

Analyzing Knowledge Graph Embedding Methods from a Multi-Embedding Interaction Perspective

Hung Nghiep Tran
SOKENDAI

(The Graduate University for Advanced Studies)
Japan
nghiepth@nii.ac.jp

Atsuhiko Takasu

National Institute of Informatics
Tokyo, Japan
takasu@nii.ac.jp

ABSTRACT

Knowledge graph is a popular format for representing knowledge, with many applications to semantic search engines, question-answering systems, and recommender systems. Real-world knowledge graphs are usually incomplete, so knowledge graph embedding methods, such as Canonical decomposition/Parallel factorization (CP), DistMult, and ComplEx, have been proposed to address this issue. These methods represent entities and relations as embedding vectors in semantic space and predict the links between them. The embedding vectors themselves contain rich semantic information and can be used in other applications such as data analysis. However, mechanisms in these models and the embedding vectors themselves vary greatly, making it difficult to understand and compare them. Given this lack of understanding, we risk using them ineffectively or incorrectly, particularly for complicated models, such as CP, with two role-based embedding vectors, or the state-of-the-art ComplEx model, with complex-valued embedding vectors. In this paper, we propose a multi-embedding interaction mechanism as a new approach to uniting and generalizing these models. We derive them theoretically via this mechanism and provide empirical analyses and comparisons between them. We also propose a new multi-embedding model based on quaternion algebra and show that it achieves promising results using popular benchmarks.

KEYWORDS

Knowledge Graph, Knowledge Graph Completion, Knowledge Graph Embedding, Multi-Embedding, Representation Learning.

1 INTRODUCTION

Knowledge graphs provide a unified format for representing knowledge about relationships between entities. A knowledge graph is a collection of triples, with each triple (h, t, r) denoting the fact that relation r exists between head entity h and tail entity t . Many large real-world knowledge graphs have been built, including WordNet [22] representing English lexical knowledge, and Freebase [3] and Wikidata [29] representing general knowledge. Moreover, knowledge graph can be used as a universal format for data from applied domains. For example, a knowledge graph for recommender systems would have triples such as $(UserA, Item1, review)$ and $(UserB, Item2, like)$.

Knowledge graphs are the cornerstones of modern semantic web technology. They have been used by large companies such as

Google to provide semantic meanings into many traditional applications, such as semantic search engines, semantic browsing, and question answering [2]. One important application of knowledge graphs is recommender systems, where they are used to unite multiple sources of data and incorporate external knowledge [5] [36]. Recently, specific methods such as knowledge graph embedding have been used to predict user interactions and provide recommendations directly [10].

Real-world knowledge graphs are usually incomplete. For example, Freebase and Wikidata are very large but they do not contain all knowledge. This is especially true for the knowledge graphs used in recommender systems. During system operation, *users review new items* or *like new items*, generating new triples for the knowledge graph, which is therefore inherently incomplete. Knowledge graph completion, or link prediction, is the task that aims to predict new triples.

This task can be undertaken by using knowledge graph embedding methods, which represent entities and relations as embedding vectors in semantic space, then model the interactions between these embedding vectors to compute matching scores that predict the validity of each triple. Knowledge graph embedding methods are not only used for knowledge graph completion, but the learned embedding vectors of entities and relations are also very useful. They contain rich semantic information similar to word embeddings [21] [20] [14], enabling them to be used in visualization or browsing for data analysis. They can also be used as extracted or pretrained feature vectors in other learning models for tasks such as classification, clustering, and ranking.

Among the many proposed knowledge graph embedding methods, the most efficient and effective involve trilinear-product-based models, such as Canonical decomposition/Parallel factorization (CP) [13] [17], DistMult [35], or the state-of-the-art ComplEx model [28]. These models solve a tensor decomposition problem with the matching score of each triple modeled as the result of a trilinear product, i.e., a multilinear map with three variables corresponding to the embedding vectors \mathbf{h} , \mathbf{t} , and \mathbf{r} of head entity h , tail entity t , and relation r , respectively. The trilinear-product-based score function for the three embedding vectors is denoted as $\langle \mathbf{h}, \mathbf{t}, \mathbf{r} \rangle$ and will be defined mathematically in Section 2.

However, the implementations of embedding vectors for the various models are very diverse. DistMult [35] uses one real-valued embedding vector for each entity or relation. The original CP [13] uses one real-valued embedding vector for each relation, but two real-valued embedding vectors for each entity when it is as head and as tail, respectively. ComplEx [28] uses one complex-valued embedding vector for each entity or relation. Moreover, a recent heuristic for CP [17], here denoted as CP_h , was proposed to augment the training data, helping CP achieve results competitive with the state-of-the-art model ComplEx. This heuristic introduces an additional embedding vector for each relation, but

First International Workshop on Data Science for Industry 4.0.
Copyright ©2019 for the individual papers by the papers' authors. Copying permitted for private and academic purposes. This volume is published and copyrighted by its editors.

the underlying mechanism is different from that in ComplEx. All of these complications make it difficult to understand and compare the various models, to know how to use them and extend them. If we were to use the embedding vectors for data analysis or as pretrained feature vectors, a good understanding would affect the way we would use the complex-valued embedding vectors from ComplEx or the different embedding vectors for head and tail roles from CP.

In this paper, we propose a multi-embedding interaction mechanism as a new approach to uniting and generalizing the above models. In the proposed mechanism, each entity e is represented by multiple embedding vectors $\{e^{(1)}, e^{(2)}, \dots\}$ and each relation r is represented by multiple embedding vectors $\{r^{(1)}, r^{(2)}, \dots\}$. In a triple (h, t, r) , all embedding vectors of h , t , and r interact with each other by trilinear products to produce multiple interaction scores. These scores are then weighted summed by a weight vector ω to produce the final matching score for the triple. We show that the above models are special cases of this mechanism. Therefore, it unifies those models and lets us compare them directly. The mechanism also enables us to develop new models by extending to additional embedding vectors.

