

# Unsupervised Post-Processing of Word Vectors via Conceptor Negation

Tianlin Liu\*

Department of Computer Science and  
Electrical Engineering  
Jacobs University Bremen  
28759 Bremen, Germany  
t.liu@jacobs-university.de

Lyle Ungar, João Sedoc

Department of Computer and  
Information Science  
University of Pennsylvania  
Philadelphia, PA 19104  
{ungar, joao}@cis.upenn.edu

## Abstract

Word vectors are at the core of many natural language processing tasks. Recently, there has been interest in post-processing word vectors to enrich their semantic information. In this paper, we introduce a novel word vector post-processing technique based on *matrix conceptors* (Jaeger 2014), a family of regularized identity maps. More concretely, we propose to use conceptors to suppress those latent features of word vectors having high variances. The proposed method is purely unsupervised: it does not rely on any corpus or external linguistic database. We evaluate the post-processed word vectors on a battery of intrinsic lexical evaluation tasks, showing that the proposed method consistently outperforms existing state-of-the-art alternatives. We also show that post-processed word vectors can be used for the downstream natural language processing task of dialogue state tracking, yielding improved results in different dialogue domains.

## Introduction

Distributional representations of words, better known as word vectors, are a cornerstone of practical natural language processing (NLP). Examples of word vectors include Word2Vec (Mikolov et al. 2013), GloVe (Pennington, Socher, and Manning 2014), Eigenwords (Dhillon, Foster, and Ungar 2015), and Fasttext (Bojanowski et al. 2017). These word vectors are usually referred to as *distributional word vectors*, as their training methods rely on the *distributional hypothesis* of semantics (Firth 1957).

Recently, there has been interest in post-processing distributional word vectors to enrich their semantic content. The post-process procedures are usually performed in a lightweight fashion, i.e., without re-training word vectors on a text corpus. In one line of study, researchers used *supervised methods* to enforce linguistic constraints (e.g., synonym relations) on word vectors (Faruqui et al. 2015; Mrksic et al. 2016; 2017), where the linguistic constraints are extracted from an external linguistic knowledge base such as WordNet (Miller 1995) and PPDB (Pavlick et al. 2015). In another line of study, researchers devised *unsupervised methods* to post-process word vectors. Spectral-

decomposition methods such as singular value decomposition (SVD) and principal component analysis (PCA) are usually used in this line of research (Caron 2001; Bullinaria and Levy 2012; Turney 2012; Levy and Goldberg 2014; Levy, Goldberg, and Dagan 2015; Mu and Viswanath 2018). The current paper is in line with the second, unsupervised, research direction.

Among different unsupervised word vector post-processing schemes, the *all-but-the-top* approach (Mu and Viswanath 2018) is a prominent example. Empirically studying the latent features encoded by principal components (PCs) of distributional word vectors, Mu and Viswanath (2018) found that the variances explained by the leading PCs “encode the frequency of the word to a significant degree”. Since word frequencies are arguably unrelated to lexical semantics, they recommend removing such leading PCs from word vectors using a PCA reconstruction.

The current work advances the findings of Mu and Viswanath (2018) and improves their post-processing scheme. Instead of discarding a fixed number of PCs, we *softly* filter word vectors using *matrix conceptors* (Jaeger 2014; 2017), which characterize the linear space of those word vector features having high variances – the features most contaminated by word frequencies according to Mu and Viswanath (2018). The proposed approach is mathematically simple and computationally efficient, as it is founded on elementary linear algebra. Besides these traits, it is also practically effective: using a standard set of lexical-level intrinsic evaluation tasks and a deep neural network-based dialogue state tracking task, we show that conceptor-based post-processing considerably enhances linguistic regularities captured by word vectors. A more detailed list of our contributions are:

1. We propose an unsupervised algorithm that leverages Boolean operations of conceptors to post-process word vectors. The resulting word vectors achieve up to 18.86% and 28.34% improvement on the SimLex-999 and SimVerb-3500 dataset relative to the original word representations.
2. A closer look at the proposed algorithm reveals commonalities across several existing post-processing techniques for neural-based word vectors and pointwise mutual information (PMI) matrix based word vectors. Unlike

\*Research done while visiting University of Pennsylvania.  
Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the existing alternatives, the proposed approach is flexible enough to remove lexically-unrelated noise, while general-purpose enough to handle word vectors induced by different learning algorithms.

The rest of the paper is organized as follows. We first briefly review the principal component nulling approach for unsupervised word vector post-processing introduced in (Mu and Viswanath 2018), upon which our work is based. We then introduce our proposed approach, *Conceptor Negation* (CN). Analytically, we reveal the links and differences between the CN approach and the existing alternatives. Finally, we showcase the effectiveness of the CN method with numerical experiments<sup>1</sup>.

## Notation

We assume a collection of words  $w \in V$ , where  $V$  is a vocabulary set. Each word  $w \in V$  is embedded as a  $n$  dimensional real valued vector  $v_w \in \mathbb{R}^n$ . An identity matrix will be denoted by  $\mathbf{I}$ . For a vector  $v$ , we denote  $\text{diag}(v)$  as the diagonal matrix with  $v$  on its diagonal. We write  $[n] = \{1, 2, \dots, n\}$  for a positive integer  $n$ .

## Post-processing word vectors by PC removal

This section is an overview of the all-but-the-top (ABTT) word vector post-processing approach introduced by Mu and Viswanath (2018). In brief, the ABTT approach is based on two key observations of distributional word vectors. First, using a PCA, Mu and Viswanath (2018) revealed that word vectors are strongly influenced by a few leading principal components (PCs). Second, they provided an interpretation of such leading PCs: they empirically demonstrated a correlation between the variances explained by the leading PCs and word frequencies. Since word frequencies are arguably unrelated to lexical semantics, they recommend eliminating top PCs from word vectors via a PCA reconstruction. This method is described in Algorithm 1.

