

Efficient Parallel Translating Embedding For Knowledge Graphs

Denghui Zhang, Manling Li, Yantao Jia, Yuanzhuo Wang, Xueqi Cheng

Institute of Computing Technology, Chinese Academy of Sciences

No.6 Kexueyuan South Road, Zhongguancun, Haidian, Beijing, China 100190

zhangdenghui@ict.ac.cn, limanling@software.ict.ac.cn, jiyantao@ict.ac.cn, wangyuanzhuo@ict.ac.cn, cxq@ict.ac.cn

ABSTRACT

Knowledge graph embedding aims to embed entities and relations of knowledge graphs into low-dimensional vector spaces. Translating embedding methods regard relations as the translation from head entities to tail entities, which achieve the state-of-the-art results among knowledge graph embedding methods. However, a major limitation of these methods is the time consuming training process, which may take several days or even weeks for large knowledge graphs, and result in great difficulty in practical applications. In this paper, we propose an efficient parallel framework for translating embedding methods, called ParTrans-X, which enables the methods to be paralleled without locks by utilizing the distinguished structures of knowledge graphs. Experiments on two datasets with three typical translating embedding methods, i.e., TransE [3], TransH [19], and a more efficient variant TransE-AdaGrad [11] validate that ParTrans-X can speed up the training process by more than an order of magnitude.

CCS CONCEPTS

•Computing methodologies → Reasoning about belief and knowledge;

KEYWORDS

Knowledge Graph Embedding, Translation-based, Parallel

ACM Reference format:

Denghui Zhang, Manling Li, Yantao Jia, Yuanzhuo Wang, Xueqi Cheng. 2017. Efficient Parallel Translating Embedding For Knowledge Graphs. In *Proceedings of WI '17, Leipzig, Germany, August 23-26, 2017*, 9 pages. DOI: 10.1145/3106426.3106447

1 INTRODUCTION

Knowledge graphs are structured graphs with various entities as nodes and relations as edges. They are usually in form of RDF-style triples (h, r, t) , where h represents a head entity, t a tail entity, and r the relation between them. In the past decades, a quantity of large scale knowledge graphs have sprung up, e.g., Freebase [2], WordNet [14], YAGO [12], OpenKN [7], and have played a pivotal role in supporting many applications, such as link prediction, question answering, etc. Although these knowledge graphs are very large, i.e., usually containing thousands of relation types, millions of

entities and billions of triples, they are still far from complete. As a result, *knowledge graph completion (KGC)* has been paid much attention to, which mainly aims to predict missing relations between entities under the supervision of existing triples.

Recent years have witnessed great advances of translating embedding methods to tackle KGC problem. The methods represent entities and relations as the embedding vectors by regarding relations as translations from head entities to tail entities, such as TransE [3], TransH [19], TransR[11], etc. However, the training procedure is time consuming, since they all employ stochastic gradient descent (SGD) to optimize a translation-based loss function, which may require days to converge for large knowledge graphs.

For instance, Table 1¹ shows the complexity of typical translating embedding methods, where T_{total} stands for the total training time with T_{epoch} for the time of each epoch, and one epoch is a single pass over all triples. n_e , n_r and n_t are the number of entities, relations and triples in the knowledge graph respectively. d is the embedding dimension which is the same for entities and relations in this case, and ep is the minimum epochs which used to be set to 1000. It can be seen that the time complexity of TransE is proportional to n_t , d and ep . When d is 100 and ep is 1000, it will take 78 minutes for TransE to learn the embeddings of FB15k², which is a subset of Freebase with 483,142 training triples, and has been widely used as experimental dataset in knowledge graph embedding methods [3, 8, 11, 19]. Nevertheless, Freebase-rdf-latest³ is the latest public available data dump of Freebase with 1.9 billion triples, which results in approximately 3932 times the training time, namely, 212 days. Furthermore, the whole Freebase contains over 3 billion triples⁴, and it will take about 357 days to learn the embeddings of it. Despite its large size, Freebase still suffers from data incomplete problem, e.g., 75% persons do not have nationalities in Freebase [5]. On top of that, most improved variants of TransE employ more complex loss function to better train the embedding vectors, thus they possess higher time complexity or model complexity, and the training time of them will be even unbearable. For example, it will take more than 59 years for Freebase-rdf-latest when employing TransR, which is one of the typical improved variants and achieves far better performance than TransE.

There have been attempts to resolve the efficiency issue of translating embedding methods for knowledge graphs. Pasquale[15] proposed TransE-AdaGrad to speed up the training process by leveraging adaptive learning rates. However, TransE-AdaGrad essentially reduces the number of epochs to converge, and still can

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WI '17, Leipzig, Germany

© 2017 ACM. 978-1-4503-4951-2/17/08...\$15.00

DOI: 10.1145/3106426.3106447

¹The experiments are conducted on a dual Intel Xeon E5-2640 CPUs (10 cores each \times 2 hyperthreading, running at 2.4 GHz) machine with 128GB of RAM. The kernel is Red Hat 4.4.7

²<https://everest.hds.utc.fr/lib/exe/fetch.php?media=en:fb15k.tgz>

³<http://commondatastorage.googleapis.com/freebase-public/>

⁴<https://github.com/nchah/freebase-triples>, there are 3,197,653,841 triples in Freebase on May 2, 2016

Table 1: Complexity Analysis of Typical translating embedding Methods, with $d = 100, ep = 1000$