In this paper, our contributions include the following.

- We introduce a multi-embedding interaction mechanism as a new approach to unifying and generalizing a class of state-of-the-art knowledge graph embedding models.
- We derive each of the above models theoretically via this mechanism. We then empirically analyze and compare these models with each other and with variants.
- We propose a new multi-embedding model by an extension to four-embedding vectors based on quaternion algebra, which is an extension of complex algebra. We show that this model achieves promising results.

2 RELATED WORK

Knowledge graph embedding methods for link prediction are actively being researched [30]. Here, we only review the works that are directly related to this paper, namely models that use only triples, not external data such as text [32] or graph structure such as relation paths [18]. Models using only triples are relatively simple and they are also the current state of the art.

2.1 General architecture

Knowledge graph embedding models take a triple of the form (h, t, r) as input and output the validity of that triple. A general model can be viewed as a three-component architecture:

- (1) *Embedding lookup*: linear mapping from one-hot vectors to embedding vectors. A one-hot vector is a sparse discrete vector representing a discrete input, e.g., the first entity could be represented as $[1, 0, \dots, 0]^T$. A triple could be represented as a tuple of three one-hot vectors representing h , t , and r , respectively. An embedding vector is a dense continuous vector of much lower dimensionality than a one-hot vector thus lead to efficient distributed representations [11] [12].
- (2) *Interaction mechanism*: modeling the interaction between embedding vectors to compute the matching score of a triple. This is the main component of a model.
- (3) *Prediction*: using the matching score to predict the validity of each triple. A higher score means that the triple is more likely to be valid.

2.2 Categorization

Based on the modeling of the second component, a knowledge graph embedding model falls into one of three categories, namely *translation-based*, *neural-network-based*, or *trilinear-product-based*, as described below.

2.2.1 Translation-based: These models translate the head entity embedding by summing with the relation embedding vector, then measuring the distance between the translated images of head entity and the tail entity embedding, usually by L^1 or L^2 distance:

$$\begin{aligned} S(h, t, r) &= -\|h + r - t\|_p \\ &= -\left(\sum_d |h_d + r_d - t_d|^p\right)^{1/p}, \end{aligned} \quad (1)$$

where

- h, t, r are embedding vectors of h, t , and r , respectively,
- p is 1 or 2 for L^1 or L^2 distance, respectively,
- D is the embedding size and d is each dimension.

TransE [4] was the first model of this type, with score function basically the same as the above equation. There have been many extensions such as TransR [19], TransH [33], and TransA [34]. Most extensions are done by linear transformation of the entities into a relation-specific space before translation [19].

These models are simple and efficient. However, their modeling capacity is generally weak because of over-strong assumptions about translation using relation embedding. Therefore, they are unable to model some forms of data [31].

2.2.2 Neural-network-based: These models use a nonlinear neural network to compute the matching score for a triple:

$$S(h, t, r) = NN(h, t, r), \quad (2)$$

where

- h, t, r are the embedding vectors of h, t , and r , respectively,
- NN is the neural network used to compute the score.

One of the simplest neural-network-based model is ER-MLP [7], which concatenates the input embedding vectors and uses a multi-layer perceptron neural network to compute the matching score. NTN [26] is an earlier model that employs nonlinear activation functions to generalize the linear model RESCAL [24]. Recent models such as ConvE [6] use convolution networks instead of fully-connected networks.

These models are complicated because of their use of neural networks as a black-box universal approximator, which usually make them difficult to understand and expensive to use.

2.2.3 Trilinear-product-based: These models compute their scores by using trilinear product between head, tail, and relation embeddings, with relation embedding playing the role of matching weights on the dimensions of head and tail embeddings:

$$\begin{aligned} S(h, t, r) &= \langle h, t, r \rangle \\ &= h^T \text{diag}(r) t \\ &= \sum_{d=1}^D (h \odot t \odot r)_d \\ &= \sum_{d=1}^D (h_d t_d r_d), \end{aligned} \quad (3)$$

where

- $\mathbf{h}, \mathbf{t}, \mathbf{r}$ are embedding vectors of $h, t,$ and $r,$ respectively,
- $\text{diag}(\mathbf{r})$ is the diagonal matrix of $\mathbf{r},$
- \odot denotes the element-wise Hadamard product,
- D is the embedding size and d is the dimension for which $\mathbf{h}_d, \mathbf{t}_d,$ and \mathbf{r}_d are the entries.

In this paper, we focus on this category, particularly on DistMult, ComplEx, CP, and CP_h with augmented data. These models are simple, efficient, and can scale linearly with respect to embedding size in both time and space. They are also very effective, as has been shown by the state-of-the-art results for ComplEx and CP_h using popular benchmarks [28] [17].

DistMult [35] embeds each entity and relation as a single real-valued vector. DistMult is the simplest model in this category. Its score function is symmetric, with the same scores for triples (h, t, r) and $(t, h, r).$ Therefore, it cannot model asymmetric data for which only one direction is valid, e.g., asymmetric triples such as (*Paper1, Paper2, cite*). Its score function is:

$$\mathcal{S}(h, t, r) = \langle \mathbf{h}, \mathbf{t}, \mathbf{r} \rangle, \quad (4)$$

where $\mathbf{h}, \mathbf{t}, \mathbf{r} \in \mathbb{R}^k.$

ComplEx [28] is an extension of DistMult that uses complex-valued embedding vectors that contain complex numbers. Each complex number c with two components, real a and imaginary $b,$ can be denoted as $c = a + bi.$ The complex conjugate \bar{c} of c is $\bar{c} = a - bi.$ The complex conjugate vector $\bar{\mathbf{t}}$ of \mathbf{t} is form from the complex conjugate of the individual entries. Complex algebra requires using the complex conjugate vector of tail embedding in the inner product and trilinear product [1]. Thus, these products can be antisymmetric, which enables ComplEx to model asymmetric data [28] [27]. Its score function is:

$$\mathcal{S}(h, t, r) = \text{Re}(\langle \mathbf{h}, \bar{\mathbf{t}}, \mathbf{r} \rangle), \quad (5)$$

where $\mathbf{h}, \mathbf{t}, \mathbf{r} \in \mathbb{C}^k$ and $\text{Re}(c)$ means taking the real component of the complex number $c.$

