However, computing a full-batch gradient in our setting requires summing up  $O(|\mathcal{I}_n|) = O(n^2)$  terms for each iteration. In practice, the computational complexity is remarkably high and non-negligible. To reduce this high computational complexity, we employ a minibatch-based stochastic algorithm.

Let  $\mathcal{W}_n$  be an index set of positive weights  $\{(i, j) \in \mathcal{I}_n \mid w_{ij} > 0\} \subset \mathcal{I}_n$ . At iteration  $t = 1, 2, \dots$ , we pick up  $m_1, m_2 \in \mathbb{N}$  elements from  $\mathcal{W}_n, \mathcal{I}_n$  uniformly at random, and denote the sets as  $\mathcal{W}_n^{(t)}, \mathcal{I}_n^{(t)}$ , respectively. Here  $\mathcal{W}_n^{(t)}$  and  $\mathcal{I}_n^{(t)}$  can overlap, but no duplication in each set. In the remaining of this section,  $E^{(t)}, E^*$  represent expectations with respect to resampling sets  $(\mathcal{W}_n^{(t)}, \mathcal{I}_n^{(t)}), \{(\mathcal{W}_n^{(t)}, \mathcal{I}_n^{(t)})\}_{t=1,2,\dots}$ , respectively.

Similarly to Okuno et al. (2018) Section 4, EMBS is stochastically approximated by

$$h^{(t)}(\boldsymbol{\theta}) := - \sum_{(i,j) \in \mathcal{W}_n^{(t)}} w_{ij} \frac{\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)^\beta - 1}{\beta} + \lambda \sum_{(i,j) \in \mathcal{I}_n^{(t)}} \frac{\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)^{1+\beta}}{1+\beta},$$

where  $\lambda > 0$  is a tuning parameter. Then, our iterative algorithm updates  $\boldsymbol{\theta}$  by

$$\boldsymbol{\theta}^{(t+1)} := \mathcal{P}_{\Theta} \left( \boldsymbol{\theta}^{(t)} - \delta^{(t)} \frac{\partial h^{(t)}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(t)}} \right) \quad (18)$$

with learning rate  $\delta^{(t)} \geq 0$ ,  $t \in \mathbb{N}_0$ , user-specified initial value  $\boldsymbol{\theta}^{(0)}$ , and projection function  $\mathcal{P}_{\Theta}(\boldsymbol{\theta}') := \arg \min_{\boldsymbol{\theta} \in \Theta} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2$ . This algorithm is a slight modification of minibatch stochastic gradient descent (minibatch SGD; see Goodfellow et al., 2016) with a flavor of negative sampling (Mikolov et al., 2013). A similar algorithm for likelihood-based GE can be found in Okuno et al. (2018).

Compared with the plain SGD that uses only one sample for computing the gradient in each step, minibatch-based stochastic algorithm is proved to be more stable, in the sense that its asymptotic variance with respect to the number of iterations is smaller (Bonakdarpour and Toulis, 2016; Toulis et al., 2017). Although our proposed algorithm (18) involves the randomness to make sets  $\mathcal{W}_n^{(t)}, \mathcal{I}_n^{(t)}$  ( $t = 1, 2, \dots$ ), its convergence limit can be established as shown in the following theorem.

**Theorem 3.2.** In order to consider the optimization locally, we redefine the parameter set as  $\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^q \mid \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq D\}$  with a constant  $D > 0$  and a fixed

