
Comprehensive Analysis of Negative Sampling in Knowledge Graph Representation Learning

Hidetaka Kamigaito¹ Katsuhiko Hayashi²

Abstract

Negative sampling (NS) loss plays an important role in learning knowledge graph embedding (KGE) to handle a huge number of entities. However, the performance of KGE degrades without hyperparameters such as the margin term and number of negative samples in NS loss being appropriately selected. Currently, empirical hyperparameter tuning addresses this problem at the cost of computational time. To solve this problem, we theoretically analyzed NS loss to assist hyperparameter tuning and understand the better use of the NS loss in KGE learning. Our theoretical analysis showed that scoring methods with restricted value ranges, such as TransE and RotatE, require appropriate adjustment of the margin term or the number of negative samples different from those without restricted value ranges, such as RESCAL, ComplEx, and DistMult. We also propose subsampling methods specialized for the NS loss in KGE studied from a theoretical aspect. Our empirical analysis on the FB15k-237, WN18RR, and YAGO3-10 datasets showed that the results of actually trained models agree with our theoretical findings.

1. Introduction

Knowledge graph completion (KGC) has attracted increasing attention as a research topic for automatically inferring new links in a KG that are likely but not yet known to be true. KG embedding (KGE) methods have been actively pursued as a scalable approach to KGC (Bordes et al., 2013; Yang et al., 2015; Trouillon et al., 2016; Sun et al., 2019). Since a KG contains a huge number of entities, negative sampling (NS) loss (Mikolov et al., 2013b) is widely used

in KGE training from the viewpoint of computational efficiency. However, NS loss was originally proposed for word-representation learning (Mikolov et al., 2013b), and there are several issues that need to be considered when using NS loss for training KGE models.

While the original NS loss was designed to be used with scoring methods on the basis of the inner product, which have no limit on the value range, many KGE models such as TransE (Bordes et al., 2013) and RotatE (Sun et al., 2019) use score functions with value ranges limited by the p -norm distance. To apply NS loss to distance-based KGE models, a margin term has been introduced as a heuristic means of adjusting the value ranges of the score functions (Bordes et al., 2013; Sun et al., 2019). This margin term is not present in the original NS loss, and there has been no discussion of how it affects the KGE model training. In certain KGE studies (Trouillon et al., 2016; Sun et al., 2019), the NS loss is normalized by the number of negative samples. This normalization term also does not exist in the original NS loss, and the validity of the term has not yet been verified.

In KG representation learning, self-adversarial NS (SANS) (Sun et al., 2019) is often used to reduce the effect of negative samples that do not carry meaningful information to the learning process. To generate negative samples, SANS uses probability distributions calculated from the current embedding model instead of using the uniform noise distribution. However, the normalization term of the probability distribution is computationally expensive and is approximated in practice by randomly sampled examples. Though certain characteristics of SANS have been investigated (Ruffinelli et al., 2020; Kamigaito & Hayashi, 2021), more detailed analysis is needed because the previous studies did not take into account the effect of the approximation of the normalization term.

Due to the difference from the original NS loss previously described, we have to consider various combinations of score functions, loss functions, and hyperparameters in KGE training. It is known that KGE performance can be significantly improved by selecting the appropriate combinations (Ruffinelli et al., 2020; Ali et al., 2021; Zhang et al., 2022), but it is unclear what combinations are theoretically appropriate. Therefore, the current state-of-the-art method

¹Nara Institute of Science and Technology (NAIST), Nara, Japan ²Hokkaido University, Hokkaido, Japan. Correspondence to: Hidetaka Kamigaito <kamigaito.h@is.naist.jp>.

relies on empirical hyperparameter tuning, which requires a huge amount of computation time to consider various combinations.

The nature of KG data also makes the NS loss in KGE unique. Since a triplet, which is a training instance of KGE, appears at most once in a KG, KGE learning suffers from the serious data-sparseness problem. Subsampling has been developed as a method for mitigating the data-sparseness problem in word-representation learning. However, there has been very little discussion on subsampling in KGE learning.

To address the aforementioned issues, we theoretically analyzed the NS loss to assist hyperparameter tuning and understand the better use of the NS loss in KGE learning. Our theoretical analysis showed that scoring methods with restricted value ranges, such as TransE, RotatE, require appropriate adjustment of the margin term or number of negative samples different from those without restricted value ranges, such as RESCAL, ComplEx, and DistMult. We present novel subsampling methods specialized for the NS loss in KGE learning from a theoretical perspective. Our empirical analysis on the FB15k-237, WN18RR, and YAGO3-10 datasets showed that the results of the models trained on the real-world datasets agree with our theoretical findings. Our code is available at <https://github.com/kamigaito/icml2022>.

2. NS Loss in KGE

In this section, we formulate KGE and explain the difference between the original NS loss (Mikolov et al., 2013b) and that used in KGE (Sun et al., 2019).

2.1. Formulation of KGE

We denote a triplet representing entities e_i , e_j and their relation r_k as (e_i, r_k, e_j) . In a typical KGC task, the model receives a query $(e_i, r_k, ?)$ or $(?, r_k, e_j)$ and predicts the entity corresponding to $?$. Let the input query be x and its answer entity be y . The probability $p_\theta(y|x)$ that y is predicted from x under the score function $s_\theta(x, y)$ on the basis of the model parameter θ is defined as follows using the softmax function:

$$p_\theta(y|x) = \frac{\exp(s_\theta(x, y))}{\sum_{y' \in Y} \exp(s_\theta(x, y'))}, \quad (1)$$

where Y is the set of entities and the size $|Y|$ can be large in KGE learning.

2.2. NS Loss Functions

The softmax function can explicitly represent probabilities, but the normalization term increases the learning time. For this reason, the original NS loss was proposed as a method

to approximate the softmax function without normalization during training (Mikolov et al., 2013b). For observables $D = \{(x_1, y_1), \dots, (x_n, y_n)\}$ that follow $p_d(x, y)$, the original NS loss is defined as

$$-\frac{1}{|D|} \sum_{(x, y) \in D} \left[\log(\sigma(s_\theta(x, y))) + \sum_{y_i \sim p_n(y_i|x)} \log(\sigma(-s_\theta(x, y_i))) \right], \quad (2)$$

where $p_n(y|x)$ is the noise distribution, σ is the sigmoid function, and ν is the number of negative samples per positive sample (x, y) .

KGE uses the following NS loss (Sun et al., 2019; Ahrabian et al., 2020) with the addition of the margin term γ and normalization term for ν to Eq. (2):

$$-\frac{1}{|D|} \sum_{(x, y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \frac{1}{\nu} \sum_{y_i \sim p_n(y_i|x)} \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right]. \quad (3)$$

Hereafter, we call Eq. (1) when the loss functions, such as Eqs. (2) and (3), are minimized, which means $s_\theta(x, y)$ reaches an optimal solution, as the *objective distribution*. In the next section, we discuss the difference between the two loss functions.

3. Theoretical Analysis

3.1. Equivalence of the Two Loss Functions

By comparing Eq. (2) with Eq. (3), we can show the following proposition.

Proposition 3.1. *Eqs. (2) and (3) have the same objective distribution:*

$$\frac{p_d(y|x)/p_n(y|x)}{\sum_{y' \in Y} (p_d(y'|x)/p_n(y'|x))}. \quad (4)$$

See Appendix A.1 for the proof. From Prop. 3.1, we can see that the existence of ν and γ has no effect on the distribution that the model will fit when the NS loss reaches the optimal solution. However, if the score function does not have the value range $(-\infty, +\infty)$, the score limitation prevents a model from reaching the optimal solution, and we need to discuss this point.

3.2. Roles of the Margin Term γ

In this section, we discuss the role of the margin term γ from a theoretical perspective.