CP [13] is similar to DistMult but embeds entities as head and as tail differently. Each entity e has two embedding vectors \mathbf{e} and $\mathbf{e}^{(2)}$ depending on its role in a triple as head or as tail, respectively. Using different role-based embedding vectors leads to an asymmetric score function, enabling CP to also model asymmetric data. However, experiments have shown that CP's performance is very poor on unseen test data [17]. Its score function is:

$$\mathcal{S}(h, t, r) = \langle \mathbf{h}, \mathbf{t}^{(2)}, \mathbf{r} \rangle, \quad (6)$$

where $\mathbf{h}, \mathbf{t}^{(2)}, \mathbf{r} \in \mathbb{R}^k.$

CP_h [17] is a direct extension of CP. Its heuristic augments the training data by making an inverse triple $(t, h, r^{(a)})$ for each existing triple $(h, t, r),$ where $r^{(a)}$ is the augmented relation corresponding to $r.$ With this heuristic, CP_h significantly improves CP, achieving results competitive with ComplEx. Its score function is:

$$\mathcal{S}(h, t, r) = \langle \mathbf{h}, \mathbf{t}^{(2)}, \mathbf{r} \rangle \text{ and } \langle \mathbf{t}, \mathbf{h}^{(2)}, \mathbf{r}^{(a)} \rangle, \quad (7)$$

where $\mathbf{h}, \mathbf{h}^{(2)}, \mathbf{t}, \mathbf{t}^{(2)}, \mathbf{r}, \mathbf{r}^{(a)} \in \mathbb{R}^k.$

In the next section, we present a new approach to analyzing these trilinear-product-based models.

3 MULTI-EMBEDDING INTERACTION

In this section, we first formally present the multi-embedding interaction mechanism. We then derive each of the above trilinear-product-based models using this mechanism, by changing the

embedding vectors and setting appropriate weight vectors. Next, we specify our attempt at learning weight vectors automatically. We also propose a four-embedding interaction model based on quaternion algebra.

3.1 Multi-embedding interaction mechanism

We globally model each entity e as the multiple embedding vectors $\{\mathbf{e}^{(1)}, \mathbf{e}^{(2)}, \dots, \mathbf{e}^{(n)}\}$ and each relation r as the multiple embedding vectors $\{\mathbf{r}^{(1)}, \mathbf{r}^{(2)}, \dots, \mathbf{r}^{(n)}\}.$ The triple (h, t, r) is therefore modeled by multiple embeddings as $\mathbf{h}^{(i)}, \mathbf{t}^{(j)}, \mathbf{r}^{(k)}, i, j, k \in \{1, \dots, n\}.$

In each triple, the embedding vectors for head, tail, and relation interact with each and every other embedding vector to produce multiple interaction scores. Each interaction is modeled by the trilinear product of corresponding embedding vectors. The interaction scores are then weighted summed by a weight vector:

$$\mathcal{S}(h, t, r; \Theta, \omega) = \sum_{i, j, k \in \{1, \dots, n\}} \omega^{(i, j, k)} \langle \mathbf{h}^{(i)}, \mathbf{t}^{(j)}, \mathbf{r}^{(k)} \rangle, \quad (8)$$

where

- Θ is the parameter denoting embedding vectors $\mathbf{h}^{(i)}, \mathbf{t}^{(j)}, \mathbf{r}^{(k)},$
- ω is the parameter denoting the weight vector used to combine the interaction scores, with $\omega^{(i, j, k)}$ being an element of $\omega.$

3.2 Deriving trilinear-product-based models

The existing trilinear-product-based models can be derived from the proposed general multi-embedding interaction score function in Eq. (8) by setting the weight vector ω as shown in Table 1.

For DistMult, we can see the equivalence directly. For ComplEx, we need to expand its score function following complex algebra [1]:

$$\begin{aligned} \mathcal{S}(h, t, r) &= \text{Re}(\langle \mathbf{h}, \bar{\mathbf{t}}, \mathbf{r} \rangle) \\ &= \langle \text{Re}(\mathbf{h}), \text{Re}(\mathbf{t}), \text{Re}(\mathbf{r}) \rangle + \langle \text{Re}(\mathbf{h}), \text{Im}(\mathbf{t}), \text{Im}(\mathbf{r}) \rangle \\ &\quad - \langle \text{Im}(\mathbf{h}), \text{Re}(\mathbf{t}), \text{Im}(\mathbf{r}) \rangle + \langle \text{Im}(\mathbf{h}), \text{Im}(\mathbf{t}), \text{Re}(\mathbf{r}) \rangle, \end{aligned} \quad (9)$$

where

- $\mathbf{h}, \mathbf{t}, \mathbf{r} \in \mathbb{C}^k,$
- $\text{Re}(c)$ and $\text{Im}(c)$ mean taking the real and imaginary components of the complex vector $c,$ respectively.