parameter value  $\boldsymbol{\theta}_0 \in \mathbb{R}^q$ . Suppose that  $\boldsymbol{\theta}_* \in \Theta$  is a solution of  $\partial h(\boldsymbol{\theta})/\partial \boldsymbol{\theta} = \mathbf{0}$  for  $h(\boldsymbol{\theta}) := E^{(1)}(h^{(1)}(\boldsymbol{\theta}))$ . We also assume that (i)  $\mu_{\boldsymbol{\theta}}(\mathbf{x}_1, \mathbf{x}_2) \in C^1(\Theta)$  for any  $(\mathbf{x}_1, \mathbf{x}_2) \in \mathcal{X}$ , (ii)  $h(\boldsymbol{\theta})$  is strongly-convex on  $\Theta$ , and (iii)  $\delta^{(t)} = O(t^{-\alpha})$  for some  $\alpha \in (0, 1]$ . Then, the estimator  $\boldsymbol{\theta}^{(t)}$  of our proposed algorithm (18) converges to  $\boldsymbol{\theta}_*$ , in the sense that  $E^*(\|\boldsymbol{\theta}^{(t)} - \boldsymbol{\theta}_*\|_2^2) \rightarrow 0$  as  $t \rightarrow \infty$ . The solution  $\boldsymbol{\theta}_*$  is the unique minimizer of  $h(\boldsymbol{\theta})$  over  $\Theta$ , and the equation  $\partial h(\boldsymbol{\theta})/\partial \boldsymbol{\theta} = \mathbf{0}$  is written as

$$\sum_{(i,j) \in \mathcal{I}_n} \left\{ \left( -vm_1 w_{ij} + \lambda m_2 \mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j) \right) \times \mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)^\beta \frac{\partial \log \mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)}{\partial \boldsymbol{\theta}} \right\} = \mathbf{0}, \quad (19)$$

where  $v := |\mathcal{I}_n|/|\mathcal{W}_n|$ ,  $m_1 := |\mathcal{W}_n^{(1)}|$ , and  $m_2 := |\mathcal{I}_n^{(1)}|$ .

As (18) is classified as a standard projected stochastic gradient descent, applying existing theorems such as Moulines and Bach (2011) Theorem 2 leads to Theorem 3.2. Proof and further explanation are given in Supplement A.3.

Theorem 3.2 indicates that the convergence limit of our proposed algorithm satisfies the estimating equation  $\partial L_{\beta,n}(\boldsymbol{\theta})/\partial \boldsymbol{\theta} = \mathbf{0}$  if  $\lambda m_2 = vm_1$ . Then it locally minimizes  $L_{\beta,n}(\boldsymbol{\theta})$ .

Moreover, even if  $\lambda m_2 \neq vm_1$ , (19) indicates that  $w_{ij}$  is approximated by  $k \mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)$  with  $k = \lambda m_2 / vm_1$ , because (19) is equivalent to replacing  $\mu_{\boldsymbol{\theta}}$  by  $k \mu_{\boldsymbol{\theta}}$  in (7). Then  $\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)$  approximates the link weight  $w_{ij}$  up to the scaling  $k$ . In practice, the scaling is not really an issue; we only need the ratio  $\mu_{\boldsymbol{\theta}}(\mathbf{x}_i, \mathbf{x}_j)/\mu_{\boldsymbol{\theta}}(\mathbf{x}_k, \mathbf{x}_l)$  to infer which of the pairs  $(\mathbf{x}_i, \mathbf{x}_j), (\mathbf{x}_k, \mathbf{x}_l)$  has a stronger relation. Thus, we can ignore the condition  $\lambda m_2 = vm_1$ , and we empirically determine  $\lambda$  in experiments in order to seek faster convergence of the algorithm.

### 3.6 Selection of the Parameter $\beta$

In this section, we discuss the selection of  $\beta > 0$ . As Theorem 3.1 shows, the bias term in (16) is  $O(\varepsilon^{\beta/\beta_0})$  if  $\boldsymbol{\theta} \in \Theta_\varepsilon$ . Thus larger  $\beta$  may make the term smaller; a larger  $\beta$  enhances the robustness. On the other hand, as  $\beta \downarrow 0$ , the  $\beta$ -estimator converges to the MLE, which is known to be asymptotically efficient (Ferguson, 1996); a smaller  $\beta$  enhances the asymptotic efficiency. To simultaneously attain high robustness and high asymptotic efficiency,  $\beta$  needs to be specified properly. In fact, specifying a proper  $\beta$  has been a central issue when considering UBSF-related methods (Durio and Isaia, 2011; Ghosh and Basu, 2015) for the same reason. However, no decisive way has been presented, even when considering simple linear regression.