	Time Complexity	Model Complexity	on FB15k		on Freebase-rdf-latest		on the whole Freebase	
			T_{epoch}	T_{total}	T_{epoch}	T_{total}	T_{epoch}	T_{total}
TransE	$O(n_t \times d \times k)$	$O(dn_e + dn_r)$	4.5s	4658s ≈ 78 minutes	17,696s ≈ 5 hours	18,323,395s ≈ 212 days	29,781s ≈ 8 hours	30,828,893s ≈ 357 days
TransH	$O(n_t \times d \times k)$	$O(dn_e + 2dn_r)$	6s	100 minutes	6.5 hours	273 days	11 hours	459 days
TransR	$O(n_t \times d^2 \times k)$	$O(dn_e + dn_r + n_r d^2)$	473s	5 days	21.5 days	59 years	36 days	99 years

not do well with large scale knowledge graphs. In fact, with more and more computation resources available, it is natural and more effective to parallel these embedding methods, which will lead to significant improvement in training efficiency and can scale to quite large knowledge graphs if given sufficient hardware resources.

However, it is challenging to parallel the translating embedding methods, since the training processes mainly employ stochastic gradient descent algorithm (SGD) or the variants of it. SGD is inherently sequential, as a dependence exists between each iteration. Parallelizing translating embedding methods straightforwardly will result in collisions between different processors. For instance, an entity embedding vector is updated by two processors at the same time, and the gradients calculated by these processors are different. In this case, the diverse gradients are called *collisions*. To avoid collisions, some methods [9] lock embedding vectors, which will slow the training process greatly as there are so many vectors. On the contrary, updating vectors without locks leads to high efficiency, but should be based on specific assumptions [4, 16]. Since the lock-free training process may result in poor convergence if adopting suboptimal strategy to resolve collisions.

Our key observation of translating embedding methods is that the update performed in one iteration of SGD is based on only one triple and its corrupted sample, which is not necessarily bound up with other embedding vectors. This gives us chance to learn the embedding vectors in parallel without being locked. In this article, we analyze the distinguished data structure of knowledge graphs, and propose an efficient parallel framework for translating embedding methods, called ParTrans-X. It enables translating methods to update the embedding vectors efficiently in shared memory without locks. Thus the training process is greatly speeded up with multi-processors, which can be more than an order of magnitude faster without lowering learning quality.

The contribution of this article is:

1. We explore the law of collisions along with increasing number of processors, by modelling the training data of knowledge graph into hypergraphs.
2. We propose ParTrans-X framework to train translating methods efficiently in parallel. It utilizes the training data sparsity of large scale knowledge graphs, and can be easily applied to many translating embedding methods.
3. We apply ParTrans-X to typical translating embedding methods, i.e., TransE [3], TransH [19], and a more efficient variant TransE-AdaGrad, and experiments validate the effectiveness of ParTrans-X on two widely used datasets.

The paper is organized as follows. Related work is in Sec.2. The collision formulation is introduced in Sec.3 and ParTrans-X is proposed based on it in Sec.4. Then, experiments demonstrate training efficiency of ParTrans-X in Sec.5, with conclusions in Sec.6.

2 RELATED WORK

In recent years, translating embedding methods have played a pivotal role in Knowledge Graph Completion, which usually employ stochastic gradient descent algorithm to optimize a translation-based loss function, i.e.,

$$L = \sum_{(h,r,t)} \sum_{(h',r,t')} \max [0, f_r(h, t) + M - f_r(h', t')], \quad (1)$$

where (h, r, t) represents the positive triple that exists in the knowledge graph, while (h', r, t') stands for the negative triple that is not in the knowledge graph. $\max [0, \cdot]$ is the hinge loss, and M is the margin between positive and negative triples. $f_r(h, t)$ is the score function to determine whether the triple (h, r, t) should exist in the knowledge graph, which varies from different translating embedding methods.

A significant work is TransE [3], which heralds the start of translating embedding methods. It looks upon a triple (h, r, t) as a translation from the head entity h to the tail entity t , i.e., $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$, and the score function is $f_r(h, t) = \|\mathbf{h} + \mathbf{r} - \mathbf{t}\|$, where $\|\cdot\|$ represents L1-similarity or L2-similarity. The boldface suggests the vectors in the embedding space, namely, $\mathbf{h}, \mathbf{t} \in \mathbb{R}^d, \mathbf{r} \in \mathbb{R}^d$, where $d = d_e = d_r$ is the dimension of embedding space, d_e the dimension for entities and d_r for relations. Moreover, TransH [19] assumes that it is the projections of entities to a relation-specific hyperplane that satisfy the translation constraint, i.e., $f_r(h, t) = \|\mathbf{h}_\perp + \mathbf{r} - \mathbf{t}_\perp\|$, where $\mathbf{h}_\perp = \mathbf{h} - \mathbf{w}_r^T \mathbf{h} \mathbf{w}_r$ and $\mathbf{t}_\perp = \mathbf{t} - \mathbf{w}_r^T \mathbf{t} \mathbf{w}_r$, with $\mathbf{w}_r \in \mathbb{R}^{d_e}$ as the normal vector of the hyperplane related to r . Furthermore, TransR [11] employs rotation transformation to project the entities to a relation-specific space, i.e., $f_r(h, t) = \|\mathbf{h}_r + \mathbf{r} - \mathbf{t}_r\|$, where $\mathbf{h}_r = \mathbf{M}_r \mathbf{h}$ and $\mathbf{t}_r = \mathbf{M}_r \mathbf{t}$, and $\mathbf{M}_r \in \mathbb{R}^{d_r \times d_e}$ is the projection matrix relation to r . Some works also involves more information to better embedding, e.g., paths [10], margins [8].