Changing $\text{Re}(\mathbf{h})$ to $\mathbf{h}^{(1)}, \text{Im}(\mathbf{h})$ to $\mathbf{h}^{(2)}, \text{Re}(\mathbf{t})$ to $\mathbf{t}^{(1)}, \text{Im}(\mathbf{t})$ to $\mathbf{t}^{(2)}, \text{Re}(\mathbf{r})$ to $\mathbf{r}^{(1)},$ and $\text{Im}(\mathbf{r})$ to $\mathbf{r}^{(2)},$ we can rewrite the score function of ComplEx as:

$$\begin{aligned} \mathcal{S}(h, t, r) &= \text{Re}(\langle \mathbf{h}, \bar{\mathbf{t}}, \mathbf{r} \rangle) \\ &= \langle \mathbf{h}^{(1)}, \mathbf{t}^{(1)}, \mathbf{r}^{(1)} \rangle + \langle \mathbf{h}^{(1)}, \mathbf{t}^{(2)}, \mathbf{r}^{(2)} \rangle \\ &\quad - \langle \mathbf{h}^{(2)}, \mathbf{t}^{(1)}, \mathbf{r}^{(2)} \rangle + \langle \mathbf{h}^{(2)}, \mathbf{t}^{(2)}, \mathbf{r}^{(1)} \rangle, \end{aligned} \quad (10)$$

which is equivalent to the weighted sum using the weight vectors in Table 1. Note that by the symmetry between h and $t,$ we can also obtain the equivalent weight vector **ComplEx equiv. 1.** By symmetry between embedding vectors of the same entity or relation, we can also obtain the equivalent weight vectors **ComplEx equiv. 2** and **ComplEx equiv. 3.**

For CP, note that the two role-based embedding vectors for each entity can be mapped to two-embedding vectors in our model and the relation embedding vector can be mapped to $\mathbf{r}^{(1)}.$ For $\text{CP}_h,$ further note that its data augmentation is equivalent to adding the score of the original triple and the inverse triple

Table 1: Weight vectors for special cases.

Weighted terms	DistMult	ComplEx	ComplEx equiv. 1	ComplEx equiv. 2	ComplEx equiv. 3	CP	CP _h	CP _h equiv.
$\langle \mathbf{h}^{(1)}, \mathbf{t}^{(1)}, \mathbf{r}^{(1)} \rangle$	1	1	1	0	0	0	0	0
$\langle \mathbf{h}^{(1)}, \mathbf{t}^{(1)}, \mathbf{r}^{(2)} \rangle$	0	0	0	1	1	0	0	0
$\langle \mathbf{h}^{(1)}, \mathbf{t}^{(2)}, \mathbf{r}^{(1)} \rangle$	0	0	0	-1	1	1	1	0
$\langle \mathbf{h}^{(1)}, \mathbf{t}^{(2)}, \mathbf{r}^{(2)} \rangle$	0	1	-1	0	0	0	0	1
$\langle \mathbf{h}^{(2)}, \mathbf{t}^{(1)}, \mathbf{r}^{(1)} \rangle$	0	0	0	1	-1	0	0	1
$\langle \mathbf{h}^{(2)}, \mathbf{t}^{(1)}, \mathbf{r}^{(2)} \rangle$	0	-1	1	0	0	0	1	0
$\langle \mathbf{h}^{(2)}, \mathbf{t}^{(2)}, \mathbf{r}^{(1)} \rangle$	0	1	1	0	0	0	0	0
$\langle \mathbf{h}^{(2)}, \mathbf{t}^{(2)}, \mathbf{r}^{(2)} \rangle$	0	0	0	1	1	0	0	0

when training using stochastic gradient descent (SGD):

$$\mathcal{S}(h, t, r) = \langle \mathbf{h}, \mathbf{t}^{(2)}, \mathbf{r} \rangle + \langle \mathbf{t}, \mathbf{h}^{(2)}, \mathbf{r}^{(a)} \rangle. \quad (11)$$

We can then map $\mathbf{r}^{(a)}$ to $\mathbf{r}^{(2)}$ to obtain the equivalence given in Table 1. By symmetry between h and t , we can also obtain the equivalent weight vector **CP_h equiv. 1**.

From this perspective, all four models DistMult, ComplEx, CP, and CP_h can be seen as special cases of the general multi-embedding interaction mechanism. This provides an intuitive perspective on using the embedding vectors in complicated models. For the ComplEx model, instead of using a complex-valued embedding vector, we can treat it as two real-valued embedding vectors. These vectors can then be used directly in common learning algorithms that take as input real-valued vectors rather than complex-valued vectors. We also see that multiple embedding vectors are a natural extension of single embedding vectors. Given this insight, multiple embedding vectors can be concatenated to form a longer vector for use in visualization and data analysis, for example.

3.3 Automatically learning weight vectors

As we have noted, the weight vector ω plays an important role in the model, because it determines how the interaction mechanism is implemented and therefore how the specific model can be derived. An interesting question is how to learn ω automatically. One approach is to let the model learn ω together with the embeddings in an end-to-end fashion. For a more detailed examination of this idea, we will test different restrictions on the range of ω by applying $\tanh(\omega)$, $\text{sigmoid}(\omega)$, and $\text{softmax}(\omega)$.

Note also that the weight vectors for related models are usually sparse. We therefore enforce a sparsity constraint on ω by an additional Dirichlet negative log-likelihood regularization loss:

$$\mathcal{L}_{dir} = -\lambda_{dir} \sum_{i,j,k \in \{1, \dots, n\}} (\alpha - 1) \log \frac{|\omega^{(i,j,k)}|}{\|\omega\|_1}, \quad (12)$$

where α is a hyperparameter controlling sparseness (a small α will make the weight vector sparser) and λ_{dir} is the regularization strength.

3.4 Quaternion-based four-embedding interaction model

Another question is whether using more embedding vectors in the multi-embedding interaction mechanism is helpful. Motivated by the derivation of ComplEx from a two-embedding interaction model, we develop a four-embedding interaction model by using quaternion algebra to determine the weight vector and the interaction mechanism.

