

# Revisiting Cosine Similarity via Normalized ICA-transformed Embeddings

Hiroaki Yamagiwa<sup>1</sup> Momose Oyama<sup>1,2</sup> Hidetoshi Shimodaira<sup>1,2</sup>

<sup>1</sup>Kyoto University <sup>2</sup>RIKEN

hiroaki.yamagiwa@sys.i.kyoto-u.ac.jp,  
oyama.momose@sys.i.kyoto-u.ac.jp, shimo@i.kyoto-u.ac.jp

## Abstract

Cosine similarity is widely used to measure the similarity between two embeddings, while interpretations based on angle and correlation coefficient are common. In this study, we focus on the interpretable axes of embeddings transformed by Independent Component Analysis (ICA), and propose a novel interpretation of cosine similarity as the sum of semantic similarities over axes. The normalized ICA-transformed embeddings exhibit sparsity, enhancing the interpretability of each axis, and the semantic similarity defined by the product of the components represents the shared meaning between the two embeddings along each axis. The effectiveness of this approach is demonstrated through intuitive numerical examples and thorough numerical experiments. By deriving the probability distributions that govern each component and the product of components, we propose a method for selecting statistically significant axes.

## 1 Introduction

Cosine similarity is widely used to measure the similarity between two embeddings (Bojanowski et al., 2017; Reimers and Gurevych, 2019; Sitikhu et al., 2019) and can be computed efficiently (Li and Han, 2013; Xia et al., 2015; Gao et al., 2021). For word embeddings, the norm represents the importance of the word and the direction represents the meaning of the word (Yokoi et al., 2020; Oyama et al., 2023). Therefore, cosine similarity, which is the inner product of the normalized embeddings, makes sense as word similarity. Studies dealing with cosine tend to focus on the angle (Deng et al., 2019; Li and Li, 2023) or interpret it as a correlation coefficient (van Dongen and Enright, 2012).

Unlike existing studies, our research introduces a novel interpretation of cosine similarity, focusing on embeddings transformed by Independent Component Analysis (ICA) (Hyvärinen and Oja,






Figure 1: Heatmaps of 300-dimensional GloVe embeddings transformed by (left) Independent Component Analysis (ICA) and (right) Principal Component Analysis (PCA), with embeddings normalized to unit length following the transformations. We select five specific axes (50th, 100th, etc.) and display the top five words by component values for each axis. For the normalized ICA-transformed embeddings, the maximum component values on the axes are substantial, highlighting significant features, while the remaining values are typically small, resulting in a sparse representation. Conversely, for the normalized PCA-transformed embeddings, even the maximum values are not large, making it difficult to interpret the meanings of the axes.

2000), which aims to maximize the independence of components. FastICA (Hyvärinen, 1999) is widely used as an implementation of ICA, where it further rotates the whitened embeddings from Principal Component Analysis (PCA) (Hotelling, 1933) to make the components of the embeddings closer to independent random variables. ICA-transformed embeddings are known to have interpretable axes (Mareček et al., 2020; Musil and Mareček, 2024; Yamagiwa et al., 2023). Specifically, Yamagiwa et al. (2023) determined the meanings of the axes by examining the top words with the highest component values in the normalized ICA-transformed embeddings. Figure 1 shows



(a) Normalized ICA-transformed GloVe embeddings of *ultraviolet* and *light* and their component-wise products.



(b) Normalized PCA-transformed GloVe embeddings of *ultraviolet* and *light* and their component-wise products.

Figure 2: For the (a) ICA and (b) PCA transformations, bar graphs are displayed for each, plotting the component values of the normalized GloVe embeddings: (left) *ultraviolet*, (middle) *light*, and (right) their component-wise products. The axes with the top five component values in the *ultraviolet* embedding are highlighted, and these same axes are consistently colored across the other two plots. For the normalized ICA-transformed embedding of *ultraviolet*, the meanings of the top five axes are *[chemistry]*, *[biology]*, *[space]*, *[spectrum]*, and *[virology]* in the order of their indices. See Table 4 in Appendix A for the top words of the axes. The component *[spectrum]* of the normalized ICA-transformed embeddings should be much more emphasized in the component-wise products than in the component values. This is because the standard deviation of the probability distribution for the component-wise products is  $1/d$ , which is smaller than the standard deviation of  $1/\sqrt{d}$  for the component values. See Appendix A for more descriptions and Appendix C for details of the distribution theory.


Axis	Meaning	Value	$p$ -value	Bonferroni	Axis	Meaning	Value	$p$ -value	Bonferroni	Axis	Meaning	Value	$p$ -value	Bonferroni
194	<i>[spectrum]</i>	0.535	$9.97 \times 10^{-21}$	$2.99 \times 10^{-18}$	194	<i>[spectrum]</i>	0.554	$4.44 \times 10^{-22}$	$1.33 \times 10^{-19}$	194	<i>[spectrum]</i>	0.296	$1.12 \times 10^{-40}$	$3.36 \times 10^{-38}$
53	<i>[chemistry]</i>	0.265	$2.17 \times 10^{-6}$	$6.51 \times 10^{-4}$	1	<i>[function words]</i>	0.379	$2.67 \times 10^{-11}$	$8.02 \times 10^{-9}$	53	<i>[chemistry]</i>	0.030	$1.38 \times 10^{-5}$	$4.15 \times 10^{-3}$
141	<i>[space]</i>	0.220	$6.97 \times 10^{-5}$	$2.09 \times 10^{-2}$	153	<i>[boxing]</i>	0.230	$3.35 \times 10^{-5}$	$1.00 \times 10^{-2}$	1	<i>[function words]</i>	0.022	$1.69 \times 10^{-4}$	$5.07 \times 10^{-2}$
68	<i>[biology]</i>	0.142	$6.88 \times 10^{-3}$	1.00	53	<i>[chemistry]</i>	0.114	$2.37 \times 10^{-2}$	1.00	153	<i>[boxing]</i>	0.014	$2.56 \times 10^{-3}$	$7.68 \times 10^{-1}$
197	<i>[virology]</i>	0.140	$7.53 \times 10^{-3}$	1.00	58	<i>[ordinal]</i>	0.114	$2.46 \times 10^{-2}$	1.00	158	<i>[police]</i>	0.007	$2.76 \times 10^{-2}$	1.00

(a) *ultraviolet*

(b) *light*

(c) *ultraviolet*  $\odot$  *light*

Table 1: (a, b) show the observed component values of normalized ICA-transformed GloVe embeddings in Fig. 2a and (c) shows their component-wise products. The  $p$ -values and their Bonferroni-corrected values are shown for the top five axes in each table. Refer to Section 4 for details on the  $p$ -value calculations. Appendix A presents the top words of these axes and Table 6 in Appendix C shows the results of normalized PCA-transformed embeddings.



(a) Components

(b) Component-wise products

(c) Component-wise products (magnified)

Figure 3: For 10,000 randomly sampled pairs of normalized ICA-transformed GloVe embeddings  $\hat{s}, \hat{s}' \in \mathbb{R}^d$ , (a) the histogram of the components  $\hat{s}^{(\ell)}$  and (b, c) the histograms of the products of the components  $\hat{s}^{(\ell)} \hat{s}'^{(\ell)}$  are displayed. The observed component values in Fig. 2a are also indicated as vertical lines. Appendix A presents additional descriptions and Fig. 12 in Appendix C shows the results for the normalized PCA-transformed GloVe embeddings. The theoretical probability density of (10) for the components and that of (11) for the component-wise products in Appendix C are almost identical to their observed histograms. The theory in Appendix C is also supported by the inverse of the observed variance,  $300.005 \approx d$  in (a) and  $89,917.992 \approx d^2$  in (b, c) for  $d = 300$ .

heatmaps of the normalized GloVe embeddings after ICA and PCA transformations. Here, as an example, the meaning of an axis is denoted as *[decimals]*, and this notation applies to other axes as well. These axes of the ICA-transformed embeddings can be interpreted as *[france]*, *[payments]*, *[guns]*, *[canada]*, and *[decimals]*, whereas those of the PCA-transformed embeddings remain uninterpretable. As demonstrated in the experiments in Section 6, ICA provides better interpretability than PCA, and normalization further enhances this interpretability. Hereafter, a word will be denoted as *paris*, as an example. For the indices of an axis and a word, we use  $\ell_{[decimals]}$  and  $i_{paris}$ , respectively, as examples. The same notation applies to other axes and words.

The inner product of normalized embeddings represents the cosine similarity. Figure 2 shows the normalized GloVe embeddings transformed by ICA and PCA for *ultraviolet* and *light*, along with their component-wise products, shown in a bar graph. The sum of the component-wise products forms the inner product, yielding an identical cosine similarity value of 0.485 for both transformations. This equivalence arises because the ICA embeddings are obtained by rotating the PCA embeddings. However, a closer examination of the component-wise products shows distinct differences: In the ICA-transformed embeddings, both *ultraviolet* and *light* exhibit a large component value in *[spectrum]*, resulting in a significant product value and sparse values elsewhere. Conversely, the PCA-transformed embeddings do not have any axes with large component values, resulting in a uniformly dense vector. These differences illustrate that, while the overall cosine similarity is the same, the underlying structural contributions to this similarity vary significantly between the two transformations.

