# Why PairDiff works? - A Mathematical Analysis of Bilinear Relational Compositional Operators for Analogy Detection 

Huda Hakami ${ }^{\dagger} \quad$ Danushka Bollegala ${ }^{\dagger} \quad$ Hayashi Kohei ${ }^{\ddagger}$<br>University of Liverpool ${ }^{\dagger}$<br>Japan Advanced Institute of Science and Technology ${ }^{\ddagger}$<br>\{H.A.Hakami, danushka\}@liverpool.ac.uk hayashi.kohei@gmail.com


#### Abstract

Representing the semantic relations that exist between two given words (or entities) is an important first step in a wide-range of NLP applications such as analogical reasoning, knowledge base completion and relational information retrieval. A simple, yet surprisingly accurate method for representing a relation between two words is to compute the vector offset (PairDiff) between their corresponding word embeddings. Despite the empirical success, it remains unclear as to whether PairDiff is the best operator for obtaining a relational representation from word embeddings. We conduct a theoretical analysis of generalised bilinear operators that can be used to measure the $\ell_{2}$ relational distance between two word-pairs. We show that, if the word embeddings are standardised and uncorrelated, such an operator will be independent of bilinear terms, and can be simplified to a linear form, where PairDiff is a special case. For numerous word embedding types, we empirically verify the uncorrelation assumption, demonstrating the general applicability of our theoretical result. Moreover, we experimentally discover PairDiff from the bilinear relation composition operator on several benchmark analogy datasets.


## 1 Introduction

Different types of semantic relations exist between words such as HYPERNYMY between ostrich and bird, or ANTONYMY between hot and cold. If we consider entities ${ }^{1}$, we can observe even a richer diversity of relations such as FOUNDER-OF between Bill Gates and Microsoft, or CAPITAL-OF between Tokyo and Japan. Identifying the relations between words

[^0]and entities is important for various Natural Language Processing (NLP) tasks such as automatic knowledge base completion (Socher et al., 2013), analogical reasoning (Turney and Littman, 2005; Bollegala et al., 2009) and relational information retrieval (Duc et al., 2010). For example, to solve a word analogy problem of the form " $a$ is to $b$ as $c$ is to ?", the relationship between the two words in the pair $(a, b)$ must be correctly identified in order to find candidates $d$ that have similar relations with $c$. For example, given the query "Bill Gates is to Microsoft as Steve Jobs is to ?", a relational search engine must retrieve Apple Inc. because the FOUNDER-OF relation exists between the first and the second entity pairs.

Two main approaches for creating relation embeddings can be identified in the literature. In the first approach, from given corpora or knowledge bases, word and relation embeddings are jointly learnt such that some objective is optimised (Guo et al., 2016; Yang et al., 2015; Nickel et al., 2016; Bordes et al., 2013; Rocktäschel et al., 2016; Minervini et al., 2017; Trouillon et al., 2016). In this approach, word and relation embeddings are considered to be independent parameters that must be learnt by the embedding method. For example, TransE (Bordes et al., 2013) learns the word and relation embeddings such that we can accurately predict relations (links) in a given knowledge base using the learnt word and relation embeddings. Because relations are learnt independently from the words, we refer to methods that are based on this approach as independent relational embedding methods.

A second approach for creating relational embeddings is to apply some operator on two word embeddings to compose the embedding for the relation that exits between those two words, if any. In contrast to the first approach, we do not have to learn relational embeddings and hence this can
be considered as an unsupervised setting, where the compositional operator is predefined. A popular operator for composing a relational embedding from two word embeddings is PairDiff, which is the vector difference (offset) of the word embeddings (Mikolov et al., 2013b; Levy and Goldberg, 2014; Vylomova et al., 2016; Bollegala et al., 2015b; Blacoe and Lapata, 2012). Specifically, given two words $a$ and $b$ represented by their word embeddings respectively $\boldsymbol{a}$ and $\boldsymbol{b}$, the relation between $a$ and $b$ is given by $\boldsymbol{a}-\boldsymbol{b}$ under the PairDiff operator. Mikolov et al. (2013b) showed that PairDiff can accurately solve analogy equations such as $\overrightarrow{\text { king }}-\overrightarrow{\mathbf{m a n}}+\overrightarrow{\text { woman }}=\overrightarrow{\text { queen }}$, where we have used the top arrows to denote the embeddings of the corresponding words. Bollegala et al. (2015a) showed that PairDiff can be used as a proxy for learning better word embeddings and Vylomova et al. (2016) conducted an extensive empirical comparison of PairDiff using a dataset containing 16 different relation types. Besides PairDiff, concatenation (Hakami and Bollegala, 2017; Yin and Schütze, 2016), circular correlation and convolution (Nickel et al., 2016) have been used in prior work for representing the relations between words. Because the relation embedding is composed using word embeddings instead of learning as a separate parameter, we refer to methods that are based on this approach as compositional relational embedding methods. Note that in this approach it is implicitly assumed that there exist only a single relation between two words.

In this paper, we focus on the operators that are used in compositional relational embedding methods. If we assume that the words and relations are represented by vectors embedded in some common space, then the operator we are seeking must be able to produce a vector representing the relation between two words, given their word embeddings as the only input. Although there have been different proposals for computing relational embeddings from word embeddings, it remains unclear as to what is the best operator for this task. The space of operators that can be used to compose relational embeddings is open and vast. A space of particular interest from a computational point-of-view is the bilinear operators that can be parametrised using tensors and matrices. Specifically, we consider operators that consider pairwise interactions between two word embeddings (second-order terms) and contributions from indi-
vidual word embeddings towards their relational embedding (first-order terms). The optimality of a relational compositional operator can be evaluated, for example, using the expected relational distance/similarity such as $\ell_{2}$ between analogous (positive) vs. nonanalogous (negative) word-pairs.

