



Statistical Analysis of Multi-Relational Network Recovery

Zhi Wang¹, Xueying Tang^{2*} and Jingchen Liu¹

¹ Department of Statistics, Columbia University, New York, NY, United States, ² Department of Mathematics, University of Arizona, Tucson, AZ, United States

In this paper, we develop asymptotic theories for a class of latent variable models for large-scale multi-relational networks. In particular, we establish consistency results and asymptotic error bounds for the (penalized) maximum likelihood estimators when the size of the network tends to infinity. The basic technique is to develop a non-asymptotic error bound for the maximum likelihood estimators through large deviations analysis of random fields. We also show that these estimators are nearly optimal in terms of minimax risk.

Keywords: multi-relational network, knowledge graph completion, tail probability, risk, asymptotic analysis, non-asymptotic analysis, maximum likelihood estimation

OPEN ACCESS

Edited by:

Yajun Mei,
Georgia Institute of Technology,
United States

Reviewed by:

Xiangyu Chang,
Xi'an Jiaotong University, China
Alex Jung,
Aalto University, Finland

*Correspondence:

Xueying Tang
xytang@math.arizona.edu

Specialty section:

This article was submitted to
Mathematics of Computation and
Data Science,
a section of the journal
Frontiers in Applied Mathematics and
Statistics

Received: 04 March 2020

Accepted: 31 August 2020

Published: 19 October 2020

Citation:

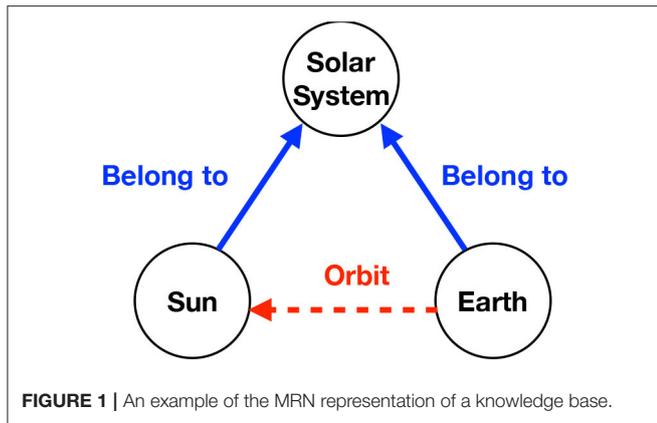
Wang Z, Tang X and Liu J (2020)
Statistical Analysis of Multi-Relational
Network Recovery.
Front. Appl. Math. Stat. 6:540225.
doi: 10.3389/fams.2020.540225

1. INTRODUCTION

A multi-relational network (MRN) describes multiple relations among a set of entities simultaneously. Our work on MRNs is mainly motivated by its applications to knowledge bases that are repositories of information. Examples of knowledge bases include WordNet [1], Unified Medical Language System [2], and Google Knowledge Graph (<https://developers.google.com/knowledge-graph>). They have been used as the information source in many natural language processing tasks, such as word-sense disambiguation and machine translation [3–5]. A knowledge base often includes knowledge on a large number of real-world objects or concepts. When a knowledge base is characterized by MRN, the objects and concepts correspond to nodes, and knowledge types are relations. **Figure 1** provides an excerpt from an MRN in which “Earth,” “Sun,” and “solar system” are three nodes. The knowledge about the orbiting patterns of celestial objects forms a relation “orbit,” and the knowledge on classification of the objects forms another relation “belong to” in the MRN.

An important task of network analysis is to recover the unobserved network based on data. In this paper, we consider a latent variable model for MRNs. The presence of an edge from node i to node j of relation type k is a Bernoulli random variable Y_{ijk} with success probability M_{ijk} . Each node is associated with a vector, θ , called the embedding of the node. The probability M_{ijk} is modeled as a function f of the embeddings, θ_i and θ_j , and a relation-specific parameter vector w_k . This is a natural generalization of the latent space model for single-relational networks [6]. Recently, it has been successfully applied to knowledge base analysis [7–14]. Various forms of f are proposed, such as distance models [7], bilinear models [12–14], and neural networks [15]. Computational algorithms are proposed to improve link prediction for knowledge bases [16, 17]. The statistical properties of the embedding-based MRN models have not been rigorously studied. It remains unknown whether and to what extent the underlying distribution of MRN can be recovered, especially when there are a large number of nodes and relations.

The results in this paper fill in the void by studying the error bounds and asymptotic behaviors of the estimators for M_{ijk} 's for a general class of models. This is a challenging problem due to the following facts. Traditional statistical inference of latent variable models often requires a (proper



or improper) prior distribution for θ_i . In such settings, one works with the marginalized likelihood with θ_i integrated out. For the analysis of MRN, the sample size and the latent dimensions are often so large that the above-mentioned inference approaches are computationally infeasible. For instance, a small-scale MRN could have a sample size as large as a few million, and the dimension of the embeddings is as large as several hundred. Therefore, in practice, the prior distribution is often dropped, and the latent variables θ_i 's are considered as additional parameters and estimated via maximizing the likelihood or penalized likelihood functions. The parameter space is thus substantially enlarged due to the addition of θ_i 's whose dimension is proportionate to the number of entities. As a result, in the asymptotic analysis, we face a double-asymptotic regime of both the sample size and the parameter dimension.

In this paper, we develop results for the (penalized) maximum likelihood estimator of such models and show that under regularity conditions the estimator is consistent. In particular, we overcome the difficulty induced by the double-asymptotic regime via non-asymptotic bounds for the error probabilities. Then, we show that the distribution of MRN can be consistently estimated in terms of average Kullback-Leibler (KL) divergence even when the latent dimension increases slowly as the sample size tends to infinity. A probability error bound is also provided together with the upper bound for the risk (expected KL divergence). We further study the lower bound and show the near-optimality of the estimator in terms of minimax risk. Besides the average KL divergence, similar results can be established for other criteria, such as link prediction accuracy.

The outline of the remaining sections is as follows. In section 2, we provide the model specification and formulate the problem. Our main results are presented in section 3. Finite sample performance is examined in section 4 through simulated and real data examples. Concluding remarks are included in section 5.