Although this category of methods achieve the state-of-the-art results, the main limitation is the computationally expensive training process when facing large scale knowledge graphs. Recently, a method TransE-AdaGrad [15] was proposed to reduce the training time of TransE by employing AdaGrad [6], an variant of SGD, to adaptively modify the learning rate. Although the training time has been reduced greatly, there is still some way to go when facing large scale knowledge graphs. With the computation resources greatly enriched, training in parallel seems to be a more reliable way to relieve this issue. Actually, there are some works, e.g., [18], to parallel some graph computation paradigms, such as online query processing, offline analytics, etc. Nevertheless, it is not easy to train translating embedding methods in parallel, since the main optimization algorithm SGD is born to run in sequence. The major obstacle to parallel SGD is the collisions between updates of different processors for the same parameter [17], to overcome which there are two main brunches of methods.

The first brunch is to design a strategy to resolve collisions according to specific data structure. For example, Hogwild! [16] is a lock-free scheme works well for sparse data, which means that there is only a small part of parameters to update by each iteration of SGD. It has been proved that processors are unlikely to overwrite each other's progress, and the method can achieve a nearly optimal rate of convergence. While the second brunch is to split the training data to reduce collisions. Downpour SGD [4] mainly employ Dist-Belief [4] framework, which divides the training data into a number of subsets, then the model replicas run independently on each of these subsets, and do not communicate with each other. Inspired by this, TensorFlow [1] splits a computation graph into a subgraph for every worker and communication takes place using Send/Receive node pairs. Motivated by training large-scale convolutional neural networks for image classification, Elastic Averaging SGD (EASGD) [20] reduces the amount of communication between local workers and the master to allow the parameters of local workers to fluctuate further from the center ones. There are also works to improve the performance in parallel settings, e.g., Delay-tolerant Algorithms for SGD [13] adapts not only to the sequence of gradients, but also to the precise update delays that occur, inspired by AdaGrad.

However, these parallel framework are based on specific assumptions, and can not directly apply to translating embedding models without exploring distinguished data structures of knowledge graphs. Therefore, we shall propose a parallel framework for translating embedding models, called ParTrans-X, as knowledge graphs are mainly in form of triples, and trained triple by triple, it will lead to particular parallel framework.

3 LAW OF COLLISIONS EMERGING IN KG

As mentioned previously, there may exist collisions between processors when they update the same embedding vector, which ends up being one of the most challenging aspects of parallelizing translating embedding methods. Hence, we explore the law of collisions emerging in this section. At first we formulate the training data of knowledge graphs into hypergraphs. Then the collisions in training process are further discussed based on this formulation.

3.1 Hypergraph Formulation

Firstly, we model the knowledge graph formally as $G = (E, R, T)$, where E is the set of entities with R the set of relations, and T is the set of triples (h, r, t) , in which $h, t \in E$ and $r \in R$. The cardinalities of E, R and T are n_e, n_r and n respectively. In this graph, nodes are entities, and edges are triples that connecting nodes with a distinguished relation. For example, the knowledge graph shown in Figure1(a), where black nodes stand for the entities in knowledge graphs and lines for relations, can be represented as $G = (E, R, T)$, where $E = \{e_1, e_2, e_3, e_4, e_5\}$, $R = \{r_1, r_2, r_3\}$ and $T = \{(e_1, r_2, e_4), (e_1, r_1, e_3), (e_2, r_3, e_5), (e_4, r_3, e_2)\}$. In this case, $n_e = 5$, $n_r = 3$ and $n = 4$.

Secondly, the training data of knowledge graphs can be looked upon as hypergraphs. Recall the loss function of translating embedding methods in Eq.(1), which means in one iteration of SGD, only one positive triple (h, r, t) and one negative triple (h', r, t') are concerned. To be more clear, the data used in one iteration, i.e., $[(h, r, t), (h', r, t')]$, is called a *sample*. Note that (h', r, t') is

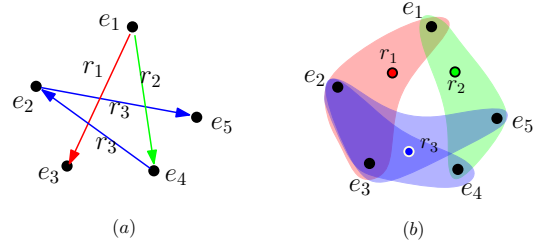


Figure 1: A Knowledge Graph (a) and one of the Hypergraphs generated by its training data (b).

constructed by substituting one entity $h' \in E$ or $t' \in E$ for h or t respectively, contributing to a corrupted triple (h', r, t) or (h, r, t') , which is just simply denoted by (h', r, t') following [3]. Consequently, a sample corresponding to three entities, i.e., h, t, h' or t' , and one relation r . As a result, the training data can be formulated in to a 4-uniform hypergraph, in which all the hyperedges have the same cardinality 4. In this hypergraph, nodes are entities and relations, and edges are training samples containing 4 nodes, i.e., three entities and one relation. More formally,

Definition 3.1. The training data to embed the knowledge graph $G = (E, R, T)$ by translating embedding methods is organized as a 4-uniform hypergraph $H = (V, S)$, where $V = \{E \cup R\}$ is the set of entities or relations, and S is the set of training samples s , where $s = \{h, r, t, h'(\text{or } t') : h, r, t, h', t' \in E, r \in R\}$.