If we assume that word embeddings are standardised, uncorrelated and word-pairs are i.i.d, then we prove in $\S 3$ that bilinear relational compositional operators are independent of bilinear pairwise interactions between the two input word embeddings. Moreover, under regularised settings (§3.1), the bilinear operator further simplifies to a linear combination of the input embeddings, and the expected loss over positive and negative instances becomes zero. In §4.1, we empirically validate the uncorrelation assumption for different pre-trained word embeddings such as the Continuous Bag-of-Words Model (CBOW) (Mikolov et al., 2013a), Skip-Gram with negative sampling (SG) (Mikolov et al., 2013a), Global Vectors (GloVe) (Pennington et al., 2014), word embeddings created using Latent Semantic Analysis (LSA) (Deerwester et al., 1990), Sparse Coding (HSC) (Faruqui et al., 2015; Yogatama et al., 2015), and Latent Dirichlet Allocation (LDA) (Blei et al., 2003a). This empirical evidence implies that our theoretical analysis is applicable to relational representations composed from a wide-range of word embedding learning methods. Moreover, our experimental results show that a bilinear operator reaches its optimal performance in two different word-analogy benchmark datasets, when it satisfies the requirements of the PairDiff operator. We hope that our theoretical analysis will expand the understanding of relational embedding methods, and inspire future research on accurate relational embedding methods using word embeddings as the input.

## 2 Related Work

As already mentioned in §1, methods for representing a relation between two words can be broadly categorised into two groups depending on whether the relational embeddings are learnt independently of the word embeddings, or they are composed from the word embeddings, in which case the relational embeddings fully depend on the input word embeddings. Next, we briefly overview the different methods that fall under each category. For a detailed survey of relation embed-
ding methods see (Nickel et al., 2015).
Given a knowledge base where an entity $h$ is linked to an entity $t$ by a relation $r$, the TransE model (Bordes et al., 2013) scores the tuple $(h, t, r)$ by the $\ell_{1}$ or $\ell_{2}$ norm of the vector $(\boldsymbol{h}+\boldsymbol{r}-\boldsymbol{t})$. (Nickel et al., 2011) proposed RESCAL, which uses $\boldsymbol{h}^{\top} \mathbf{M}_{r} t$ as the scoring function, where $\mathbf{M}_{r}$ is a matrix embedding of the relation $r$. Similar to RESCAL, Neural Tensor Network (Socher et al., 2013) also models a relation by a matrix. However, compared to vector embeddings of relations, matrix embeddings increase the number of parameters to be estimated, resulting in an increase in computational time/space and likely to overfit. To overcome these limitations, DistMult (Yang et al., 2015) models relations by vectors and use elementwise multilinear dot product $\boldsymbol{r} \odot \boldsymbol{h} \odot \boldsymbol{t}$. Unfortunately, DistMult cannot capture directionality of a relation. Complex Embeddings (Trouillon et al., 2016) overcome this limitation of DistMult by using complex embeddings and defining the score to be the real part of $\boldsymbol{r} \odot \boldsymbol{h} \odot \overline{\boldsymbol{t}}$, where $\overline{\boldsymbol{t}}$ denotes the complex conjugate of $\boldsymbol{t}$.

The observation made by Mikolov et al. (2013b) that the relation between two words can be represented by the difference between their word embeddings sparked a renewed interest in methods that compose relational embeddings using word embeddings. Word analogy datasets such as Google dataset (Mikolov et al., 2013b), SemEval 2012 Task2 dataset (Jurgens et al., 2012), BATS (Drozd et al., 2016) etc. have established as benchmarks for evaluating word embedding learning methods.

Different methods have been proposed to measure the similarity between the relations that exist between two given word pairs such as CosMult, CosAdd and PairDiff (Levy and Goldberg, 2014; Bollegala et al., 2015a). Vylomova et al. (2016) studied as to what extent the vectors generated using simple PairDiff encode different relation types. Under supervised classification settings, they conclude that PairDiff can cover a wide range of semantic relation types. Holographic embeddings proposed by Nickel et al. (2016) use circular convolution to mix the embeddings of two words to create an embedding for the relation that exist between those words. It can be showed that circular correlation is indeed an elementwise product in the Fourier space and is mathematically equivalent to
complex embeddings (Hayashi and Shinbo, 2017).
Although PairDiff operator has been widely used in prior work for computing relation embeddings from word embeddings, to the best of our knowledge, no theoretical analysis has been conducted so far explaining why and under what conditions PairDiff is optimal, which is the focus of this paper.

## 3 Bilinear Relation Representations

Let us consider the problem of representing the semantic relation $\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})$ between two given words $h$ and $t$. We assume that $h$ and $t$ are already represented in some $d$-dimensional space respectively by their word embeddings $\boldsymbol{h}, \boldsymbol{t} \in \mathbb{R}^{d}$. The relation between two words can be represented using different linear algebraic structures. Two popular alternatives are vectors (Nickel et al., 2016; Bordes et al., 2013; Minervini et al., 2017; Trouillon et al., 2016) and matrices (Socher et al., 2013; Bollegala et al., 2015b). Vector representations are preferred over matrix representations because of the smaller number of parameters to be learnt (Nickel et al., 2015).