---

**Algorithm 1:** The all-but-the-top (ABTT) algorithm for word vector post-processing.

---

**Input :** (i)  $\{v_w \in \mathbb{R}^n : w \in V\}$ : word vectors with a vocabulary  $V$ ; (ii)  $d$ : the number of PCs to be removed.

- 1 Center the word vectors: Let  $\bar{v}_w := v_w - \mu$  for all  $w \in V$ , where  $\mu$  is the mean of the input word vectors.
- 2 Compute the first  $d$  PCs  $\{u_i \in \mathbb{R}^n\}_{i \in [d]}$  of the column-wisely stacked centered word vectors  $[\bar{v}_w]_{w \in V} \in \mathbb{R}^{n \times |V|}$  via a PCA.

- 3 Process the word vectors:

$$\tilde{v}_w^{\text{ABTT}} := \bar{v}_w - \sum_{i=1}^d u_i^\top u_i \bar{v}_w, \forall w \in V.$$

**Output:**  $\{\tilde{v}_w^{\text{ABTT}}, w \in V\}$

---

In practice, Mu and Viswanath (2018) found that the improvements yielded by ABTT are particularly impressive for

<sup>1</sup>Our codes are available at <https://github.com/liutianlin0121/Conceptor-Negation-WV>

word similarity tasks. Here, we provide a straightforward interpretation of the effects. Concretely, consider two arbitrary words  $w_1$  and  $w_2$  with word vectors  $v_{w_1}$  and  $v_{w_2}$ . Without loss of generality, we assume  $v_{w_1}$  and  $v_{w_2}$  are normalized, i.e.,  $\|v_{w_1}\|_2 = \|v_{w_2}\|_2 = 1$ . Given PCs  $\{u_1, \dots, u_n\}$  of the word vectors  $\{v_w : w \in V\}$ , we re-write  $v_{w_1}$  and  $v_{w_2}$  via linear combinations with respect to the basis  $\{u_1, \dots, u_n\}$ :  $v_{w_1} := \sum_{i=1}^n \beta_i u_i$  and  $v_{w_2} := \sum_{i=1}^n \beta'_i u_i$ , for some  $\beta_i, \beta'_i \in \mathbb{R}$  and for all  $i \in [n]$ . We see

$$\text{cosine}(v_{w_1}, v_{w_2}) \stackrel{(*)}{=} v_{w_1}^\top v_{w_2} \quad (1)$$

$$= \left( \sum_{i=1}^n \beta_i u_i \right)^\top \left( \sum_{i=1}^n \beta'_i u_i \right)$$

$$\stackrel{(**)}{=} \sum_{i=1}^n \beta_i \beta'_i \quad (2)$$

where  $(*)$  holds because the word vectors were assumed to be normalized and  $(**)$  holds because  $\{u_1, \dots, u_n\}$  is an orthonormal basis of  $\mathbb{R}^n$ . Via Equation 2, the similarity between word  $w_1$  and  $w_2$  can be seen as the overall “compatibility” of their measurements  $\beta_i$  and  $\beta'_i$  with respect to each latent feature  $u_i$ . If leading PCs encode the word frequencies, removing the leading PCs, in theory, help the word vectors capture semantic similarities, and consequently improve the experiment results of word similarity tasks.

## Post-processing word vectors via Conceptor Negation

Removing the leading PCs of word vectors using the ABTT algorithm described above is effective in practice, as seen in the elaborate experiments conducted by Mu and Viswanath (2018). However, the method comes with a potential limitation: for each latent feature taking form as a PC of the word vectors, ABTT either completely removes the feature or keeps it intact. For this reason, Khodak et al. (2018) argued that ABTT is liable either to not remove enough noise or to cause too much information loss.

The objective of this paper is to address the limitations of ABTT. More concretely, we propose to use *matrix conceptors* (Jaeger 2017) to gate away variances explained by the leading PCs of word vectors. As will be seen later, the proposed *Conceptor Negation* method removes noise in a “softer” manner when compared to ABTT. We show that it shares the spirit of an eigenvalue weighting approach for PMI-based word vector post-processing. We proceed by providing the technical background of conceptors.

## Conceptors

Conceptors are a family of regularized identity maps introduced by Jaeger (2014). We present a sketch of conceptors by heavily re-using (Jaeger 2014; He and Jaeger 2018) sometimes verbatim. In brief, a *matrix conceptor*  $C$  for some vector-valued random variable  $x$  taking values in  $\mathbb{R}^N$  is defined as a linear transformation that minimizes the following loss function.

$$\mathbb{E} [\|x - Cx\|_2^2] + \alpha^{-2} \|C\|_F^2 \quad (3)$$

where  $\alpha$  is a control parameter called *aperture*,  $\|\cdot\|_2$  is the  $\ell_2$  norm, and  $\|\cdot\|_F$  is the Frobenius norm. This optimization problem has a closed-form solution

$$C = R(R + \alpha^{-2} \mathbf{I})^{-1} \quad (4)$$

where  $R = \mathbb{E}[xx^\top]$  and  $\mathbf{I}$  are  $N \times N$  matrices. If  $R = \Psi T \Psi^\top$  is the SVD of  $R$ , then the SVD of  $C$  is given as  $\Psi S \Psi^\top$ , where the singular values  $s_i$  of  $C$  can be written in terms of the singular values  $t_i$  of  $R$ :  $s_i = t_i / (t_i + \alpha^{-2}) \in (0, 1)$  for  $\alpha \in (0, \infty)$ . In intuitive terms,  $C$  is a soft projection matrix on the linear subspace where the samples of  $x$  lie, such that for a vector  $y$  in this subspace,  $C$  acts like the identity:  $Cy \approx y$ , and when some  $\epsilon$  orthogonal to the subspace is added to  $y$ ,  $C$  reconstructs  $y$ :  $C(y + \epsilon) \approx y$ .