For example, the hypergraph in Figure1(b) is one of the hypergraphs generated by Figure1(a), where black nodes are entities and colored nodes are relations, and the colored blocks represent hyperedges. Here, different colors are related to different relations. For instance, for triple (e_1, r_1, e_3) , the negative triple sampled in Figure1(b) is (e_1, r_1, e_2) , which contributes to a sample $s_1 = \{e_1, r_1, e_2, e_3\}$, thus the hyperedge colored by red contains e_1, e_2, e_3 and r_1 . Note that many other negative triples can be constructed, e.g., (e_1, r_1, e_5) for triple (e_1, r_1, e_3) , and the hypergraph generated in Figure1(b) is just an example. Similarly, the other samples in Figure1(b) are $s_2 = \{e_1, r_2, e_4, e_5\}$, $s_3 = \{e_2, r_3, e_3, e_5\}$ and $s_4 = \{e_4, r_3, e_2, e_3\}$.

To better analyze the collisions between processors, we define the following statistics of the hypergraph H . Given a hyperedge s ,

$$\sigma(s) = \{s' : \exists r \in s \cap s', r \in R\} \quad (2)$$

denotes the set of hyperedges containing the same relations with hyperedge s .

$$\hat{\sigma} = \max_{s \in S} |\sigma(s)| \quad (3)$$

denotes the maximal number of hyperedges containing same relations, where $|\cdot|$ denotes the cardinality.

$$\rho(s) = \{s' : \exists e \in s \cap s', e \in E\} \quad (4)$$

denotes the set of hyperedges containing one or more same entities with hyperedge s .

$$\hat{\rho} := \max_{s \in S} |\rho(s)|. \quad (5)$$

denotes the maximal number of hyperedges containing same entities, where $|\cdot|$ denotes the cardinality the same as before.

3.2 Collision Formulation

In this section, we will verify that it is highly possible that few collisions happen when training by p processors for large and sparse knowledge graphs. Let X_{samp} represent the event that p processors select p different samples. X_{rel} represents the event that there are collisions between relations, i.e., different processors updates a same relation vector, and X_{ent} between entities similarly. The verification is decomposed into two steps, 1) to prove it is quite likely that the processors handle different samples, i.e., $P(X_{samp} = 1) \approx 1$, which is the prerequisite to no collisions; 2) to prove it is unlikely that these different samples correspond to the same relations or entities, i.e., $P(X_{rel} = 0) \approx 1$ and $P(X_{ent} = 0) \approx 1$.

Supposing that for embedding methods and the knowledge graph, the training samples $S = \{s_1, s_2, \dots, s_i, \dots, s_n\}$ of size n is drawn independent and identically distributed (i.e., i.i.d.) from some unknown distribution \mathcal{D} . Therefore, the probability of s_i being selected P_i is supposed to be

$$P_i = \frac{1}{n}. \quad (6)$$

Moreover, according to i.i.d., it is reasonable to assume that the sample selecting process by p processors is an observation from a Multinomial Distribution, i.e., selecting one sample from n samples and repeated p times. Let x_i denote the number of processors that select s_i during the same iteration of SGD, then the possibility of s_i being selected by c_1 processors, \dots , s_{n_e} being selected by c_{n_e} processors is as follows,

$$P(x_1 = c_1, \dots, x_{n_e} = c_{n_e}) = \begin{cases} \frac{p!}{c_1! c_2! \dots c_{n_e}!} P_1^{c_1} \dots P_{n_e}^{c_{n_e}}, & \sum_{i=1}^{n_e} c_i = p \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where $\sum_{i=1}^{n_e} c_i = p$ indicates that there are p and only p samples being selected in the same iteration of SGD.

THEOREM 3.2. *For a knowledge graph with n triples and training by p processors in parallel, when n is large and p is relatively small, the possibility that p processors select p different samples is*

$$P(X_{samp} = 1) \approx 1 \quad (8)$$

with probability at least γ , where

$$\gamma = \prod_{i=1}^{p-1} \left(1 - \frac{i}{n}\right). \quad (9)$$

PROOF. Provided that samples selected by processors are different, it can be easily derived that $\forall s_i \in S, x_i \leq 1$. Then there are only $\binom{n}{p}$ sampling circumstances satisfying no collisions between samples, where p distinct samples are selected once, and other $n-p$ samples are not selected, e.g., $x_1 = 1, x_2 = 1, \dots, x_p = 1, x_{p+1} = 0, \dots, x_n = 0$. Therefore, according to Eq.(7) and Eq.(6),

$$\begin{aligned} P(X_{samp} = 1) &= \binom{n}{p} \frac{p!}{\prod_{i=1}^p 1!} \prod_{i=1}^p (P_i)^1 = \frac{n!}{(n-p)!} \left(\frac{1}{n}\right)^p \\ &= \frac{n(n-1)(n-2) \dots (n-p+1)}{n^p} \end{aligned}$$

$$\begin{aligned} &= \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \dots \left(1 - \frac{p-1}{n}\right) \\ &= \prod_{i=1}^{p-1} \left(1 - \frac{i}{n}\right) \end{aligned}$$