Let us assume that the relation $r$ is represented by a vector $\boldsymbol{r} \in \mathbb{R}^{\delta}$ in some $\delta$-dimensional space. Therefore, we can write $\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})$ as a function that takes two vectors (corresponding to the embeddings of the two words) as the input and returns a single vector (representing the relation between the two words) as given in (1).

$$
\begin{equation*}
r: \mathbb{R}^{d} \times \mathbb{R}^{d} \rightarrow \mathbb{R}^{\delta} \tag{1}
\end{equation*}
$$

Having both words and relations represented in the same $\delta=d$ dimensional space is useful for performing linear algebraic operations using those representations in that space. For example, in TransE (Bordes et al., 2013), the strength of a relation $r$ that exists between two words $h$ and $t$ is computed as the $\ell_{1,2}$ norm of the vector $(\boldsymbol{h}+\boldsymbol{r}-\boldsymbol{t})$ using the word and relation embeddings. Such direct comparisons between word and relation embeddings would not be possible if words and relations were not embedded in the same vector space. If $\delta<d$, we can first project word embeddings to a lower $\delta$-dimensional space using some dimensionality reduction method such as SVD, whereas if $\delta>d$ we can learn higher $\delta$-dimensional overcomplete word representations (Faruqui et al., 2015) from the original $d$-dimensional word embeddings. Therefore, we will limit our theoretical
analysis to the $\delta=d$ case for ease of description.
Different functions can be used as $\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})$ that satisfy the domain and range requirements specified by (1). If we limit ourselves to bilinear functions, the most general functional form is given by (2).

$$
\begin{equation*}
\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})=\boldsymbol{h}^{\top} \underline{\mathbf{A} t} \boldsymbol{t}+\mathbf{P} \boldsymbol{h}+\mathbf{Q} t \tag{2}
\end{equation*}
$$

Here, $\underline{\mathbf{A}} \in \mathbb{R}^{d \times d \times d}$ is a 3-way tensor in which each slice is a $d \times d$ real matrix. Let us denote the $k$-th slice of $\underline{\mathbf{A}}$ by $\mathbf{A}^{(k)}$ and its $(i, j)$ element by $A_{i j}^{(k)}$. The first term in (2) corresponds to the pairwise interactions between $\boldsymbol{h}$ and $\boldsymbol{t} . \mathbf{P}, \mathbf{Q} \in \mathbb{R}^{d \times d}$ are the nonsingular ${ }^{2}$ projection matrices involving first-order contributions respectively of $\boldsymbol{h}$ and $\boldsymbol{t}$ towards $r$.
Let us consider the problem of learning the simplest bilinear functional form according to (2) from a given dataset of analogous word-pairs $\mathcal{D}_{+}=\left\{\left((h, t),\left(h^{\prime}, t^{\prime}\right)\right)\right\}$. Specifically, we would like to learn the parameters $\mathbf{A}, \mathbf{P}$ and $\mathbf{Q}$ such that some distance (loss) between analogous wordpairs is minimised. As a concrete example of a distance function, let us consider the popularly used Euclidean distance ${ }^{3}$ ( $\ell_{2}$ loss) for two word pairs given by (3).

$$
\begin{equation*}
J\left((h, t),\left(h^{\prime}, t^{\prime}\right)\right)=\left\|\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})-\boldsymbol{r}\left(\boldsymbol{h}^{\prime}, \boldsymbol{t}^{\prime}\right)\right\|_{2}^{2} \tag{3}
\end{equation*}
$$

If we were provided only analogous word-pairs (i.e. positive examples), then this task could be trivially achieved by setting all parameters to zero. However, such a trivial solution would not generalise to unseen test data. Therefore, in addition to $\mathcal{D}_{+}$we would require a set of nonanalogous word-pairs $\mathcal{D}_{-}$as negative examples. Such negative examples are often generated in prior work by randomly corrupting positive relational tuples (Nickel et al., 2016; Bordes et al., 2013; Trouillon et al., 2016) or by training an adversarial generator (Minervini et al., 2017).

The total loss $J$ over both positive and negative training data can be written as follows:

$$
\begin{align*}
J= & \sum_{\left((h, t),\left(h^{\prime}, t^{\prime}\right)\right) \in \mathcal{D}_{+}}\left\|\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})-\boldsymbol{r}\left(\boldsymbol{h}^{\prime}, \boldsymbol{t}^{\prime}\right)\right\|_{2}^{2} \\
& -\sum_{\left((h, t),\left(h^{\prime}, t^{\prime}\right)\right) \in \mathcal{D}_{-}}\left\|\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})-\boldsymbol{r}\left(\boldsymbol{h}^{\prime}, \boldsymbol{t}^{\prime}\right)\right\|_{2}^{2} \tag{4}
\end{align*}
$$

[^1]Assuming that the training word-pairs are randomly sampled from $\mathcal{D}_{+}$and $\mathcal{D}_{-}$according to two distributions respectively $p_{+}$and $p_{-}$, we can compute the total expected loss, $\mathbb{E}_{p}[J]$, as follows:

$$
\begin{align*}
\mathbb{E}_{p}[J]= & \mathbb{E}_{p_{+}}\left[\left\|\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})-\boldsymbol{r}\left(\boldsymbol{h}^{\prime}, \boldsymbol{t}^{\prime}\right)\right\|_{2}^{2}\right]- \\
& \mathbb{E}_{p_{-}}\left[\left\|\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})-\boldsymbol{r}\left(\boldsymbol{h}^{\prime}, \boldsymbol{t}^{\prime}\right)\right\|_{2}^{2}\right] \tag{5}
\end{align*}
$$

We make the following assumptions to further analyse the properties of relational embeddings.

Uncorrelation: The correlation between any two distinct dimensions of a word embedding is zero. One might think that the uncorrelation of word embedding dimensions to be a strong assumption, but we later show its validity empirically in $\S 4.1$ for a wide range of word embeddings.

Standerdisation: Word embeddings are standerdised to zero mean and unit variance. This is a linear transformation in the word embedding space and does not affect the topology of the embedding space. In particular, translating word embeddings such that they have a zero mean has shown to improve performance in similarity tasks (Mu et al., 2017).