Referring to existing studies on a tuning parameter selection of the  $\beta$ -score, an idea worth considering is applying cross-validation (CV) with EMBS using a fixed parameter  $\beta_0 > 0$  (independent of  $\beta$ ). By virtue of the robustness of EMBS, this EMBS-based CV is expected to be robust against noisy link weights whereas a negative log-likelihood-based CV may not be. A similar idea can be found in Mollah et al. (2007) and Kawashima and Fujisawa (2017).

However, CV requires significant computational resources, and existing studies on the UBSF in fact empirically demonstrate that the UBSF with  $\beta$  heuristically chosen from  $[0.1, 1.0]$  even demonstrates a good performance in practice (Jian and Vemuri, 2005; Dessein et al., 2010).

### 3.7 Extension to the Multi-View Setting

Here, we consider the *multi-view* setting, that there exist  $D$  different types of data vectors, whereas the above-mentioned methods consider only one type. We denote the type of  $i$ th vector  $\mathbf{x}_i$  as  $d_i \in \{1, 2, \dots, D\}$ , and the dimension  $p_{d_i}$  of  $\mathbf{x}_i$  depends on the type  $d_i$ . A typical example of multi-view data is images and their text tags, whose associations are represented by links.

To employ such multi-view data vectors, (2) can be simply extended to the multi-view setting as

$$\log \mu_{\theta}(\mathbf{x}_i, \mathbf{x}_j) = \langle f_{\psi^{(d_i)}}^{(d_i)}(\mathbf{x}_i), f_{\psi^{(d_j)}}^{(d_j)}(\mathbf{x}_j) \rangle - \gamma^{(d_i, d_j)}, \quad (20)$$

where  $f_{\psi^{(d)}}^{(d)} : \mathbb{R}^{p_d} \rightarrow \mathbb{R}^K$ , ( $d = 1, 2, \dots, D$ ) are different NNs whose parameters are  $\{\psi^{(d)}\}_{d=1}^D$ , and  $K$  is specified by users such that  $K \leq p_1, p_2, \dots, p_d$ . As we also consider the scaling parameter  $\{\gamma^{(d,e)}\}_{1 \leq d < e \leq D}$ , the full parameter vector is  $\theta = (\{\psi^{(d)}\}_{d=1}^D, \{\gamma^{(d,e)}\}_{1 \leq d < e \leq D})$ . This multi-view extension of likelihood-based GE using (20) is proposed recently and it is called probabilistic multi-view GE (PMvGE; see Okuno et al., 2018).

By specifying linear transformations for  $f_{\psi^{(d)}}^{(d)}$ , PMvGE approximates CDMCA (Shimodaira, 2016) that generalizes various multivariate analysis methods such as principal component analysis (PCA), canonical correlation analysis (CCA; see Hotelling, 1936), LPP, HIM-FAC (Nori et al., 2012); see Section 3.6 of Okuno et al. (2018). Thus, PMvGE extends various multivariate analysis methods.

The multi-view setting of PMvGE is easily incorporated into  $\beta$ -GE by (20). This multi-view  $\beta$ -GE corresponds to robustification of PMvGE, and therefore it indeed robustifies various existing methods for multi-view analysis.

## 4 EXPERIMENTS

In order to confirm the robustness of  $\beta$ -GE, we conduct numerical experiments on a synthetic dataset in Section 4.1 and a real-world dataset in Section 4.2. Feature vectors obtained by graph embedding methods are evaluated by clustering task.