Quaternion numbers are extension of complex numbers to four components [15] [8]. Each quaternion number q , with one real component a and three imaginary components b, c, d , could be written as $q = a + bi + cj + dk$ where i, j, k are fundamental quaternion units, similar to the imaginary number i in complex algebra. As for complex conjugates, we also have a quaternion conjugate $\bar{q} = a - bi - cj - dk$.

An intuitive view of quaternion algebra is that each quaternion number represents a 4-dimensional vector (or 3-dimensional vector when the real component $a = 0$) and quaternion multiplication is rotation of this vector in 4- (or 3-)dimensional space. Compared to complex algebra, each complex number represents a 2-dimensional vector and complex multiplication is rotation of this vector in 2-dimensional plane [1].

Several works have shown the benefit of using complex, quaternion, or other hyper-complex numbers in the hidden layers of deep neural networks [9] [23] [25]. To the best of our knowledge, this paper is the first to motivate and use quaternion numbers for the embedding vectors of knowledge graph embedding.

Quaternion multiplication is noncommutative, thus there are multiple ways to multiply three quaternion numbers in the trilinear product. Here, we choose to write the score function of the quaternion-based four-embedding interaction model as:

$$\mathcal{S}(h, t, r) = \text{Re}(\langle \mathbf{h}, \bar{\mathbf{t}}, \mathbf{r} \rangle), \quad (13)$$

where $\mathbf{h}, \mathbf{t}, \mathbf{r} \in \mathbb{H}^k$.

By expanding this formula using quaternion algebra [15] and mapping the four components of a quaternion number to four embeddings in the multi-embedding interaction model, respectively, we can write the score function in the notation of the multi-embedding interaction model as:

$$\begin{aligned} \mathcal{S}(h, t, r) &= \text{Re}(\langle \mathbf{h}, \bar{\mathbf{t}}, \mathbf{r} \rangle) \\ &= \langle \mathbf{h}^{(1)}, \mathbf{t}^{(1)}, \mathbf{r}^{(1)} \rangle + \langle \mathbf{h}^{(2)}, \mathbf{t}^{(2)}, \mathbf{r}^{(1)} \rangle \\ &\quad + \langle \mathbf{h}^{(3)}, \mathbf{t}^{(3)}, \mathbf{r}^{(1)} \rangle + \langle \mathbf{h}^{(4)}, \mathbf{t}^{(4)}, \mathbf{r}^{(1)} \rangle \\ &\quad + \langle \mathbf{h}^{(1)}, \mathbf{t}^{(2)}, \mathbf{r}^{(2)} \rangle - \langle \mathbf{h}^{(2)}, \mathbf{t}^{(1)}, \mathbf{r}^{(2)} \rangle \\ &\quad + \langle \mathbf{h}^{(3)}, \mathbf{t}^{(4)}, \mathbf{r}^{(2)} \rangle - \langle \mathbf{h}^{(4)}, \mathbf{t}^{(3)}, \mathbf{r}^{(2)} \rangle \\ &\quad + \langle \mathbf{h}^{(1)}, \mathbf{t}^{(3)}, \mathbf{r}^{(3)} \rangle - \langle \mathbf{h}^{(2)}, \mathbf{t}^{(4)}, \mathbf{r}^{(3)} \rangle \\ &\quad - \langle \mathbf{h}^{(3)}, \mathbf{t}^{(1)}, \mathbf{r}^{(3)} \rangle + \langle \mathbf{h}^{(4)}, \mathbf{t}^{(2)}, \mathbf{r}^{(3)} \rangle \\ &\quad + \langle \mathbf{h}^{(1)}, \mathbf{t}^{(4)}, \mathbf{r}^{(4)} \rangle + \langle \mathbf{h}^{(2)}, \mathbf{t}^{(3)}, \mathbf{r}^{(4)} \rangle \\ &\quad - \langle \mathbf{h}^{(3)}, \mathbf{t}^{(2)}, \mathbf{r}^{(4)} \rangle - \langle \mathbf{h}^{(4)}, \mathbf{t}^{(1)}, \mathbf{r}^{(4)} \rangle, \end{aligned} \quad (14)$$

where $\mathbf{h}, \mathbf{t}, \mathbf{r} \in \mathbb{H}^k$.

4 LOSS FUNCTION AND OPTIMIZATION

The learning problem in knowledge graph embedding methods can be modeled as the binary classification of valid and invalid triples. Because knowledge graphs do not contain invalid triples, we generate them by negative sampling [20]. For each valid triple (h, t, r) , we replace the h or t entities in each training triple with other random entities to obtain the invalid triples (h', t, r) and (h, t', r) [4].

We can then learn the model parameters by minimizing the negative log-likelihood loss for the training data with the predicted probability modeled by the logistic sigmoid function $\sigma(\cdot)$ on the matching score. This loss is the cross-entropy:

$$\begin{aligned} \mathcal{L}(\mathcal{D}, \mathcal{D}'; \Theta, \omega) = & - \sum_{(h,t,r) \in \mathcal{D}} \log \sigma(S(h, t, r; \Theta, \omega)) \\ & - \sum_{(h',t',r) \in \mathcal{D}'} \log \sigma(1 - S(h', t', r; \Theta, \omega)), \end{aligned} \quad (15)$$

where \mathcal{D} is true data ($\hat{p} = 1$), \mathcal{D}' is negative sampled data ($\hat{p} = 0$), and \hat{p} is the empirical probability.