3.2.1. VALUE RANGES OF SCORING METHODS

Distance-based scoring methods, such as TransE (Bordes et al., 2013) and RotatE (Sun et al., 2019), have score functions with a restricted value range. Distance-based scoring methods are generally expressed using the p -norm as follows:

$$-||f_\theta(x, y)||_p, \quad (5)$$

where $f_\theta(x, y)$ is a function that returns a vector value for (x, y) . From Eq. (5), when $f_\theta(x, y)$ can represent any vector, the value range of a distance-based scoring method is $(-\infty, 0]$. Due to the value range limitation, the following proposition holds for distance-based scoring methods.

Proposition 3.2. *In Eq. (3), the distance-based scoring function cannot reach the optimal solution when there exists (x, y) that satisfies $\exp(\gamma)p_n(y|x) < p_d(y|x)$.*

See Appendix A.2 for the proof. Since the value range restriction causes this problem in Eq. (3), the smaller $\exp(\gamma)p_n(y|x)$ than $p_d(y|x)$ is, the less able to approach the optimal solution the distance-based scoring function is. It is clear from Prop. 3.2 that we need to choose the noise distribution and γ appropriately when learning a distance-based scoring method with the NS loss in Eq. (3). To satisfy $\exp(\gamma)p_n(y|x) \geq p_d(y|x)$ in the NS loss in Eq. (3), we can simply set a sufficiently large γ . Since uniform distribution is generally used for $p_n(y|x)$ in KGE and $p_d(y|x)$ does not exceed 1 from the definition of probability, $\exp(\gamma)$ should be larger than the number of labels $|Y|$. To make a distance-based method capable to reach the optimal solution, it is appropriate to use γ satisfying

$$\gamma \geq \log(|Y|). \quad (6)$$

Note that scoring methods with unlimited value ranges, such as RESCAL (Bordes et al., 2011), ComplEx (Trouillon et al., 2016), and DistMult (Yang et al., 2015) are not related to the discussion of Prop. 3.2 and Eq. (6) that indicate the importance of adjusting γ .

3.2.2. GRADIENT CHANGES

From the discussion in §3.2.1, it seems that setting a large value for γ will facilitate the learning of the model without facing any disadvantages. However, the following proposition shows that γ cannot be set freely.

Proposition 3.3. *The margin term γ affects the gradient of the NS loss in Eq. (3).*

See Appendix A.3 for the proof. From Prop. 3.3, we can see that when we change γ , we also have to set the other hyperparameters related to the gradient appropriately. Thus, when using a distance-based scoring method, we need to set a sufficiently large γ and adjust the hyperparameters related to the gradient appropriately.

3.3. Roles of the Number of Negative Samples ν

Next, we discuss the effect of ν on learning from a theoretical perspective.

3.3.1. VALUE RANGES OF SCORING METHODS

Similar to Prop. 3.2, the following proposition holds:

Proposition 3.4. *In Eq. (2), the distance-based scoring function cannot reach the optimal solution when there exists (x, y) that satisfies $\nu p_n(y|x) < p_d(y|x)$.*

See Appendix A.2 for the proof. Similar to Prop. 3.2, in Eq. (2), the smaller $\nu p_n(y|x)$ than $p_d(y|x)$ is, the less able to approach the optimal solution the distance-based scoring function is. From Prop. 3.4, a distance-based scoring method should satisfy the condition $\nu p_n(y|x) \geq p_d(y|x)$ to reach the optimal solution in the NS loss of Eq. (2). When using a uniform distribution for $p_n(y|x)$, then

$$\nu \geq |Y|, \quad (7)$$

must be satisfied. However, since the goal of the NS loss is to reduce the computational cost by setting $\nu < |Y|$, this condition would greatly degrade the advantage of the NS loss in Eq. (2). Note that similar to γ , the discussion of Prop. 3.4 and Eq. (7) is not related to scoring methods that have no limited value range.

3.3.2. GRADIENT CHANGES

Furthermore, we can induce the following proposition by the Monte Carlo method on the basis of the law of large numbers.

Proposition 3.5. *When ν is sufficiently large, ν affects the gradient of the NS loss in Eq. (2), whereas ν does not affect the gradient of the NS loss in Eq. (3).*

See Appendix A.4 for the proof. When using the NS loss in Eq. (3), if ν is sufficiently large, it is unnecessary to tune ν in detail. However, for the NS loss in Eq. (2), when increasing ν for distance-based scoring methods, it is necessary to adjust hyperparameters related to the gradient.

3.4. Effects of Noise Distributions

Instead of uniform noise distributions, SANS presented by Sun et al. (2019) uses the output of the training model as the noise distribution for sampling negative entities. Though it is difficult to apply the theoretical analysis conducted on the loss functions in Eqs. (2) and (3) directly to SANS, the Monte Carlo method on the basis of the law of large numbers enables us to derive the following proposition.

Proposition 3.6. *When ν is sufficiently large, SANS becomes equivalent to the NS loss set to $p_n(y|x) = p_\theta(y|x)$ in Eq. (3).*

See Appendix A.5 for the proof. Prop. 3.6 means that if ν is sufficiently large, SANS can be treated as the NS loss expressed in the same form as in Eq. (3). This enables us to apply to SANS to what we have discussed for Eq. (3).

3.5. Subsampling for KGE

The discussion thus far has been on the assumption that the NS loss function fits the model to the distribution $p_d(y|x)$ defined from the observed data. However, what the NS loss actually does is to fit the model to the true distribution $p'_d(y|x)$ that exists behind the observed data. To fill in the gap between $p_d(y|x)$ and $p'_d(y|x)$, we reformulate the NS loss in Eq. (3) by the Monte Carlo method as

$$(3) \approx - \sum_{x,y} \left[\log(\sigma(s_\theta(x,y) + \gamma)) p_d(x,y) + p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) p_d(x) \right], \quad (8)$$

and consider replacements of $p_d(x,y)$ with $p'_d(x,y)$ and $p_d(x)$ with $p'_d(x)$. By assuming two functions, $A(x,y)$ and $B(x)$, that convert $p_d(x,y)$ into $p'_d(x,y)$ and $p_d(x)$ into $p'_d(x)$, we further reformulate Eq. (8) as follows.

$$\begin{aligned} & - \sum_{x,y} \left[\log(\sigma(s_\theta(x,y) + \gamma)) p'_d(x,y) + p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) p'_d(x) \right] \\ = & - \sum_{x,y} \left[\log(\sigma(s_\theta(x,y) + \gamma)) A(x,y) p_d(x,y) + p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) B(x) p_d(x) \right] \\ \approx & - \frac{1}{|D|} \sum_{(x,y) \in D} \left[A(x,y) \log(\sigma(s_\theta(x,y) + \gamma)) + \frac{1}{\nu} \sum_{y_i \sim p_n(y|x)} B(x) \log(\sigma(-s_\theta(x,y_i) - \gamma)) \right]. \quad (9) \end{aligned}$$

See Appendix B for the detailed derivation. Mikolov et al. (2013b) proposed subsampling for the NS loss in Eq. (2) to balance the appearance probability of words and word pairs by discounting their frequency. In Eq. (2), we can consider a word as x and a word pair as (x,y) . We denote the appearance probability of x as $p_d(x)$ and (x,y) as $p_d(x,y)$ in Eq. (9). Since $A(x,y)$ and $B(x)$ adjust $p_d(x,y)$ and $p_d(x)$ in Eq. (9), we can understand that subsampling has the same role with $A(x,y)$ and $B(x)$.