When n is large and p is relatively small, $P(X_{samp} = 1) \approx 1$. \square

THEOREM 3.3. *For a knowledge graph with n triples and training in p processors in parallel, when $\frac{\hat{\sigma}}{n}$ is relatively small and $p < \frac{n}{\hat{\sigma}} + 1$, we have the possibility of no relation in a collision is*

$$P(X_{rel} = 0) \approx 1 \quad (10)$$

with probability at least γ , where

$$\gamma = \prod_{i=1}^{p-1} \left(1 - \frac{i\hat{\sigma}}{n}\right). \quad (11)$$

PROOF. Given that p processors select p different samples, the possibility of relations in a collision can be deduced according to conditional probability as follows,

$$P(X_{rel} = 0) = P(X_{rel} = 0 | X_{samp} = 1) \cdot P(X_{samp} = 1), \quad (12)$$

where $P(X_{rel} = 0 | X_{samp} = 1)$ is the possibility of p samples containing distinct relations being selected, which is supposed to be similar to sampling without replacement. More precisely, assuming a sample s is selected randomly, then the next sample selected s' should be from $S - \sigma(s)$, and the third sample s'' should be selected from in $S - \sigma(s) \cup \sigma(s')$. Accordingly, $P(X_{rel} = 0 | X_{samp} = 1)$ is deduced as follows when $p < \frac{n}{\hat{\sigma}} + 1$ is satisfied,

$$\begin{aligned} P(X_{rel} = 0 | X_{samp} = 1) &= \sum_{s_1, s_2, \dots, s_p \in S} \frac{1}{n} \cdot \frac{|S - \sigma(s_1)|}{n-1} \cdot \frac{|S - \sigma(s_1) \cup \sigma(s_2)|}{n-2} \dots \\ &\quad \frac{|S - \sigma(s_1) \cup \sigma(s_2) \cup \dots \cup \sigma(s_{p-1})|}{n - (p-1)} \\ &\geq \sum_{s_1, s_2, \dots, s_p \in S} \frac{1}{n} \cdot \frac{n - \hat{\sigma}}{n-1} \cdot \frac{n - 2\hat{\sigma}}{n-2} \dots \frac{n - (p-1)\hat{\sigma}}{n - (p-1)} \\ &= \left(1 - \frac{\hat{\sigma} - 1}{n-1}\right) \left(1 - \frac{2(\hat{\sigma} - 1)}{n-2}\right) \dots \left(1 - \frac{(p-1)(\hat{\sigma} - 1)}{n - (p-1)}\right) \\ &= \prod_{i=1}^{p-1} \left(1 - \frac{i(\hat{\sigma} - 1)}{n-i}\right) \end{aligned}$$

By Eq.(12), the possibility of no collisions between relations in different processors is

$$P(X_{rel} = 0) = \prod_{i=1}^{p-1} \left(1 - \frac{i(\hat{\sigma} - 1)}{n-i}\right) \cdot \prod_{i=1}^{p-1} \left(1 - \frac{i}{n}\right) = \prod_{i=1}^{p-1} \left(1 - \frac{i\hat{\sigma}}{n}\right). \quad (13)$$

Note that $p > \frac{n}{\hat{\sigma}} + 1$ results in $1 - \frac{(p-1)(\hat{\sigma}-1)}{n-(p-1)} < 0$, which means $\hat{\sigma}$ is so large that one or more processors will definitely select the same relation among p processors, namely, $P(X_{rel} = 0) = 0$. Furthermore, when $\frac{\hat{\sigma}}{n}$ is relatively small, $P(X_{rel} = 0) \approx 1$. \square

According to $P[i]^{(j)}$, a negative triple $N[i]^{(j)} = (h', r, t')$ is generated by sampling a corrupted entity h' (or t') from shared memory, where i and j are the same as before. That is to say, a sample $S[i]^{(j)}$ is constructed by $P[i]^{(j)}$ and $N[i]^{(j)}$, which then be used to calculate the gradient $\nabla[f_r(h, t) + M - f_r(h', t')]$ according to Eq.(1), and update the embeddings of entities and relations $(h, r, t, h', t')^{(j+1)} \leftarrow (h, r, t, h', t')^{(j)} - \eta \nabla[i]^{(j)}$.

Algorithm 1 ParTrans-X

Require:

Training triples $T = \{(h, r, t)\}$, entities and relations set E and R , embedding dimension d , margin M , training epochs ep , the number of processors p ;

Ensure:

Embeddings of entities and relations;

- 1: Initialize $r \in R$ and $e \in E$ by uniform distribution and persist them in the shared memory
 - 2: **for** $i \leftarrow 0$ **to** p **do**
 - 3: $ep' \leftarrow \frac{ep}{p}$ ▷ In Parallel
 - 4: $seed[i] \leftarrow SEED_RAND(i)$
 - 5: **for** $j \leftarrow 0$ **to** ep' **do**
 - 6: **loop**
 - 7: $e := \frac{e}{\|e\|}$ for each entity $e \in E$
 - 8: $P[i]^{(j)} \leftarrow SAMPLE(T, seed[i])$
 - 9: $N[i]^{(j)} \leftarrow SAMPLE_NEG(P[i]^{(j)}, T, E, seed[i])$
 - 10: $S[i]^{(j)} \leftarrow [P[i]^{(j)}, N[i]^{(j)}]$
 - 11: $\nabla[i]^{(j)} \leftarrow \nabla[f_r(h, t) + M - f_r(h', t')]$, where $h, r, t, h', t' \in S[i]^{(j)}$
 - 12: $(h, r, t, h', t')^{(j+1)} \leftarrow (h, r, t, h', t')^{(j)} - \eta \nabla[i]^{(j)}$
 - 13: **end loop**
 - 14: **end for**
 - 15: **end for**
 - 16: Generate embeddings of E and R after all processors finish
-

4.2 Application to Typical translating embedding Methods

The framework can be applied to many translating embedding methods, which employ SGD or its variants to optimize the hinge loss with similar algorithm framework, and are only different in the score function $f_r(h, t)$ as mentioned in Sec.2, e.g., TransE, TransH and so on. Hence, the parallel algorithm of them can be obtained by applying the corresponding score function in Lines 11-12 of the pseudocode in Algorithm 1.