2. PROBLEM SETUP

2.1. Notation

Let $|\cdot|$ be the cardinality of a set and \times be the Cartesian product. Set $\{1, \dots, N\}$ is denoted by $[N]$. The sign function $\text{sgn}(x)$ is

defined to be 1 for $x \geq 0$ and 0 otherwise. The logistic function is denoted by $\sigma(x) = e^x / (1 + e^x)$. Let 1_A be the indicator function on event A . We use $U[a, b]$ to denote the uniform distribution on $[a, b]$ and $\text{Ber}(p)$ to denote the Bernoulli distribution with probability p . The KL divergence between $\text{Ber}(p)$ and $\text{Ber}(q)$ is written as $D(p||q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$. We use $\|\cdot\|$ to denote the Euclidean norm for vectors and the Frobenius norm for matrices.

For two real positive sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n = O(b_n)$ if $\limsup_{n \rightarrow \infty} a_n/b_n < \infty$. Similarly, we write $a_n = \Omega(b_n)$ if $\limsup_{n \rightarrow \infty} b_n/a_n < \infty$ and $a_n = o(b_n)$ if $\lim_{n \rightarrow \infty} a_n/b_n = 0$. We denote $a_n \lesssim b_n$ if $\limsup_{n \rightarrow \infty} a_n/b_n \leq 1$. When $\{a_n\}$ and $\{b_n\}$ are negative sequences, $a_n \lesssim b_n$ means $\liminf_{n \rightarrow \infty} a_n/b_n \geq 1$. In some places, we use $b_n \gtrsim a_n$ as an interchangeable notation of $a_n \lesssim b_n$. Finally, if $\lim_{n \rightarrow \infty} a_n/b_n = 1$, we write $a_n \sim b_n$.

2.2. Model

Consider an MRN with N entities and K relations. Given $i, j \in [N]$ and $k \in [K]$, the triple $\lambda = (i, j, k)$ corresponds to the edge from entity i to entity j of relation k . Let $\Lambda = [N] \times [N] \times [K]$ denote the set of all edges. We assume in this paper that an edge can be either present or absent in a network and use $Y_\lambda \in \{0, 1\}$ to indicate the presence of edge λ . In some scenarios, the status of an edge may have more than two types. Our analysis can be generalized to accommodate these cases.

We associate each entity i with a vector θ_i of dimension d_E and each relation k with a vector w_k of dimension d_R . Let $\mathcal{E} \subseteq \mathbb{R}^{d_E}$ be a compact domain where the embeddings $\theta_1, \dots, \theta_N$ live. We call \mathcal{E} the entity space. Similarly, we define a compact relation space $\mathcal{R} \subseteq \mathbb{R}^{d_R}$ for the relation-specific parameters w_1, \dots, w_K . Let $\mathbf{x} = (\theta_1, \dots, \theta_N, w_1, \dots, w_K)$ be a vector in the product space $\Theta = \mathcal{E}^N \times \mathcal{R}^K$. The parameters associated with edge $\lambda = (i, j, k)$ is then $\mathbf{x}_\lambda = (\theta_i, \theta_j, w_k)$. We assume that given \mathbf{x} , elements in $\{Y_\lambda \mid \lambda \in \Lambda\}$ are independent with each other and that the log odds of $Y_\lambda = 1$ is

$$\log \frac{P(Y_\lambda = 1 | \mathbf{x})}{P(Y_\lambda = 0 | \mathbf{x})} = \phi(\mathbf{x}_\lambda), \text{ for } \lambda \in \Lambda. \tag{1}$$

Here ϕ is defined on $\mathcal{E}^2 \times \mathcal{R}$, and $\phi(\mathbf{x}_\lambda)$ is often called the score of edge λ .

We will use Y to represent the $N \times N \times K$ tensor formed by $\{Y_\lambda \mid \lambda \in \Lambda\}$ and $M(\mathbf{x})$ to represent the corresponding probability tensor $\{P(Y_\lambda = 1 | \mathbf{x}) \mid \lambda \in \Lambda\}$. Our model is given by

$$Y_\lambda \sim \text{Ber}(M_\lambda(\mathbf{x}^*)), \tag{2}$$

$$M_\lambda(\mathbf{x}) = \sigma(\phi(\mathbf{x}_\lambda)), \lambda \in \Lambda, \tag{3}$$

where \mathbf{x}^* stands for the true value of \mathbf{x} and Y_λ 's are independent. In the above model, the probability of the presence of an edge is entirely determined by the embeddings of the corresponding entities and the relation-specific parameters. This imposes a low-dimensional latent structure on the probability tensor $M^* = M(\mathbf{x}^*)$.

We specify our model using a generic function ϕ . It includes various existing models as special cases. Below are two examples of ϕ .

1. Distance model [7].

$$\phi(\theta_i, \theta_j, \mathbf{w}_k) = b_k - \|\theta_i + \mathbf{a}_k - \theta_j\|^2, \quad (4)$$

where $\theta_i, \theta_j, \mathbf{a}_k \in \mathbb{R}^d$, $b_k \in \mathbb{R}$ and $\mathbf{w}_k = (\mathbf{a}_k, b_k)$. In the distance model, relation k from node i to node j is more likely to exist if θ_i shifted by \mathbf{a}_k is closer to θ_j under the Euclidean norm.

2. Bilinear model [9].

$$\phi(\theta_i, \theta_j, \mathbf{w}_k) = \theta_i^T \text{diag}(\mathbf{w}_k) \theta_j, \quad (5)$$

where $\theta_i, \theta_j, \mathbf{w}_k \in \mathbb{R}^d$ and $\text{diag}(\mathbf{w}_k)$ is a diagonal matrix with \mathbf{w}_k as the diagonal elements. Model (5) is a special case of the more general model $\phi(\theta_i, \theta_j, \mathbf{w}_k) = \theta_i^T W_k \theta_j$, where $W_k \in \mathbb{R}^{d \times d}$ is a matrix parameterized by $\mathbf{w}_k \in \mathbb{R}^{dR}$. Trouillon et al. [12], Nickel et al. [13], and Liu et al. [14] explored different ways of constructing W_k .