Defining the class label $Y_{(h,t,r)} = 2\hat{p}_{(h,t,r)} - 1$, i.e., the labels of positive triples are 1 and negative triples are -1 , the above loss can be written more concisely. In cluding the L^2 regularization of embedding vectors, this loss can be written as:

$$\begin{aligned} \mathcal{L}(\mathcal{D}, \mathcal{D}'; \Theta, \omega) = & \sum_{(h,t,r) \in \mathcal{D} \cup \mathcal{D}'} \left(\log(1 + e^{-Y_{(h,t,r)} S(h,t,r; \Theta, \omega)}) \right. \\ & \left. + \frac{\lambda}{nD} \|\Theta\|_2^2 \right), \end{aligned} \quad (16)$$

where \mathcal{D} is true data, \mathcal{D}' is negative sampled data, Θ are the embedding vectors corresponding to specific current triples, n is the number of multi-embedding, D is the embedding size, and λ is the regularization strength.

5 EXPERIMENTAL SETTINGS

5.1 Datasets

For our empirical analysis, we used the WN18 dataset, the most popular of the benchmark datasets built on WordNet [22] by Bordes et al. [4]. This dataset has 40,943 *entities*, 18 *relations*, 141,442 *training triples*, 5,000 *validation triples*, 5,000 *test triples*. In our preliminary experiments, the relative performance on all datasets was quite consistent, therefore choosing the WN18 dataset is appropriate for our analysis. We will consider the use of other datasets in in future work.

5.2 Evaluation protocols

Knowledge graph embedding methods are usually evaluated on link prediction task [4]. In this task, for each true triple (h, t, r) in the test set, we replace h and t by every other entity to generate corrupted triples (h', t, r) and (h, t', r) , respectively [4]. The goal of the model now is to rank the true triple (h, t, r) before the corrupted triples based on the predicted score \mathcal{S} .

For each true triple in the test set, we compute its rank, then we can compute popular evaluation metrics including *MRR* (mean reciprocal rank) and *Hit@k* for $k \in \{1, 3, 10\}$ (how many true triples are correctly ranked in the top k) [28].

To avoid false negative error, i.e., corrupted triples are accidentally valid triples, we follow the protocols used in other works for filtered metrics [4]. In this protocol, all valid triples in the

training, validation, and test sets are removed from the corrupted triples set before computing the rank of the true triple.

5.3 Training

We trained the models using SGD with learning rates auto-tuned by Adam [16], that makes the choice of initial learning rate more robust. For all models, we found good hyperparameters with grid search on learning rates $\in \{10^{-3}, 10^{-4}\}$, embedding regularization strengths $\in \{10^{-2}, 3 \times 10^{-3}, 10^{-3}, 3 \times 10^{-4}, 10^{-4}, 0.0\}$, and batch sizes $\in \{2^{12}, 2^{14}\}$. For a fair comparison, we fixed the embedding sizes so that numbers of parameters for all models are comparable. In particular, we use embedding sizes of 400 for one-embedding models such as DistMult, 200 for two-embedding models such as ComplEx, CP, and CP_h , and 100 for four-embedding models. We also fixed the number of negative samples at 1 because, although using more negative samples is beneficial for all models, it is also more expensive and not necessary for this comparative analysis.

We constrained entity embedding vectors to have unit L^2 -norm after each training iteration. All training runs were stopped early by checking the filtered MRR on the validation set after every 50 epochs, with 100 epochs patient.

6 RESULTS AND DISCUSSION

In this section, we present experimental results and analyses for the models described in Section 3. We report results for derived weight vectors and their variants, auto-learned weight vectors, and the quaternion-based four-embedding interaction model.

6.1 Derived weight vectors and variants

6.1.1 Comparison of derived weight vectors. We evaluated the multi-embedding interaction model with the score function in Eq. (8), using the derived weight vectors in Table 1. The results are shown in Table 2. They are consistent with the results reported in other works [28]. Note that ComplEx and CP_h achieved good results, whereas DistMult performed less well. CP performed very poorly in comparison to the other models, even though it is a classical model for the tensor decomposition task [13].

For a more detailed comparison, we report the performance on training data. Note that ComplEx and CP_h can accurately predict the training data, whereas DistMult did not. This is evidence that ComplEx and CP_h are fully expressive while DistMult cannot model asymmetric data effectively.

The most surprising result was that CP can also accurately predict the training data at a comparable level to ComplEx and CP_h , despite its very poor result on the test data. This suggests that the problem with CP is not its modeling capacity, but in its generalization performance to new test data. In other words, CP is severely overfitting to the training data. However, standard regularization techniques such as L^2 regularization did not appear to help. CP_h can be seen as a regularization technique that does help CP generalize well to unseen data.

6.1.2 Comparison with other variants of weight vectors. In Table 2, we show the results for two bad examples and two good examples of weight vector variants. Note that *bad example 1* performed similarly to CP and *bad example 2* performed similarly to DistMult. *Good example 1* was similar to CP_h and *good example 2* was similar to ComplEx.

This shows that the problem of bad weight vectors is not unique to some specific models. Moreover, it shows that there

Table 2: Results for the derived weight vectors on WN18.

Weight setting	MRR	Hit@1	Hit@3	Hit@10
DistMult (1, 0, 0, 0, 0, 0, 0)	0.796	0.674	0.915	0.945
ComplEx (1, 0, 0, 1, 0, -1, 1, 0)	0.937	0.928	0.946	0.951
CP (0, 0, 1, 0, 0, 0, 0, 0)	0.086	0.059	0.093	0.139
CP _h (0, 0, 1, 0, 0, 1, 0, 0)	0.937	0.929	0.944	0.949
DistMult <i>on train</i>	0.917	0.848	0.985	0.997
ComplEx <i>on train</i>	0.996	0.994	0.998	0.999
CP <i>on train</i>	0.994	0.994	0.996	0.999
CP _h <i>on train</i>	0.995	0.994	0.998	0.999
Bad example 1 (0, 0, 20, 0, 0, 1, 0, 0)	0.107	0.079	0.116	0.159
Bad example 2 (0, 0, 1, 1, 1, 1, 0, 0)	0.794	0.666	0.917	0.947
Good example 1 (0, 0, 20, 1, 1, 20, 0, 0)	0.938	0.934	0.942	0.946
Good example 2 (1, 1, -1, 1, 1, -1, 1, 1)	0.938	0.930	0.944	0.950

Table 3: Results for the auto-learned weight vectors on WN18.