For example, for **TransE**, the gradient updating procedure in Lines 11 is performed according to Eq.(16). For **TransH**, which employs the score function $f_r(h, t) = \|\mathbf{h}_\perp + \mathbf{r} - \mathbf{t}_\perp\|$, the gradient updating procedure of \mathbf{h} in Lines 11 is as follows,

$$\mathbf{h}_k := \mathbf{h}_k - \eta \cdot 2 \left[(\mathbf{h}_k - \mathbf{w}_r^\top \mathbf{h} \mathbf{w}_{rk}) + \mathbf{r}_k - (\mathbf{t}_k - \mathbf{w}_r^\top \mathbf{t} \mathbf{w}_{rk}) \right]. \quad (17)$$

Namely, ParTrans-X has the flexibility to parallel many translating embedding methods, since they possess similar training process.

Moreover, ParTrans-X can be directly applied to the improved variant **TransE-AdaGrad**, since the training data sparsity of knowledge graph still holds. In one iteration of AdaGrad, it updates the embedding vectors according to the gradient from the previous iteration. Highly similar to SGD, AdaGrad can be easily paralleled using our framework by only performing a learning rate calculation procedure during the gradient update procedure, i.e., Line 12 of the pseudocode in Algorithm 1. For example, to parallel TransE-AdaGrad, the learning rate is determined adaptively by adding

$$\eta^{(j)} := \frac{\nabla^{(j)}}{\sqrt{\sum_{k=1}^j (\nabla^{(k)})^2}} \eta^* \quad (18)$$

before Line 12 in Algorithm 1, where j is the current epoch, with $\eta^{(j)}$ the learning rate of j -th epoch. $\nabla^{(k)}$, $k < j$ represents all the previous gradient before j -th epoch. η^* is the initial learning rate.

5 EXPERIMENT

Firstly, we apply ParTrans-X to TransE, TransH and TransE-Adagrad in Sec.5.1. In Sec.5.2, experiment results demonstrate excessive decline in training time by ParTrans-X, with scaling performance along with increasing number of processors shown in Sec.5.3.

5.1 Experimental Settings

The datasets employed are two representative datasets WN18 and FB15k, which are subsets of well-known knowledge graphs WordNet and Freebase respectively, and have been widely used by translating embedding methods [3, 8, 11, 19]. Table2 shows the statistics of them. Without loss of generality, $\frac{\rho}{n}$ and $\frac{\bar{\rho}}{n}$ are also shown, and they are both small on WN18 and FB15k. Furthermore, it can be seen that the two datasets possess different characteristics. Namely, WN18 possesses only 18 relations, which results in large possibility of collisions between relations. On the contrary, FB15k is less unbalanced in the number of entities and relations.

Table 2: Two widely used datasets in KGs.

Data	# Rel	# Ent	#Train	#Valid	#Test	$\bar{\sigma}/n$	$\bar{\rho}/n$
WN18	18	40,943	141,442	5,000	5,000	1.7e-1	5.6e-4
FB15k	1,345	14,951	483,142	50,000	59,071	9.3e-3	2.3e-3

To tackle the *KGC* problem, experiments are conducted on the link prediction task which aims to predict the missing entities h or t for a triple (h, r, t) . Namely, it predicts t given (h, r) or predict h given (r, t) . Similar to the setting in [3], the task returns a list of candidate entities from the knowledge graph.

To evaluate the performance of link prediction, we adopt Mean Rank and Hits@10 under ‘‘Raw’’ and ‘‘Filter’’ settings as evaluation measure following [3]. Mean Rank is the average rank of the correct entities, and Hits@10 is proportion of correct entities ranked in top-10. It is clear that a good predictor has low mean rank and high Hits@10. This is called ‘‘Raw’’ setting, and ‘‘Filter’’ setting filters out the corrupted triples which are correct.

To evaluate the speed up performance, we adopt Training Time and Speed-up Ratio as evaluation measures, where Training Time is measured using wall-clock seconds. Speed-up Ratio is

$$Speedup \quad Ratio = \frac{t_{serial}}{t_{parallel}}, \quad (19)$$

Table 3: Link prediction performance with all time measured in wall-clock seconds.

Metric	WN18						FB15k					
	Mean Rank		Hits@10		Training Time(s)	Speedup Ratio	Mean Rank		Hits@10		Training Time(s)	Speedup Ratio
	Raw	Filter	Raw	Filter			Raw	Filter				
TransE	214	203	58.2	65.9	473	-	184	73	44.5	60.7	4658	-
ParTransE	217	206	55.7	63.1	54	9	185	74	44.4	60.5	364	13
TransE-AdaGrad	209	197	68.9	77.7	100	4.7	185	69	45.3	62.3	496	9
ParTransE-AdaGrad	219	208	67.7	76.2	17	28 (4.7×6)	186	70	44.9	61.9	42	111 (9×12)
TransH	227	216	66.5	75.9	637	-	183	60	46.6	65.5	6066	-
ParTransH	215	203	66.8	76.6	134	4.8	183	60	46.8	65.7	474	13

where t_{serial} is the training time in serial, and $t_{parallel}$ is the training time under parallel methods.