Moreover, operations that satisfy most laws of Boolean logic such as NOT  $\neg$ , OR  $\vee$ , and AND  $\wedge$ , can be defined on matrix conceptors. These operations all have interpretation on the data level, i.e., on the distribution of the random variable  $x$  (details in (Jaeger 2014, Section 3.9)). Among these operations, the negation operation NOT  $\neg$  is relevant for the current paper:

$$\neg C := \mathbf{I} - C. \quad (5)$$

Intuitively, the negated conceceptor,  $\neg C$ , softly projects the data onto a linear subspace that can be roughly understood as the orthogonal complement of the subspace characterized by  $C$ .

### Post-processing word vectors with Conceceptor Negation

This subsection explains how conceptors can be used to post-process word vectors. The intuition behind our approach is simple. Consider a random variable  $x$  taking values on word vectors  $\{v_w \in \mathbb{R}^n : w \in V\}$ . We can estimate a conceceptor  $C$  that describes the distribution of  $x$  using Equation 4. Recall that (Mu and Viswanath 2018) found that the directions with which  $x$  has the highest variances encode word frequencies, which are unrelated to word semantics. To suppress such word-frequency related features, we can simply pass all word vectors through the negated conceceptor  $\neg C$ , so that  $\neg C$  dampens the directions with which  $x$  has the highest variances. This simple method is summarized in Algorithm 2.

---

**Algorithm 2:** The conceceptor negation (CN) algorithm for word vector post-processing.

---

**Input :** (i)  $\{v_w \in \mathbb{R}^n : w \in V\}$ : word vectors of a vocabulary  $V$ ; (ii)  $\alpha \in \mathbb{R}$ : a hyper-parameter

- 1 Compute the conceceptor  $C$  from word vectors:

$$C = R(R + \alpha^{-2} \mathbf{I})^{-1}, \text{ where } R \text{ is estimated by } \frac{1}{|V|} \sum_w v_w v_w^\top$$

- 2 Compute  $\neg C := \mathbf{I} - C$

- 3 Process the word vectors:  $\tilde{v}_w^{\text{CN}} := \neg C v_w, \forall w \in V$

**Output:**  $\{\tilde{v}_w^{\text{CN}} : w \in V\}$

---

The hyper-parameter  $\alpha$  of Algorithm 2 governs the “sharpness” of the suppressing effects on word vectors em-

ployed by  $\neg C$ . Although in this work we are mostly interested in  $\alpha \in (0, \infty)$ , it is nonetheless illustrative to consider the extreme cases where  $\alpha = 0$  or  $\infty$ : for  $\alpha = 0$ ,  $\neg C$  will be an identity matrix, meaning that word vectors will be kept intact; for  $\alpha = \infty$ ,  $\neg C$  will be a zero matrix, meaning that all word vectors will be nulled to zero vectors. The computational costs of the Algorithm 2 are dominated by its step 1: one needs to calculate the matrix product  $R = \frac{1}{|V|} [v_w]_{w \in V} [v_w]_{w \in V}^\top$  for  $[v_w]_{w \in V} \in \mathbb{R}^{n \times |V|}$  being the matrix whose columns are word vectors. Since modern word vectors usually come with a vocabulary of some millions of words (e.g., Google News Word2Vec contains 3 million tokens), performing a matrix product on such large matrices  $[v_w]_{w \in V}$  is computationally laborious. But considering that there are many uninteresting words in the vast vocabulary, we find it is empirically beneficial to only use a subset of the vocabulary, whose words are not too peculiar<sup>2</sup>. Specifically, borrowing the word list provided by Arora, Liang, and Ma (2017)<sup>3</sup>, we use the words that appear at least 200 times in a Wikipedia dump 2015 to estimate  $R$ . This greatly boosts the computation speed. Somewhat surprisingly, the trick also improves the performance of Algorithm 2. This might due to the higher quality of word vectors of common words compared with infrequent ones.

### Analytic comparison with other methods

Since most of the existing unsupervised word vector post-processing methods are ultimately based on linear data transformations, we hypothesize that there should be commonalities between the methods. In this section, we show CN resembles ABTT in that both methods can be interpreted as “spectral encode-decode processes”; when applied to word vectors induced by a pointwise mutual information (PMI) matrix, CN shares the spirit with the *eigenvalue weighting* (EW) post-processing (Caron 2001; Levy, Goldberg, and Dagan 2015): they both assign weights on singular vectors of a PMI matrix. A key distinction of CN is that it does *soft* noise removal (unlike ABTT) and that it is not restricted to post-processing PMI-matrix induced word vectors (unlike EW).

### Relation to ABTT

In this subsection, we reveal the connection between CN and ABTT. To do this, we will re-write the last step of both algorithms into different formats. For the convenience of comparison, throughout this section, we will assume that the word vectors  $\{v_w\}_{w \in V}$  in Algorithm 1 and Algorithm 2 possess a zero mean, although this is not a necessary requirement in general.

We first re-write the equation in step 3 of Algorithm 1. We let  $U$  be the matrix whose columns are the PCs estimated from the word vectors. Let  $U_{:,1:d}$  be the first  $d$  columns of  $U$ . It is clear that step 2 of Algorithm 1, under the assumption that word vectors possess zero mean, can be re-written as

<sup>2</sup>This trick has also been used for ABTT by Mu and Viswanath (2018) (personal communications).