Weight setting	MRR	Hit@1	Hit@3	Hit@10
Uniform weight (1, 1, 1, 1, 1, 1, 1, 1)	0.787	0.658	0.915	0.944
Auto weight no restriction	0.774	0.636	0.911	0.944
Auto weight $\in (-1, 1)$ by tanh	0.765	0.625	0.908	0.943
Auto weight $\in (0, 1)$ by sigmoid	0.789	0.661	0.915	0.946
Auto weight $\in (0, 1)$ by softmax	0.802	0.685	0.915	0.944
Auto weight no restriction, sparse	0.792	0.685	0.892	0.935
Auto weight $\in (-1, 1)$ by tanh, sparse	0.763	0.613	0.910	0.943
Auto weight $\in (0, 1)$ by sigmoid, sparse	0.793	0.667	0.915	0.945
Auto weight $\in (0, 1)$ by softmax, sparse	0.803	0.688	0.915	0.944

Table 4: Results for the quaternion-based four-embedding interaction model on WN18.

Weight setting	MRR	Hit@1	Hit@3	Hit@10
Quaternion-based four-embedding	0.941	0.931	0.950	0.956
Quaternion-based four-embedding <i>on train</i>	0.997	0.995	0.999	1.000

are other good weight vectors, besides those for ComplEx and CP_h, that can achieve very good results.

We note that the good weight vectors exhibit the following properties.

- *Completeness*: all embedding vectors in a triple should be involved in the weighted-sum matching score.
- *Stability*: all embedding vectors for the same entity or relation should contribute equally to the weighted-sum matching score.
- *Distinguishability*: the weighted-sum matching scores for different triples should be *distinguishable*. For example, the score $\langle \mathbf{h}^{(1)}, \mathbf{t}^{(2)}, \mathbf{r}^{(1)} \rangle + \langle \mathbf{h}^{(2)}, \mathbf{t}^{(1)}, \mathbf{r}^{(2)} \rangle$ is *indistinguishable* because switching h and t forms a symmetric group.

As an example, consider the ComplEx model, where the multiplication of two complex numbers written in polar coordinate format, $c_1 = |c_1|e^{-i\theta_1}$ and $c_2 = |c_2|e^{-i\theta_2}$, can be written as $c_1c_2 = |c_1||c_2|e^{-i(\theta_1+\theta_2)}$ [1]. This is a rotation in the complex plane, which intuitively satisfies the above properties.

6.2 Automatically learned weight vectors

We let the models learn ω together with the embeddings in an end-to-end fashion, aiming to learn good weight vectors automatically. The results are shown in Table 3.

We first set uniform weight vector as a baseline. The results were similar to those for DistMult because the weighted-sum

matching score is also symmetric. However, other automatically learned weight vectors also performed similarly to DistMult. Different restrictions by applying $\tanh(\omega)$, $\text{sigmoid}(\omega)$, and $\text{softmax}(\omega)$ did not help. We noticed that the learned weight vectors were almost uniform, making them *indistinguishable*, suggesting that the use of sparse weight vectors might help.

We enforced a sparsity constraint by an additional Dirichlet negative log-likelihood regularization loss on ω , with α tuned to $\frac{1}{16}$ and λ_{dir} tuned to 10^{-2} . However, the results did not improve. Tracking of weight vectors value showed that the sparsity constraint seemed to amplify the initial differences between the weight values instead of learning useful sparseness. This suggests that the gradient information is too symmetric that the model cannot break the symmetry of ω and escape the local optima.

In general, these experiments show that learning good weight vectors automatically is a particularly difficult task.

6.3 Quaternion-based four-embedding interaction model

In Table 4, we present the evaluation results for the proposed quaternion-based four-embedding interaction model. The results were generally positive, with most metrics higher than those in Table 2 for state-of-the-art models such as ComplEx and CP_h. Especially, H@10 performance was much better than other models.

Note that this model needs more extensive evaluation. One potential problem is its being prone to overfitting, as seen in the *on train* results, with H@10 at absolute 1.000. This might mean that better regularization methods may be needed. However, the general results suggest that extending to more embedding vectors for multi-embedding interaction models is a promising approach.

7 CONCLUSION

This paper proposes a multi-embedding interaction mechanism as a new approach to analyzing state-of-the-art knowledge graph embedding models such as DistMult, ComplEx, CP, and CP_h. We show that these models can be unified and generalized under the new approach to provide an intuitive perspective on using the models and their embedding vectors effectively. We analyzed and compared the models and their variants empirically to better understand their properties, such as the severe overfitting problem of the CP model. In addition, we propose and have evaluated a new multi-embedding interaction model based on quaternion algebra, which showed some promising results.

There are several promising future directions. One direction is to find new methods of modeling the interaction mechanism between multi-embedding vectors and the effective extension to additional embedding vectors. Another direction is to evaluate multi-embedding models such as the proposed quaternion-based four-embedding interaction model more extensively.

ACKNOWLEDGMENTS

This work was supported by a JSPS Grant-in-Aid for Scientific Research (B) (15H02789).