Very often, only a small portion of the network is observed [18]. We assume that each edge in the MRN is observed independently with probability γ and that the observation of an edge is independent of Y . Let $S \subset \Lambda$ be the set of observed edges. Then the elements in S are independent draws from Λ . For convenience, we use n to represent the expected number of observed edges, namely, $n = E[|S|] = \gamma|\Lambda| = \gamma N^2 K$. Our goal is to recover the underlying probability tensor M^* based on the observed edges $\{Y_\lambda \mid \lambda \in S\}$.

REMARK 1. Ideally, if there exists \mathbf{x}^* such that $Y_\lambda = \text{sgn}(M_\lambda(\mathbf{x}^*) - \frac{1}{2})$ for all $\lambda \in \Lambda$, then Y can be recovered with no error under \mathbf{x}^* . This is, however, a rare case in practice, especially for large-scale MRN. A relaxed assumption is that Y can be recovered with some low dimensional \mathbf{x}^* and noise $\{\epsilon_\lambda\}$ such that

$$Y_\lambda = \text{sgn}\left(M_\lambda(\mathbf{x}^*) + \epsilon_\lambda - \frac{1}{2}\right), \quad \epsilon_\lambda \stackrel{i.i.d}{\sim} U\left[-\frac{1}{2}, \frac{1}{2}\right], \quad \forall \lambda \in \Lambda. \quad (6)$$

By introducing the noise term, we formulate the deterministic MRN as a random graph. The model described in (2) is an equivalent but simpler form of (6).

2.3. Estimation

According to (2), the log-likelihood function of our model is

$$l(\mathbf{x}; Y_S) = \sum_{\lambda \in S} Y_\lambda \log M_\lambda(\mathbf{x}) + (1 - Y_\lambda) \log(1 - M_\lambda(\mathbf{x})). \quad (7)$$

We omit the terms $\sum_{\lambda \in S} \log \gamma + \sum_{\lambda \notin S} \log(1 - \gamma)$ in (7) since γ is not the parameter of interest. To obtain an estimator of M^* , we take the following steps.

1. Obtain the maximum likelihood estimator (MLE) of \mathbf{x}^* ,

$$\hat{\mathbf{x}} = \underset{\mathbf{x} \in \Theta}{\text{argmax}} l(\mathbf{x}; Y_S). \quad (8)$$

2. Use the plug-in estimator

$$\hat{M} = M(\hat{\mathbf{x}}) \quad (9)$$

as an estimator of M^* .

In (8), the estimator $\hat{\mathbf{x}}$ is a maximizer over the compact parameter space $\Theta = \mathcal{E}^N \times \mathcal{R}^K$. The dimension of Θ is

$$m = Nd_E + Kd_R,$$

which grows linearly in the number of entities N and the number of relations K .

2.4. Evaluation Criteria

We consider the following criteria to measure the error of the above-mentioned estimator. They will be used in both the main results and numerical studies.

1. Average KL divergence of the predictive distribution from the true distribution

$$L(\hat{M}, M^*) = \frac{1}{|\Lambda|} \sum_{\lambda \in \Lambda} D(M_\lambda^* \parallel \hat{M}_\lambda). \quad (10)$$

2. Mean squared error of the predicted scores

$$MSE_\phi = \frac{1}{|\Lambda|} \sum_{\lambda \in \Lambda} (\phi(\hat{\mathbf{x}}_\lambda) - \phi(\mathbf{x}_\lambda^*))^2. \quad (11)$$

3. Link prediction error

$$\widehat{err} = \frac{1}{|\Lambda|} \sum_{\lambda \in \Lambda} 1_{\hat{Y}_\lambda \neq Y_\lambda^*}, \quad (12)$$

where $\hat{Y}_\lambda = \text{sgn}(\hat{M}_\lambda - \frac{1}{2})$ and $Y_\lambda^* = \text{sgn}(M_\lambda^* - \frac{1}{2})$.

REMARK 2. The latent attributes of entities and relations are often not identifiable, so the MLE $\hat{\mathbf{x}}$ is not unique. For instance, in (4), the values of ϕ and $M(\mathbf{x})$ remain the same if we replace θ_i and \mathbf{a}_k , respectively by $\Gamma\theta_i + \mathbf{t}$ and $\Gamma\mathbf{a}_k$, where \mathbf{t} is an arbitrary vector in \mathbb{R}^{dE} and Γ is an orthonormal matrix. Therefore, we consider the mean squared error of scores, which are identifiable.

3. MAIN RESULTS

We first provide results of the MLE in terms of KL divergence between the estimated and the true model. Specifically, we investigate the tail probability $P(L(\hat{M}, M^*) > t)$ and the expected loss $E[L(\hat{M}, M^*)]$. In section 3.1, we discuss upper bounds for the two quantities. The lower bounds are provided in section 3.2. In section 3.3, we extend the results to penalized maximum likelihood estimators (pMLE) and other loss functions. All proofs are included in the **Supplementary Materials**.

3.1. Upper Bounds

We first present an upper bound for the tail probability $P(L(\hat{M}, M^*) > t)$ in Lemma 1. The result depends on the tensor size, the number of observed edges, the functional form of ϕ , and the geometry of parameter space Θ . The lemma explicitly quantifies the impact of these element on the error probability. It is key to the subsequent analyses. Lemma 2 gives a non-asymptotic upper bound for the expected loss (risk). We then establish the consistency of \hat{M} and the asymptotic error bounds in Theorem 1.

We will make the following assumptions throughout this section.

ASSUMPTION 1. $\mathbf{x}^* \in \Theta = \mathcal{E}^N \times \mathcal{R}^K$, where \mathcal{E} and \mathcal{R} are Euclidean balls of radius U .

ASSUMPTION 2. The function ϕ is Lipschitz continuous under the Euclidean norm,

$$|\phi(\mathbf{u}) - \phi(\mathbf{v})| \leq \alpha \|\mathbf{u} - \mathbf{v}\|, \quad \forall \mathbf{u}, \mathbf{v} \in \mathcal{E}^2 \times \mathcal{R}, \quad (13)$$

where α is a Lipschitz constant.