Based on these observations, we define the semantic similarity on the  $\ell$ -th axis for words  $w_i$  and  $w_j$  as the component-wise product of the normalized ICA-transformed embeddings, denoted as  $\text{sem}_\ell(w_i, w_j)$ . For example, in Fig. 2a,

$$\text{sem}_{\ell_{[spectrum]}}(\textit{ultraviolet}, \textit{light}) = 0.296.$$

A large value of  $\text{sem}_\ell(w_i, w_j)$  indicates that both  $w_i$  and  $w_j$  have the meaning represented by  $\ell$ -th axis. The cosine similarity can be interpreted as the sum of the semantic similarities across all axes:

$$\cos(w_i, w_j) = \sum_{\ell=1}^d \text{sem}_\ell(w_i, w_j), \quad (1)$$

which represents the ‘‘additive compositionality’’ of semantic similarities, decomposing overall similarity into the component-wise similarities.

As shown in parts (a, b) of Table 1, selecting the axes with large component values for each word allows us to interpret the meaning of that word. Similarly, as shown in part (c) of Table 1, selecting the axes with large  $\text{sem}_\ell(w_i, w_j)$  values for a word pair allows us to interpret the meaning shared by both words.

The number of selected axes can be determined by statistical methods. As shown in Fig. 3, we theoretically derived the probability distributions for both the component values and their products. Therefore, as will be discussed in Section 4, we can calculate the  $p$ -value to determine whether the observed values are significantly greater than zero. We select the axes where the Bonferroni-corrected  $p$ -value, accounting for the multiplicity of hypothesis testing, is smaller than the significance level  $\alpha$ . For example, with  $\alpha = 0.05$ , *[spectrum]*, *[chemistry]*, and *[space]* are selected for *ultraviolet*, while *[spectrum]* and *[chemistry]* are selected for the pair *ultraviolet* and *light*.

## 2 Background: Independent Components in Embeddings

In this section, we explain PCA and ICA transformations for embeddings based on Hyvärinen et al. (2001); Yamagiwa et al. (2023). Let  $\mathbf{X} \in \mathbb{R}^{n \times d}$  be pre-trained embeddings with vocabulary size  $n$  and dimension  $d$ . We assume that  $\mathbf{X}$  is centered, i.e., the mean of components in each column is zero.

### 2.1 PCA-transformed embeddings

PCA, typically implemented using algorithms such as SVD, transforms the embeddings so that their components align with the directions of maximum variance. The PCA-transformed embeddings  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  of  $\mathbf{X}$  are given by the transformation matrix  $\mathbf{A} \in \mathbb{R}^{d \times d}$  as

$$\mathbf{Z} = \mathbf{X}\mathbf{A}. \quad (2)$$

The columns of  $\mathbf{Z}$  are called principal components (PC), and the matrix  $\mathbf{Z}$  is whitened; i.e., the variance of components in each column is 1 and the columns are uncorrelated with each other. Whitening generally improves the quality of the embeddings (Su et al., 2021; Sasaki et al., 2023).

## 2.2 ICA-transformed embeddings

ICA transforms the embeddings so that their components are as statistically independent as possible. Statistical independence, a stronger property than uncorrelatedness or whitening, ensures that random variables are not only uncorrelated but also independent in a probabilistic sense (Hyvärinen et al., 2001). The ICA-transformed embeddings  $\mathbf{S} \in \mathbb{R}^{d \times d}$  of  $\mathbf{X}$  are given by the transformation matrix  $\mathbf{B} \in \mathbb{R}^{d \times d}$  as

$$\mathbf{S} = \mathbf{X}\mathbf{B}. \quad (3)$$

The columns of  $\mathbf{S}$  are called independent components (IC).

In particular, FastICA (Hyvärinen, 1999) uses  $\mathbf{Z}$  in (2) to compute  $\mathbf{S}$  as follows:

$$\mathbf{S} = \mathbf{Z}\mathbf{R}_{\text{ica}}, \quad (4)$$

where  $\mathbf{R}_{\text{ica}} \in \mathbb{R}^{d \times d}$  is an orthogonal matrix that maximizes the statistical independence of the columns of  $\mathbf{S}$ . Similar to  $\mathbf{Z}$ , the matrix  $\mathbf{S}$  is also whitened.

## 2.3 Normalized ICA-transformed embeddings

The ICA-transformed embedding of a word  $w_i$ , denoted by  $\mathbf{s}_i \in \mathbb{R}^d$ , is normalized to  $\hat{\mathbf{s}}_i \in \mathbb{R}^d$ :

$$\hat{\mathbf{s}}_i := \mathbf{s}_i / \|\mathbf{s}_i\| = (\hat{s}_i^{(1)}, \dots, \hat{s}_i^{(\ell)}, \dots, \hat{s}_i^{(d)}). \quad (5)$$

The  $\ell$ -th component  $\hat{s}_i^{(\ell)}$  of the normalized ICA-transformed embedding  $\hat{\mathbf{s}}_i$  for the word  $w_i$  can be interpreted as the semantic component of  $w_i$  along the  $\ell$ -th axis. For example, as shown in Fig. 2a, the normalized ICA-transformed embedding of *ultraviolet* has large semantic components of [chemistry], [biology], [space], [spectrum], and [virology].

## 3 Decomposition and Interpretation of Cosine Similarity

We define the semantic similarity mentioned in Section 1 using the notation introduced in Section 2 and explain how ICA improves interpretability compared to PCA.

### 3.1 Semantic similarity on axes

Cosine similarity is widely used to measure the similarity between words. The cosine similarity between words  $w_i$  and  $w_j$  can be expressed as the

inner product of their normalized ICA-transformed embeddings  $\hat{\mathbf{s}}_i$  and  $\hat{\mathbf{s}}_j$ :

$$\cos(w_i, w_j) = \hat{\mathbf{s}}_i^\top \hat{\mathbf{s}}_j = \sum_{\ell=1}^d \hat{s}_i^{(\ell)} \hat{s}_j^{(\ell)}. \quad (6)$$

As seen in Section 2,  $\hat{s}_i^{(\ell)}$  can be interpreted as the semantic component of a word  $w_i$  on the  $\ell$ -th axis. Therefore, the semantic similarity for words  $w_i$  and  $w_j$  on the  $\ell$ -th axis,  $\text{sem}_\ell(w_i, w_j)$ , is defined as:

$$\text{sem}_\ell(w_i, w_j) := \hat{s}_i^{(\ell)} \hat{s}_j^{(\ell)}. \quad (7)$$

Using the element-wise product (i.e., Hadamard product),  $\text{sem}_\ell(w_i, w_j)$ ,  $\ell = 1, \dots, d$ , can also be interpreted as the  $\ell$ -th component of the vector<sup>1</sup>

$$\hat{\mathbf{s}}_i \odot \hat{\mathbf{s}}_j = (\hat{s}_i^{(1)} \hat{s}_j^{(1)}, \dots, \hat{s}_i^{(\ell)} \hat{s}_j^{(\ell)}, \dots, \hat{s}_i^{(d)} \hat{s}_j^{(d)}).$$

The expression for the cosine similarity in (6) can be rewritten as in (1) with the definition (7). Thus, the cosine similarity can be interpreted as the sum of the semantic similarities over all axes.

Figure 4 shows the cosine similarity computations for *ultraviolet* with *salts*, *proteins*, *spacecraft*, *light*, and *virus*, and semantic similarities on the axes of [chemistry], [biology], [space], [spectrum], and [virology]. The cosine similarity values, which is the sum of the all semantic similarities, can be interpreted from these semantic similarity values.

### 3.2 ICA improves interpretability

Note that the cosine similarity between two embeddings is the same for the ICA-transformed embeddings and the PCA-transformed embeddings. As seen in (4), since  $\mathbf{S}$  is  $\mathbf{Z}$  multiplied by the orthogonal matrix  $\mathbf{R}_{\text{ica}}$ ,  $\hat{\mathbf{z}}_i = \mathbf{R}_{\text{ica}} \hat{\mathbf{s}}_i$ , where  $\hat{\mathbf{z}}_i$  is the normalized PCA-transformed embedding of  $w_i$ . Then  $\sum_{\ell=1}^d \hat{z}_i^{(\ell)} \hat{z}_j^{(\ell)} = \hat{\mathbf{z}}_i^\top \hat{\mathbf{z}}_j = \hat{\mathbf{s}}_i^\top \hat{\mathbf{s}}_j = \cos(w_i, w_j)$ , meaning that  $\hat{z}_i^{(\ell)} \hat{z}_j^{(\ell)}$  can also be interpreted as the semantic similarity for the normalized PCA-transformed embeddings.

However, the PCA-based semantic similarity lacks interpretability. Figure 2 shows bar graphs for the normalized ICA-transformed and PCA-transformed GloVe embeddings of *ultraviolet*, *light*, and their component-wise products. The sum of the component-wise products is equal to the cosine similarity value (0.485) for both transformations. In the normalized ICA-transformed embeddings, the semantic components

<sup>1</sup>For words  $w_i$  and  $w_j$ , we abbreviate  $\hat{\mathbf{s}}_i \odot \hat{\mathbf{s}}_j$  as  $w_i \odot w_j$ .