<sup>3</sup>[https://github.com/PrincetonML/SIF/tree/master/auxiliary\\_data](https://github.com/PrincetonML/SIF/tree/master/auxiliary_data)

$$\begin{aligned}\tilde{v}_w^{\text{ABTT}} &:= (I - U_{:,1:d}U_{:,1:d}^\top) v_w \\ &= U \text{diag}(\underbrace{[0, \dots, 0]}_{d \text{ copies of } 0}, 1, \dots, 1) U^\top v_w.\end{aligned}\quad (6)$$

Next, we re-write step 3 of the Conceptor Negation (CN) method of algorithm 2. Note that for word vectors with zero mean, the estimation for  $R$  is a (sample) covariance matrix of a random variable taking values as word vectors, and therefore the singular vectors of  $R$  are PCs of word vectors. Letting  $R = U\Sigma U^\top$  be the SVD of  $R$ , the equation in step 3 of Algorithm 2 can be re-written via elementary linear algebraic operations:

$$\begin{aligned}\tilde{v}_w^{\text{CN}} &:= -C\bar{v}_w \\ &= (\mathbf{I} - C) v_w \\ &= (\mathbf{I} - R(R + \alpha^{-2}I)^{-1}) v_w \\ &= (\mathbf{I} - U\Sigma U^\top (U\Sigma U^\top + \alpha^{-2}UU^\top)^{-1}) v_w \\ &= \left( \mathbf{I} - U \text{diag}\left(\left[\frac{\sigma_1}{\sigma_1 + \alpha^{-2}}, \dots, \frac{\sigma_n}{\sigma_n + \alpha^{-2}}\right]\right) U^\top \right) v_w \\ &= U \text{diag}\left(\left[\frac{\alpha^{-2}}{\sigma_1 + \alpha^{-2}}, \dots, \frac{\alpha^{-2}}{\sigma_n + \alpha^{-2}}\right]\right) U^\top v_w,\end{aligned}\quad (7)$$

where  $\sigma_1, \dots, \sigma_n$  are diagonal entries of  $\Sigma$ .

Examining Equations 6 and 7, we see ABTT and CN share some similarities. In particular, they both can be unified into “spectral encode-decode processes,” which contain the following three steps:

1. **PC encoding.** Load word vectors on PCs by multiplying  $U^\top$  with  $v_w$ .
2. **Variance gating.** Pass the PC-encoded data through the variance gating matrices  $\text{diag}([0, \dots, 0, 1, \dots, 1])$  and  $\text{diag}([\frac{\alpha^{-2}}{\sigma_1 + \alpha^{-2}}, \dots, \frac{\alpha^{-2}}{\sigma_n + \alpha^{-2}}])$  respectively for ABTT and CN.
3. **PC decoding.** Transform the data back to the usual coordinates using the matrix  $U$ .

With the above encode-decode interpretation, we see CN differ from ABTT is its variance gating step. In particular, ABTT does a *hard* gating, in the sense that the diagonal entries of the variance gating matrix (call them variance gating coefficients) take values in the set  $\{0, 1\}$ . The CN approach, on the other hand, does a softer gating as the entries take values in  $(0, 1)$ :

$$0 < \frac{\alpha^{-2}}{\sigma_i + \alpha^{-2}} \leq \frac{\alpha^{-2}}{\sigma_j + \alpha^{-2}} < 1,$$

for all  $1 \leq i < j \leq n$  and  $\alpha \in (0, \infty)$ . To illustrate the gating effects, we plot the variance gating coefficients for ABTT and CN for Word2Vec in Figure 1.

### Relation with eigenvalue weighting

We relate the conceptor approach to the eigenvalue weighting approach for post-processing PMI-based word vectors. This effort is in line with the ongoing research in the NLP

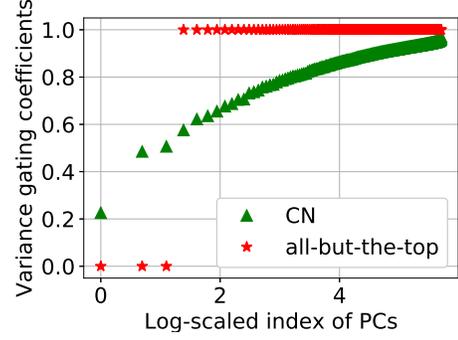


Figure 1: The variance gating coefficients of ABTT and CN for Word2Vec. Hyper-parameters:  $d = 3$  for ABTT and  $\alpha = 2$  for CN.

community that envisages a connection between “neural word embedding” and PMI-matrix factorization based word embedding (Levy and Goldberg 2014; Pennington, Socher, and Manning 2014; Levy, Goldberg, and Dagan 2015).

In the PMI approach for word association modeling, for each word  $w$  and each context (i.e., sequences of words)  $q$ , the PMI matrix  $M$  assigns a value for the pair  $(w, q)$ :  $M(w, q) = \log \frac{\mathbb{P}(w, q)}{\mathbb{P}(w)\mathbb{P}(q)}$ . In practical NLP tasks, the sets of words and contexts tend to be large, and therefore, directly working with  $M$  is inconvenient. To lift the problem, one way is to perform a truncated SVD on  $M$ , factorizing  $M$  into the product of three smaller matrices  $M \approx \Theta_{:,1:n} D_{1:n,1:n} \Gamma_{:,1:n}^\top$ , where  $\Theta_{:,1:n}$  is the first  $n$  left singular vectors of the matrix  $M$ ,  $D_{1:n,1:n}$  is the diagonal matrix containing  $n$  leading singular values of  $M$ , and  $\Gamma_{:,1:n}$  are the first  $n$  right singular vectors of the matrix  $M$ . A generic way to induce word vectors from  $M$  is to let

$$E := \Theta_{:,1:n} D_{1:n,1:n} \in \mathbb{R}^{|V| \times n},$$

which is a matrix containing word vectors as rows. Coined by Levy, Goldberg, and Dagan (2015), the term *eigenvalue weighting*<sup>4</sup> (EW) refers to a post-processing technique for PMI-matrix-induced word vectors. This technique has its root in Latent Semantic Analysis (LSA): Caron (2001) first propose to define the post-processed version of  $E$  as

$$\tilde{E}^{\text{EW}} := \Theta_{:,1:n} D_{1:n,1:n}^p,$$

where  $p$  is the *weighting exponent* determining the relative weights assigned to each singular vector of  $\Theta_{:,1:n}$ . While an optimal  $p$  depends on specific task demands, previous research suggests that  $p < 1$  is generally preferred, i.e., the contributions of the initial singular vectors of  $M$  should be suppressed. For instance,  $p = 0, 0.25$ , and  $0.5$  are recommended in (Caron 2001; Bullinaria and Levy 2012; Levy,

<sup>4</sup>It seems to us that a term “singular value weighting” is more appropriate because the weighting is based on singular values of a PMI matrix  $M$  but not eigenvalues of  $M$ . The term “eigenvalue” is relevant here only because the singular values of  $M$  are also the square roots of eigenvalues of  $M^\top M$ .