Relational Independence Word pairs in the training data are assumed to be i.i.d. For example, whether a particular semantic relation $r$ exists between $h$ and $t$, is assumed to be independent of any other relation $r^{\prime}$ that exists between $h^{\prime}$ and $t^{\prime}$ in a different pair.

For relation representations given by (2), Theorem 1 holds:
Theorem 1. Consider the bilinear relational embedding defined by (2) computed using uncorrelated word embeddings. If the word embeddings are standerdised, then the expected loss given by (5) over a relationally independent set of word pairs is independent of $\underline{\mathbf{A}}$.

Proof. Let us consider the bilinear term in (2), because $i$ and $j(\neq i)$ dimensions of word embeddings are uncorrelated by the assumption (i.e. $\operatorname{corr}\left(u_{i}, u_{j}\right)=0$ ), from the definition of correlation we have,

$$
\begin{align*}
\operatorname{corr}\left(u_{i}, u_{j}\right) & =\mathbb{E}\left[u_{i} u_{j}\right]-\mathbb{E}\left[u_{i}\right] \mathbb{E}\left[u_{j}\right]=0  \tag{6}\\
\mathbb{E}\left[u_{i} u_{j}\right] & =\mathbb{E}\left[u_{i}\right] \mathbb{E}\left[u_{j}\right] \tag{7}
\end{align*}
$$

Moreover, from the standerdisation assumption we have, $\mathbb{E}\left[u_{i}\right]=0, \quad \forall_{i=1 \ldots n}$. From (7) it follows that:

$$
\begin{equation*}
\mathbb{E}\left[u_{i} u_{j}\right]=0 \tag{8}
\end{equation*}
$$

for $i \neq j$ dimensions.
We will next show that (5) is independent of $\mathbf{A}$. For this purpose, let us consider the $\mathbb{E}_{p_{+}}$term first and write the $k$-th dimension of $\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})$ using $\mathbf{A}^{(k)}$, $\mathbf{P}$ and $\mathbf{Q}$ as follows:

$$
\begin{equation*}
\sum_{i, j}\left(A_{i j}^{(k)} h_{i} t_{j}\right)+\sum_{n} P_{k n} h_{n}+\sum_{n} Q_{k n} t_{n} \tag{9}
\end{equation*}
$$

Plugging (9) in (5) and computing the loss over all positive training instances we get,

$$
\begin{align*}
& \mathbb{E}_{p_{+}}\left[\sum _ { k } \left(\sum_{i, j}\left(A_{i j}^{(k)}\left(h_{i} t_{j}-h_{i}^{\prime} t_{j}^{\prime}\right)\right)+\right.\right. \\
& \left.\left.\sum_{n} P_{k n}\left(h_{n}-h_{n}^{\prime}\right)+\sum_{n} Q_{k n}\left(t_{n}-t_{n}^{\prime}\right)\right)^{2}\right] \tag{10}
\end{align*}
$$

Terms that involve only elements in $\mathbf{A}^{(k)}$ take the form:

$$
\begin{align*}
& \sum_{i, j} \sum_{l, m} \mathbb{E}_{p_{+}}\left[A_{i j}^{(k)} A_{l m}^{(k)}\left(h_{i} t_{j}-h_{i}^{\prime} t_{j}^{\prime}\right)\left(h_{l} t_{m}-h_{l}^{\prime} t_{m}^{\prime}\right)\right] \\
& =\sum_{i, j} \sum_{l, m} A_{i j}^{(k)} A_{l m}^{(k)}\left(\mathbb{E}_{p_{+}}\left[h_{i} t_{j} h_{l} t_{m}\right]-\mathbb{E}_{p_{+}}\left[h_{i} t_{j} h_{l} t_{m}^{\prime}\right]-\right. \\
& \left.\mathbb{E}_{p_{+}}\left[h_{i} t_{j}^{\prime} h_{l} t_{m}\right]+\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{j}^{\prime} h_{l}^{\prime} t_{m}^{\prime}\right]\right) \tag{11}
\end{align*}
$$

In cases where $i \neq j$ and $l \neq m$, each of the four expectations in (11) contains the product of different dimensionalities, which is zero from (8). For $i=j=l=m$ case we have,

$$
\begin{equation*}
A_{i j}^{(k)^{2}}\left(\mathbb{E}_{p_{+}}\left[h_{i}^{2} t_{i}^{2}\right]-2 \mathbb{E}_{p_{+}}\left[h_{i} t_{i} h_{i}^{\prime} t_{i}^{\prime}\right]+\mathbb{E}_{p_{+}}\left[h_{i}^{\prime 2} t_{i}^{\prime 2}\right]\right) \tag{12}
\end{equation*}
$$

From the relational independence we have $\mathbb{E}_{p_{+}}\left[h_{i} t_{i} h_{i}^{\prime} t_{i}^{\prime}\right]=\mathbb{E}_{p_{+}}\left[h_{i} t_{i}\right] \mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{i}^{\prime}\right]$. Moreover, because the word embeddings are assumed to be standerdised to unit variance we have $\mathbb{E}_{p_{+}}\left[h_{i} t_{i}\right]=$ $\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{i}^{\prime}\right]=1$ and $\mathbb{E}_{p_{+}}\left[h_{i}^{2} t_{i}^{2}\right]=\mathbb{E}_{p_{+}}\left[h_{i}^{\prime 2} t^{\prime 2}\right]=1$. Therefore, (12) evaluates to zero and none of the terms arising purely from $\underline{\mathbf{A}}$ will remain in the expected loss over positive examples.
Next, lets consider the $A_{i j}^{(k)} P_{k n}$ terms in the expansion of (10) given by,

$$
\begin{equation*}
2 \sum_{i, j} \sum_{n} A_{i j}^{(k)} P_{k n}\left(h_{i} t_{j}-h_{i}^{\prime} t_{j}^{\prime}\right)\left(h_{n}-h_{n}^{\prime}\right) . \tag{13}
\end{equation*}
$$