To the best of our knowledge, no paper discusses the use of subsampling in KGE. However, Sun et al. (2019) used subsampling which follows word2vec in their implementation¹. Regarding $B(x)$ as $B(x,y)$, we can understand that their

subsampling fills the gap between $p_d(y|x)$ and $p'_d(y|x)$ in Eq. (9) as follows:

$$A(x,y) = B(x,y) = \frac{\frac{1}{\sqrt{\#(x,y)}}}{\sum_{(x',y') \in D} \frac{1}{\sqrt{\#(x',y')}}}. \quad (10)$$

Here, $\#$ is the symbol for frequency, and $\#(x,y)$ represents the frequency of (x,y) . Note that the actual (x,y) occurs at most once in the KG, so when $(x,y) = (e_i, r_k, e_j)$, they approximate the frequency of (x,y) as follows:

$$\#(x,y) \approx \#(e_i, r_k) + \#(r_k, e_j). \quad (11)$$

We can interpret this approximation as the backoff (Katz, 1987) of $\#(x,y)$ to $\#(e_i, r_k)$ and $\#(r_k, e_j)$.

We derived our proposed KGE-specific subsampling methods from a theoretical perspective on the basis of Eq. (9) that discounts by frequency as in Eq. (10). This derivation depends on how we make assumptions about $p'_d(y|x)$. First, the derivation is based on the assumption that in $p'_d(y|x)$, (x,y) originally has a frequency, but the observed one is at most 1. In this case, we cannot actually calculate the frequency of $\#(x,y)$, as in Eq. (10). Thus, we need to use the approximation in Eq. (11). Since $A(x,y)$ needs to discount the frequency of (x,y) , and $B(x)$ needs to discount that of x , we can derive the following subsampling method:

$$\begin{aligned} A(x,y) &= \frac{\frac{1}{\sqrt{\#(x,y)}}}{\sum_{(x',y') \in D} \frac{1}{\sqrt{\#(x',y')}}}, \\ B(x) &= \frac{\frac{1}{\sqrt{\#x}}}{\sum_{x' \in D} \frac{1}{\sqrt{\#x'}}}. \quad (12) \end{aligned}$$

In $p'_d(y|x)$, however, if we assume that (x,y) has frequency 1 at most, as in the observation, then $p'_d(y|x) = p'_d(x,y)/p'_d(x) \propto 1/p'_d(x)$, so $p'_d(y|x)$ is the same for an x independent from y . Therefore, under this assumption, we only consider a discount for $p_d(x)$ and derive the following subsampling method:

$$A(x,y) = B(x) = \frac{\frac{1}{\sqrt{\#x}}}{\sum_{x' \in D} \frac{1}{\sqrt{\#x'}}}. \quad (13)$$

Although we derive our methods using Eqs. (12) and (13) from a theoretical perspective, $p'_d(y|x)$, which is the target of generalization in the actual task, varies depending on the dataset. Therefore, we cannot discuss the superiority or inferiority of our subsampling methods only from a theoretical perspective, and it is necessary to verify the methods using the development data in actual use.

4. Empirical Analysis

We examined whether the theoretical analysis discussed in §3 is valid for actual datasets and models.

¹<https://github.com/DeepGraphLearning/KnowledgeGraphEmbedding>

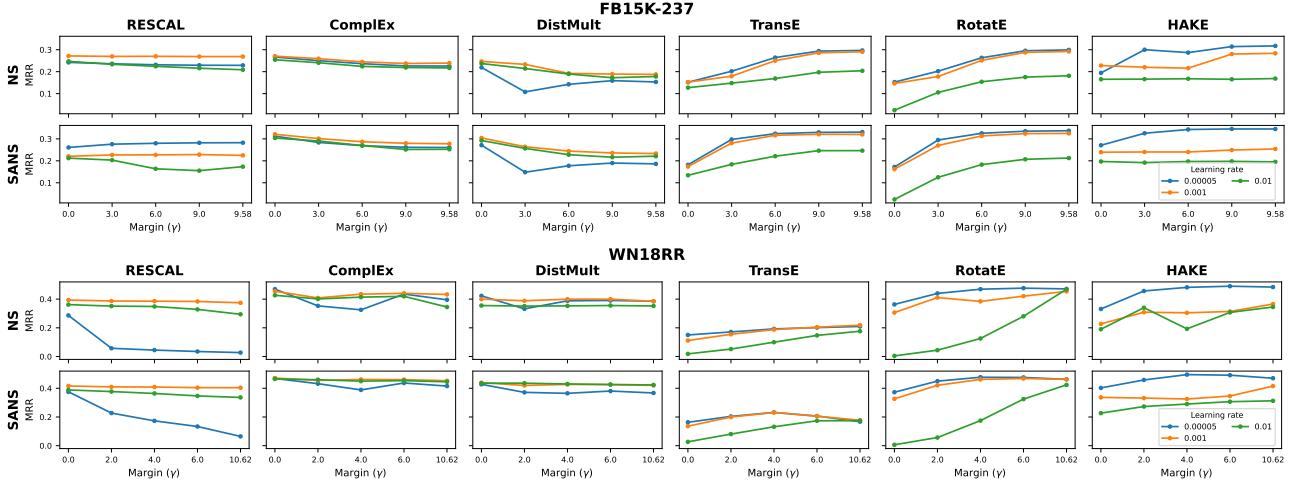

 Figure 1: MRR scores for each margin term γ in each learning rate on FB15k-237 and WN18RR.

Table 1: Statistics for each dataset.

Dataset	Entities	Relations	Tuples		
			Train	Valid	Test
FB15k-237	14,541	237	272,115	17,535	20,466
WN18RR	40,943	11	86,835	3,034	3,134
YAGO3-10	123,182	37	1,079,040	4,978	4,982

 Table 2: Scoring functions. d denotes a dimension size.

Model	Score Function	Parameters
RESCAL	$\mathbf{h}^T \mathbf{M}_r \mathbf{t}$	$\mathbf{h}, \mathbf{t} \in \mathbb{R}^d$, $\mathbf{M} \in \mathbb{R}^{d \times d}$
DistMult	$\mathbf{h}^T \text{diag}(\mathbf{r}) \mathbf{t}$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{R}^d$
ComplEx	$\text{Re}(\mathbf{h}^T \text{diag}(\mathbf{r}) \mathbf{t})$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{C}^d$
TransE	$-\ \mathbf{h} + \mathbf{r} - \mathbf{t}\ _p$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{R}^d$
RotatE	$-\ \mathbf{h} \circ \mathbf{r} - \mathbf{t}\ _p$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{C}^d, r_i = 1$, $\mathbf{h}, \mathbf{t} \in \mathbb{R}^d, \mathbf{r} \in \mathbb{R}_+^d$
HAKE	$-\ \mathbf{h} \circ \mathbf{r} - \mathbf{t}\ _p$ $-\lambda \ \sin((\mathbf{h}' + \mathbf{r}' - \mathbf{t}')/2)\ _1$	$\mathbf{h}', \mathbf{r}', \mathbf{t}' \in [0, 2\pi]^d$, $\lambda \in \mathbb{R}$

4.1. Settings

Datasets We used FB15k-237 (Toutanova & Chen, 2015), WN18RR, and YAGO3-10 (Dettmers et al., 2018) as the datasets. Table 1 shows the statistics for each dataset.

Models We chose RESCAL, ComplEx, DistMult, TransE, RotatE, and HAKE (Zhang et al., 2020) for our experiments. Table 2 shows the scoring functions for each model. From the table, we can see that the value range of RESCAL, ComplEx, and DistMult is $(-\infty, \infty)$, whereas that of TransE and RotatE is $(-\infty, 0]$. HAKE’s value range is close to $(-\infty, \infty)$ when λ is close to $-\infty$, whereas it is close to $(-\infty, 0]$ when λ is close to 0. Since λ is initialized near 0, a large learning rate is required to make λ close to $-\infty$. Thus, we can expect that HAKE pretends like TransE and RotatE

in a small learning rate.

Implementation & Hyperparameters We used the implementations and hyperparameters reported by Sun et al. (2019)² for ComplEx, DistMult, TransE, and RotatE and the one reported by Zhang et al. (2020)³ for HAKE in our experiments unless otherwise noted in the next subsections. We implemented RESCAL by modifying DistMult by basically inheriting the original hyperparameters. We chose uniform noise as $p_n(y|x)$ for the NS loss in Eqs. (2) and (3). We also applied SANS for model training. We used Adam (Kingma & Ba, 2015) as our optimizer.