Assumption 1 is imposed for technical convenience. The results can be easily extended to general compact parameter spaces. Let $C = \sup_{\mathbf{u} \in \mathcal{E}^2 \times \mathcal{R}} |\phi(\mathbf{u})|$. Without loss of generality, we assume that $C \geq 2$.

LEMMA 1. Consider \hat{M} defined in (9) and the average KL divergence L in (10). Under Assumptions 1 and 2, for every $t > 0$, $\beta > 0$ and $0 < s < nt$,

$$P(L(\hat{M}, M^*) \geq t) \leq \exp \left\{ -\frac{nt-s}{C} h \left(\frac{1}{2} - \frac{s}{2nt} \right) \right\} \left(1 + \frac{2\sqrt{3}\alpha U n(1+\beta)}{s} \right)^m + \exp \{-n\beta h(\beta)\}, \quad (14)$$

where $m = Nd_E + Kd_R$ is the dimension of Θ , $n = \gamma N^2 K$ is the expected number of observations, and $h(u) = (1 + \frac{1}{u}) \log(1+u) - 1$.

In the proof of Lemma 1, we use Bennett’s inequality to develop a uniform bound that does not depend on the true parameters. It is sufficient for the current analysis. If the readers need sharper bounds, they can read through the proof and replace the Bennett’s bound by the usual large deviation rate function which provides a sharp exponential bound that depends on the true parameters. We don’t pursue this direction in this paper.

Lemma 2 below gives an upper bound of risk $E[L(\hat{M}, M^*)]$, which follows from Lemma 1.

LEMMA 2. Consider \hat{M} defined in (9) and loss function L in (10). Let $C_1 = 18C$, $C_2 = 8\sqrt{3}\alpha U$ and $C_3 = 2 \max \{C_1, C_2\}$. If Assumptions 1 and 2 hold and $\frac{n}{m} \geq C_2 + e$, then

$$E[L(\hat{M}, M^*)] \leq C_3 \frac{m}{n} \log \frac{n}{m} + \frac{C_1}{n} \exp \left\{ -m \log \frac{n}{m} \right\} + \frac{3}{n} \exp \left\{ -\frac{1}{3} \left(n + C_3 m \log \frac{n}{m} \right) \right\}. \quad (15)$$

We are interested in the asymptotic behavior of the tail probability in two scenarios: (i) t is a fixed constant and (ii) t decays to zero as the number of entities N tends to infinity. The following theorem gives an asymptotic upper bound for the tail probability and the risk.

THEOREM 1. Consider \hat{M} defined in (9) and the loss function L in (10). Let the number of entities $N \rightarrow \infty$ and $C, K, U, d_E, d_R, \alpha$, and γ be fixed constants. If Assumptions 1 and 2 hold, we have the following asymptotic inequalities.

When t is a fixed constant,

$$\log P(L(\hat{M}, M^*) \geq t) \lesssim -\frac{t}{5C} n. \quad (16)$$

When $t = 10C \frac{m}{n} \log \frac{n}{m}$,

$$\log P(L(\hat{M}, M^*) \geq t) \lesssim -m \log \frac{n}{m}. \quad (17)$$

Furthermore,

$$E[L(\hat{M}, M^*)] \lesssim 10C \frac{m}{n} \log \frac{n}{m}. \quad (18)$$

The consistency of \hat{M} is implied by (16) and the rate of convergence is $|\log P(L(\hat{M}, M^*) \geq t)| = \Omega(N^2)$ if t is a fixed constant. The rate decreases to $\Omega(N \log N)$ for the choice of t producing (17). It is also implied by (17) that $L(\hat{M}, M^*) = O(\frac{1}{N} \log N)$ with high probability. We show in the next section that this upper bound is reasonably sharp.

The condition that K, U, d_E, d_R , and α are fixed constants can be relaxed. For instance, we can let U, d_E, d_R , and α go to infinity slowly at the rate $O(\log N)$ and K at the rate $O(N)$. We can let γ go to zero provided that $\frac{m}{n} \log \frac{n}{m} = o(1)$.

3.2. Lower Bounds

We show in Theorem 2 that the order of the minimax risk is $\Omega(\frac{m}{n})$, which implies the near optimality of \hat{M} in (9) and the upper bound $O(\frac{m}{n} \log \frac{n}{m})$ in Theorem 1. To begin with, we introduce the following definition and assumption.

DEFINITION 1. For $\mathbf{u} = (\boldsymbol{\theta}, \boldsymbol{\theta}', \mathbf{w}) \in \mathcal{E}^2 \times \mathcal{R}$, the r -neighborhood of \mathbf{u} is

$$\mathcal{N}_r(\mathbf{u}) = \{(\boldsymbol{\eta}, \boldsymbol{\eta}', \boldsymbol{\zeta}) \in \mathcal{E}^2 \times \mathcal{R} \mid \|\boldsymbol{\eta} - \boldsymbol{\theta}\| \leq r, \|\boldsymbol{\eta}' - \boldsymbol{\theta}'\| \leq r, \|\boldsymbol{\zeta} - \mathbf{w}\| \leq r\}.$$

Similarly, for $\mathbf{x} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{w}_1, \dots, \mathbf{w}_K) \in \mathcal{E}^N \times \mathcal{R}^K$, the r -neighborhood of \mathbf{x} is

$$\mathcal{N}_r(\mathbf{x}) = \{(\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_N, \boldsymbol{\zeta}_1, \dots, \boldsymbol{\zeta}_K) \in \mathcal{E}^N \times \mathcal{R}^K \mid \|\boldsymbol{\eta}_i - \boldsymbol{\theta}_i\| \leq r, \|\boldsymbol{\zeta}_k - \mathbf{w}_k\| \leq r, \forall i \in [N], k \in [K]\}.$$