### 4.1 Experiment on a Synthetic Dataset

**Synthetic dataset:** For generating  $n = 200$  data vectors  $\{\mathbf{x}_i\}_{i=1}^{200}$  of  $p = 20$  dimensions, we first prepared centers of clusters  $\mathbf{x}_k^{(0)} \in \mathbb{R}^5$ ,  $k = 1, \dots, 4$ , and a linear transformation  $\mathbf{A} \in \mathbb{R}^{20 \times 5}$ , where all the elements are generated from  $N(0, 1)$  independently. Then data vectors are generated by  $\mathbf{x}_i \stackrel{\text{i.i.d.}}{\sim} N(\mathbf{A}\mathbf{x}_{k_i}^{(0)}, \mathbf{I}_{20})$  for  $i = 1, 2, \dots, 200$ , where  $k_i := \lceil i/50 \rceil$  represents cluster index to which  $\mathbf{x}_i$  belongs.  $\{\mathbf{x}_i\}_{i=1}^{50}$  are finally rescaled such that  $\sum_{i=1}^{200} \|\mathbf{x}_i\|_2 / 200 = 4$ . Then binary link weights  $\{w_{ij}\}_{i,j=1}^{200}$  are generated from Bernoulli distribution. For  $k_i = k_j$ ,  $w_{ij} = w_{ji} \sim B(0.05)$  ( $i \neq j$ ) and  $w_{ii} = 0$ . For  $k_i \neq k_j$ ,  $w_{ij} = w_{ji} \sim B(\xi)$  with a specified parameter  $\xi \in [0, 1]$ .

**Estimation:** Feature vectors  $\{\mathbf{y}_i\}$  are computed by applying likelihood-based GE (corresponding to  $\beta = 0$ ) and  $\beta$ -GE ( $\beta = 0.1, 0.5, 1$ ) to the generated data vectors  $\{\mathbf{x}_i\}$  and link weights  $\{w_{ij}\}$ . Loss functions are ridge-regularized. For optimization, we utilize the BFGS algorithm (Fletcher, 2013) with random initialization.

**Evaluation:** We apply  $k$ -means clustering to obtained feature vectors, and evaluate the result using the purity score (Manning et al., 2008).

**Results:** Sample average and standard deviation of purity scores over 10 experiments are listed in Table 1 along with noisy link probability  $\xi$ . The purity score takes a value in  $[0, 1]$  and a higher value is better. When  $\xi$  is small ( $\xi = 0.01$ ), likelihood-based GE achieves the best score, however, if the probability increases ( $\xi = 0.02, 0.03$ ),  $\beta$ -GE with larger  $\beta$  achieves a better score than likelihood-based GE.

Table 1: Sample average and standard error of purity scores over 10 experiments. Higher score is better.

	$\xi = 0.01$	$\xi = 0.02$	$\xi = 0.03$
Likelihood	<b>0.72</b> $\pm$ 0.01	0.66 $\pm$ 0.01	0.58 $\pm$ 0.01
$\beta = 0.1$	0.71 $\pm$ 0.01	<b>0.72</b> $\pm$ 0.01	0.60 $\pm$ 0.01
$\beta = 0.5$	0.71 $\pm$ 0.01	0.69 $\pm$ 0.01	<b>0.64</b> $\pm$ 0.00
$\beta = 1$	0.70 $\pm$ 0.00	0.64 $\pm$ 0.01	<b>0.64</b> $\pm$ 0.00

**Discussion:** Although Table 1 demonstrates the robustness of  $\beta$ -GE, the difference between  $\beta$ -GE and existing GE is not drastic. This may be because the expected noise  $\eta_*$  is not very different from the correct

expected weight  $\mu_*$ . Although  $\beta$ -GE may improve the likelihood-based GE drastically if some  $w_{ij}$  take extremely large values as in the classical setting of robust statistics, we do not investigate such a setting in this paper.

**Visualization in Figure 1:** Using  $\xi = 0.03$ , feature vectors obtained by the likelihood-based GE ( $\Leftrightarrow \beta = 0$ ) and  $\beta$ -GE with  $\beta = 0.5, 1.0$  are plotted in Figure 1 as an illustrative example. In each figure, nodes are colored by the four clusters, and links are shown as gray segments. Our proposed  $\beta$ -GE distinguishes the four clusters even when noisy links are included in the synthetic dataset, whereas the clusters are unclear for the likelihood-based GE.