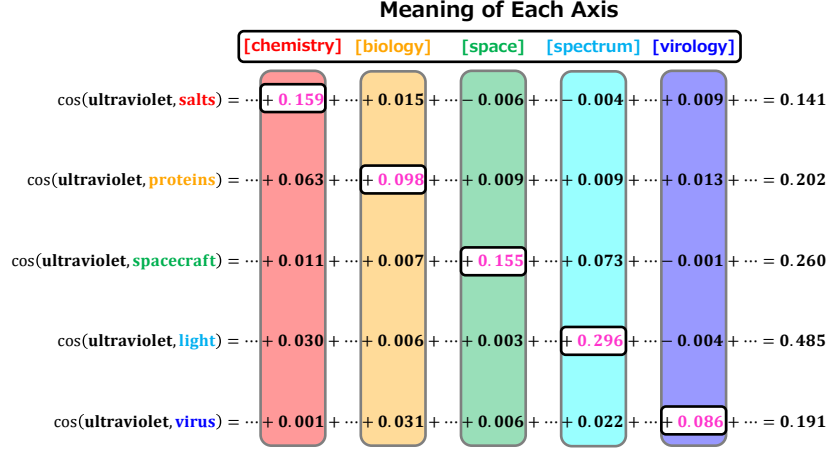


Figure 4: Cosine similarity interpretation. For the normalized ICA-transformed GloVe embedding of *ultraviolet*, the meanings of the axes of the top five components are [chemistry], [biology], [space], [spectrum], and [virology]. For the top words on these axes, see Table 4 in Appendix A. The cosine similarities are computed between *ultraviolet* and the top word of each axis: *salts*, *proteins*, *spacecraft*, *light*, and *virus*, respectively. The inner products (i.e., cosine similarities) of their normalized ICA-transformed GloVe embeddings are computed, and the semantic similarities of these five axes are displayed. For example, the semantic similarity on the [space] axis between *ultraviolet* and *spacecraft* is 0.155, which is more than half of the cosine similarity of 0.260.

of [spectrum] are large (0.535 for *ultraviolet* and 0.554 for *light*), and the semantic similarity is  $\text{sem}_{\ell_{[\text{spectrum}]}}(\text{ultraviolet}, \text{light}) = 0.296$ . Other semantic similarities are close to zero. However, in the normalized PCA-transformed embeddings, although no axis has a component-wise product as large as in the normalized ICA-transformed embeddings, the sum of the component-wise products is still equal to the cosine similarity. Therefore, the component-wise products are not close to zero compared to ICA, resulting in a dense vector.

#### 4 Statistical Analysis of Axis Selection

For the normalized ICA-transformed embeddings of words  $w_i$  and  $w_j$ , the criteria for selecting axes  $\ell$  where the component values  $\hat{s}_i^{(\ell)}$  or the component-wise product values  $\hat{s}_i^{(\ell)}\hat{s}_j^{(\ell)}$  become large can be determined based on the theory of statistical hypothesis testing. Using the probability distribution provided in Appendix C, the observed component values can be converted into  $p$ -values, allowing for a probabilistic interpretation. This is illustrated with numerical examples in Fig. 3 and Table 1.

##### 4.1 Component values

As given by equation (10) in Appendix C.2, the probability distribution that  $\hat{s}_i^{(\ell)}$  follows is a normal distribution with mean 0 and variance  $1/d$ ,

$$\hat{s}_i^{(\ell)} \sim \mathcal{N}(0, 1/d).$$

Looking at Fig. 3a, we can see that the histogram of the observed  $\hat{s}_i^{(\ell)}$  fits well with the theoretical curve of the normal distribution. The  $p$ -value is given by the upper tail probability of this probability distribution, and if it is smaller than the pre-determined significance level  $\alpha$  (for example, 0.05), that axis is selected. The explicit formula for the  $p$ -value is expressed using the cumulative distribution function<sup>2</sup>  $\Phi(x)$  of  $\mathcal{N}(0, 1)$  as  $p_1 = \Phi(-\sqrt{d}\hat{s}_i^{(\ell)})$ .

The  $p$ -value described above is correct for a single predetermined axis, but in practice, since we are selecting from  $d$  axes, it is necessary to account for the multiplicity of tests to avoid false positives. To address this, the Bonferroni correction can be applied by using  $p_d = \min\{dp_1, 1\}$ . Thus, as a conservative and safe approach, we select axes where  $p_d < \alpha$ .

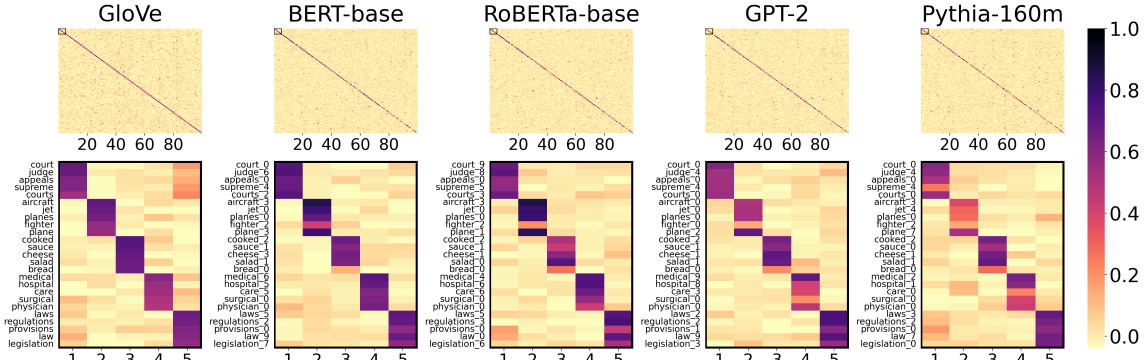
##### 4.2 Product of two component values

As given by equation (11) in Appendix C.3, the probability density function of  $\hat{s}_i^{(\ell)}\hat{s}_j^{(\ell)}$  is


$$\hat{s}_i^{(\ell)}\hat{s}_j^{(\ell)} \sim (d/\pi)K_0(d|\hat{s}_i^{(\ell)}\hat{s}_j^{(\ell)}|),$$

where  $K_0(\cdot)$  is the modified Bessel function of the second kind of order zero. This probability distribution is particularly intriguing from the perspective of mathematical statistics, and since it is a relatively novel result, the probability distribution has not yet

<sup>2</sup>`pnorm` in R or `norm.cdf` in `scipy.stats`.



(a) Normalized ICA-transformed embeddings for GloVe and contextualized embedding models.



(b) Normalized PCA-transformed embeddings for GloVe and contextualized embedding models.

Figure 5: Heatmaps of normalized (a) ICA-transformed and (b) PCA-transformed embeddings for GloVe and four contextualized embedding models, with the axes aligned by the procedure of Yamagiwa et al. (2023). (a) In the top panel, the components of the first 100 axes for 500 embeddings are shown, while the bottom panel magnifies the components of the first 5 axes. For each GloVe axis, we selected five top words for which corresponding contextualized embeddings exist, and displayed the embeddings along with the contextualized embeddings corresponding to each top word. The embeddings selected for each axis have large component values on that axis, confirming the existence of common semantic axes across the five models. (b) In contrast to the ICA-transformed embeddings, the PCA-transformed embeddings do not reveal such common semantic axes across the five models.

been given a specific name. From Figs. 3b and 3c, we can see that the histogram of the observed  $\hat{s}_i^{(\ell)} \hat{s}_j^{(\ell)}$  closely fits the theoretical curve. The upper tail probability of this distribution, denoted as  $p_1$ , was numerically calculated by integrating the probability density function<sup>3</sup>. As a conservative approach, similar to Section 4.1, we apply the Bonferroni correction using  $p_d = \min\{dp_1, 1\}$ .

## 5 Consistency of ICA Transformation for Contextualized Embeddings

In this section, we show the consistency of the interpretability of ICA components and their products across different contextualized embeddings. See Appendix D for the details of the experiments.

<sup>3</sup>Using Mathematica, we first derived a formula involving the modified Bessel function and the modified Struve function, and then numerically evaluated it for computing  $p_1$ .

**Consistency of component values.** For contextualized embeddings, we used BERT-base (Devlin et al., 2019), RoBERTa-base (Radford et al., 2019), GPT-2 (Radford et al., 2019), and Pythia-160m (Biderman et al., 2023), and computed 50,000 embeddings for each model.

Figure 5a shows the result of axis matching using correlation coefficients for these ICA-transformed embeddings. Despite applying ICA transformations to the embeddings of different models individually, we observed over 100 axes representing common semantic content across the five models. In contrast, Fig. 5b shows that PCA-transformed embeddings do not exhibit such common semantic axes among the models. To improve clarity, prefixes such as ## and  $\tilde{G}$  were removed in Fig. 5.