Goldberg, and Dagan 2015). Bullinaria and Levy (2012) argue that the initial singular vectors of  $M$  tend to be contaminated most by aspects other than lexical semantics.

We now show that applying CN on the PMI-matrix-based word embedding  $E := \Theta_{:,1:n} D_{1:n,1:n}$  has a tantamount effect with “suppressing initial singular vectors” of EW. Acting the negated  $-C$  on word vectors of  $E$  (i.e., rows of  $E$ ), we get the post-processed word vectors as rows of the  $\tilde{E}^{\text{CN}}$ :

$$\begin{aligned} \tilde{E}^{\text{CN}} &:= (-CE^{\top})^{\top} \\ &= E(I - R(R + \alpha^{-2}I)^{-1}) \\ &= E(I - \frac{1}{|V|} E^{\top} E (\frac{1}{|V|} E^{\top} E + \alpha^{-2}I)^{-1}) \\ &= \Theta_{:,1:n} D_{1:n,1:n} (I - \frac{1}{|V|} D_{1:n,1:n}^2 (\frac{1}{|V|} D_{1:n,1:n}^2 + \alpha^{-2}I)^{-1}) \\ &= \Theta_{:,1:n} D_{1:n,1:n} \text{diag}([\frac{|V| \cdot \alpha^{-2}}{\lambda_1^2 + |V| \cdot \alpha^{-2}}, \dots, \frac{|V| \cdot \alpha^{-2}}{\lambda_n^2 + |V| \cdot \alpha^{-2}}]) \end{aligned}$$

Since

$$0 < \frac{|V| \cdot \alpha^{-2}}{\lambda_i^2 + |V| \cdot \alpha^{-2}} \leq \frac{|V| \cdot \alpha^{-2}}{\lambda_j^2 + |V| \cdot \alpha^{-2}} < 1,$$

for all  $1 \leq i < j \leq n$  and  $\alpha \in (0, \infty)$ , these weights suppress the contribution of the initial singular vectors, similar to what has been done in EW.

## Experiments

We evaluate the post-processed word vectors on a variety of lexical-level intrinsic tasks and a down-stream deep learning task. We use the publicly available pre-trained Google News Word2Vec (Mikolov et al. 2013)<sup>5</sup> and Common Crawl GloVe<sup>6</sup> (Pennington, Socher, and Manning 2014) to perform lexical-level experiments. For CN, we fix  $\alpha = 2$  for Word2Vec and GloVe throughout the experiments<sup>7</sup>. For ABTT, we set  $d = 3$  for Word2Vec and  $d = 2$  for GloVe, as what has been suggested by Mu and Viswanath (2018).

**Word similarity** We test the performance of CN on seven benchmarks that have been widely used to measure word similarity: the RG65 (Rubenstein and Goodenough 1965), the WordSim-353 (WS) (Finkelstein et al. 2002), the rare-words (RW) (Luong, Socher, and Manning 2013), the MEN dataset (Bruni, Tran, and Baroni 2014), the MTurk (Radinsky et al. 2011), the SimLex-999 (SimLex) (Hill, Reichart, and Korhonen 2015), and the SimVerb-3500 (Gerz et al. 2016). To evaluate the word similarity, we calculate the cosine distance between vectors of two words using Equation 1. We report the Spearman’s rank correlation coefficient (Myers and Well 1995) of the estimated rankings against the rankings by humans in Table 1. We see that the proposed CN method consistently outperforms the original word embedding (orig.) and the post-processed word embedding by ABTT for most of the benchmarks.

<sup>5</sup><https://code.google.com/archive/p/word2vec/>

<sup>6</sup><https://nlp.stanford.edu/projects/glove/>

<sup>7</sup>Analytical optimization methods for the aperture  $\alpha$  are available from (Jaeger 2014), remaining to be connected with the word vector post-processing scheme in the future.

	WORD2VEC			GLOVE		
	orig.	ABTT	CN	orig.	ABTT	CN
RG65	76.08	78.34	<b>78.92</b>	76.96	74.36	<b>78.40</b>
WS	68.29	69.05	<b>69.30</b>	73.79	76.79	<b>79.08</b>
RW	53.74	54.33	<b>58.04</b>	46.41	52.04	<b>58.98</b>
MEN	78.20	<b>79.08</b>	78.67	80.49	81.78	<b>83.38</b>
MTurk	68.23	<b>69.35</b>	66.81	69.29	70.85	<b>71.07</b>
SimLex	44.20	45.10	<b>46.82</b>	40.83	44.97	<b>48.53</b>
SimVerb	36.35	36.50	<b>38.30</b>	28.33	32.23	<b>36.36</b>

Table 1: Post-processing results (Spearman’s rank correlation coefficient  $\times 100$ ) under seven word similarity benchmarks. The baseline results (orig. and ABTT) are collected from (Mu and Viswanath 2018).