Taking the expectation of (13) w.r.t. $p_{+}$we get,

$$
\begin{align*}
& 2 \sum_{i, j} \sum_{n} A_{i j}^{(k)} P_{k n}\left(\mathbb{E}_{p_{+}}\left[h_{i} t_{j} h_{n}\right]-\mathbb{E}_{p_{+}}\left[h_{i} t_{j} h_{n}^{\prime}\right]-\right. \\
& \left.\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{j}^{\prime} h_{n}\right]+\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{j}^{\prime} h_{n}^{\prime}\right]\right) . \tag{14}
\end{align*}
$$

Likewise, from the uncorrelation assumption and relational independence it follows that all the expectations in (14) are zero. A similar argument can be used to show that terms that involve $A_{i j}^{(k)} Q_{k n}$ disappear from (10). Therefore, $\underline{\mathbf{A}}$ does not play any part in the expected loss over positive examples. Similarly, we can show that $\underline{\mathbf{A}}$ is independent of the expected loss over negative examples. Therefore, from (5) we see that the expected loss over the entire training dataset is independent of A.

### 3.1 Regularised $\ell_{2}$ loss

As a special case, if we attempt to minimise the expected loss under some regularisation on $\mathbf{A}$ such as the Frobenius norm regularisation, then this can be achieved by sending $\underline{\mathbf{A}}$ to zero tensor because according to Theorem 1 (2) is independent from A.

With $\underline{\mathbf{A}}=\underline{\mathbf{0}}$, the relation between $h$ and $t$ can be simplified to:

$$
\begin{equation*}
\boldsymbol{r}(\boldsymbol{h}, \boldsymbol{t})=\mathbf{P} \boldsymbol{h}+\mathbf{Q} t \tag{15}
\end{equation*}
$$

Then the expected loss over the positive instances is given by (16).
$\mathbb{E}_{p_{+}}\left[\left\|\mathbf{P}\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)+\mathbf{Q}\left(\boldsymbol{t}-\boldsymbol{t}^{\prime}\right)\right\|_{2}^{2}\right]=$
$\mathbb{E}_{p_{+}}\left[\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)^{\top} \mathbf{P}^{\top} \mathbf{P}\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)\right]+\mathbb{E}_{p_{+}}\left[\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)^{\top} \mathbf{P}^{\top} \mathbf{Q}\left(\boldsymbol{t}-\boldsymbol{t}^{\prime}\right)\right]+$
$\mathbb{E}_{p_{+}}\left[\left(\boldsymbol{t}-\boldsymbol{t}^{\prime}\right)^{\top} \mathbf{Q}^{\top} \mathbf{P}\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)\right]+\mathbb{E}_{p_{+}}\left[\left(\boldsymbol{t}-\boldsymbol{t}^{\prime}\right)^{\top} \mathbf{Q}^{\top} \mathbf{Q}\left(\boldsymbol{t}-\boldsymbol{t}^{\prime}\right)\right]$

The second expectation term in RHS of (16) can be computed as follows:

$$
\begin{align*}
& \mathbb{E}_{p_{+}}\left[\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)^{\top} \mathbf{P}^{\top} \mathbf{Q}\left(\boldsymbol{t}-\boldsymbol{t}^{\prime}\right)\right] \\
& =\sum_{i, j}\left(\mathbf{P}^{\top} \mathbf{Q}\right)_{i j} \mathbb{E}_{p_{+}}\left[\left(h_{i}-h_{i}^{\prime}\right)\left(t_{j}-t_{j}^{\prime}\right)\right] \\
& =\sum_{i, j}\left(\mathbf{P}^{\top} \mathbf{Q}\right)_{i j}\left(\mathbb{E}_{p_{+}}\left[h_{i} t_{j}\right]-\mathbb{E}_{p_{+}}\left[h_{i} t_{j}^{\prime}\right]-\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{j}\right]+\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} t_{j}^{\prime}\right]\right) \tag{17}
\end{align*}
$$

When $i \neq j$, each of the four expectations in the RHS of (19) are zero from the uncorrelation assumption. When $i=j$, each term will be equal to one from the standeridisation assumption (unit variance) and cancel each other out. A similar argument can be used to show that the third expectation term in the RHS of (16) vanishes.

Now lets consider the first expectation term in the RHS of (16), which can be computed as follows:

$$
\begin{align*}
& \mathbb{E}_{p_{+}}\left[\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)^{\top} \mathbf{P}^{\top} \mathbf{P}\left(\boldsymbol{h}-\boldsymbol{h}^{\prime}\right)\right] \\
& =\sum_{i, j}\left(\mathbf{P}^{\top} \mathbf{P}\right)_{i j} \mathbb{E}_{p_{+}}\left[\left(h_{i}-h_{i}^{\prime}\right)\left(h_{j}-h_{j}^{\prime}\right)\right] \\
& =\sum_{i, j}\left(\mathbf{P}^{\top} \mathbf{P}\right)_{i j}\left(\mathbb{E}_{p_{+}+}\left[h_{i} h_{j}\right]-\mathbb{E}_{p_{+}}\left[h_{i} h_{j}^{\prime}\right]\right. \\
& \left.\quad-\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} h_{j}\right]+\mathbb{E}_{p_{+}}\left[h_{i}^{\prime} h_{j}^{\prime}\right]\right) \tag{18}
\end{align*}
$$