Baselines include typical translating embedding methods, TransE, TransH and TransE-Adagrad, which can all be trained in parallel using the ParTrans-X framework, denoted by ParTransE, ParTransH and ParTransE-Adagrad respectively in Table3. Note that TransE and TransH adopt the programs publicly available⁵, which are the most efficient serial versions to our knowledge, and TransE-Adagrad is implemented based on TransE.

Each experiment is conducted 10 times and the average is taken as results, with all time measured in wall-clock seconds. Our experiments are carried out on dual Intel Xeon E5-2640 CPUs, and each of them possesses 10 physical cores 20 logical cores and running at 2.4 GHz. The machine has 128 GB RAM and runs Red Hat 4.4.7. The language used is C++ and the program is compiled with the gcc compiler version 6.3.0. We use OpenMP for multithreading, each thread binds a processor.

5.2 Link Prediction Performance of ParTrans-X

Experiments on each baseline and its parallel implementation in ParTrans-X employ the same hyper-parameters, which are decided on the validation set. The learning rate η during the stochastic gradient descent process is selected among $\{0.1, 0.01, 0.001\}$, the embedding dimension d_e and d_r are selected in $\{20, 50, 100\}$, the margin M between positive and negative triples is selected among $\{1, 2, 3, 4\}$. For TransE and ParTransE, the parameters are $\eta = 0.01$, $d_e = d_r = 20$, $M = 3$ on WN18, and $\eta = 0.001$, $d_e = d_r = 100$, $M = 4$ on Fb15k. For TransH and ParTransH, the parameters are $\eta = 0.01$, $d_e = d_r = 20$, $M = 3$ on WN18, and $\eta = 0.001$, $d_e = d_r = 100$, $M = 3$ on Fb15k. For TransEAdaGrad and ParTransE-AdaGrad, the parameters are $\eta^* = 0.3$, $d_e = d_r = 50$, $M = 4$ on WN18, and $\eta^* = 0.1$, $d_e = d_r = 100$, $M = 3$ on Fb15k. All the experiments employ L1-similarity. ParTransE, ParTransH and ParTransE-AdaGrad all run in 20 processors for both datasets.

It can be observed from Table3 that:

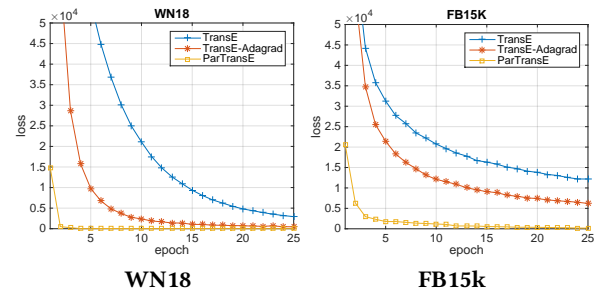
1. Link prediction performance in parallel is as good as the serial counterparts on both WN18 and FB15k, which demonstrates that ParTrans-X will not affect embedding performance.

2. The training time is greatly reduced by ParTrans-X. On WN18, TransE-AdaGrad only speeds up TransE by 4.7 times, compared to

our 28 times. On FB15k, the training time of TransE is reduced from more than 1 hour to less than 1 minute by ParTransE-AdaGrad.

3. ParTrans-X achieves higher speedup ratio on FB15k than on WN18. Since FB15k has far more training triples than WN18, the time of each epoch on FB15k is much longer than WN18. As a result, the overhead of multi-threading is less important compared to the whole training time on FB15k, which leads to a higher speedup ratio. It further validates the superiority of ParTrans-X to handle the data with large size.

4. ParTrans-X achieves enormous improvement on training time when applying to TransEAdaGrad, especially on FB15k, where the speedup ratio has been improve to 111 from 9. Since AdaGrad decreases the total epochs needed by making the convergence come earlier, and ParTrans-X reduce training time by running in parallel, the two different strategies can achieve higher speedup ratio when combined.

**Figure 3: The descent process of loss.**

Moreover, the descent process of loss for the three algorithms on WN18 and FB15k is shown in Figure3. It can be seen that, for both datasets, the loss optimizing by ParTrans-X has already fallen sharply in the preceding epochs, and it yields sensibly lower values of the loss than TransE-AdaGrad and TransE even after a few iterations (< 5 epoches). Still, ParTrans-X performs better on FB15k than WN18, shows that it is more effective on large data size.

5.3 Scaling Results for Multi-Processors

Furthermore, we carry out a number of experiments to test if the implementations scale with increasing number of processors. We mainly analyze two aspects of experiment results, i.e., the training time and the link prediction performance.

⁵<https://github.com/thunlp/KB2E>

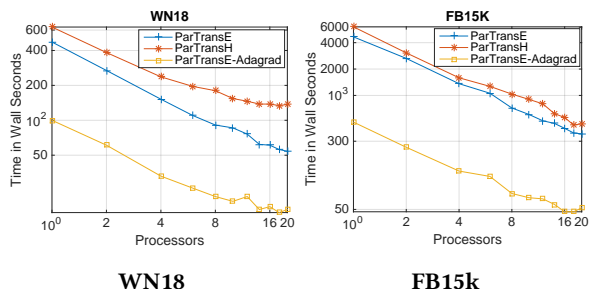


Figure 4: Log-log plot of Training Time along with number of processors

Figure 4 shows the log-log plot of the training time in wall-clock seconds for different number of processors. We can observe that the training time continue to decrease along with the increasing number of mutli-processors on both WN18 and FB15k. While the absolute training time of ParTransE-AdaGrad is better than ParTransE, which is better than ParTransH, consistent with the previous result. Moreover, the total training time of ParTransE-AdaGrad drops sharply when processor number is less than four, it is because the training time of ParTransE-AdaGrad with few processors is fairly short, the increase of communication time cost with the more processors has larger effect on the total training time compared with other methods, which leads to small decline.