REFERENCES

- [1] Lars V. Ahlfors. 1953. *Complex Analysis: An Introduction to the Theory of Analytic Functions of One Complex Variable*. New York, London (1953), 177.
- [2] Amit Singhal. 2012. Official Google Blog: Introducing the Knowledge Graph: Things, Not Strings. <https://googleblog.blogspot.com/2012/05/introducing-knowledge-graph-things-not.html>.
- [3] Kurt Bollacker, Colin Evans, Praven Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge. In *In SIGMOD Conference*. 1247–1250.
- [4] 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.
- [5] Walter Carrer-Neto, María Luisa Hernández-Alcaraz, Rafael Valencia-García, and Francisco García-Sánchez. 2012. Social Knowledge-Based Recommender System. Application to the Movies Domain. *Expert Systems with Applications* 39, 12 (Sept. 2012), 10990–11000. <https://doi.org/10.1016/j.eswa.2012.03.025>
- [6] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d Knowledge Graph Embeddings. In *In Thirty-Second AAAI Conference on Artificial Intelligence*.
- [7] 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. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '14*. ACM Press, New York, New York, USA, 601–610. <https://doi.org/10.1145/2623330.2623623>
- [8] Ron Goldman. 2010. Rethinking Quaternions. *Synthesis Lectures on Computer Graphics and Animation* 4, 1 (Oct. 2010), 1–157. <https://doi.org/10.2200/S00292ED1V01Y201008CGR013>
- [9] Nitzan Guberman. 2016. On Complex Valued Convolutional Neural Networks. *arXiv:1602.09046 [cs.NE]* (Feb. 2016). [arXiv:cs.NE/1602.09046](https://arxiv.org/abs/1602.09046)
- [10] Ruining He, Wang-Cheng Kang, and Julian McAuley. 2017. Translation-Based Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*. ACM, New York, NY, USA, 161–169. <https://doi.org/10.1145/3109859.3109882>
- [11] Geoffrey E. Hinton. 1986. Learning Distributed Representations of Concepts. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, Vol. 1. Amherst, MA, 12.
- [12] G E Hinton, J L McClelland, and D E Rumelhart. 1984. Distributed Representations. In *Parallel Distributed Processing*. Carnegie-Mellon University, Pittsburgh, PA, 33.
- [13] Frank L. Hitchcock. 1927. The Expression of a Tensor or a Polyadic as a Sum of Products. *Journal of Mathematics and Physics* 6, 1–4 (April 1927), 164–189. <https://doi.org/10.1002/sapm192761164>
- [14] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543.
- [15] Isai Lvovich Kantor and Aleksandr Samuilovich Solodovnikov. 1989. *Hypercomplex Numbers: An Elementary Introduction to Algebras*. Springer.
- [16] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- [17] Timothée Lacroix, Nicolas Usunier, and Guillaume Obozinski. 2018. Canonical Tensor Decomposition for Knowledge Base Completion. In *Proceedings of the 35th International Conference on Machine Learning (ICML '18)*.
- [18] Yankai Lin, Zhiyuan Liu, Huanbo Luan, Maosong Sun, Siwei Rao, and Song Liu. 2015. Modeling Relation Paths for Representation Learning of Knowledge Bases. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- [19] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*. 2181–2187.
- [20] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR'13 Workshop*.
- [21] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Advances in Neural Information Processing Systems*. 3111–3119.
- [22] Miller, George A. 1995. WordNet: A Lexical Database for English. *Commun. ACM* (1995), 39–41.
- [23] Toshifumi Minemoto, Tejiro Isokawa, Haruhiko Nishimura, and Nobuyuki Matsui. 2017. Feed Forward Neural Network with Random Quaternionic Neurons. *Signal Processing C*, 136 (2017), 59–68. <https://doi.org/10.1016/j.sigpro.2016.11.008>
- [24] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *Proceedings of the 28th International Conference on Machine Learning*. 809–816.
- [25] Titouan Parcollet, Mirco Ravanelli, Mohamed Morchid, Georges Linares, Chih-Heb Trabelsi, Renato De Mori, and Yoshua Bengio. 2019. Quaternion Recurrent Neural Networks. In *Proceedings of the International Conference on Learning Representations (ICLR'19)*.
- [26] Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Y Ng. 2013. Reasoning With Neural Tensor Networks for Knowledge Base Completion. In *Advances in Neural Information Processing Systems*. 926–934.
- [27] Théo Trouillon, Christopher R. Dance, Éric Gaussier, Johannes Welbl, Sebastian Riedel, and Guillaume Bouchard. 2017. Knowledge Graph Completion via Complex Tensor Factorization. *The Journal of Machine Learning Research* 18, 1 (2017), 4735–4772.
- [28] Theo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. 2016. Complex Embeddings for Simple Link Prediction. In *International Conference on Machine Learning (ICML'16)*. 2071–2080.
- [29] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM* 57, 10 (Sept. 2014), 78–85. <https://doi.org/10.1145/2629489>
- [30] Q. Wang, Z. Mao, B. Wang, and L. Guo. 2017. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Transactions on Knowledge and Data Engineering* 29, 12 (Dec. 2017), 2724–2743. <https://doi.org/10.1109/TKDE.2017.2754499>
- [31] Yanjie Wang, Rainer Gemulla, and Hui Li. 2018. On Multi-Relational Link Prediction with Bilinear Models. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [32] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge Graph and Text Jointly Embedding. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*. 1591–1601.
- [33] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. In *AAAI Conference on Artificial Intelligence*. Citeseer, 1112–1119.
- [34] Han Xiao, Minlie Huang, Yu Hao, and Xiaoyan Zhu. 2015. TransA: An Adaptive Approach for Knowledge Graph Embedding. In *AAAI Conference on Artificial Intelligence*. [arXiv:1509.05490](https://arxiv.org/abs/1509.05490)
- [35] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *International Conference on Learning Representations*.
- [36] Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. 2016. Collaborative Knowledge Base Embedding for Recommender Systems. ACM Press, 353–362. <https://doi.org/10.1145/2939672.2939673>