When $i \neq j$, it follows from the uncorrelation assumption that each of the four expectation terms in the RHS of (18) will be zero. For $i=j$ case we have,

$$
\begin{align*}
& \sum_{i, j}\left(\mathbf{P}^{\top} \mathbf{P}\right)_{i i}\left(\mathbb{E}_{p_{+}}\left[h_{i}^{2}\right]-2 \mathbb{E}_{p_{+}}\left[h_{i} h_{h^{\prime}}^{\prime}\right]+\mathbb{E}_{p_{+}}\left[h_{i}^{\prime 2}\right]\right) \\
& =2 \sum_{i, j}\left(\mathbf{P}^{\top} \mathbf{P}\right)_{i i} \tag{19}
\end{align*}
$$

Note that from the relational independence between $h$ and $h^{\prime}$ we have $\mathbb{E}_{p_{+}}\left[h_{i} h_{i}^{\prime}\right]=$ $\mathbb{E}_{p_{+}}\left[h_{i}\right] \mathbb{E}_{p_{+}}\left[h_{i}^{\prime}\right]$. From the standerdidation (zero mean) assumption this term is zero. On the other hand $\mathbb{E}_{p_{+}}\left[h_{i}^{2}\right]=\mathbb{E}_{p_{+}}\left[h_{i}^{\prime 2}\right]=1$ from the standerdidation (unit variance) assumption, which gives the result in (19).

Similarly, the fourth expectation term in the RHS of (16) evaluates to $2 \sum_{i, j}\left(\mathbf{Q}^{\top} \mathbf{Q}\right)_{i i}$, which shows that (16) evaluates to $2 \sum_{i, j}\left(\left(\mathbf{P}^{\top} \mathbf{P}\right)_{i i}+\left(\mathbf{Q}^{\top} \mathbf{Q}\right)_{i i}\right) . \quad$ Note that this is independent of the positive instances and will be equal to the expected loss over negative instances, which gives $\mathbb{E}_{p}[J]=0$ for the relational embedding given by (15).
It is interesting to note that PairDiff is a special case of (15), where $\mathbf{P}=\mathbf{I}$ and $\mathbf{Q}=-\mathbf{I}$. In the general case where word embeddings are nonstanderdised to unit variance, we can set $\mathbf{P}$ to be the diagonal matrix where $\mathbf{P}_{i i}=1 / \sigma_{i}$, where $\sigma_{i}$ is the variance of the $i$-th dimension of the word embedding space, to enforce standerdisation. Considering that $\mathbf{P}, \mathbf{Q}$ are parameters of the relational embedding, this is analogous to batch normalisation (Ioffe and Szegedy, 2015), where the appropriate parameters for the normalisation are learnt during training.

## 4 Experimental Results

### 4.1 Corss-dimensional Correlations

A key assumption in our theoretical analysis is the uncorrelations between different dimensions
in word embeddings. Here, we empirically verify the uncorrelation assumption for different input word embeddings. For this purpose, we create SG, CBOW and GloVe embeddings from the $u k W a C$ corpus ${ }^{4}$. We use a context window of 5 tokens and select words that occur at least 6 times in the corpus. We use the publicly available implementations for those methods by the original authors and set the parameters to the recommended values in (Levy et al., 2015) to create 50dimensional word embeddings. As a representative of counting-based word embeddings, we create a word co-occurrence matrix weighted by the positive pointwise mutual information (PPMI) and apply singular value decomposition (SVD) to obtain 50 -dimensional embeddings, which we refer to as the Latent Semantic Analysis (LSA) embeddings.

We use Latent Dirichlet Allocation (LDA) (Blei et al., 2003b) to create a topic model, and represent each word by its distribution over the set of topics. Ideally, each topic will capture some semantic category and the topic distribution provides a semantic representation for a word. We use gensim ${ }^{5}$ to extract 50 topics from a 2017 January dump of English Wikipedia. In contrast to the abovementioned word embeddings, which are dense and flat structured, we used Hierarchical Sparse Coding $^{6}$ (HSC) (Yogatama et al., 2015) to produce sparse and hierarchical word embeddings.

Given a word embedding matrix $\mathbf{W} \in \mathbb{R}^{m \times d}$, where each row correspond to the $d$-dimensional embedding of a word in a vocabulary containing $m$ words, we compute a correlation matrix $\mathbf{C} \in \mathbb{R}^{d \times d}$, where the $(i, j)$ element, $C_{i j}$, denotes the Pearson correlation coefficient between the $i$ th and $j$-th dimensions in the word embeddings over the $m$ words. By construction $C_{i i}=1$ and the histograms of the cross-dimensional correlations ( $i \neq j$ ) are shown in Figure 1 for 50 dimensional word embeddings obtained from the six methods described above. The mean of the absolute pairwise correlations for each embedding type and the standard deviation (sd) are indicated in the figure.

From Figure 1, irrespective of the word em-

[^2]

Figure 1: Cross-dimensional correlations for six word embeddings.
bedding learning method used, we see that crossdimensional correlations are distributed in a narrow range with an almost zero mean. This result empirically validates the uncorrelation assumption we used in our theoretical analysis. Moreover, this result indicates that Theorem 1 can be applied to a wide-range of existing word embeddings.

### 4.2 Learning Relation Representations

Our theoretical analysis in $\S 3$ claims that the performance of the bilinear relational embedding is independent of the tensor operator $\underline{\mathbf{A}}$. To empirically verify this claim, we conduct the following experiment. For this purpose, we use the BATS dataset (Gladkova et al., 2016) that contains of 40 semantic and syntactic relation types ${ }^{7}$, and generate positive examples by pairing word-pairs that have the same relation types. Approximately each relation type has 1,225 word-pairs, which enables us to generate a total of 48 k positive training instances (analogous word-pairs) of the form $\left((h, t),\left(h^{\prime}, t^{\prime}\right)\right)$. For each pair $(h, t)$ related by a relation $r$, we randomly select pairs $\left(h^{\prime}, t^{\prime}\right)$ with a different relation type $r^{\prime}$, according to the $\ell_{2}$ distance between the two pairs to create negative (nonanalogous) instances. ${ }^{8}$ We collectively refer both positive and negative training instances as the training dataset.