**Consistency of the product of two components.** Figure 6 shows the component values of the normalized ICA-transformed embeddings for GPT-2




Figure 6: Similar to Fig. 2a, the component values and their component-wise products for the normalized ICA-transformed contextualized embeddings of *ultraviolet\_3* and *light\_1* are shown in bar graphs. The cosine similarity is 0.576 for GPT-2 and 0.699 for Pythia-160m. In Appendix D, Fig. 13 shows results for BERT and RoBERTa, and Fig. 14 for normalized PCA-transformed contextualized embeddings. Table 9 in Appendix D shows the top words of the axes corresponding to the top 5 components of these normalized embeddings.

Axis	Meaning	Value	$p$ -value	Bonferroni
349	[light]	0.413	$1.28 \times 10^{-30}$	$9.86 \times 10^{-28}$
273	[elements]	0.275	$1.26 \times 10^{-14}$	$9.70 \times 10^{-12}$
28	[biology]	0.156	$7.62 \times 10^{-6}$	$5.85 \times 10^{-3}$
619	[occupations]	0.124	$3.02 \times 10^{-4}$	$2.32 \times 10^{-1}$
512	[technology]	0.117	$5.85 \times 10^{-4}$	$4.49 \times 10^{-1}$

(a) *ultraviolet\_3*

Axis	Meaning	Value	$p$ -value	Bonferroni
349	[light]	0.450	$5.88 \times 10^{-36}$	$4.51 \times 10^{-33}$
273	[elements]	0.322	$2.31 \times 10^{-19}$	$1.78 \times 10^{-16}$
619	[occupations]	0.145	$2.97 \times 10^{-5}$	$2.28 \times 10^{-2}$
402	[rainstorm]	0.128	$1.89 \times 10^{-4}$	$1.45 \times 10^{-1}$
284	[challenges]	0.122	$3.58 \times 10^{-4}$	$2.75 \times 10^{-1}$

(b) *light\_1*

Axis	Meaning	Value	$p$ -value	Bonferroni
349	[light]	0.186	$3.84 \times 10^{-64}$	$2.95 \times 10^{-61}$
273	[elements]	0.089	$1.43 \times 10^{-31}$	$1.10 \times 10^{-28}$
28	[biology]	0.019	$4.60 \times 10^{-8}$	$3.53 \times 10^{-5}$
619	[occupations]	0.018	$1.08 \times 10^{-7}$	$8.27 \times 10^{-5}$
402	[rainstorm]	0.013	$6.89 \times 10^{-6}$	$5.29 \times 10^{-3}$

(c) *ultraviolet\_3*  $\odot$  *light\_1*

Table 2: (a, b) show the observed component values of normalized ICA-transformed GPT-2 embeddings in Fig. 6 and (c) shows their component-wise products with the  $p$ -values and Bonferroni-corrected values in each table.

Axis	Meaning	Value	$p$ -value	Bonferroni
262	[light]	0.601	$1.32 \times 10^{-62}$	$1.02 \times 10^{-59}$
619	[business]	0.149	$1.78 \times 10^{-5}$	$1.36 \times 10^{-2}$
660	[exceptional]	0.116	$6.80 \times 10^{-4}$	$5.22 \times 10^{-1}$
485	[function words]	0.087	$7.70 \times 10^{-3}$	1.00
371	[subwords]	0.086	$8.35 \times 10^{-3}$	1.00

(a) *ultraviolet\_3*

Axis	Meaning	Value	$p$ -value	Bonferroni
262	[light]	0.645	$7.92 \times 10^{-72}$	$6.09 \times 10^{-69}$
619	[business]	0.151	$1.51 \times 10^{-5}$	$1.16 \times 10^{-2}$
264	[elements]	0.122	$3.54 \times 10^{-4}$	$2.72 \times 10^{-1}$
485	[function words]	0.116	$6.40 \times 10^{-4}$	$4.92 \times 10^{-1}$
548	[subwords]	0.114	$7.82 \times 10^{-4}$	$6.00 \times 10^{-1}$

(b) *light\_1*

Axis	Meaning	Value	$p$ -value	Bonferroni
262	[light]	0.388	$9.66 \times 10^{-132}$	$7.42 \times 10^{-129}$
619	[business]	0.022	$2.99 \times 10^{-9}$	$2.29 \times 10^{-6}$
485	[function words]	0.010	$5.44 \times 10^{-5}$	$4.18 \times 10^{-2}$
264	[elements]	0.008	$2.68 \times 10^{-4}$	$2.06 \times 10^{-1}$
336	[creation]	0.006	$1.23 \times 10^{-3}$	$9.42 \times 10^{-1}$

(c) *ultraviolet\_3*  $\odot$  *light\_1*

Table 3: (a, b) show the observed component values of normalized ICA-transformed Pythia-160m embeddings in Fig. 6 and (c) shows their component-wise products with the  $p$ -values and Bonferroni-corrected values in each table.

and Pythia-160m, as well as the product of those components, for *ultraviolet\_3* and *light\_1* in the same sentence. Similar to Fig. 2a, the component values of the embeddings are sparse overall, but when looking at the component-wise products, the common semantic components between the two embeddings (e.g., [light]) stand out more clearly. Additionally, Tables 2 and 3 present the top 5 values for bar graphs in Fig. 6, along with the  $p$ -values and Bonferroni-corrected values. As in Table 1, our method identifies statistically significant axes.

## 6 Quantitative Experiments

We quantitatively evaluate that ICA provides better interpretability than PCA, and that normalization further enhances interpretability (Sec. 6.2). We confirm that ICA tends to yield larger component values compared to PCA, resulting in more dimen-

sions that can be interpreted with specific meanings (Sec. 6.3). Then, we quantitatively verify the sparsity of the semantic similarities on axes in the normalized ICA-transformed embeddings (Sec. 6.4).

### 6.1 General settings

**Word embedding.** We use GloVe (Pennington et al., 2014) embeddings with  $d = 300$  dimensions and vocabulary size  $n = 400,000$ .

**ICA-transformation.** To compute the ICA transformation, we utilized the FastICA implementation available in scikit-learn (Pedregosa et al., 2011). The hyperparameters for FastICA were configured with a maximum number of iterations set to 10,000 and a convergence tolerance of  $1 \times 10^{-10}$ . As post-processing, we flip the signs of the axes in  $\mathbf{S}$  if necessary to ensure positive skewness, and sort the axes in descending order of their skewness.




Figure 7: Result of word intrusion task. A large value indicates high consistency in the semantic components represented by each axis of the embeddings. We used the GloVe embeddings transformed by ICA and PCA, and their respective normalized versions. For each dimension, we select the top  $k = 5$  words with the largest component values and compute the average score of 10 different intruder choices. The final result shown is the median of these scores across  $d = 300$  dimensions.

## 6.2 Normalization enhances interpretability


We conducted a word intrusion task (Chang et al., 2009; Sun et al., 2016) to quantitatively evaluate that ICA provides better interpretability than PCA and that normalization further enhances this interpretability. The word intrusion task assesses the semantic coherence of a set of the top  $k$  words by measuring the ability to identify an intruder.

**Settings.** For the four embedding types, we compute a semantic coherence score for each dimension  $\ell = 1, \dots, d$ . The interpretability score for each embedding type is the median of these  $d$  scores. Details of the word intrusion task and scoring method are provided in Appendix E.

**Results and discussion.** Figure 7 shows that ICA consistently outperforms PCA in interpretability scores, with normalized embeddings achieving higher scores for both methods. This indicates that ICA components offer better interpretability than PCA, and that normalization further enhances the interpretability. As an intuitive explanation, normalizing embeddings emphasizes each word’s association with specific ICA-defined semantic axes. This allows words with focused meanings to stand out along relevant dimensions, while words with mixed semantics exhibit less pronounced values across axes. This improvement in interpretability may be related to the experimental results showing that the fidelity of rankings along axes and embeddings increases with normalization (Appendix B).

## 6.3 ICA: Larger component values

By sorting the components, we investigate whether normalized ICA-transformed or PCA-transformed



(a) Sorted along embeddings (b) Sorted along axes

Figure 8: Comparison of component values of the normalized GloVe embeddings transformed by ICA and PCA. The component values are averaged after being sorted in descending order (a) along embeddings for each axis, and (b) along axes for each embedding. The range of  $\pm 1\sigma$  is shown, where  $\sigma$  is the standard deviation of the component values. See Fig. 16 in Appendix D.5 for the results of contextualized embeddings.




Figure 9: Performance comparison between ICA and PCA for the GloVe embeddings when reducing non-zero normalized component-wise products and computing cosine similarity. Each value represents the average of 8 word similarity tasks or 30 analogy tasks.

embeddings give greater emphasis on larger component values to improve interpretability.

**Settings.** To compare the ICA and PCA transformations, we employed two methods: (a) We sorted the  $n$  component values **along embeddings** for each axis in descending order and averaged them over the  $d$  axes. (b) We sorted the  $d$  component values **along axes** for each embedding in descending order and averaged them over the  $n$  embeddings.

**Results and discussion.** The sorted component values are shown in Fig. 8. In Fig. 8a, ICA component values are larger than those of PCA up to approximately the 10,000th embedding, and in Fig. 8b, ICA component values are larger for the first few axes. These results demonstrate that ICA tends to emphasize larger component values compared to PCA.