The improvement of results by CN are particularly impressive for two “modern” word similarity benchmarks SimLex and SimVerb – these two benchmarks carefully distinguish *genuine word similarity* from *conceptual association* (Hill, Reichart, and Korhonen 2015). For instance, *coffee* is associated with *cup* but by no means similar to *cup*, a confusion often made by earlier benchmarks. In particular, SimLex has been heavily used to evaluate word vectors yielded by *supervised* word vector fine-tuning algorithms, which perform gradient descent on word vectors with respect to linguistic constraints such as synonym and antonym relationships of words extracted from WordNet and/or PPDB. When compared to a recent supervised approach of counter-fitting. Our results on SimLex are comparable to those reported by Mrksic et al. (2016), as shown in Table 2.

	Post-processing method	WORD2VEC	GLOVE
supervised	Counter-Fitting + syn.	0.45	0.46
	Counter-Fitting + ant.	0.33	0.43
	Counter-Fitting + syn. + ant.	<b>0.47</b>	<b>0.50</b>
unsupervised	CN	<b>0.47</b>	0.49

Table 2: Comparing the testing results (Spearman’s rank correlation coefficient) on SimLex with those of Counter-Fitting approach (results collected from (Mrksic et al. 2016, Table 2) and (Mrksic et al. 2017, Table 3)). The linguistic constraints for Counter-Fitting are synonym (syn.) and/or antonym (ant.) relationships extracted from English PPDB.

**Semantic Textual Similarity** In this subsection, we showcase the effectiveness of the proposed post-processing method using semantic textual similarity (STS) benchmarks, which are designed to test the semantic similarities of sentences. We use 2012-2015 SemEval STS tasks (Agirre et al. 2012; 2013; 2014; 2015) and the 2012 SemEval Semantic Related task (SICK) (Marelli et al. 2014).

Concretely, for each pair of sentences,  $s_1$  and  $s_2$ , we computed  $v_{s_1}$  and  $v_{s_2}$  by averaging their constituent word vectors. We then calculated the cosine distance between two sentence vectors  $v_{s_1}$  and  $v_{s_2}$ . This naive method has been shown to be a strong baseline for STS tasks (Wieting et al. 2016). As in Agirre et al. (2012), we used Pearson correlation of the estimated rankings of sentence similarity against

the rankings by humans to assess model performance.

In Table 3, we report the average result for the STS tasks each year (detailed results are in supplementary material). Again, our CN method consistently outperforms the alternatives.

	WORD2VEC			GLOVE		
	orig.	ABTT	CN	orig.	ABTT	CN
STS 2012	57.22	<b>57.67</b>	54.31	48.27	54.06	<b>54.38</b>
STS 2013	56.81	57.98	<b>59.17</b>	44.83	51.71	<b>55.51</b>
STS 2014	62.89	63.30	<b>66.22</b>	51.11	59.23	<b>62.66</b>
STS 2015	62.74	63.35	<b>67.15</b>	47.23	57.29	<b>63.74</b>
SICK	70.10	70.20	<b>72.71</b>	65.14	<b>67.85</b>	66.42

Table 3: Post-processing results ( $\times 100$ ) on the semantic textual similarity tasks. The baseline results (orig. and ABTT) are collected from (Mu and Viswanath 2018).

**Concept Categorization** In the concept categorization task, we used  $k$ -means to cluster words into concept categories based on their vector representations (for example, “bear” and “cat” belong to the concept category of animals). We use three standard datasets: (i) a rather small dataset ESSLLI 2008 (Baroni, Evert, and Lenci 2008) that contains 44 concepts in 9 categories; (ii) the Almuhareb-Poesio (AP) (Poesio and Almuhareb 2005), which contains 402 concepts divided into 21 categories; and (iii) the BM dataset (Battig and Montague 1969) that 5321 concepts divided into 56 categories. Note that the datasets of ESSLLI, AP, and BM are increasingly challenging for clustering algorithms, due to the increasing numbers of words and categories.

Following (Baroni, Dinu, and Kruszewski 2014; Schnabel et al. 2015; Mu and Viswanath 2018), we used “purity” of clusters (Manning, Raghavan, and Schütze 2008, Section 16.4) as the evaluation criterion. That the results of  $k$ -means heavily depend on two hyper-parameters: (i) the number of clusters and (ii) the initial centroids of clusters. We follow previous research (Baroni, Dinu, and Kruszewski 2014; Schnabel et al. 2015; Mu and Viswanath 2018) to set  $k$  as the ground-truth number of categories. The settings of the initial centroids of clusters, however, are less well-documented in previous work – it is not clear how many initial centroids have been sampled, or if different centroids have been sampled at all. To avoid the influences of initial centroids in  $k$ -means (which are particularly undesirable for this case because word vectors live in  $\mathbb{R}^{300}$ ), in this work, we simply fixed the initial centroids as the average of original, ABTT-processed, and CN-processed word vectors respectively from ground-truth categories. This initialization is fair because all post-processing methods make use of the ground-truth information equally, similar to the usage of the ground-truth numbers of clusters. We report the experiment results in Table 4.

The performance of the proposed methods and the baseline methods performed equally well for the smallest dataset ESSLLI. As the dataset got larger, the results differed and the proposed CN approaches outperformed the baselines by a margin.

	WORD2VEC			GLOVE		
	orig.	ABTT	CN.	orig.	ABTT	CN
ESSLLI	100.0	100.0	100.0	100.0	100.0	100.0
AP	87.28	88.3	<b>89.31</b>	86.43	87.19	<b>90.95</b>
BM	58.15	59.24	<b>60.19</b>	65.34	67.35	<b>67.63</b>

Table 4: Purity ( $\times 100$ ) of the clusters in concept categorization task with fixed centroids.