ASSUMPTION 3. There exists $\mathbf{u}_0 \in \mathcal{E}^2 \times \mathcal{R}$ and $r, \kappa > 0$ such that $\mathcal{N}_r(\mathbf{u}_0) \subset \mathcal{E}^2 \times \mathcal{R}$ and

$$|\sigma(\phi(\mathbf{u})) - \sigma(\phi(\mathbf{v}))| \geq \kappa \|\mathbf{u} - \mathbf{v}\|, \quad \forall \mathbf{u}, \mathbf{v} \in \mathcal{N}_r(\mathbf{u}_0). \quad (19)$$

THEOREM 2. Let $b = \sup_{\mathbf{u} \in \mathcal{N}_r(\mathbf{u}_0)} \sigma(\phi(\mathbf{u}))$. Under Assumptions 2 and 3, if $\tau^2 \geq \frac{(m/16-1)b(1-b)}{12\alpha^2 n}$, then for any estimator \hat{M} , there exists $\mathbf{x}^* \in \Theta$ such that

$$P\left(L(\hat{M}, M^*) > \tilde{C} \frac{m/16-1}{n}\right) \geq \frac{1}{2}, \quad (20)$$

where $\tilde{C} = \frac{\kappa^2 b(1-b)}{108\alpha^2}$. Consequently, the minimax risk

$$\min_{\hat{M}} \max_{M^*} E[L(\hat{M}, M^*)] \geq \tilde{C} \frac{m/16-1}{2n}. \quad (21)$$

3.3. Extensions

3.3.1. Regularization

In this section, we extend our asymptotic results in Theorem 1 to regularized estimators. In practice, regularization is often considered to prevent overfitting. We consider a regularization similar to elastic net [19]

$$l_\rho(\mathbf{x}; Y_S) = l(\mathbf{x}; Y_S) - \rho_1 \|\mathbf{x}\|_1 - \rho_2 \|\mathbf{x}\|^2, \quad (22)$$

where $\|\cdot\|_1$ stands for L_1 norm and $\rho_1, \rho_2 \geq 0$ are regularization parameters. The pMLE is

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in \Theta} l_\rho(\mathbf{x}; Y_S). \quad (23)$$

Note that the MLE in (8) is a special case of the pMLE above with $\rho_1 = \rho_2 = 0$. Since $\hat{\mathbf{x}}$ is shrunk toward $\mathbf{0}$, without loss of generality, we assume that \mathcal{E} and \mathcal{R} are centered at $\mathbf{0}$. We generalize Theorem 1 to pMLE in the following theorem.

THEOREM 3. Consider the estimator \hat{M} given by (23) and (9) and the loss function L in (10). Let the number of entities $N \rightarrow \infty$ and $C, K, U, d_E, d_R, \alpha, \gamma$ be absolute constants. If Assumptions 1 and 2 hold and $\rho_1 + \rho_2 = o(\log N)$, then asymptotic inequalities (16), (17), and (18) in Theorem 1 hold.

3.3.2. Other Loss Functions

We present some results for the mean squared error loss MSE_ϕ defined in (11) and the link prediction error \widehat{err} defined in (12). Corollaries 1 and 2 give upper and lower bounds for MSE_ϕ , and Corollary 3 gives an upper bound for \widehat{err} under an additional assumption.

COROLLARY 1. Under the setting of Theorem 3 with the loss function replaced by MSE_ϕ , we have the following asymptotic results.

If t is a fixed constant,

$$\log P(MSE_\phi \geq t) \lesssim -\frac{5\sigma(C)(1-\sigma(C))t}{2C} n. \quad (24)$$

If $t = \frac{20C}{\sigma(C)(1-\sigma(C))} \frac{m}{n} \log \frac{n}{m}$,

$$\log P(MSE_\phi \geq t) \lesssim -m \log \frac{n}{m}. \quad (25)$$

Furthermore,

$$E[MSE_\phi] \lesssim \frac{20C}{\sigma(C)(1-\sigma(C))} \frac{m}{n} \log \frac{n}{m}. \quad (26)$$

COROLLARY 2. Under the setting of Theorem 2 with the loss function replaced by MSE_ϕ , we have

$$P\left(MSE_\phi > \tilde{C} \frac{m/16-1}{8n}\right) \geq \frac{1}{2}, \quad (27)$$

and

$$\min_{\hat{M}} \max_{M^*} E[MSE_\phi] \geq \tilde{C} \frac{m/16-1}{16n}. \quad (28)$$

ASSUMPTION 4. There exists $\varepsilon > 0$ such that $|M_\lambda^* - \frac{1}{2}| \geq \varepsilon$ for every $\lambda \in \Lambda$.

COROLLARY 3. Under the setting of Theorem 3 with the loss function replaced by \widehat{err} and Assumption 4 added, we have the following asymptotic results.

If t is a fixed constant,

$$\log P(\widehat{err} \geq t) \lesssim -\frac{2\varepsilon^2 t}{5C} n. \quad (29)$$

If $t = \frac{5C}{\varepsilon^2} \frac{m}{n} \log \frac{n}{m}$,

$$\log P(\widehat{err} \geq t) \lesssim -m \log \frac{n}{m}. \quad (30)$$

Furthermore,

$$E[\widehat{err}] \lesssim \frac{5C}{\varepsilon^2} \frac{m}{n} \log \frac{n}{m}. \quad (31)$$

3.3.3. Sparse Representations

We are interested in sparse entity embeddings and relation parameters. Let $\|\cdot\|_0$ be the number of non-zero elements of a vector and τ be a pre-specified sparsity level of \mathbf{x} (i.e., the proportion of non-zero elements). Let $m_\tau = m\tau$ be the upper bound of non-zero parameters, that is, $\|\mathbf{x}^*\|_0 \leq m_\tau$. Consider the following estimator

$$\hat{\mathbf{x}} = \operatorname{argmax}_{\mathbf{x} \in \Theta} l(\mathbf{x}; Y_S) \quad \text{subject to} \quad \|\mathbf{x}\|_0 \leq m_\tau. \quad (32)$$

The estimator defined above maximizes the L_0 -penalized log-likelihood.

THEOREM 4. Consider \hat{M} defined in (32) and (9) and the loss function L in (10). Let the number of entities $N \rightarrow \infty$ and $\tau, C, K, U, d_E, d_R, \alpha$ be absolute constants. Under Assumptions 1 and 2, the following asymptotic inequalities hold.