Using the $d=50$ dimensional word embeddings from CBOW, SG, GloVe, LSA, LDA, and

[^3]

Figure 2: The learnt model parameters for different word embeddings of 50 dimensions.

HSC methods created in $\S 4.1$, we learn relational embeddings according to (2) by minimising the $\ell_{2}$ loss, (4). To avoid overfitting, we perform $\ell_{2}$ regularisation on $\underline{\mathbf{A}}, \mathbf{P}$ and $\mathbf{Q}$ are regularised to diagonal matrices $p \mathbf{I}$ and $q \mathbf{I}$, for $p, q \in \mathbb{R}$. We initialise all parameters by uniformly sampling from $[-1,+1]$ and use AdaGrad (Duchi et al., 2011) with initial learning rate set to 0.01 .

Figure 2 shows the Frobenius norm of the tensor $\underline{\mathbf{A}}$ (on the left vertical axis) and the values of $p$ and $q$ (on the right vertical axis) for the six word embeddings. In all cases, we see that as the training progresses, $\underline{\mathbf{A}}$ goes to zero as predicted by Theorem 1 under regularisation. Moreover, we see that approximately $p \approx-q=c$ is reached for some $c \in \mathbb{R}$ in all cases, which implies that $\mathbf{P} \approx-\mathbf{Q}=c \mathbf{I}$, which is the PairDiff operator. Among the six input word embeddings compared in Figure 1, HSC has the highest mean correlation (0.082), which implies that its dimensions are correlated more than in the other word embeddings. This is to be expected by design because a hierarchical structure is imposed on the dimensions of the word embedding during training. However, HSC embeddings also satisfy the $\underline{\mathbf{A}} \approx \underline{\mathbf{0}}$ and $p \approx-q=c$ requirements, as expected by the PairDiff. This result shows that the claim of Theorem 1 is empirically true even when the uncorrelation assumption is mildly violated.


Figure 3: The training loss and test performance on SAT and SemEval benchmarks for relational embeddings. When the loss converges to zero, the performance on both benchmarks reaches that of the PairDiff.

### 4.3 Generalisation Performance on Analogy Detection

So far we have seen that the bilinear relational representation given by (2) does indeed converge to the form predicted by our theoretical analysis for different types of word embeddings. However, it remains unclear whether the parameters learnt from the training instances generated from the BATS dataset accurately generalise to other benchmark datasets for analogy detection. To emphasise, our focus here is not to outperform relational representation methods proposed in previous works, but rather to empirically show that the learnt operator converges to the popular PairDiff for the analogy detection task.

To measure the generalisation capability of the learnt relational embeddings from BATS, we measure their performance on two other benchmark datasets: SAT (Turney and Bigham, 2003) and SemEval 2012-Task2 ${ }^{9}$. Note that we do not retrain A, $\mathbf{P}$ and $\mathbf{Q}$ in (2) on SAT nor SemEval, but simply to use their values learnt from BATS because the purpose here to evaluate the generalisation of the learnt operator.

In SAT analogical questions, given a stem wordpair $(a, b)$ with five candidate word-pairs $(c, d)$, the task is to select the word-pair that is relationally similar to the the stem word-pair. The relational similarity between two word-pairs $(a, b)$ and $(c, d)$ is computed by the cosine similarity

[^4]between the corresponding relational embeddings $\boldsymbol{r}(\boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{r}(\boldsymbol{c}, \boldsymbol{d})$. The candidate word-pair that has the highest relational similarity with the stem word-pair is selected as the correct answer to a word analogy question. The reported accuracy is the ratio of the correctly answered questions to the total number of questions. On the other hand, SemEval dataset has 79 semantic relations, with each relation having ca. 41 word-pairs and four prototypical examples. The task is to assign a score for each word pair which is the average of the relational similarity between the given word-pair and prototypical word-pairs in a relation. Maximum difference scaling (MaxDiff) is used as the evaluation measure in this task.

Figure 3 shows the performance of the relational embeddings composed from 50dimensional CBOW embeddings. ${ }^{10}$ The level of performance reported by PairDiff on SAT and SemEval datasets are respectively $35.16 \%$ and $41.94 \%$, and are shown by horizontal dashed lines. From Figure 3, we see that the training loss gradually decreases with the number of training epochs and the performance of the relational embeddings on SAT and SemEval datasets reach that of the PairDiff operator. This result indicates that the relational embeddings learnt not only converge to PairDiff operator on training data but also generalise to unseen relation types in SAT and SemEval test datasets.

## 5 Conclusion

We showed that, if the word embeddings are standardised and uncorrelated, then the expected $\ell_{2}$ distance between analogous and non-analogous word-pairs is independent of bilinear terms, and the relation embedding further simplifies to the popular PairDiff operator under regularised settings. Moreover, we provided empirical evidence showing the uncorrelation in word embedding dimensions, where their cross-dimensional correlations are narrowly distributed around a mean close to zero. An interesting future research direction of this work is to extend the theoretical analysis to nonlinear relation composition operators, such as for nonlinear neural networks.

[^5]
## References

William Blacoe and Mirella Lapata. 2012. A comparison of vector-based representations for semantic composition. In Proc. of EMNLP. pages 546-556.

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003a. Latent dirichlet allocation. Journal of Machine Learning Research 3:993-1022.

David M Blei, Andrew Y Ng, and Michael I Jordan. 2003b. Latent dirichlet allocation. Machine Learning research 3:993-1022.

Danushka Bollegala, Takanori Maehara, and Ken ichi Kawarabayashi. 2015a. Embedding semantic relationas into word representations. In Proc. of IJCAI. pages 1222-1228.