4.2. Effects of the Margin Term γ

To investigate the performance when changing γ , we compared the mean reciprocal rank (MRR). We first examined whether the theoretical conclusion that γ is necessary to properly learn p -norm-based scoring methods as discussed in §3.2 is also valid in practice.

Settings For this investigation, we compared the MRRs of the models introduced in §4.1 on FB15k-237 and WN18RR. We chose the NS loss in Eq. (3) with uniform noise and the SANS loss to train the models. We chose margin term γ from $\{0.0, 3.0, 6.0, 9.0, 9.58\}$ on FB15k-237 and $\{0.0, 2.0, 4.0, 6.0, 10.62\}$ on WN18RR for each initial learning rate in $\{e^{-2}, e^{-3}, 5e^{-5}\}$. Note that we decided the set for margin values on the basis of previous studies (Sun et al., 2019; Zhang et al., 2020) and Eq. (6). We list the detailed hyperparameters for each model in Appendix C.1.

²<https://github.com/DeepGraphLearning/KnowledgeGraphEmbedding>

³<https://github.com/zyjcs/KGE-HAE>

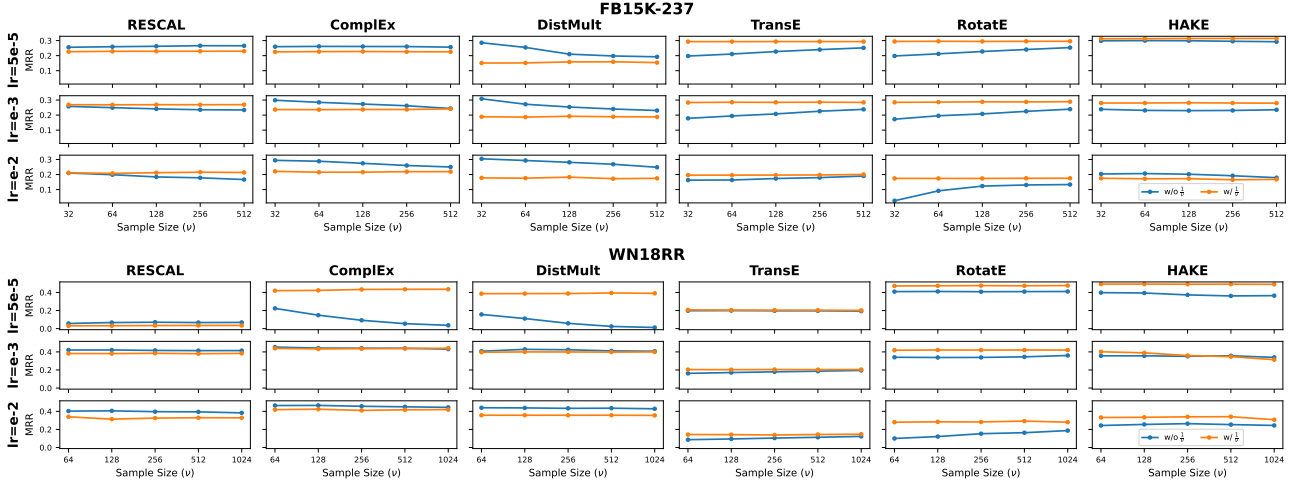


Figure 2: MRR scores using zero as margin term γ with and without $\frac{1}{\nu}$ normalizing for negative samples in Eq. (3).

Results Figure 1 shows the results. The results indicate that there was no performance gain for RESCAL, ComplEx, and DistMult by setting non-zero values on γ , which again confirms no necessity of adjusting γ in these scoring methods. We believe the gradient changes by different γ caused these performance degradations as shown in Prop. 3.3. For TransE and RotatE based on the p -norm, there is a significant performance degradation when γ decreases, which confirms the necessity of adjusting γ with these scoring methods. HAKE also follows the tendencies of TransE and RotatE, especially when the learning rate is small. These tendencies are consistent regardless of the NS loss with uniform noise or SANS loss.

Since the lowest γ satisfying Eq. (6) is 9.58 for FB15k-237 and 10.62 for WN18RR, the model cannot reach the optimal solution with γ reported in the previous study, 9.0 for FB15k-237 and 6.0 for WN18RR. However, the performance gains in MRR when changing γ from 9.0 to 9.58 on FB15k-237 and from 6.0 to 10.62 on WN18RR are slight. The goal of learning is to generalize the model to not fit perfectly to the training data. From this viewpoint, even if the model does not reach the optimal solution, it is not a large problem if it can reach the vicinity of the optimal solution. In fact, the tendencies of losses and MRRs are not perfectly the same.

From the results, we can conclude that in actual training, we should use γ to train distance-based scoring methods; however, Eq. (6) does not have to be strictly followed.

4.3. Effects of the Number of Negative Samples ν

Next, we investigated the effect of ν in learning.

4.3.1. VALUE RANGES OF SCORING METHODS

From the theoretical perspective, ν helps the distance-based scoring method reach the optimal solution in the original NS loss in Eq. (2) as shown in Prop. 3.4. We conducted an experiment to confirm this theoretical finding.

Settings We compared the MRRs of the models introduced in §4.1 on FB15k-237 and WN18RR. We used the NS loss in Eq. (2) (w/o $\frac{1}{\nu}$). We varied the number of negative samples ν from $\{32, 64, 128, 256, 512\}$ on FB15k-237 and $\{64, 128, 256, 512, 1024\}$ on WN18RR for each initial learning rate in $\{e-2, e-3, 5e-5\}$. We list the detailed hyperparameters for each model in Appendix C.2.

Results The blue lines of w/o $\frac{1}{\nu}$ in Figure 2 show the result. This result indicates that increasing ν assists in learning distance-based scoring methods TransE and RotatE on FB15k-237. However, the performance improvement on WN18RR by increasing ν is slight or none. Because WN18RR has a larger number of entities than FB15k-237, we consider that ν was insufficient for WN18RR. RESCAL, ComplEx, and DistMult have no performance gains by increasing ν . These observations indicate that Prop. 3.4 holds in actual training.

4.3.2. GRADIENT CHANGES

As shown in Prop. 3.5, changing ν affects the gradient of the NS loss in Eq. (2) but not in Eq. (3). We investigated the effect of the gradients in the actual training.

Settings We inherited the setting in §4.3.1 considering the NS loss in Eq. (3) (w/ $\frac{1}{\nu}$). We set γ as 9.0 on FB15k-237 and 6.0 on WN18RR.

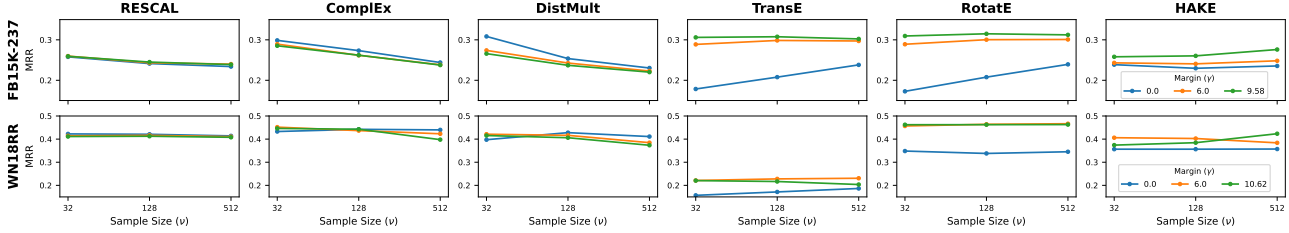


Figure 3: MRR scores for each model of each dataset when we vary negative sample size ν and margin term γ .