## 6.4 ICA: Sparsity of semantic similarities

We investigate whether the sparsity of component-wise products observed in Fig. 2 generalizes to



other word pairs. To this end, we perform word similarity and analogy tasks, comparing ICA-transformed and PCA-transformed embeddings.

**Settings.** We evaluate the performance degradation as component-wise products are replaced by zero in ascending order until only  $p$  non-zero products remain<sup>4</sup>. For this analysis, we use the Word Embedding Benchmark (Jastrzebski et al., 2017). See Appendix F for more details.

**Results and discussion.** In Fig. 9, for both word similarity and analogy tasks, the ICA-transformed embeddings consistently outperform the PCA-transformed embeddings, even when the number of non-zero products is small. This indicates that ICA can represent cosine similarity with fewer dimensions than PCA.

**Applications.** Based on these results, we present potential applications of ICA-transformed embeddings. In Appendix G, we consider sparse embeddings that contain only the semantic components of [food], [animals], or [plants] and show that word retrieval works effectively. Additionally, in Appendix H, we show that ablating the [female] semantic component from the ICA-transformed embedding of *woman* results in greater similarity to embeddings such as *man* and *boy*. Furthermore, Appendix I shows case studies on the noun and verb senses of *shore*, as well as multilingual analyses.

## 7 Related Work

**Cosine similarity.** Cosine similarity is widely used to measure the similarity between two embeddings, such as word (Mikolov et al., 2013b,c), token (Zhang et al., 2020; Bommasani et al., 2020), and sentence embeddings (Reimers and Gurevych, 2019; Gao et al., 2021). Cross-lingual alignment methods based on cosine similarity have also been proposed (Xing et al., 2015; Alvarez-Melis and Jaakkola, 2018; Lampe et al., 2018).

There are studies that question the effectiveness of cosine similarity. For example, Schnabel et al. (2015) showed that word frequency can affect cosine similarity. Steck et al. (2024) used linear models to show cases where cosine similarity fails.

**Interpretability of embeddings.** Research on the interpretability of embeddings has used various

<sup>4</sup>In Section 6.2 of Yamagiwa et al. (2023), small component values are zeroed out before computing cosine similarity. In contrast, we first compute element-wise products between normalized embeddings and then zero out small components.

methods, including non-negative matrix factorization (Murphy et al., 2012), sparse coding (Faruqui et al., 2015), methods for learning interpretable embeddings (Luo et al., 2015; Sun et al., 2016), rotation of embeddings (Park et al., 2017), Singular Value Decomposition (SVD) (Shin et al., 2018), autoencoders (Subramanian et al., 2018; Huben et al., 2024), and PCA (Musil, 2019).

Some research has applied ICA to embeddings. For example, Chagnaa et al. (2007) showed that the features of similar verbs have a large same independent component through bar graphs. The axes of the ICA-transformed embeddings are known to be interpretable (Mareček et al., 2020; Musil and Mareček, 2024), and such axes are also observed in embeddings for other languages, dynamic models, and images (Yamagiwa et al., 2023). Li et al. (2024) statistically validated their existence. The dependencies between axes are explored via semantic continuity optimization (Yamagiwa et al., 2024) and higher-order correlations (Oyama et al., 2024).

## 8 Conclusion

Inspired by previous studies showing that ICA transformation can produce sparse and interpretable embeddings from dense distributed representations, we made the following contributions:

(1) We experimentally demonstrated that normalized ICA-transformed embeddings exhibit sparsity, enhancing the interpretability of each axis. In particular, ICA provides better interpretability than PCA, and normalization further improves this interpretability. These findings were tested using static embeddings (GloVe) and four contextualized embeddings (BERT, RoBERTa, GPT-2, and Pythia).

(2) We proposed interpreting cosine similarity as the sum of semantic similarities across axes, where the component-wise product of two embeddings represents the semantic similarity on each axis. The experiments showed that ICA results in greater sparsity in the component-wise products compared to PCA, allowing cosine similarity to be represented with fewer dimensions.

(3) By deriving the probability distribution that governs each component in the normalized ICA-transformed embeddings, we proposed a method to statistically select significant axes from an embedding. Similarly, by newly deriving the probability distribution that governs the component-wise products, we proposed a method to select statistically significant axes shared between two embeddings.

## Limitations

- This study explains the interpretation of cosine similarity using centered and whitened embeddings. These embeddings are different from the original embeddings.
- The meanings of the axes of the ICA-transformed embeddings are manually interpreted based on their top words after normalization. In addition, it is not always possible to interpret the meaning of an axis from its top words. Note that the top words used to interpret the meanings of the axes in this study can all be found in the Appendix sections.
- We need to pay attention to the signs of the axes of the ICA-transformed embeddings. In this study, following Yamagiwa et al. (2023), we ensure that all axes have positive skewness by flipping their signs when necessary. We then assume that larger component values are more representative of the meanings of the axes.
- To explain our interpretation of cosine similarity, we mainly use the GloVe<sup>5</sup> embeddings for which cosine similarity works well. Although cosine similarity may not be effective for all embeddings (Steck et al., 2024), this study does not cover which specific types of embeddings are suitable for cosine similarity.
- The Bonferroni correction, used to adjust for multiple testing when selecting statistically significant axes from the  $d$  axes, is safe but conservative, and tends to result in fewer axes being selected. Other approaches, such as the False Discovery Rate (FDR), should also be considered.
- The axes selected from the component-wise product of the two embeddings do not necessarily have logical consistency with the axes selected from the components of each embedding. For example, in the numerical example in Table 1, where axes are selected with  $\alpha = 0.05$ , *[spectrum]*, *[function words]*, and *[boxing]* are selected for *light*, and *[chemistry]* is not selected. However, as shown in Section 1, *[chemistry]* is selected for the pair *ultraviolet* and *light*.

<sup>5</sup><https://nlp.stanford.edu/data/glove.6B.zip>.

- ICA is a more complex method than PCA; the computation of the orthogonal matrix  $\mathbf{R}_{\text{ica}}$  in (4) for ICA takes considerably more time than that of the matrix  $\mathbf{A}$  in (2) for PCA. The computation time for FastICA using `scikit-learn` depends mainly on the embedding dimension  $d$ , vocabulary size  $n$ , the maximum number of iterations, and the convergence tolerance. For example, it takes several hours for GloVe embeddings with the settings in Section 6.1.

## Ethics Statement

This study complies with the [ACL Ethics Policy](#).

## Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments and suggestions. This study was partially supported by JSPS KAKENHI 22H05106, 23H03355, JST CREST JPMJCR21N3, JST SPRING JPMJSP2110, JST BOOST JPMJBS2407.

## Code availability

Code is available at <https://github.com/ymgw55/Cosine-Similarity-via-ICA>.

## References

- David Alvarez-Melis and Tommi S. Jaakkola. 2018. [Gromov-wasserstein alignment of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1881–1890. Association for Computational Linguistics.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 2397–2430. PMLR.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomáš Mikolov. 2017. [Enriching word vectors with subword information](#). *Trans. Assoc. Comput. Linguistics*, 5:135–146.
- Rishi Bommasani, Kelly Davis, and Claire Cardie. 2020. [Interpreting pretrained contextualized representations via reductions to static embeddings](#). In *Proceedings*

- of the 58th Annual Meeting of the Association for Computational Linguistics, *ACL 2020, Online, July 5-10, 2020*, pages 4758–4781. Association for Computational Linguistics.
- Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49:1–47.
- Altangerel Chagnaa, Cheol-Young Ock, Chang-Beom Lee, and Purev Jaimai. 2007. Feature extraction of concepts by independent component analysis. *Journal of Information Processing Systems*, 3(1):33–37.
- Jonathan Chang, Sean Gerrish, Chong Wang, Jordan Boyd-graber, and David Blei. 2009. **Reading tea leaves: How humans interpret topic models**. In *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc.
- Ciprian Chelba, Tomás Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and Tony Robinson. 2014. **One billion word benchmark for measuring progress in statistical language modeling**. In *INTER-SPEECH 2014, 15th Annual Conference of the International Speech Communication Association, Singapore, September 14-18, 2014*, pages 2635–2639. ISCA.
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. **Arcface: Additive angular margin loss for deep face recognition**. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4690–4699. Computer Vision Foundation / IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah A. Smith. 2015. **Sparse overcomplete word vector representations**. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1491–1500. The Association for Computer Linguistics.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2002. Placing search in context: The concept revisited. *ACM Transactions on information systems*, 20(1):116–131.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. **Simcse: Simple contrastive learning of sentence embeddings**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6894–6910. Association for Computational Linguistics.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- Harold Hotelling. 1933. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417.
- Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, and Lee Sharkey. 2024. **Sparse autoencoders find highly interpretable features in language models**. In *The Twelfth International Conference on Learning Representations*.
- Aapo Hyvärinen. 1999. **Fast and robust fixed-point algorithms for independent component analysis**. *IEEE Trans. Neural Networks*, 10(3):626–634.
- Aapo Hyvärinen, Juha Karhunen, and Erkki Oja. 2001. *Independent Component Analysis*. Wiley.
- Aapo Hyvärinen and Erkki Oja. 2000. **Independent component analysis: algorithms and applications**. *Neural Networks*, 13(4-5):411–430.
- Yoichi Ishibashi, Katsuhito Sudoh, Koichiro Yoshino, and Satoshi Nakamura. 2020. **Reflection-based word attribute transfer**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, ACL 2020, Online, July 5-10, 2020*, pages 51–58. Association for Computational Linguistics.
- Stanislaw Jastrzebski, Damian Lesniak, and Wojciech Marian Czarnecki. 2017. **How to evaluate word embeddings? on importance of data efficiency and simple supervised tasks**. *CoRR*, abs/1702.02170.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. **Word translation without parallel data**. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Baoli Li and Liping Han. 2013. **Distance weighted cosine similarity measure for text classification**. In *Intelligent Data Engineering and Automated Learning - IDEAL 2013 - 14th International Conference, IDEAL 2013, Hefei, China, October 20-23, 2013. Proceedings*, volume 8206 of *Lecture Notes in Computer Science*, pages 611–618. Springer.
- Rongzhi Li, Takeru Matsuda, and Hitomi Yanaka. 2024. **Exploring intra and inter-language consistency in embeddings with ICA**. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language*