Danushka Bollegala, Takanori Maehara, Yuichi Yoshida, and Ken ichi Kawarabayashi. 2015b. Learning word representations from relational graphs. In Proc. of AAAI. pages 2146-2152.

Danushka Bollegala, Yutaka Matsuo, and Mitsuru Ishizuka. 2009. A relational model of semantic similarity between words using automatically extracted lexical pattern clusters from the web. In Proc. of $E M N L P$. pages 803-812.

Antoine Bordes, Nicolas Usunier, Alberto GarciaDurán, Jason Weston, and Oksana Yakhenko. 2013. Translating embeddings for modeling multirelational data. In Proc. of NIPS. pages 2787-2795.

Scott Deerwester, Susan T. Dumais, George W. Furnas, Thomas K. Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. American Society for Information Science 41(6):391-407.

Aleksandr Drozd, Anna Gladkova, and Satoshi Matsuoka. 2016. Word embeddings, analogies, and machine learning: Beyond king - man + woman $=$ queen. In Proc. of COLING. pages 3519-3530.

Nguyen Tuan Duc, Danushka Bollegala, and Mitsuru Ishizuka. 2010. Using relational similarity between word pairs for latent relational search on the web. In Proc. of WI. pages 196-199.

John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. Journal of Machine Learning Research 12:2121-2159.

Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah A. Smith. 2015. Sparse overcomplete word vector representations. In Proc. of ACL. pages 1491-1500.

Anna Gladkova, Aleksandr Drozd, and Satoshi Matsuoka. 2016. Analogy-based detection of morphological and semantic relations with word embeddings: what works and what doesn't. In Proc. of SRW@ HLT-NAACL. pages 8-15.

Shu Guo, Quan Wang, Lihong Wang, Bin Wang, and Li Guo. 2016. Jointly embedding knowledge graphs and logical rules. In Proc. of EMNLP. pages 192202.

Huda Hakami and Danushka Bollegala. 2017. Compositional approaches for representing relations between words: A comparative study. KnowledgeBased Systems 136:172-182.

Katsuhiko Hayashi and Masashi Shinbo. 2017. On the equivalence of holograpic and complex embeddings for link prediction. In Proc. of ACL. pages 554-559.

Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Proc. of Machine Learning Research. pages 448-456.

David A. Jurgens, Saif Mohammad, Peter D. Turney, and Keith J. Holyoak. 2012. Semeval-2012 task 2: Measuring degrees of relational similarity. In Proc. of $* S E M$. pages $356-364$.

Omer Levy and Yoav Goldberg. 2014. Linguistic regularities in sparse and explicit word representations. In Proc. of CoNLL. pages 171-180.

Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. Transactions of Association for Computational Linguistics 3:211-225.

Tomas Mikolov, Kai Chen, and Jeffrey Dean. 2013a. Efficient estimation of word representation in vector space. In Proc. of ICLR.

Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013b. Linguistic regularities in continuous space word representations. In Proc. of hlt-Naacl. pages 746-751.

Pasquale Minervini, Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. 2017. Adversarial sets for regularising neural link predictors. In Proc of UAI.
J. Mu, S. Bhat, and P. Viswanath. 2017. All-but-theTop: Simple and Effective Postprocessing for Word Representations. ArXiv e-prints .

Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2015. A review of relational machine learning for knowledge graphs. In Proc. of the IEEE. volume 104, pages 11-33.

Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016. Holographic embeddings of knowledge graphs. In Proc. of AAAI.

Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In Proc. of ICML. pages 809-816.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In Proc. of EMNLP. volume 14, pages 1532-1543.

Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2016. Reasoning about entailment with neural attention. In Proc. of ICLR.

Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In Proc. of NIPS. pages 926-934.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In Proc. of ICML.

Peter D Turney and Jeffrey Bigham. 2003. Combining independent modules to solve multiple-choice synonym and analogy. In Proc. of RANLP.

Peter D Turney and Michael L Littman. 2005. Corpusbased learning of analogies and semantic relations. Machine Learning 60(1):251-278.

Ekaterina Vylomova, Laura Rimell, Trevor Cohn, and Timothy Baldwin. 2016. Take and took, gaggle and goose, book and read: Evaluating the utility of vector differences for lexical relational learning. In Proc. of ACL. pages 1671-1682.

Bishan Yang, Wen tau Yih, Xiadong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. In Proc. of ICLR.

Wenpeng Yin and Hinrich Schütze. 2016. Learning meta-embeddings by using ensembles of embedding sets. In Proc. of ACL. pages 1351-1360.

Dani Yogatama, Manaal Faruqui, Chris Dyer, and Noah A. Smith. 2015. Learning word representations with hierarchical sparse coding. In Proc. of ICML. pages $87-96$.


[^0]:    ${ }^{1}$ We interchangeably use the terms word and entity to represent both unigrams as well as a multi-word expressions including named entities.

[^1]:    ${ }^{2}$ If the projection matrix is nonsingular, then the inverse projection exists, which preserves the dimensionality of the embedding space.
    ${ }^{3}$ For $\ell_{2}$ normalised vectors, their Euclidean distance is a monotonously decreasing function of their cosine similarity.

[^2]:    ${ }^{4}$ http://wacky.sslmit.unibo.it/doku. php?id=corpora
    ${ }^{5}$ https://radimrehurek.com/gensim/wiki. html
    ${ }^{6}$ http://www.cs.cmu.edu/~ark/dyogatam/ wordvecs/

[^3]:    ${ }^{7}$ http://vsm.blackbird.pw/bats
    ${ }^{8} 10$ negative instances are generated from each word-pair in our experiments.

[^4]:    ${ }^{9}$ https://sites.google.com/site/ semeval2012task2/

[^5]:    ${ }^{10}$ Similar trends were observed for all six word embedding types but not shown here due to space limitations.