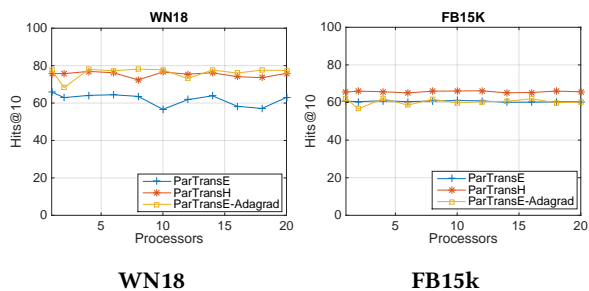


Figure 5: Hits@10 performance along with number of processors

The predictive performance measured by Hits@10 along with increasing number of processors is shown in Figure 5. It can be seen that ParTransE, ParTransE-AdaGrad and ParTransH always maintain good performance, which validates the applicability and superiority of ParTrans-X. Note that the performance on FB15k is more stable than WN18, since there are more training triples in FB15k, and the model will learn more sufficient so that the stability of predictive performance is better on FB15k, which validates the superiority of ParTrans-X on large data size.

6 CONCLUSION

In this paper, we explore the law of collisions emerging in knowledge graphs by modelling training data to hypergraphs. Our key observation is that one learning iteration only concerns few embeddings, which is not necessarily bound up with others, thus the probability of collisions between different processors can be negligible. Based on this assumption, we propose an efficient parallel framework for translating embedding methods, called ParTrans-X.

It employs the intrinsic sparsity of training data in large knowledge graphs, which enables the embedding vectors to be learnt without locks and not inducing errors. Experiments validate that ParTrans-X can speed up the training process by more than an order of magnitude, without degrading embedding performance. The source code of this paper can be obtained from here⁶.

7 ACKNOWLEDGE

We thank Jun Xu and the anonymous reviewers for valuable suggestions. The work was funded by National Natural Science Foundation of China (No. 61572469, 61402442, 91646120, 61572473, 61402022), the National Key R&D Program of China (No. 2016QY02D0405, 2016YFB1000902), and National Grand Fundamental Research 973 Program of China (No. 2013CB329602, 2014CB340401).

REFERENCES

- [1] Martn Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, and Matthieu Devin. 2016. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. (2016).
- [2] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. ACM, 1247–1250.
- [3] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*. 2787–2795.
- [4] Jeffrey Dean, Greg S Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V Le, Mark Z Mao, Marc'Aurelio Ranzato, Andrew Senior, and Paul Tucker. 2012. Large scale distributed deep networks. In *International Conference on Neural Information Processing Systems*. 1223–1231.
- [5] Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014. *Knowledge vault: a web-scale approach to probabilistic knowledge fusion*. ACM. 601–610 pages.
- [6] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12, Jul (2011), 2121–2159.
- [7] Yantao Jia, Yuanzhuo Wang, Xueqi Cheng, Xiaolong Jin, and Jiafeng Guo. 2014. OpenKN: An open knowledge computational engine for network big data. In *Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 657–664.
- [8] Yantao Jia, Yuanzhuo Wang, Hailun Lin, Xiaolong Jin, and Xueqi Cheng. 2016. Locally Adaptive Translation for Knowledge Graph Embedding. In *AAAI*.
- [9] John Langford, Alexander J Smola, and Martin Zinkevich. 2009. Slow learners are fast. In *International Conference on Neural Information Processing Systems*. 2331–2339.
- [10] Yankai Lin, Zhiyuan Liu, Huanbo Luan, Maosong Sun, Siwei Rao, and Song Liu. 2015. Modeling Relation Paths for Representation Learning of Knowledge Bases. *Computer Science* (2015).
- [11] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*. 2181–2187.
- [12] Farzaneh Mahdisoltani, Joanna Biega, and Fabian Suchanek. 2014. YAGO3: A Knowledge Base from Multilingual Wikipedias. (2014).
- [13] H. B. McMahan and M. Streeter. 2014. Delay-tolerant algorithms for asynchronous distributed online learning. *Advances in Neural Information Processing Systems* 4 (2014), 2915–2923.
- [14] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.
- [15] Pasquale Minervini, Claudia d'Amato, Nicola Fanizzi, and Floriana Esposito. 2015. Efficient Learning of Entity and Predicate Embeddings for Link Prediction in Knowledge Graphs. In *URSW@ ISWC*. 26–37.
- [16] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. 2011. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*. 693–701.
- [17] Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. (2016).
- [18] Bin Shao, Yatao Li, and Haixun Wang. 2015. Parallel Processing of Graphs. (2015).

⁶<https://github.com/zdh2292390/ParTrans-X>

- [19] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. *AAAI - Association for the Advancement of Artificial Intelligence* (2014).
- [20] Sixin Zhang, Anna Choromanska, and Yann Lecun. 2015. Deep learning with Elastic Averaging SGD. *Computer Science* (2015), 685–693.