- Processing, *EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 19104–19111. Association for Computational Linguistics.
- Xianming Li and Jing Li. 2023. [Angle-optimized text embeddings](#). *CoRR*, abs/2309.12871.
- Hongyin Luo, Zhiyuan Liu, Huan-Bo Luan, and Maosong Sun. 2015. [Online learning of interpretable word embeddings](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1687–1692. The Association for Computational Linguistics.
- Thang Luong, Richard Socher, and Christopher Manning. 2013. Better word representations with recursive neural networks for morphology. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 104–113.
- David Mareček, Jindřich Libovický, Tomáš Musil, Rudolf Rosa, and Tomasz Limisiewicz. 2020. *Hidden in the Layers: Interpretation of Neural Networks for Natural Language Processing*, volume 20 of *Studies in Computational and Theoretical Linguistics*. Institute of Formal and Applied Linguistics, Prague, Czechia.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Tomás Mikolov, Quoc V. Le, and Ilya Sutskever. 2013b. [Exploiting similarities among languages for machine translation](#). *CoRR*, abs/1309.4168.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013c. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Tomás Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013d. [Linguistic regularities in continuous space word representations](#). In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*, pages 746–751. The Association for Computational Linguistics.
- Brian Murphy, Partha Pratim Talukdar, and Tom M. Mitchell. 2012. [Learning effective and interpretable semantic models using non-negative sparse embedding](#). In *COLING 2012, 24th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, 8-15 December 2012, Mumbai, India*, pages 1933–1950. Indian Institute of Technology Bombay.
- Tomás Musil. 2019. [Examining structure of word embeddings with PCA](#). In *Text, Speech, and Dialogue - 22nd International Conference, TSD 2019, Ljubljana, Slovenia, September 11-13, 2019, Proceedings*, volume 11697 of *Lecture Notes in Computer Science*, pages 211–223. Springer.
- Tomás Musil and David Mareček. 2024. [Exploring interpretability of independent components of word embeddings with automated word intruder test](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pages 6922–6928. ELRA and ICCL.
- Saralees Nadarajah and Tibor K. Pogány. 2016. [On the distribution of the product of correlated normal random variables](#). *Comptes Rendus Mathématique*, 354(2):201–204.
- Momose Oyama, Hiroaki Yamagiwa, and Hidetoshi Shimodaira. 2024. [Understanding higher-order correlations among semantic components in embeddings](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 2883–2899. Association for Computational Linguistics.
- Momose Oyama, Sho Yokoi, and Hidetoshi Shimodaira. 2023. [Norm of word embedding encodes information gain](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 2108–2130. Association for Computational Linguistics.
- Sungjoon Park, JinYeong Bak, and Alice Oh. 2017. [Rotated word vector representations and their interpretability](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 401–411. Association for Computational Linguistics.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake VanderPlas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. 2011. [Scikit-learn: Machine learning in python](#). *J. Mach. Learn. Res.*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

- Kira Radinsky, Eugene Agichtein, Evgeniy Gabrilovich, and Shaul Markovitch. 2011. A word at a time: Computing word relatedness using temporal semantic analysis. In *Proceedings of the 20th International Conference on World Wide Web*, page 337–346.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics.
- Herbert Rubenstein and John B. Goodenough. 1965. [Contextual correlates of synonymy](#). *Commun. ACM*, 8(10):627–633.
- Shota Sasaki, Benjamin Heinzerling, Jun Suzuki, and Kentaro Inui. 2023. [Examining the effect of whitening on static and contextualized word embeddings](#). *Inf. Process. Manag.*, 60(3):103272.
- Tobias Schnabel, Igor Labutov, David M. Mimno, and Thorsten Joachims. 2015. [Evaluation methods for unsupervised word embeddings](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 298–307. The Association for Computational Linguistics.
- Jamin Shin, Andrea Madotto, and Pascale Fung. 2018. Interpreting word embeddings with eigenvector analysis. In *32nd Conference on Neural Information Processing Systems (NIPS 2018), IRASL workshop*, pages 73–81.
- Pinky Sitikhu, Kritish Pahi, Pujan Thapa, and Subarna Shakya. 2019. A comparison of semantic similarity methods for maximum human interpretability. In *2019 artificial intelligence for transforming business and society (AITB)*, volume 1, pages 1–4. IEEE.
- Harald Steck, Chaitanya Ekanadham, and Nathan Kallus. 2024. [Is cosine-similarity of embeddings really about similarity?](#) *CoRR*, abs/2403.05440.
- Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. 2021. [Whitening sentence representations for better semantics and faster retrieval](#). *CoRR*, abs/2103.15316.
- Anant Subramanian, Danish Pruthi, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Eduard H. Hovy. 2018. [SPINE: sparse interpretable neural embeddings](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 4921–4928. AAAI Press.
- Fei Sun, Jiafeng Guo, Yanyan Lan, Jun Xu, and Xueqi Cheng. 2016. [Sparse word embeddings using  \$\ell\_1\$  regularized online learning](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 2915–2921. IJCAI/AAAI Press.
- Stijn van Dongen and Anton J. Enright. 2012. [Metric distances derived from cosine similarity and pearson and spearman correlations](#). *CoRR*, abs/1208.3145.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 38–45. Association for Computational Linguistics.
- Peipei Xia, Li Zhang, and Fanzhang Li. 2015. [Learning similarity with cosine similarity ensemble](#). *Inf. Sci.*, 307:39–52.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. [Normalized word embedding and orthogonal transform for bilingual word translation](#). In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 1006–1011. The Association for Computational Linguistics.
- Hiroaki Yamagiwa, Momose Oyama, and Hidetoshi Shimodaira. 2023. [Discovering universal geometry in embeddings with ICA](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 4647–4675. Association for Computational Linguistics.
- Hiroaki Yamagiwa, Yusuke Takase, and Hidetoshi Shimodaira. 2024. [Axis tour: Word tour determines the order of axes in ica-transformed embeddings](#). *CoRR*, abs/2401.06112.
- Sho Yokoi, Ryo Takahashi, Reina Akama, Jun Suzuki, and Kentaro Inui. 2020. [Word rotator’s distance](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2944–2960. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

	Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning
Normalized ICA	53	salts	solvents	chlorine	hydrogen	inorganic	ammonia	chloride	flammable	sulfide	sulfur	[chemistry]
	68	proteins	protein	genes	gene	mrna	receptor	transcription	activation	p53	rna	[biology]
	141	spacecraft	astronauts	orbit	nasa	astronaut	orbiter	space	orbiting	mars	atlantis	[space]
	194	light	ultraviolet	infrared	sunlight	uv	shadows	bright	illumination	glow	illuminate	[spectrum]
	197	virus	h5n1	influenza	flu	contagious	outbreak	swine	avian	viruses	pandemic	[virology]
Normalized PCA	80	zhongshan	solar	optical	shyh	electrics	wafers	nerpa	selden	wbut	reutemann	[PC80]
	92	woodcuts	natrum	eriboll	linocuts	shamva	cellblock	hafslund	g2	heidelberg	venter	[PC92]
	152	sidebar	maternal	smoker	customizer	non-qualified	sufia	foundresses	frosting	traudl	romm	[PC152]
	153	monogamous	lifespan	necessitate	loyals	supplementation	skrall	zng	gietzen	remnant	well-meaning	[PC153]
	222	replication	isley	guaporé	bobos	pawhuska	foss	rigoberto	angara	laporta	200-250	[PC222]

Table 4: For the normalized ICA-transformed and PCA-transformed GloVe embeddings of *ultraviolet* in Fig. 2, the axes of the top 5 component values are focused on, and their top 10 words are shown. For ICA, the meanings of the axes are interpreted from these listed words and labeled such as [chemistry]. For PCA, however, since it is difficult to interpret the meanings of the axes, they are simply labeled such as [PC80].

	Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning
light	194	light	ultraviolet	infrared	sunlight	uv	shadows	bright	illumination	glow	illuminate	[spectrum]
	1	.	but	though	although	,	even	that	instance	both	however	[function words]
	153	welterweight	heavyweight	middleweight	featherweight	ibf	wbc	holyfield	wba	cruiserweight	bout	[boxing]
	53	salts	solvents	chlorine	hydrogen	inorganic	ammonia	chloride	flammable	sulfide	sulfur	[chemistry]
	58	23rd	35th	27th	22nd	26th	39th	24th	36th	31st	37th	[ordinal]
ultraviolet $\odot$ light	194	light	ultraviolet	infrared	sunlight	uv	shadows	bright	illumination	glow	illuminate	[spectrum]
	53	salts	solvents	chlorine	hydrogen	inorganic	ammonia	chloride	flammable	sulfide	sulfur	[chemistry]
	1	.	but	though	although	,	even	that	instance	both	however	[function words]
	153	welterweight	heavyweight	middleweight	featherweight	ibf	wbc	holyfield	wba	cruiserweight	bout	[boxing]
	158	gendarmierie	force	contingent	contingents	detachments	500-strong	detachment	constabulary	5,000-strong	constables	[police]

Table 5: For the normalized ICA-transformed GloVe embeddings of *light* and the component-wise products *ultraviolet*  $\odot$  *light* in Table 1, the axes of the top 5 values are focused on, and their top 10 words are shown. Axes are sorted by component values. The meanings of the axes are interpreted from these listed words.

## A Interpretation of ICA Components for *ultraviolet*, *light*, and *ultraviolet* $\odot$ *light*

***ultraviolet*.** For the normalized ICA-transformed and PCA-transformed GloVe embeddings of *ultraviolet* in Fig. 2, Table 4 shows the top 10 words for the axes of the top 5 component values. The semantic interpretations of the axes are provided for the ICA-transformed embeddings, but such interpretations are challenging for the PCA-transformed embeddings.

Based on Table 4, Fig. 4 illustrates our interpretation of cosine similarity between *ultraviolet* and the top word of each axis: *salts*, *proteins*, *spacecraft*, *light*, and *virus*.

***light* and *ultraviolet*  $\odot$  *light*.** Similar to Table 4, Table 5 shows the interpretation of ICA components for *light* and *ultraviolet*  $\odot$  *light*. The meaning of each axis in Table 1 is derived from Tables 4 and 5.

## B Normalization of Embeddings: Fidelity of Component Rankings

As shown in Fig. 1, although the meanings of the axes of the ICA-transformed and PCA-transformed

embeddings can be interpreted from the top words, the rankings of component values along the embeddings can change before and after normalization. Therefore, we compare the rankings along the embeddings with those along the axes, both before and after normalization.

Figure 10 shows scatterplots of these two rankings for the component values of the 150th axis of the ICA-transformed and PCA-transformed GloVe embeddings, both before and after normalization. Before normalization, as shown in Fig. 10a, embeddings with high rankings along the embeddings tend to have lower rankings along the axes as their norms increase. After normalization, as shown in Fig. 10b, the norm-derived artifacts observed in Fig. 10a disappear. These results support the improvement in interpretability through normalization, as discussed in Section 6.2.

## C Distribution of Cosine Similarity, Component Values and Their Products

In this section, vectors are denoted simply as  $X$  instead of in boldface as  $\mathbf{x}$ , and the elements of the vector  $X$  are denoted using the subscript  $X_\ell$  rather than the superscript  $X^{(\ell)}$ .

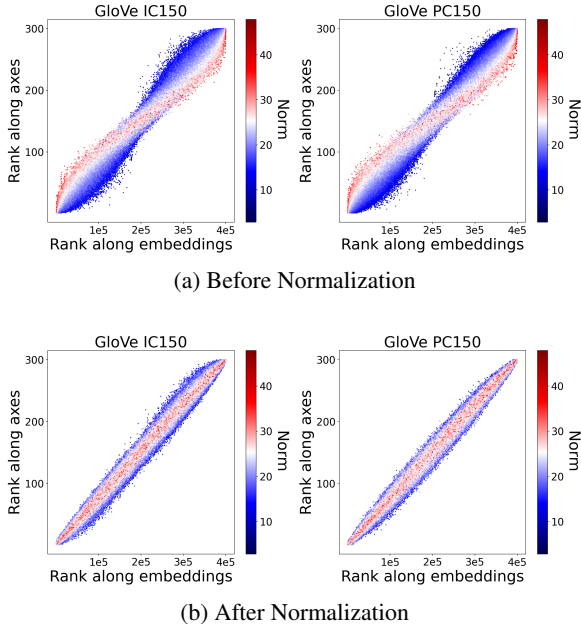


Figure 10: Scatterplots for the 150th axis of the (left) ICA-transformed and (right) PCA-transformed GloVe embeddings, (a) before and (b) after normalization, showing the rankings of component values along the embeddings and those along the axes, colored by the norms. The larger the norm of an embedding, the more it is plotted in the foreground. See Appendix D.4 for contextualized embeddings.

### C.1 Cosine similarity

Let us consider two random vectors  $X = (X_1, \dots, X_d), Y = (Y_1, \dots, Y_d) \in \mathbb{R}^d$  with elements of mean zero  $\mathbb{E}(X_\ell) = \mathbb{E}(Y_\ell) = 0$  and variance one  $\mathbb{E}(X_\ell^2) = \mathbb{E}(Y_\ell^2) = 1$ . We assume that the elements  $X_1, \dots, X_d$  and  $Y_1, \dots, Y_d$  are independent. Then, for sufficiently large  $d$ , the cosine similarity  $\cos(X, Y)$  asymptotically follows  $\mathcal{N}(0, 1/d)$ , the normal distribution with mean 0 and variance  $1/d$ ,

$$\cos(X, Y) \sim \mathcal{N}(0, 1/d). \quad (8)$$

This is easily shown as follows; a more general argument can be found in Appendix C of Yamagiwa et al. (2024). First note that  $\mathbb{E}(X_\ell Y_\ell) = \mathbb{E}(X_\ell)\mathbb{E}(Y_\ell) = 0$ ,  $\mathbb{E}(X_\ell^2 Y_\ell^2) = \mathbb{E}(X_\ell^2)\mathbb{E}(Y_\ell^2) = 1$ . Thus the inner product, if scaled by dimension,  $d^{-1/2}\langle X, Y \rangle = d^{-1/2} \sum_{\ell=1}^d X_\ell Y_\ell$  has mean zero and variance one. Thus, according to the central limit theorem,

$$d^{-1/2}\langle X, Y \rangle \sim \mathcal{N}(0, 1) \quad (9)$$

for sufficiently large  $d$ . It also follows from the law of large numbers that  $d^{-1}\|X\|^2 = d^{-1} \sum_{\ell=1}^d X_\ell^2$




Figure 11: For 10,000 randomly sampled pairs of ICA-transformed embeddings  $s, s' \in \mathbb{R}^d$ , the histogram of the cosine similarity  $\cos(s, s')$  is displayed. Since cosine similarities are invariant under the orthogonal transformation, exactly the same plot is also obtained from PCA-transformed embeddings. The theoretical probability density of (8) is almost identical to the observed histogram. The theory in Appendix C is also supported by the inverse of the variance,  $309.663 \approx d$  for  $d = 300$ . The observed cosine similarities in Fig. 4 are also indicated as vertical lines.

converges in probability to  $\mathbb{E}(X_\ell^2) = 1$ , and similarly  $d^{-1}\|Y\|^2 \rightarrow 1$  in probability. Therefore,

$$\sqrt{d} \cos(X, Y) = \frac{d^{-1/2}\langle X, Y \rangle}{\sqrt{d^{-1}\|X\|^2} \sqrt{d^{-1}\|Y\|^2}}$$

converges to (9), thereby concluding (8).

### C.2 Component values

Let  $S = (S_1, \dots, S_d) \in \mathbb{R}^d$  be a random vector representing ICA-transformed embeddings, and let  $e_\ell = (0, \dots, 0, 1, 0, \dots, 0) \in \mathbb{R}^d$  be the one-hot vector with a one at the  $\ell$ -th element. Then,  $\hat{S}_\ell = \cos(S, e_\ell)$  represents the  $\ell$ -th component of the normalized ICA-transformed embeddings. Although  $e_\ell$  is not a random vector, formally letting  $X = S$  and  $Y = e_\ell$  in (8) gives

$$\hat{S}_\ell \sim \mathcal{N}(0, 1/d) \quad (10)$$

for sufficiently large  $d$ . Considering  $S$  and  $e_\ell$  in the original coordinate system before the ICA-transformation, they can be regarded as almost random, which suggests that the formal argument above is valid.

### C.3 Product of two component values

Let  $\hat{S} = (\hat{S}_1, \dots, \hat{S}_d), \hat{S}' = (\hat{S}'_1, \dots, \hat{S}'_d) \in \mathbb{R}^d$  be two independent random vectors representing




Figure 12: The plots for the normalized PCA-transformed GloVe embeddings are displayed in the same manner as in Fig. 3. The inverse of the variance is  $300.000 \approx d$  in (a) and  $90,108.283 \approx d^2$  in (b, c) for  $d = 300$ .