If t is a fixed constant,

$$\log P(L(\hat{M}, M^*) \geq t) \lesssim -\frac{t}{5C} n. \quad (33)$$

If $t = 10C \frac{m_\tau}{n} \log \frac{n}{m_\tau}$,

$$\log P(L(\hat{M}, M^*) \geq t) \lesssim -m_\tau \log \frac{n}{m_\tau}. \quad (34)$$

Furthermore,

$$E[L(\hat{M}, M^*)] \lesssim 10C \frac{m_\tau}{n} \log \frac{n}{m_\tau}. \quad (35)$$

We omit the results for other loss functions as well as the lower bounds since they can be analogously obtained.

4. NUMERICAL EXAMPLES

In this section, we demonstrate the finite sample performance of \hat{M} through simulated and real data examples. Throughout the numerical experiments, AdaGrad algorithm [20] is used to compute \hat{x} in (8) or (23). It is a first-order optimization method that combines stochastic gradient descent (SGD) [21] with adaptive step sizes for finding the local optima. Since the objective function in (8) is non-convex, a global maximizer is not guaranteed. Our objective function usually has many global maximizers, but, empirically, we found the algorithm works well on MRN recovery and the recovery performance is insensitive to the choice of the starting point of SGD. Computationally, SGD is also more appealing to handle large-scale MRNs than those more expensive global optimization methods.

4.1. Simulated Examples

In the simulated examples, we fix $K = 20$, $d = 20$ and consider various choices of N ranging from 100 to 10,000 to investigate the estimation performance as N grows. The function ϕ we consider is a combination of the distance model (4) and the bilinear model (5),

$$\phi(\theta_i, \theta_j, \mathbf{w}_k) = (\theta_i + \mathbf{a}_k - \theta_j)^T \text{diag}(\mathbf{b}_k) (\theta_i + \mathbf{a}_k - \theta_j), \quad (36)$$

where $\theta_i, \theta_j, \mathbf{a}_k, \mathbf{b}_k \in \mathbb{R}^d$ and $\mathbf{w}_k = (\mathbf{a}_k, \mathbf{b}_k)$. We independently generate the elements of θ_i^* , \mathbf{a}_k^* , and \mathbf{b}_k^* from normal distributions $N(0, 1)$, $N(0, 1)$, and $N(0, 0.25)$, respectively, where $N(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 . To guarantee that the parameters are from a compact set, the normal distributions are truncated to the interval $[-20, 20]$. Given the latent attributes, each Y_{ijk} is generated from the Bernoulli distribution with success probability $M_{ijk}^* = \sigma(\phi(\theta_i^*, \theta_j^*, \mathbf{w}_k^*))$. The observation probability γ takes value from $\{0.005, 0.01, 0.02\}$. For each combination of γ and N , 100 independent datasets are generated. For each dataset, we compute \hat{x} and \hat{M} in (8) and (9) with AdaGrad algorithm and then calculate $L(\hat{M}, M^*)$ defined in (10) as well as the link prediction error \widehat{err} defined in (12). The two types of losses are averaged over the 100 datasets for each combination of N and γ to approximate the theoretical risks $E[L(\hat{M}, M^*)]$ and $E[\widehat{err}]$. These quantities are plotted against N in log scale in **Figure 2**. As the figure shows, in general, both risks decrease as N increases. When N is small, n/m is not large enough to satisfy the condition $n/m \geq C_2 + e$ in Lemma 2 and the expected KL risk increases at the beginning. After N gets sufficiently large, the trend agrees with our asymptotic analysis.

4.2. Real Data Example: Knowledge Base Completion

WordNet [1] is a large lexical knowledge base for English. It has been used in word sense disambiguation, text classification, question answering, and many other tasks in natural language processing [3, 5]. The basic components of WordNet are groups of words. Each group, called a synset, describes a distinct concept. In WordNet, synsets are linked by conceptual-semantic and lexical relations, such as super-subordinate relation and

antonym. We model WordNet as an MRN with the synsets as entities and the links between synsets as relations.

Following Bordes et al. [7], we use a subset of WordNet for analysis. The dataset contains 40,943 synsets and 18 types of relations. A triple (i, j, k) is called valid if relation k from entity i to entity j exists, i.e., $Y_{ijk} = 1$. All the other triples are called invalid triples. Among more than 3.0×10^{10} possible triples in WordNet, only 151,442 triples are valid. We assume that 141,442 valid triples and the same proportion of invalid triples are observed. The goal of our analysis is to recover the unobserved part of the knowledge base. We adopt the ranking procedure, which is commonly used in knowledge graph embedding literature, to evaluate link predictions. Given a valid triple $\lambda = (i, j, k)$, we rank estimated scores for all the invalid triples inside $\Lambda_{jk} = \{(i', j, k) \mid i' \in [N]\}$ in descending order and call the rank of $\phi(\hat{x}_\lambda)$ as the head rank of λ , denoted by H_λ . Similarly, we can define the tail rank T_λ and the relation rank R_λ by ranking $\phi(\hat{x}_\lambda)$ among the estimated scores of invalid triples in Λ_{ij} and $\Lambda_{i,k}$, respectively. For a set V of valid triples, the prediction performance can be evaluated by rank-based criteria, mean rank (MR), mean reciprocal rank (MRR), and hits at q (Hits@ q), which are defined as

$$\text{MR}_E = \frac{1}{2|V|} \sum_{\lambda \in V} H_\lambda + T_\lambda, \text{MR}_R = \frac{1}{|V|} \sum_{\lambda \in V} R_\lambda,$$

$$\text{MRR}_E = \frac{1}{2|V|} \sum_{\lambda \in V} \frac{1}{H_\lambda} + \frac{1}{T_\lambda}, \text{MRR}_R = \frac{1}{|V|} \sum_{\lambda \in V} \frac{1}{R_\lambda},$$

and

$$\text{Hits}_E@q = \frac{1}{2|V|} \sum_{\lambda \in V} \mathbf{1}_{\{H_\lambda \leq q\}} + \mathbf{1}_{\{T_\lambda \leq q\}},$$

$$\text{Hits}_R@q = \frac{1}{|V|} \sum_{\lambda \in V} \mathbf{1}_{\{R_\lambda \leq q\}}.$$

The subscripts E and R represent the criteria for predicting entities and relations, respectively. Models with higher MRRs, Hits@ q 's or lower MRs are more preferable. In addition, MRR is more robust to outliers than MR.