## 4.2 Experiment on the Cora Citation Network Dataset

**Dataset:** Cora citation network (Sen et al., 2008) consists of 2,708 nodes and 5,278 ordered edges. Each node  $v_i$  represents a document, which has 1,433-dimensional (bag-of-words) data vector  $\mathbf{x}_i \in \{0, 1\}^{1,433}$  and a class label of 7 classes. Each directed edge represents citation from a document  $v_i$  to another document  $v_j$ . We set the link weight as  $w_{ij} = w_{ji} = 1$  by ignoring the direction, and  $w_{ij} = 0$  otherwise. We divide the dataset into a training set consisting of 2,166 nodes (80%) with their edges, and a test set consisting of the remaining 542 nodes (20%) with their edges. Hyper-parameters are tuned by a validation set consisting of 20% of the training set.

**Neural network architecture:** We employ a one-hidden-layer fully connected network, which consists of 3,000 tanh hidden units and 100 tanh output units. The dimension of the feature vector is  $K = 100$ . The minibatch-based stochastic algorithm shown in Section 3.5 is used for optimization with batch normalization and dropout ( $p = 0.5$ ). The learning rate and momentum are tuned on the validation set.

**Evaluation:**  $k$ -means clustering is applied to the obtained feature vectors of nodes representing documents. The number of clusters is set as 7, and the clustering result is evaluated by normalized mutual information (NMI; see Manning et al., 2008). We compare the result with the stochastic block model (SBM; see Holland et al., 1983), ISOMAP (Tenenbaum et al., 2000), locally linear embedding (LLE; see Roweis and Saul, 2000), SGE, multi-dimensional scaling (MDS; see Kruskal, 1964), DeepWalk (Perozzi et al., 2014), and GraphSAGE (Hamilton et al., 2017).

**Results:** The quality of clustering is evaluated by NMI. Sample averages and standard errors over 10 experiments are listed in Table 2. In experiment (A), feature vectors are computed from both the training

set and test set, and they are evaluated on the test set. In experiment (B), feature vectors are computed from only the training set, and they are evaluated on the test set. SGE, MDS, and DeepWalk are not inductive, and they cannot be applied to unseen data vectors in experiment (B).

Table 2: Sample average and standard error of NMIs (multiplied by 100) over 10 times experiments on Cora citation dataset. Higher is better.

	(A)	(B)
SBM <sup>†</sup>	4.37 ± 0.46	2.81 ± 0.03
ISOMAP <sup>†</sup>	13.0 ± 0.11	14.3 ± 0.63
LLE <sup>†</sup>	7.40 ± 1.08	9.47 ± 0.95
SGE <sup>†</sup>	1.41 ± 0.11	-
MDS <sup>†</sup>	2.81 ± 0.03	-
DeepWalk <sup>†</sup>	16.7 ± 0.33	-
GraphSAGE <sup>†</sup>	19.6 ± 0.29	12.4 ± 0.95
Likelihood-based <sup>†</sup>	35.9 ± 0.28	30.5 ± 1.23
$\beta$ -GE ( $\beta = 0.1$ )	<b>36.2 ± 0.28</b>	<b>30.7 ± 1.11</b>
$\beta$ -GE ( $\beta = 0.5$ )	<b>36.1 ± 0.28</b>	<b>31.0 ± 0.94</b>

<sup>†</sup> scores are referred to Okuno et al. (2018).

## 5 CONCLUSION

We have proposed  $\beta$ -GE, by incorporating the newly proposed EMBS into existing likelihood-based GE. We have proved that  $\beta$ -GE is robust against noise in link weights, and is free from the distributional specification. We have also proposed an efficient minibatch-based stochastic algorithm that is theoretically proven to exactly locally minimize EMBS. Although robustification of GE is very challenging in practice, numerical experiments on synthetic and real-world datasets demonstrated the promising performance of  $\beta$ -GE compared with existing methods.

## ACKNOWLEDGEMENT

We would like to thank Takayuki Kawashima for helpful discussions. This work was partially supported by JSPS KAKENHI grant 16H02789 to HS and 17J03623 to AO.