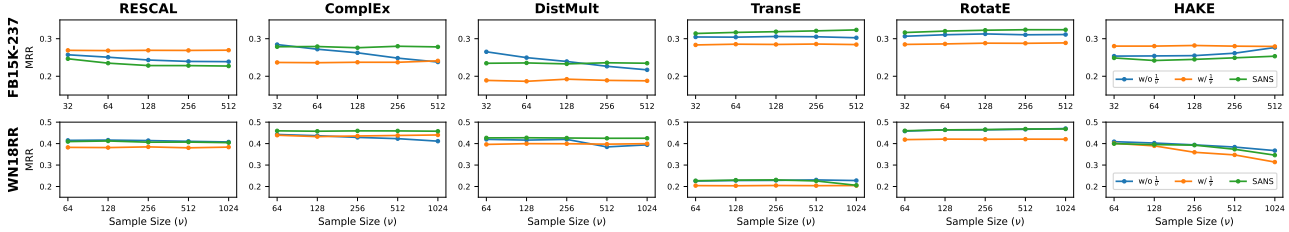


Figure 4: MRR scores for each model of each dataset when we vary negative sample size ν . $w/o \frac{1}{\nu}$ and $w/ \frac{1}{\nu}$ represent NS loss in Eqs. (2) and (3), respectively.

Results Figure 2 shows the result. By changing ν , we can see the stable performance of the NS loss in Eq. (3) compared with that in Eq. (2) for each learning rate. From the result, we can understand that Prop. 3.5 holds in the actual training.

4.4. Relationship between the Margin Term γ and the Number of Negative Samples ν

Props. 3.2 and 3.4 indicate that margin term γ and number of negative samples ν have a similar role. Since many factors affect the model performance, we investigated whether this relationship still holds in the actual training.

Settings In this investigation, we compared the MRRs of the models introduced in §4.1 on FB15k-237 and WN18RR. To consider the effect of both ν and γ , we used the NS loss in Eq. (3) by canceling its normalization term $\frac{1}{\nu}$. We varied ν from $\{32, 128, 512\}$ and γ from $\{0.0, 6.0, 9.58\}$ on FB15k-237 and $\{0.0, 6.0, 10.62\}$ on WN18RR with the commonly used initial learning rate $e - 3$ in Adam. We list the detailed hyperparameters for each model in Appendix C.3.

Results Figure 3 shows the result. In p -norm-based scoring methods TransE and RotatE, we can observe that large γ can be a replacement of large ν , whereas the results in the other models do not have such relationships. Therefore, we can consider that γ and ν have a similar role for p -norm-based scoring methods in the actual training. Compared with increasing γ , performance improvements by increasing ν are slight. This is because γ has an exponential effect on

the NS loss, while ν has only a linear effect and is small for reducing the computation time. Thus, using γ is more efficient than tuning ν for training a distance-based scoring method.

4.5. Effects of Gradient Changes in SANS

As discussed in Prop. 3.6, when ν is sufficiently large to approximate model probabilities, we can assume that ν does not affect the gradient during training in SANS similar to the NS loss in Eq. (3). On the other hand, as discussed in Prop. 3.5, ν affects the gradient during training of the NS loss in Eq. (2). Since ν is not the only factor that affects the gradient, it is not clear whether this difference actually affects model’s performance. We investigated these aspects through this experiment.

Settings Similar to the experiment in §4.3, we compared the MRRs of the models introduced in §4.1 on the basis of the various number of negative samples ν on FB15k-237 and WN18RR. We used the NS losses in Eqs. (2) and (3) and the SANS loss to train the models. To ignore the effect discussed in Prop. 3.4, on the basis of the best setting of the original implementation, we set γ as 9.0 on FB15k-237 and 6.0 on WN18RR. Furthermore, we also consider margin term γ for the NS loss in Eq. (2), which is in a similar form to the NS loss in Eq. (3) without normalization term $\frac{1}{\nu}$. We varied ν from $\{32, 64, 128, 256, 512\}$ on FB15k-237 and $\{64, 128, 256, 512, 1024\}$ on WN18RR with the commonly used initial learning rate $e - 3$ in Adam. We list the detailed hyperparameters for each model in Appendix C.4.

Table 3: Evaluation results for each subsampling method on FB15k-237. *Sub.* denotes subsampling, *None* denotes model that did not use subsampling, *Base* denotes Eq. (10), *Freq* denotes Eq. (12), and *Uniq* denotes Eq. (13).

FB15k-237					
Model	Sub.	MRR	Hits@1	Hits@3	Hits@10
RESCAL	None	17.2	9.9	18.1	31.8
	Base	22.3	13.9	24.2	39.8
	Freq	26.6	17.4	29.4	45.1
	Uniq	26.6	17.6	29.3	44.9
ComplEx	None	22.4	14.0	24.2	39.5
	Base	32.2	23.0	35.1	51.0
	Freq	32.8	23.6	36.1	51.2
	Uniq	32.7	23.5	35.8	51.3
DistMult	None	22.2	14.0	24.0	39.4
	Base	30.8	22.1	33.6	48.4
	Freq	29.9	21.2	32.7	47.5
	Uniq	29.1	20.3	31.8	46.6
TransE	None	33.0	22.8	37.2	53.0
	Base	32.9	23.0	36.8	52.7
	Freq	33.6	24.0	37.3	52.9
	Uniq	33.5	23.9	37.3	52.8
RotatE	None	33.1	23.1	37.1	53.1
	Base	33.6	23.9	37.4	53.2
	Freq	34.0	24.5	37.6	53.2
	Uniq	34.0	24.5	37.6	53.0
HAKE	None	32.3	21.6	36.9	53.2
	Base	34.5	24.7	38.2	54.3
	Freq	34.9	25.2	38.6	54.2
	Uniq	35.4	25.8	38.9	54.7

Results Figure 4 shows the result. We can see the stable MRRs in the NS loss in Eq. (3), whereas the NS loss in Eq. (2) shows the large changes of MRRs in RESCAL, ComplEx, and DistMult. In the SANS loss, we can also see the stable MRRs when ν is 64 or more. This means that Props. 3.5 and 3.6 can be satisfied only using small ν . In conclusion, as well as our theoretical interpretation, we found that the effect of ν on the actual learning is large for the original NS loss in Eq. (2), whereas it is small for the NS loss in Eq. (3) and SANS loss.

4.6. Effects of Subsampling Methods

We evaluated the subsampling methods discussed in §3.5. We compared the currently used subsampling method derived from Eq. (10): *Base* and the proposed methods derived from Eq. (12): *Freq* and (13): *Uniq*. We also included models trained without subsampling in this comparison.

Settings To train the models, we used the best configuration reported by Sun et al. (2019) for ComplEx, DistMult, TransE, and RotatE and Zhang et al. (2020) for HAKE since

Table 4: Evaluation results for each subsampling method in WN18RR. The notations are the same as those in Table 3.