**A Downstream NLP task: Neural Belief Tracker** The experiments we have reported so far are all intrinsic lexical evaluation benchmarks. Only evaluating the post-processed word vectors using these benchmarks, however, invites an obvious critique: the success of intrinsic evaluation tasks may not transfer to downstream NLP tasks, as suggested by previous research (Schnabel et al. 2015). Indeed, when supervised learning tasks are performed, the post-processing methods such as ABTT and CN can *in principle* be absorbed into a classifier such as a neural network. Nevertheless, good initialization for classifiers is crucial. We hypothesize that the post-processed word vectors serve as a good initialization for those downstream NLP tasks that semantic knowledge contained in word vectors is needed.

To validate this hypothesis, we conducted an experiment using Neural Belief Tracker (NBT), a deep neural network based dialogue state tracking (DST) model (Mrksic et al. 2017; Mrkšić and Vulić 2018). As a concrete example to illustrate the purpose of the task, consider a dialogue system designed to help users find restaurants. When a user wants to find a Sushi restaurant, the system is expected to know that Japanese restaurants have a higher probability to be a good recommendation than Italian restaurants or Thai restaurants. Word vectors are important for this task because NBT needs to absorb useful semantic knowledge from word vectors using a neural network.

In our experiment with NBT, we used the model specified in (Mrkšić and Vulić 2018) with default hyper-parameter settings<sup>8</sup>. We report the *goal accuracy*, a default DST performance measure, defined as the proportion of dialogue turns where all the user’s search goal constraints match with the model predictions. The test data was Wizard-of-Oz (WOZ) 2.0 (Wen et al. 2017), where the goal constraints of users were divided into three domains: *food*, *price range*, and *area*. The experiment results are reported in Table 5.

	WORD2VEC			GLOVE		
	orig.	ABTT	CN.	orig.	ABTT	CN
Food	48.6	<b>84.7</b>	78.5	86.4	83.7	<b>88.8</b>
Price range	90.2	88.1	<b>92.2</b>	91.0	93.9	<b>94.7</b>
Area	83.1	82.4	<b>86.1</b>	93.5	<b>94.9</b>	93.7
Average	74.0	85.1	<b>85.6</b>	90.3	90.8	<b>92.4</b>

Table 5: The goal accuracy of food, price range, and area.

**Further discussions** Besides the NBT task, we have also tested ABTT and CN methods on other downstream NLP

<sup>8</sup><https://github.com/nmrksic/neural-belief-tracker>

tasks such as text classification (not reported). We found that ABTT and CN yield equivalent results in such tasks. One explanation is that the ABTT and CN post-processed word vectors are different only up to a small perturbation. With a sufficient amount of training data and an appropriate regularization method, a neural network should generalize over such a perturbation. With a relatively small training data (e.g., the 600 dialogues for training NBT task), however, we found that word vectors as initializations matters, and in such cases, CN post-processed word vectors yield favorable results. Another interesting finding is that, having tested ABTT and CN on Fasttext (Bojanowski et al. 2017), we found that neither post-processing method provides visible gain. We hypothesize that this might be because Fasttext includes subword (character-level) information in its word representation during training, which suppresses the word frequency features contained in word vectors. It remains for future work to validate this hypothesis.

## Conclusion

We propose a simple yet effective method for post-processing word vectors via the negation operation of conceptors. With a battery of intrinsic evaluation tasks and a down-stream deep-learning empowered dialogue state tracking task, the proposed method enhances linguistic regularities captured by word vectors and consistently improves performance over existing alternatives.

There are several possibilities for future work. We envisage that the logical operations and abstract ordering admitted by conceptors can be used in other NLP tasks. As concrete examples, the  $AND \wedge$  operation can be potentially applied to induce and fine-tune bi-lingual word vectors, by mapping word representations of individual languages into a shared linear space; the  $OR \vee$  together with  $NOT \neg$  operation can be used to study the vector representations of polysemous words, by joining and deleting sense-specific vector representations of words; the abstraction ordering  $\leq$  is a natural tool to study graded lexical entailment of words.

**Acknowledgement** We appreciate the anonymous reviewers for their constructive comments. We thank Xu He, Jordan Rodu, and Daphne Ippolito, and Chris Callison-Burch for helpful discussions.

## References

Agirre, E.; Diab, M.; Cer, D.; and Gonzalez-Agirre, A. 2012. Semeval-2012 task 6: A pilot on semantic textual similarity. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, SemEval '12, 385–393. Stroudsburg, PA, USA: Association for Computational Linguistics.

Agirre, E.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; and Guo, W. 2013. Sem 2013 shared task: Semantic textual similarity. In *Second Joint Conference on Lexical and Computational Semantics*, volume 1, 32–43.

Agirre, E.; Banea, C.; Cardie, C.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; Guo, W.; Mihalcea, R.; Rigau, G.; and

Wiebe, J. 2014. Semeval-2014 task 10: Multilingual semantic textual similarity. In *Proceedings of the 8th international workshop on semantic evaluation*, 81–91.

Agirre, E.; Banea, C.; Cardie, C.; Cer, D.; Diab, M.; Gonzalez-Agirre, A.; Guo, W.; Lopez-Gazpio, I.; Maritxalar, M.; Mihalcea, R.; Rigau, G.; Uriaa, L.; and Wiebe, J. 2015. Semeval-2015 task 2: Semantic textual similarity, English, Spanish and pilot on interpretability. In *Proceedings of the 9th international workshop on semantic evaluation*, 252–263.

Arora, S.; Liang, Y.; and Ma, T. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *International Conference on Learning Representations*.

Baroni, M.; Dinu, G.; and Kruszewski, G. 2014. Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings the ACL 2014*, 238–247. Association for Computational Linguistics.

Baroni, M.; Evert, S.; and Lenci, A. 2008. Verb categorization: Shared tasks from the ESSLLI 2008 workshop.

Battig, W. F., and Montague, W. E. 1969. Category norms of verbal items in 56 categories a replication and extension of the connecticut category norms. *Journal of Experimental Psychology* 80(3, Pt.2):1–46.

Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5:135–146.