The three models described in (4), (5), and (36) are considered in our data analysis and we refer to them as Model 1, 2, and 3, respectively. For each model, the latent dimension d takes value from $\{50, 100, 150, 200, 250\}$. Due to the high dimensionality of the parameter space, L_2 penalized MLE is used to obtain the estimated latent attributes \hat{x} , with tuning parameters $\rho_1 = 0$ and ρ_2 chosen from $\{0, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$ in (22). Since information criteria based dimension and tuning parameter selection is computationally intensive for dataset of this scale, we set aside 5,000 of the unobserved valid triples as a validation set and select the d and ρ_2 that produce the smallest MRR_E on this validation set. The model with the selected d and ρ_2 is then evaluated on the test set consisting of the rest 5,000 unobserved valid triples.

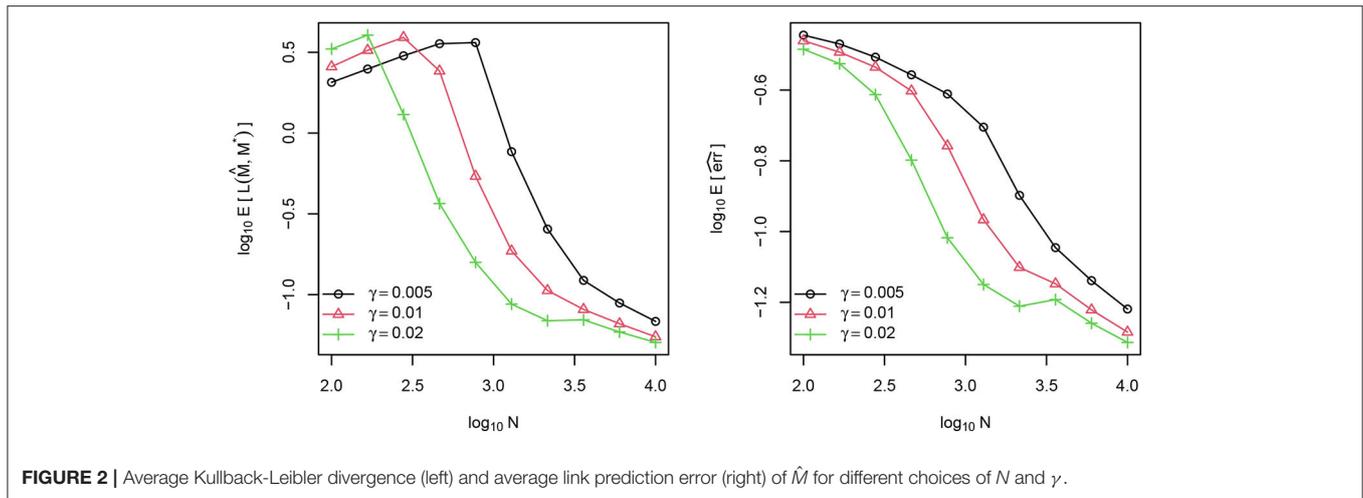


FIGURE 2 | Average Kullback-Leibler divergence (left) and average link prediction error (right) of \hat{M} for different choices of N and γ .

TABLE 1 | Results for WordNet data analysis.

Method	(d, ρ_2)	MR_E	MRR_E	$Hits_E@10$	MR_R	MRR_R	$Hits_R@1$
Model 1	$(100, 10^{-5})$	385	0.64	0.888	1.41	0.896	0.817
Model 2	$(250, 10^{-4})$	769	0.94	0.945	1.31	0.968	0.959
Model 3	$(200, 10^{-4})$	499	0.94	0.947	1.13	0.978	0.967
CP	–	–	0.075	0.125	–	–	–
RESCAL	–	–	0.890	0.928	–	–	–

The results for CP and RESCAL are extracted from Trouillon et al. [12] and Nickel et al. [13].

The computed evaluation criteria on the test set are listed in Table 1. The table also includes the selected d and ρ_2 for each of the three score models. Models 2 and 3 generate similar performance. The MRRs for the two models are very close to 1, and the Hits@ q 's are higher than 90%, suggesting that the two models can identify the valid triples very well. Although Model 1 is inferior to the other two models in terms of most of the criteria, it outperforms them in MR_E . The results imply that Model 2 and Model 3 could perform extremely bad for a few triples.

In addition to Models 1–3, we also display the performance of the Canonical Polyadic (CP) decomposition [22] and a tensor factorization approach, RESCAL [23]. Their MRR_E and $Hits_E@10$ results on the WordNet dataset are extracted from Trouillon et al. [12] and Nickel et al. [13], respectively. Both methods, especially CP, are outperformed by Model 3.

5. CONCLUDING REMARKS

In this article, we focused on the recovery of large-scale MRNs with a small portion of observations. We studied a generalized latent space model where entities and relations are associated with latent attribute vectors and conducted statistical analysis on the error of recovery. MLEs and pMLEs over a compact space are considered to estimate the latent attributes and the edge probabilities. We established non-asymptotic upper bounds for estimation error in terms of tail probability and risk, based on which we then studied the asymptotic properties when the size

of MRN and latent dimension go to infinity simultaneously. A matching lower bound up to a log factor is also provided.

We kept ϕ generic for theoretical development. The choice of ϕ is usually problem-specific in practice. How to develop a data-driven method for selecting an appropriate ϕ is an interesting problem to investigate in future works.

Besides the latent space models, sparsity [24] or clustering assumptions [25] have been used to impose low-dimensional structures in single-relational networks. An MRN can be seen as a combination of several heterogeneous single-relational networks. The distribution of edges may vary dramatically across relations. Therefore, it is challenging to impose appropriate sparsity or cluster structures on MRNs. More empirical and theoretical studies are needed to quantify the impact of heterogeneous relations and to incorporate the information for recovering MRNs.