## References

- Basu, A., Harris, I. R., Hjort, N. L., and Jones, M. (1998). Robust and Efficient Estimation by Minimising a Density Power Divergence. *Biometrika*, 85(3):549–559.
- Belkin, M. and Niyogi, P. (2001). Laplacian Eigenmaps and Spectral Techniques for Embedding and



- Clustering. In *Advances in Neural Information Processing Systems (NIPS)*, pages 585–591.
- Bonakdarpour, M. and Toulis, P. P. (2016). Statistical Perspectives of Stochastic Optimization. In *NIPS Workshop*.
- Bregman, L. M. (1967). The Relaxation Method of Finding the Common Point of Convex Sets and Its Application to the Solution of Problems in Convex Programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217.
- Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., and Shah, R. (1994). Signature Verification using a "Siamese" Time Delay Neural Network. In *Advances in Neural Information Processing Systems (NIPS)*, pages 737–744.
- Chung, F. R. (1997). *Spectral Graph Theory*. Number 92. American Mathematical Society.
- Cichocki, A., and Amari, S. (2010). Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6):1532-1568.
- Dai, Q., Li, Q., Tang, J., and Wang, D. (2018). AdversarialNetwork Embedding. In *Proceedings of the conference on Artificial Intelligence (AAAI)*.
- Dessein, A., Cont, A., and Lemaitre, G. (2010). Real-time polyphonic music transcription with non-negative matrix factorization and beta-divergence. In *Proceedings of the International Conference of Society for Music Information Retrieval (ISMIR)*, pages 489–494.
- Durio, A. and Isaia, E. D. (2011). The Minimum Density Power Divergence Approach in Building Robust Regression Models. *Informatica*, 22(1):43–56.
- Ferguson, T. (1996). *A Course in Large Sample Theory*. Chapman & Hall Texts in Statistical Science Series. Taylor & Francis.
- Fletcher, R. (2013). *Practical Methods of Optimization*. John Wiley & Sons.
- Fujisawa, H. and Eguchi, S. (2008). Robust parameter estimation with a small bias against heavy contamination. *Journal of Multivariate Analysis*, 99(9):2053–2081.
- Futami, F., Sato, I., and Sugiyama, M. (2018). Variational Inference based on Robust Divergences. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Ghosh, A., and Basu, Ayanendranath. (2013). Robust estimation for independent non-homogeneous observations using density power divergence with applications to linear regression. *Electronic Journal of statistics*, 7:2420-2456.
- Ghosh, A. and Basu, A. (2015). Robust estimation for non-homogeneous data and the selection of the optimal tuning parameter: the density power divergence approach. *Journal of Applied Statistics*, 42(9):2056–2072.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Goyal, P. and Ferrara, E. (2018). Graph Embedding Techniques, Applications, and Performance: A Survey. *Knowledge-Based Systems*, 151:78–94.
- Grover, A. and Leskovec, J. (2016). node2vec: Scalable Feature Learning for Networks. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 855–864. ACM.
- Hamilton, W. L., Ying, Z., and Leskovec, J. (2017). Inductive Representation Learning on Large Graphs. *Advances in Neural Information Processing Systems (NIPS)*, pages 1025–1035.
- He, X. and Niyogi, P. (2004). Locality Preserving Projections. In *Advances in Neural Information Processing Systems (NIPS)*, pages 153–160.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. (1983). Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137.
- Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, 28(3/4):321–377.
- Huber, P. J. (2011). Robust Statistics. In *International Encyclopedia of Statistical Science*, pages 1248–1251. Springer.
- Huber, P. J. et al. (1964). Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101.
- Jian, B. and Vemuri, B. C. (2005). A Robust Algorithm for Point Set Registration using Mixture of Gaussians. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 1246–1251. IEEE.
- Jones, M. C., Hjort, N. L., Harris, I. R., and Basu, A. (2001). A comparison of related density-based minimum divergence estimators. *Biometrika*, 88(3):865–873.
- Kanamori, T. and Fujisawa, H. (2014). Affine invariant divergences associated with proper composite scoring rules and their applications. *Bernoulli*, 20(4):2278–2304.
- Kanamori, T. and Fujisawa, H. (2015). Robust Estimation under Heavy Contamination using Unnormalized Models. *Biometrika*, 102(3):559–572.