WN18RR					
Model	Sub.	MRR	Hits@1	Hits@3	Hits@10
RESCAL	None	41.5	39.0	42.3	45.9
	Base	43.3	40.7	44.5	48.2
	Freq	44.1	41.1	45.6	49.5
	Uniq	44.1	41.4	45.5	49.5
ComplEx	None	45.0	40.9	46.6	53.4
	Base	47.1	42.8	48.9	55.7
	Freq	47.6	43.3	49.3	56.3
	Uniq	47.6	43.2	49.5	56.3
DistMult	None	42.4	38.3	43.6	51.0
	Base	43.9	39.4	45.2	53.3
	Freq	44.6	40.0	45.9	54.4
	Uniq	44.6	39.9	46.2	54.3
TransE	None	22.6	1.8	40.1	52.3
	Base	22.4	1.3	40.1	53.0
	Freq	23.0	1.9	40.7	53.7
	Uniq	23.2	2.2	41.0	53.4
RotatE	None	47.3	42.6	49.1	56.7
	Base	47.6	43.1	49.5	56.6
	Freq	47.8	42.9	49.8	57.4
	Uniq	47.9	43.5	49.6	56.7
HAKE	None	49.1	44.5	51.1	57.8
	Base	49.8	45.3	51.6	58.2
	Freq	49.7	45.2	51.4	58.5
	Uniq	49.8	45.4	51.5	58.3

they used the subsampling method in Eq. (10). In RESCAL, which was not included in the original implementation, we inherited the setting of DistMult by reducing its dimension size since RESCAL uses weight matrices instead of vectors for representing each relation. Note that on YAGO3-10, they only reported the configuration of TransE, RotatE, and HAKE. Therefore, we compared only TransE, RotatE, and HAKE on YAGO3-10. Subsampling is a smoothing method, and Kamigaito & Hayashi (2021) showed that the SANS loss also has a smoothing effect. Therefore, to confirm whether the subsampling method is effective in actual training, it is necessary to apply it to the SANS loss. Since the subsampling method’s implementation in the previous studies also adjusts the hyperparameters on the basis of the SANS loss, we used the SANS loss in this experiment. We used MRR, Hits@1, Hits@3, and Hits@10 as the evaluation metrics.

Results Tables 3, 4, and 5 show the results. We can see that for MRR, Hits@1, and Hits@3, subsampling was effective for all datasets and all models used in this experiment. Also, for Hits@10, the score was higher when we used subsampling for all models except RotatE. For RotatE, Freq and Uniq achieved a comparable or higher performance in

Table 5: Evaluation results for each subsampling method in YAGO3-10. The notations are the same as those in Table 3.

YAGO3-10					
Model	Sub.	MRR	Hits@1	Hits@3	Hits@10
TransE	None	50.6	40.9	56.6	67.7
	Base	51.2	41.5	57.6	68.3
	Freq Uniq	51.3 51.4	41.9 42.0	57.2 57.6	68.1 67.9
RotatE	None	50.6	41.1	56.5	67.8
	Base	50.8	41.8	56.5	67.6
	Freq Uniq	51.0 51.5	41.9 42.5	56.5 56.8	67.8 68.3
HAKE	None	53.4	44.9	58.7	68.4
	Base	54.3	46.1	59.5	69.2
	Freq Uniq	54.0 55.0	45.5 46.6	59.4 60.1	69.1 69.8

Hits@10 as those without subsampling. This observation suggest that the use of subsampling basically improves KGC performance.

When comparing the results of the subsampling methods, *Freq* and *Uniq* performed as well as or better than *Base* in many cases. This result shows the validity of deriving *Freq* and *Uniq* from the theoretical perspective. However, the effect of the subsampling method varies depending on the combined dataset and model. Considering the theoretical perspective, since we derived *Freq* and *Uniq* by assuming the true distribution behind the dataset, the fact that the results change depending on the model selected is considered contrary to the original intention. This is because several models have already taken into account the training data bias to be corrected by the subsampling method. Therefore, when using a subsampling method, it is important to check the performance of the combination of the model, subsampling method, and dataset by using a development dataset.

5. Related Work

Mikolov et al. (2013b) originally proposed the NS loss to approximate the softmax cross-entropy loss function for reducing the computational cost in word-representation learning (Mikolov et al., 2013a). Initially, due to its similarity to noise contrastive estimation (NCE) (Gutmann & Hyvärinen, 2010), Dyer (2014) conducted a theoretical analysis to determine the differences between the NS loss and NCE loss, but the distinct properties of the NS loss have not been clarified. In particular, the work of Levy & Goldberg (2014) attracted attention because it showed that the NS loss in word2vec is equivalent to decomposing the pointwise mutual information (PMI) matrix by using matrix factorization. However, this equivalence is only valid when the unigram distribu-

tion is used as the noise distribution; thus, this finding is restricted to word representation.

The NS loss has been widely used in KGE (Trouillon et al., 2016). Sun et al. (2019) proposed SANS for efficiently training KGE models. Since various scoring methods and loss functions are used in combination for KGE, the properties of the loss functions are not clear. To solve this problem, Ruffinelli et al. (2020) validated these combinations through large-scale experiments and empirically derived which combinations are best. Ali et al. (2021) also investigated the best configuration of the current KGE models. Zhang et al. (2022) proposed fast hyperparameter tuning in KGE.

However, these empirical approaches still require much computational time for hyperparameter tuning. Kamigaito & Hayashi (2021) clarified the properties of the softmax cross-entropy (SCE) loss function and NS loss function by theoretically identifying the differences between them. While their discussion covers a wide range of applications from word2vec to KGE, they did not include the NS loss specifically used for KGE in their analysis.

Our work is novel in that we analyzed the NS loss theoretically considering γ and the normalization of ν used in KGE in previous studies. These points have not been discussed in those studies. In addition, our study included the SANS loss, specifically proposed for KGE learning, in the discussion. Thus, our discussion is useful for training scoring methods using SANS. We also proposed subsampling methods specialized for KGE that can actually improve KGC performance.

6. Conclusion

We conducted a theoretical analysis for the NS loss used in KGE learning and derived theoretical facts. On the basis of these facts, we also proposed subsampling methods suitable for the NS loss function of KGE learning. To verify the validity of these theoretical facts and the performance of the proposed subsampling methods, we conducted experiments on the FB15k-237, WN18RR, and YAGO3-10 datasets. The experimental results indicate that the theoretical facts we derived are also observed in the real-world datasets.

Acknowledgements

The authors thank anonymous reviewers for insightful comments and suggestions. This work was partially supported by JSPS Kakenhi Grant Nos. 21H03491, 21K17801.

References

Ahrabian, K., Feizi, A., Salehi, Y., Hamilton, W. L., and Bose, A. J. Structure aware negative sampling in

- knowledge graphs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 6093–6101, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.492. URL <https://aclanthology.org/2020.emnlp-main.492>.
- Ali, M., Berrendorf, M., Hoyt, C. T., Vermue, L., Galkin, M., Sharifzadeh, S., Fischer, A., Tresp, V., and Lehmann, J. Bringing light into the dark: A large-scale evaluation of knowledge graph embedding models under a unified framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021. doi: 10.1109/TPAMI.2021.3124805.
- Bordes, A., Weston, J., Collobert, R., and Bengio, Y. Learning structured embeddings of knowledge bases. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI'11*, pp. 301–306. AAAI Press, 2011.
- Bordes, A., Usunier, N., García-Durán, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013*, pp. 2787–2795, 2013. URL <https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html>.
- Dettmers, T., Minervini, P., Stenetorp, P., and Riedel, S. Convolutional 2d knowledge graph embeddings. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pp. 1811–1818, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17366>.
- Dyer, C. Notes on noise contrastive estimation and negative sampling. *arXiv preprint arXiv:1410.8251*, 2014.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 297–304, 2010.
- Kamigaito, H. and Hayashi, K. Unified interpretation of softmax cross-entropy and negative sampling: With case study for knowledge graph embedding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021*, pp. 5517–5531, 2021. doi: 10.18653/v1/2021.acl-long.429. URL <https://doi.org/10.18653/v1/2021.acl-long.429>.
- Katz, S. Estimation of probabilities from sparse data for the language model component of a speech recognizer. *IEEE transactions on acoustics, speech, and signal processing*, 35(3):400–401, 1987.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kloek, T. and Van Dijk, H. K. Bayesian estimates of equation system parameters: an application of integration by monte carlo. *Econometrica: Journal of the Econometric Society*, pp. 1–19, 1978.
- Levy, O. and Goldberg, Y. Neural word embedding as implicit matrix factorization. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/feab05aa91085b7a8012516bc3533958-Paper.pdf>.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient estimation of word representations in vector space. In *Proceedings of the 1st International Conference on Learning Representations, ICLR 2013*, 2013a. URL <http://arxiv.org/abs/1301.3781>.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. Distributed representations of words and phrases and their compositionality. *CoRR*, abs/1310.4546, 2013b. URL <http://arxiv.org/abs/1310.4546>.
- Ruffinelli, D., Broscheit, S., and Gemulla, R. You CAN teach an old dog new tricks! on training knowledge graph embeddings. In *Proceedings of the 8th International Conference on Learning Representations, ICLR 2020*, 2020. URL <https://openreview.net/forum?id=BkxSmlBFvr>.
- Sun, Z., Deng, Z., Nie, J., and Tang, J. Rotate: Knowledge graph embedding by relational rotation in complex space. In *Proceedings of the 7th International Conference on Learning Representations, ICLR 2019*, 2019. URL <https://openreview.net/forum?id=HkgEQnRqYQ>.
- Toutanova, K. and Chen, D. Observed versus latent features for knowledge base and text inference. In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*, pp. 57–66, Beijing,

China, July 2015. Association for Computational Linguistics. doi: 10.18653/v1/W15-4007. URL <https://www.aclweb.org/anthology/W15-4007>.

Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., and Bouchard, G. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 2071–2080. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/trouillon16.html>.

Yang, B., Yih, W., He, X., Gao, J., and Deng, L. Embedding entities and relations for learning and inference in knowledge bases. In *Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015*, 2015. URL <http://arxiv.org/abs/1412.6575>.

Zhang, Y., Zhou, Z., Yao, Q., and Li, Y. Efficient hyperparameter search for knowledge graph embedding. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2715–2735, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.194. URL <https://aclanthology.org/2022.acl-long.194>.

Zhang, Z., Cai, J., Zhang, Y., and Wang, J. Learning hierarchy-aware knowledge graph embeddings for link prediction. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence, (AAAI20)*, pp. 3065–3072, 2020.

A. The Proofs of the Propositions

A.1. The Proof for Prop. 3.1

The previous studies (Levy & Goldberg, 2014; Kamigaito & Hayashi, 2021) show that the minimization for Eq. (2) induces the following equation:

$$\exp(s_\theta(x, y)) = \frac{p_d(y|x)}{\nu p_n(y|x)}. \quad (14)$$

Now we consider the minimization for the NS loss in Eq. (3). The normalization to ν in the NS loss in Eq. (3) eliminates ν in Eq. (14). Furthermore, due to the existence of the margin term γ , the minimization for the NS loss in Eq. (3) induces the following equation:

$$\exp(s_\theta(x, y)) = \frac{p_d(y|x)}{\exp(\gamma)p_n(y|x)}. \quad (15)$$

Here, substituting Eq. (14) into Eq. (1) and substituting Eq. (15) into Eq. (1) both result in the following formula:

$$\frac{p_d(y|x)/p_n(y|x)}{\sum_{y' \in Y} (p_d(y'|x)/p_n(y'|x))}. \quad (16)$$

Therefore, from Eq. (16), we can understand that Prop. 3.1 holds.

A.2. The Proof for Props. 3.2 and 3.4

When $s_\theta(x, y) \leq 0$, then $\exp(s_\theta(x, y)) \leq 1$ holds. If the right-hand side of Eq. (14) do not have the same value range as $\exp(s_\theta(x, y))$, then the NS loss in Eq. (2) can have an unreachable point. Therefore, we can derive the condition that the NS loss in Eq. (2) does not have an unreachable point as follows:

$$\begin{aligned} \frac{p_d(y|x)}{\nu p_n(y|x)} &\leq 1 \\ \Leftrightarrow p_d(y|x) &\leq \nu p_n(y|x). \end{aligned} \quad (17)$$

We can similarly induce the condition that the NS loss in Eq. (3) has no unreachable point as follows:

$$\begin{aligned} \frac{p_d(y|x)}{\exp(\gamma)p_n(y|x)} &\leq 1 \\ \Leftrightarrow p_d(y|x) &\leq \exp(\gamma)p_n(y|x). \end{aligned} \quad (18)$$

Here, from Eq. (17) and (18), we can understand that Prop. 3.2 and 3.4 hold.

A.3. The Proof for Prop. 3.3

From the derivative of a composite function, it is obvious that the result of differentiating the NS loss in Eq. (3) contains γ . Thus, we can understand that Prop. 3.3 holds.

A.4. The Proof for Prop. 3.5

In the NS loss of Eq. (2), we can consider the following approximation by utilizing the Monte Carlo method:

$$\sum_y p_n(y|x) \log(\sigma(-s_\theta(x, y))) \approx \frac{1}{\nu} \sum_{i=1, y_i \sim p_n(y_i|x)}^\nu \log(\sigma(-s_\theta(x, y_i))). \quad (19)$$

Based on Eq. (19), we can reformulate Eq. (2) as follows:

$$\begin{aligned} & -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y))) + \sum_{y_i \sim p_n(y_i|x)}^\nu \log(\sigma(-s_\theta(x, y_i))) \right] \\ & \approx -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y))) + \nu \sum_{y_i} p_n(y_i|x) \log(\sigma(-s_\theta(x, y_i))) \right]. \end{aligned} \quad (20)$$

In the NS loss of Eq. (3), we can similarly consider the following approximation by utilizing the Monte Carlo method:

$$\sum_y p_n(y|x) \log(\sigma(-s_\theta(x, y) - \gamma)) \approx \frac{1}{\nu} \sum_{i=1, y_i \sim p_n(y_i|x)}^\nu \log(\sigma(-s_\theta(x, y_i) - \gamma)). \quad (21)$$

Based on Eq. (21), we can reformulate Eq. (3) as follows:

$$\begin{aligned} & -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \frac{1}{\nu} \sum_{y_i \sim p_n(y_i|x)}^\nu \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right] \\ & \approx -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \sum_{y_i} p_n(y_i|x) \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right]. \end{aligned} \quad (22)$$

Since the reformulation of Eq. (20) and (22) is valid, when ν is enough large, we can understand that Prop. 3.5 holds.

A.5. The Proof for Prop. 3.6

Self-adversarial negative sampling (SANS) is defined as follows:

$$-\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \sum_{i=1, y_i \sim \text{uniform}}^\nu \tilde{p}_\theta(y_i|x) \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right], \quad (23)$$

where $\tilde{p}_\theta(y_i|x)$ is an approximated $p_\theta(y_i|x)$ with sampled y_j , defined as follows:

$$\tilde{p}_\theta(y_i|x) = \frac{\exp(s_\theta(x, y_i))}{\sum_{j=1}^\nu \exp(s_\theta(x, y_j))}. \quad (24)$$

Here, we can reformulate Eq. (23) as follows:

$$\begin{aligned} & -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \sum_{i=1, y_i \sim \text{uniform}}^\nu \tilde{p}_\theta(y_i|x) \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right] \\ & = -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \frac{|Y|}{\nu} \sum_{i=1, y_i \sim \text{uniform}}^\nu \frac{\exp(s_\theta(x, y_i))}{\frac{|Y|}{\nu} \sum_{j=1}^\nu \exp(s_\theta(x, y_j))} \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right]. \end{aligned} \quad (25)$$

In Eq. (25), we can consider $\frac{|Y|}{\nu} \sum_{j=1}^\nu \exp(s_\theta(x, y_j))$ as a Monte Carlo expectation to approximate $\sum_{j=1}^{|Y|} \exp(s_\theta(x, y_j))$ in $p_\theta(y_i|x)$. Thus, we can reformulate Eq. (25) as follows:

$$\begin{aligned} (25) & \approx -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \frac{|Y|}{\nu} \sum_{i=1, y_i \sim \text{uniform}}^\nu p_\theta(y_i|x) \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right] \\ & = -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x, y) + \gamma)) + \frac{1}{\nu} \sum_{i=1, y_i \sim \text{uniform}}^\nu \frac{p_\theta(y_i|x)}{\frac{1}{|Y|}} \log(\sigma(-s_\theta(x, y_i) - \gamma)) \right]. \end{aligned} \quad (26)$$