DATA AVAILABILITY STATEMENT

The datasets generated for this study are available on request to the corresponding author.

AUTHOR CONTRIBUTIONS

ZW provided the theoretical proofs. ZW and XT conducted the simulation and real data analysis. All authors contributed to manuscript writing and revision.

FUNDING

This work was supported by National Science Foundation grants SES-1826540 and IIS-1633360.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fams.2020.540225/full#supplementary-material>

REFERENCES

1. Miller GA. WordNet: a lexical database for English. *Commun ACM*. (1995) **38**:39–41. doi: 10.1145/219717.219748
2. McCray AT. An upper-level ontology for the biomedical domain. *Compar Funct Genomics*. (2003) **4**:80–4. doi: 10.1002/cfg.255
3. Gabrilovich E, Markovitch S. Wikipedia-based semantic interpretation for natural language processing. *J Artif Intell Res*. (2009) **34**:443–98. doi: 10.1613/jair.2669
4. Scott S, Matwin S. Feature engineering for text classification. In: *ICML Vol. 99 Bled* (1999). p. 379–88.
5. Ferrucci D, Brown E, Chu-Carroll J, Fan J, Gondek D, Kalyanpur AA, et al. Building Watson: an overview of the DeepQA project. *AI Mag*. (2010) **31**:59–79. doi: 10.1609/aimag.v31i3.2303
6. Hoff PD, Raftery AE, Handcock MS. Latent space approaches to social network analysis. *J Am Stat Assoc*. (2002) **97**:1090–8. doi: 10.1198/016214502388618906
7. Bordes A, Usunier N, Garcia-Duran A, Weston J, Yakhnenko O. Translating embeddings for modeling multi-relational data. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, editors. *Advances in Neural Information Processing Systems 26*. Curran Associates, Inc. (2013). p. 2787–95. Available online at: <http://papers.nips.cc/paper/5071-translating-embeddings-for-modeling-multi-relational-data.pdf>
8. Wang Z, Zhang J, Feng J, Chen Z. Knowledge graph embedding by translating on hyperplanes. In: *AAAI*. Québec City, QC (2014). p. 1112–9.
9. Yang B, Yih SWt, He X, Gao J, Deng L. Embedding entities and relations for learning and inference in knowledge bases. In: *Proceedings of the International Conference on Learning Representations*. San Diego, CA (2015).
10. Lin Y, Liu Z, Sun M, Liu Y, Zhu X. Learning entity and relation embeddings for knowledge graph completion. In: *AAAI*. Austin, TX (2015). p. 2181–7.
11. Garcia-Duran A, Bordes A, Usunier N, Grandvalet Y. Combining two and three-way embedding models for link prediction in knowledge bases. *J Artif Intell Res*. (2016) **55**:715–42. doi: 10.1613/jair.5013
12. Trouillon T, Welbl J, Riedel S, Gaussier É, Bouchard G. Complex embeddings for simple link prediction. In: *International Conference on Machine Learning*. (2016). p. 2071–80.
13. Nickel M, Rosasco L, Poggio TA. Holographic embeddings of knowledge graphs. In: *AAAI*. Phoenix, AZ (2016). p. 1955–61.
14. Liu H, Wu Y, Yang Y. Analogical inference for multi-relational embeddings. In: *Proceedings of the 34th International Conference on Machine Learning*. Vol. **70**. Sydney, NSW (2017). p. 2168–78.
15. Socher R, Chen D, Manning CD, Ng A. Reasoning with neural tensor networks for knowledge base completion. In: Burges CJC, Bottou L, Welling M, Ghahramani Z, Weinberger KQ, editors. *Advances in Neural Information Processing Systems*. Lake Tahoe, NV: Curran Associates, Inc. (2013) p. 926–34.
16. Kotnis B, Nastase V. Analysis of the impact of negative sampling on link prediction in knowledge graphs. *arXiv*. (2017) 170806816.
17. Kanojia V, Maeda H, Togashi R, Fujita S. Enhancing knowledge graph embedding with probabilistic negative sampling. In: *Proceedings of the 26th International Conference on World Wide Web Companion*. Perth, WA (2017) p. 801–2. doi: 10.1145/3041021.3054238
18. Min B, Grishman R, Wan L, Wang C, Gondek D. Distant supervision for relation extraction with an incomplete knowledge base. In: *HLT-NAACL*. Atlanta, GA (2013). p. 777–82.
19. Zou H, Hastie T. Regularization and variable selection via the elastic net. *J R Stat Soc B*. (2005) **67**:301–20. doi: 10.1111/j.1467-9868.2005.00503.x
20. Duchi J, Hazan E, Singer Y. Adaptive subgradient methods for online learning and stochastic optimization. *J Mach Learn Res*. (2011) **12**:2121–59. doi: 10.5555/1953048.2021068
21. Robbins H, Monro S. A stochastic approximation method. *Ann Math Stat*. (1951) **22**:400–7. doi: 10.1214/aoms/1177729586
22. Hitchcock FL. The expression of a tensor or a polyadic as a sum of products. *J Math Phys*. (1927) **6**:164–89. doi: 10.1002/sapm192761164
23. Nickel M, Tresp V, Krieger HP. A three-way model for collective learning on multi-relational data. In: *Proceedings of the 28th International Conference on International Conference on Machine Learning*. *ICML'11*. Madison, WI: Omnipress (2011). p. 809–16.
24. Tran N, Abramenko O, Jung A. On the sample complexity of graphical model selection from non-stationary samples. *IEEE Trans Signal Process*. (2020) **68**:17–32. doi: 10.1109/TSP.2019.2956687
25. Jung A, Hero AO III, Mara AC, Jahromi S, Heimowitz A, Eldar YC, et al. Semi-supervised learning in network-structured data via total variation minimization. *IEEE Trans Signal Process*. (2019) **67**:6256–69. doi: 10.1109/TSP.2019.2953593

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Wang, Tang and Liu. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.