- Kawashima, T. and Fujisawa, H. (2017). Robust and Sparse Regression via  $\gamma$ -divergence. *Entropy*, 19(11):608.
- Kawashima, T. and Fujisawa, H.. (2018). Robust and Sparse Regression in GLM by Stochastic Optimization. *arXiv preprint arXiv:1802.03127*.
- Kipf, T. N. and Welling, M. (2016). Variational Graph Auto-Encoders. *NIPS Workshop on Bayesian Deep Learning*.
- Kipf, T. N. and Welling, M. (2017). Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*. ACM.
- Kruskal, J. B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29(1):1–27.
- Manning, C. D., Raghavan, P., and Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed Representations of Words and Phrases and Their Compositionality. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3111–3119.
- Mollah, M. N. H., Eguchi, S., and Minami, M. (2007). Robust Prewhitening for ICA by Minimizing  $\beta$ -Divergence and Its Application to Fast ICA. *Neural Processing Letters*, 25(2):91–110.
- Moulines, E. and Bach, F. R. (2011). Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 451–459.
- Nori, N., Bollegala, D., and Kashima, H. (2012). Multinomial Relation Prediction in Social Data: A Dimension Reduction Approach. In *Proceedings of the conference on Artificial Intelligence (AAAI)*, volume 12, pages 115–121.
- Okuno, A., Hada, T., and Shimodaira, H. (2018). A probabilistic framework for multi-view feature learning with many-to-many associations via neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Okuno, A. and Shimodaira, H. (2019). Graph Embedding with Shifted Inner Product Similarity and Its Improved Approximation Capability. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- Pang, Y. and Yuan, Y. (2010). Outlier-resisting graph embedding. *Neurocomputing*, 73(4-6):968–974.
- Perozzi, B., Al-Rfou, R., and Skiena, S. (2014). DeepWalk: Online Learning of Social Representations. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 701–710. ACM.
- Robbins, H. and Monro, S. (1951). A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407.
- Roweis, S. T. and Saul, L. K. (2000). Nonlinear Dimensionality Reduction by Locally Linear Embedding. *Science*, 290(5500):2323–2326.
- Ruppert, David. (1985). A Newton-Raphson version of the multivariate Robbins-Monro procedure. *The Annals of Statistics*, 13(1):236–245.
- Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. (2008). Collective Classification in Network Data. *AI magazine*, 29(3):93.
- Shimodaira, H. (2016). Cross-validation of matching correlation analysis by resampling matching weights. *Neural Networks*, 75:126–140.
- Takenouchi, T., and Kanamori, T. (2017). Statistical inference with unnormalized discrete models and localized homogeneous divergences. *The Journal of Machine Learning Research*, 18(1):1804-1829.
- Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., and Mei, Q. (2015). LINE: Large-scale Information Network Embedding. In *Proceedings of the International Conference on World Wide Web (WWW)*, pages 1067–1077.
- Tenenbaum, J. B., De Silva, V., and Langford, J. C. (2000). A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323.
- Toulis, P., Airoldi, E. M., et al. (2017). Asymptotic and finite-sample properties of estimators based on stochastic gradients. *The Annals of Statistics*, 45(4):1694–1727.
- Wang, D., Cui, P., and Zhu, W. (2016). Structural Deep Network Embedding. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 1225–1234. ACM.
- Yan, S., Xu, D., Zhang, B., Zhang, H.-J., Yang, Q., and Lin, S. (2007). Graph Embedding and Extensions: A General Framework for Dimensionality Reduction. *IEEE transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(1):40–51.
- Zhang, H., Zha, Z.-J., Yang, Y., Yan, S., and Chua, T.-S. (2014). Robust (Semi-) Nonnegative Graph Embedding. *IEEE Transactions on Image Processing*, 23(7):2996–3012.