By using importance sampling (Kloek & Van Dijk, 1978), we can consider the following approximation:

$$\begin{aligned} & \frac{1}{\nu} \sum_{i=1, y_i \sim \text{uniform}}^\nu \frac{p_\theta(y_i|x)}{\frac{1}{|Y|}} \log(\sigma(-s_\theta(x, y_i) - \gamma)) \\ & \approx \sum_{y_i} \frac{1}{|Y|} \frac{p_\theta(y_i|x)}{\frac{1}{|Y|}} \log(\sigma(-s_\theta(x, y_i) - \gamma)) \\ & = \sum_{y_i} p_\theta(y_i|x) \log(\sigma(-s_\theta(x, y_i) - \gamma)) \end{aligned} \quad (27)$$

Based on Eq. (27), we can further reformulate Eq. (26) as follows:

$$(26) \approx -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x,y) + \gamma)) + \sum_{y_i} p_\theta(y_i|x) \log(\sigma(-s_\theta(x,y_i) - \gamma)) \right]. \quad (28)$$

Since Eq. (28) is the same when $p_n(y|x) = p_\theta(y|x)$ in Eq. (22), we can understand that Prop. 3.6 holds.

B. The Detailed Derivation of Eq. (9)

We can reformulate the NS loss in Eq. (3) as follows:

$$\begin{aligned} (3) &= -\frac{1}{|D|} \sum_{(x,y) \in D} \left[\log(\sigma(s_\theta(x,y) + \gamma)) + \frac{1}{\nu} \sum_{y_i \sim p_n(y_i|x)} \log(\sigma(-s_\theta(x,y_i) - \gamma)) \right] \\ &= -\frac{1}{|D|} \sum_{(x,y) \in D} \log(\sigma(s_\theta(x,y) + \gamma)) - \frac{1}{|D|} \sum_{(x,y) \in D} \frac{1}{\nu} \sum_{y_i \sim p_n(y_i|x)} \log(\sigma(-s_\theta(x,y_i) - \gamma)). \end{aligned} \quad (29)$$

Here, we can consider the following approximation based on the Monte Carlo method:

$$\frac{1}{\nu} \sum_{y_i \sim p_n(y_i|x)} \log(\sigma(-s_\theta(x,y_i) - \gamma)) \approx \sum_y p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)). \quad (30)$$

Using Eq. (30), we can reformulate Eq. (29) as follows:

$$(29) \approx -\frac{1}{|D|} \sum_{(x,y) \in D} \log(\sigma(s_\theta(x,y) + \gamma)) - \frac{1}{|D|} \sum_{(x,y) \in D} \sum_{y'} p_n(y'|x) \log(\sigma(-s_\theta(x,y') - \gamma)). \quad (31)$$

Similr to Eq. (30), we can consider the following approximation by the the Monte Carlo method:

$$\begin{aligned} -\frac{1}{|D|} \sum_{(x,y) \in D} \log(\sigma(s_\theta(x,y) + \gamma)) &\approx -\sum_{x,y} \log(\sigma(s_\theta(x,y) + \gamma)) p_d(x,y), \\ -\frac{1}{|D|} \sum_{(x,y) \in D} \sum_{y'} p_n(y'|x) \log(\sigma(-s_\theta(x,y') - \gamma)) &\approx -\sum_x \sum_{y'} p_n(y'|x) \log(\sigma(-s_\theta(x,y') - \gamma)) p_d(x). \end{aligned} \quad (32)$$

Using Eq. (32), we can reformulate Eq. (31) as follows:

$$\begin{aligned} (31) &\approx -\sum_{x,y} \log(\sigma(s_\theta(x,y) + \gamma)) p_d(x,y) - \sum_x \sum_{y'} p_n(y'|x) \log(\sigma(-s_\theta(x,y') - \gamma)) p_d(x) \\ &= -\sum_{x,y} \log(\sigma(s_\theta(x,y) + \gamma)) p_d(x,y) - \sum_{x,y} p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) p_d(x) \\ &= -\sum_{x,y} \left[\log(\sigma(s_\theta(x,y) + \gamma)) p_d(x,y) + p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) p_d(x) \right]. \end{aligned} \quad (33)$$

Next, we consider replacements of $p_d(x,y)$ with $p'_d(x,y)$ and $p_d(x)$ with $p'_d(x)$. By assuming two functions, $A(x,y)$ and $B(x)$, that convert $p_d(x,y)$ into $p'_d(x,y)$ and $p_d(x)$ into $p'_d(x)$, we further reformulate Eq. (33) as follows:

$$\begin{aligned} &-\sum_{x,y} \left[\log(\sigma(s_\theta(x,y) + \gamma)) p'_d(x,y) + p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) p'_d(x) \right] \\ &= -\sum_{x,y} \left[\log(\sigma(s_\theta(x,y) + \gamma)) A(x,y) p_d(x,y) + p_n(y|x) \log(\sigma(-s_\theta(x,y) - \gamma)) B(x) p_d(x) \right]. \end{aligned} \quad (34)$$

Based on the similar derivation from Eq. (3) to Eq. (33), we can reformulate Eq. (34) as follows:

$$(34) \approx -\frac{1}{|D|} \sum_{(x,y) \in D} \left[A(x,y) \log(\sigma(s_\theta(x,y) + \gamma)) + \frac{1}{\nu} \sum_{y_i \sim p_n(y_i|x)} B(x) \log(\sigma(-s_\theta(x,y_i) - \gamma)) \right]. \quad (35)$$

Table 6: Hyperparameters for each model. *Batch* denotes the batch size, *Dim* denotes the hidden dimension size, α denotes the temperature parameter for SANS, and *Step* denotes the max steps in training.

Model	FB15k-237				WN18RR				YAGO3-10			
	Batch	Dim.	α	Steps	Batch	Dim.	α	Steps	Batch	Dim.	α	Step
RESCAL	1024	500	1.0	100000	512	500	1.0	80000	-	-	-	-
ComplEx	1024	1000	1.0	100000	512	500	1.0	80000	-	-	-	-
DistMult	1024	2000	1.0	100000	512	1000	1.0	80000	-	-	-	-
TransE	1024	1000	1.0	100000	512	500	0.5	80000	1024	500	1.0	200000
RotatE	1024	1000	1.0	100000	512	500	0.5	80000	1024	500	1.0	200000
HAKE	1024	1000	1.0	100000	512	500	0.5	80000	1024	500	1.0	200000

Table 7: Hyperparameters for each model in §4.5. *LR* denotes the learning rate.

Model	FB15k-237			WN18RR			YAGO3-10		
	ν	γ	LR	ν	γ	LR	ν	γ	LR
RESCAL	256	200	0.001	1024	200	0.002	-	-	-
ComplEx	256	200	0.001	1024	200	0.002	-	-	-
DistMult	256	200	0.001	1024	200	0.002	-	-	-
TransE	256	9.0	0.00005	1024	6.0	0.00005	400	24.0	0.0002
RotatE	256	9.0	0.00005	1024	6.0	0.00005	400	24.0	0.0002
HAKE	256	9.0	0.00005	1024	6.0	0.00005	500	24.0	0.0002

C. Experimental Details

C.1. Detailed Hyperparameters for §4.2

For each model, we used the hyperparameters of FB15k-237 and WN18RR in Table 6. We explained other hyperparameters and settings in §4.2.

C.2. Detailed Hyperparameters for §4.3

For each model, we used the hyperparameters of FB15k-237 and WN18RR in Table 6. We explained other hyperparameters and settings in §4.3.

C.3. Detailed Hyperparameters for §4.4

For each model, we used the hyperparameters of FB15k-237 and WN18RR in Table 6. We explained other hyperparameters and settings in §4.4.

C.4. Detailed Hyperparameters for §4.5

For each model, we used the hyperparameters listed in Tables 6 and 7. We described other settings in §4.5.