Bruni, E.; Tran, N. K.; and Baroni, M. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research* 49(1):1–47.

Bullinaria, J. A., and Levy, P. J. 2012. Extracting semantic representations from word co-occurrence statistics: stop-lists, stemming, and SVD. *Behavior Research Methods* 44(3):890–907.

Caron, J. 2001. Experiments with LSA scoring: Optimal rank and basis. *Computational Information Retrieval* 157–169.

Dhillon, P. S.; Foster, D. P.; and Ungar, L. H. 2015. Eigenwords: Spectral word embeddings. *Journal of Machine Learning Research* 16:3035–3078.

Faruqui, M.; Dodge, J.; Jauhar, S. K.; Dyer, C.; Hovy, E. H.; and Smith, N. A. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of the NAACL HLT 2015*, 1606–1615.

Finkelstein, L.; Gabrilovich, E.; Matias, Y.; Rivlin, E.; Solan, Z.; Wolfman, G.; and Ruppin, E. 2002. Placing search in context: the concept revisited. *ACM Transactions on Information Systems* 20(1):116–131.

Firth, J. R. 1957. A synopsis of linguistic theory 1930–55. In *Studies in Linguistic Analysis (special volume of the Philological Society)*, volume 1952–59, 1–32. Oxford: The Philological Society.

Gerz, D.; Vulic, I.; Hill, F.; Reichart, R.; and Korhonen, A. 2016. SimVerb-3500: a large-scale evaluation set of verb similarity. In *Proceedings of the EMNLP 2016*, 2173–2182.

- He, X., and Jaeger, H. 2018. Overcoming catastrophic interference using conceptor-aided backpropagation. In *International Conference on Learning Representations*.
- Hill, F.; Reichart, R.; and Korhonen, A. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics* 41(4):665–695.
- Jaeger, H. 2014. Controlling recurrent neural networks by conceptors. Technical report, Jacobs University Bremen.
- Jaeger, H. 2017. Using conceptors to manage neural long-term memories for temporal patterns. *Journal of Machine Learning Research* 18(13):1–43.
- Khodak, M.; Saunshi, N.; Liang, Y.; Ma, T.; Stewart, B.; and Arora, S. 2018. A la carte embedding: Cheap but effective induction of semantic feature vectors. In *the Proceedings of ACL*.
- Levy, O., and Goldberg, Y. 2014. Neural word embedding as implicit matrix factorization. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N. D.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems* 27. Curran Associates, Inc. 2177–2185.
- Levy, O.; Goldberg, Y.; and Dagan, I. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics* 3:211–225.
- Luong, M.; Socher, R.; and Manning, C. D. 2013. Better word representations with recursive neural networks for morphology. In *Proceedings of the CoNLL 2013*.
- Manning, C. D.; Raghavan, P.; and Schütze, H. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- Marelli, M.; Menini, S.; Baroni, M.; Bentivogli, L.; Bernardi, R.; and Zamparelli, R. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation*. European Language Resources Association (ELRA).
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In Burges, C. J. C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems* 26. Curran Associates, Inc. 3111–3119.
- Miller, G. A. 1995. Wordnet: A lexical database for English. *Communications of the ACM* 38(11):39–41.
- Mrkšić, N., and Vulić, I. 2018. Fully statistical neural belief tracking. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 108–113. Association for Computational Linguistics.
- Mrksic, N.; Séaghdha, D.; Thomson, B.; Gasic, M.; Rojas-Barahona, L. M.; Su, P.; Vandyke, D.; Wen, T.; and Young, S. J. 2016. Counter-fitting word vectors to linguistic constraints. In *Proceedings of the NAACL HLT 2016*, 142–148.
- Mrksic, N.; Vulić, I.; Séaghdha, D. Ó.; Leviant, I.; Reichart, R.; Gasic, M.; Korhonen, A.; and Young, S. J. 2017. Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *TACL* 5:309–324.
- Mu, J., and Viswanath, P. 2018. All-but-the-top: Simple and effective postprocessing for word representations. In *International Conference on Learning Representations*.
- Myers, J. L., and Well, A. D. 1995. *Research Design & Statistical Analysis*. Routledge, 1 edition.
- Pavlick, E.; Rastogi, P.; Ganitkevitch, J.; Durme, B. V.; and Callison-Burch, C. 2015. PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification. In *Proceedings of the ACL 2015 (Volume 2: Short Papers)*, 425–430. Beijing, China: Association for Computational Linguistics.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP*, 1532–1543.
- Poesio, M., and Almuhareb, A. 2005. Identifying concept attributes using a classifier. In *Proceedings of the ACL-SIGLEX Workshop on Deep Lexical Acquisition*, 18–27. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Radinsky, K.; Agichtein, E.; Gabrilovich, E.; and Markovitch, S. 2011. A word at a time: Computing word relatedness using temporal semantic analysis. In *Proceedings of the 20th International World Wide Web Conference*, 337–346.
- Rubenstein, H., and Goodenough, J. B. 1965. Contextual correlates of synonymy. *Communications of the ACM* 8(10):627–633.
- Schnabel, T.; Labutov, I.; Mimno, D. M.; and Joachims, T. 2015. Evaluation methods for unsupervised word embeddings. In *Proceedings of EMNLP 2015*, 298–307.
- Turney, P. D. 2012. Domain and function: A dual-space model of semantic relations and compositions. *Journal of Artificial Intelligence Research* 44(1):533–585.
- Wen, T.; Vandyke, D.; Mrkšić, N.; Milica, M.; Rojas-Barahona, L. M.; Su, P.; Ultes, S.; and Young, S. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*, 438–449. Valencia, Spain: Association for Computational Linguistics.
- Wieting, J.; Bansal, M.; Gimpel, K.; and Livescu, K. 2016. Towards universal paraphrastic sentence embeddings. In *International Conference on Learning Representations*.