Axis	Value	$p$ -value	Bonferroni
152	0.197	$3.21 \times 10^{-4}$	$9.62 \times 10^{-2}$
80	0.172	$1.41 \times 10^{-3}$	$4.22 \times 10^{-1}$
153	0.134	$1.01 \times 10^{-2}$	1.00
92	0.129	$1.30 \times 10^{-2}$	1.00
222	0.113	$2.52 \times 10^{-2}$	1.00

(a) ultraviolet

Axis	Value	$p$ -value	Bonferroni
1	0.247	$9.12 \times 10^{-6}$	$2.74 \times 10^{-3}$
5	0.216	$9.38 \times 10^{-5}$	$2.81 \times 10^{-2}$
61	0.176	$1.15 \times 10^{-3}$	$3.44 \times 10^{-1}$
118	0.176	$1.16 \times 10^{-3}$	$3.47 \times 10^{-1}$
127	0.146	$5.76 \times 10^{-3}$	1.00

(b) light

Axis	Value	$p$ -value	Bonferroni
1	0.024	$1.09 \times 10^{-4}$	$3.26 \times 10^{-2}$
80	0.022	$2.21 \times 10^{-4}$	$6.62 \times 10^{-2}$
131	0.020	$3.80 \times 10^{-4}$	$1.14 \times 10^{-1}$
152	0.019	$4.58 \times 10^{-4}$	$1.38 \times 10^{-1}$
60	0.017	$9.65 \times 10^{-4}$	$2.90 \times 10^{-1}$

(c) ultraviolet  $\odot$  light

Table 6: (a, b) show the observed component values of normalized PCA-transformed GloVe embeddings in Fig. 2b, while (c) shows their component-wise products. The  $p$ -values and their Bonferroni-corrected values are shown for the top five axes in each table.

normalized ICA-transformed embeddings. We consider  $Z_\ell = \hat{S}_\ell \hat{S}'_\ell$ , the  $\ell$ -th element of the component-wise product of  $\hat{S}$  and  $\hat{S}'$ . Then, for sufficiently large  $d$ , the probability density function of  $Z_\ell$  is

$$Z_\ell \sim (d/\pi)K_0(d|Z_\ell|), \quad (11)$$

where  $K_0(\cdot)$  is the modified Bessel function of the second kind of order zero. This result follows directly from Theorem 2.1 of Nadarajah and Pogány (2016), assuming that  $\hat{S}_\ell$  and  $\hat{S}'_\ell$  are independently distributed as  $\mathcal{N}(0, 1/d)$ . The mean of  $Z_\ell$  is  $\mathbb{E}(Z_\ell) = \mathbb{E}(\hat{S}_\ell \hat{S}'_\ell) = \mathbb{E}(\hat{S}_\ell)\mathbb{E}(\hat{S}'_\ell) = 0$  and the variance of  $Z_\ell$  is  $\mathbb{E}(Z_\ell^2) = \mathbb{E}(\hat{S}_\ell^2 \hat{S}'_\ell^2) = \mathbb{E}(\hat{S}_\ell^2)\mathbb{E}(\hat{S}'_\ell^2) = 1/d^2$ .

#### C.4 Histogram of cosine similarity, component values and their products

**Histogram of cosine similarity.** The observed and theoretical distributions of the cosine similarities are exhibited in Fig. 11. The histogram and the theoretical curve of the probability density function clearly suggest that the distribution theory in Appendix C.1 is strongly supported and that the cosine similarities are normally distributed with mean zero and variance  $1/d$ . However, when embeddings are centered but not whitened, which violates the assumption in Appendix C.1, the variance should tend to be larger than  $1/d$ . In such a case, the in-

verse of the variance would provide an effective dimensionality of the embeddings.

**Histogram of component values and component-wise products.** The observed and theoretical distributions of the components and the component-wise products are exhibited in Fig. 3 for ICA and Fig. 12 for PCA. The histograms and the theoretical curves of the probability density functions clearly suggest that the distribution theory in Appendices C.2 and C.3 is strongly supported. The components are normally distributed with mean zero and variance  $1/d$ , and the probability density function of the component-wise products is expressed by the modified Bessel function, where the mean is zero and the variance is  $1/d^2$ . Interestingly, while the theory was originally provided for ICA-transformed embeddings, the experiments suggest that the theory also applies to the PCA-transformed embeddings.

**Comparison between observed values and theoretical distributions.** In Figs. 11, 3, and 12, the observed values, indicated as vertical lines, are compared to their respective distributions to identify the significance of the values. By multiplying the values by the inverse of the standard deviations, these values can be easily interpreted. Thus,  $\sqrt{d} \cos(\mathbf{s}, \mathbf{s}')$ ,  $\sqrt{d} \hat{s}^{(\ell)}$ , and  $d \hat{s}^{(\ell)} \hat{s}'^{(\ell)}$  may be used



Model	Layers	Dimensions	Parameters
bert-base-uncased			110M
roberta-base	12	768	125M
gpt2			117M
pythia-160m			160M

Table 7: The number of layers, the dimensions, and the parameter size for each model.

for numerical comparisons.

Similar to Table 1, Table 6 presents the top five values from the bar graphs in Fig. 2b, along with their corresponding  $p$ -values and Bonferroni-corrected values. As shown in Fig. 12, fewer axes are statistically significant with  $\alpha = 0.05$  compared to those in Table 1.

## D Details of ICA Consistency for Contextualized Embeddings

### D.1 Experimental settings

To conduct experiments with contextualized embedding models, we used the `transformers` library (Wolf et al., 2020). The models used are listed in Table 7.

Following Yamagiwa et al. (2023), we selected the One Billion Word Benchmark (Chelba et al., 2014) as the corpus for our experiments. To compute the embeddings, we used the sentences in the first file<sup>6</sup> of the dataset’s training data, converted to lowercase. It is important to note that, for computing the embeddings of *ultraviolet* in Fig. 6, we prioritized sentences in the file that contain the word *ultraviolet*<sup>7</sup>. Then, embeddings were computed sequentially from the remaining sentences. To avoid the overrepresentation of high-frequency words during the ICA transformation, we computed up to 10 embeddings for each token, resulting in a total of 50,000 embeddings. Each token was distinguished as *ultraviolet\_0*, *ultraviolet\_1*, . . . , and so on.

For the ICA transformation of the contextualized embeddings of each model, we set the hyperparameters to the same values used for the GloVe embeddings: a maximum number of iterations of 10,000 and a convergence tolerance of  $1 \times 10^{-10}$ .

### D.2 Details of the heatmaps in Fig. 5

We explain the details of the heatmaps in Fig. 5 using the ICA-transformed embeddings, and ap-

<sup>6</sup>`news.en-00001-of-00100`

<sup>7</sup>The word *light* were included without prioritizing specific sentences.

ply the same procedure to the PCA-transformed embeddings.

**Axis matching.** As a preprocessing step, for each model’s 50,000 tokens, the prefix `##` is removed for BERT, and the prefix `Ġ` is removed for other models, if present. Then, if a word in the GloVe vocabulary is included in the token sets of each model, the word and its corresponding tokens are paired. If there are multiple corresponding tokens, multiple pairs are created. Using these pairs, the correlation coefficients between the axes of the ICA-transformed GloVe embeddings and those of the ICA-transformed contextualized embeddings are computed. The greedy algorithm by Yamagiwa et al. (2023) for matching these axes is then applied using these correlation coefficients. For further details, refer to Yamagiwa et al. (2023).

**Selection of words and tokens.** After matching these axes, for the axes of GloVe, we selected the top five words with the largest components among the words that are paired with the tokens of all other models. For each model, if multiple pairs existed, the token with the largest component on the corresponding axis was selected.

### D.3 Details of the bar graphs by the embeddings of *ultraviolet\_3* and *light\_1*

#### D.3.1 Details of *ultraviolet\_3* and *light\_1*

Table 8 shows the sentence that contains both *ultraviolet\_3* and *light\_1* used in Fig. 6 in Section 5. Among all the pairs of *ultraviolet* and *light* in the corpus (see Appendix D.1 for details), the pair of *ultraviolet\_3* and *light\_1* had the highest average cosine similarity computed by the ICA-transformed embeddings of the four models.

#### D.3.2 Interpretation of ICA Components for *ultraviolet\_3*, *light\_1*, and *ultraviolet\_3* $\odot$ *light\_1*

***ultraviolet\_3*.** Table 9 presents the top 10 words for the axes of the top 5 component values of the normalized ICA-transformed contextualized embeddings of *ultraviolet\_3*, along with the semantic interpretations of these axes.

***light\_1* and *ultraviolet\_3*  $\odot$  *light\_1*.** Similar to Table 9, for the normalized ICA-transformed embeddings of GPT-2 and Pythia-160m, Tables 10 and 11 provide the interpretation of ICA components for *light\_1* and *ultraviolet\_3*  $\odot$  *light\_1*. The meaning of each axis in Tables 2 and 3 for  $p$ -values

The *ultraviolet light* used in tanning salons are now considered as carcinogenic as tobacco and ...

Table 8: The sentence that contains both *ultraviolet\_3* and *light\_1* used in Fig. 6 in Section 5. Note that the sentence is lowercased when the embeddings are computed.

Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning	
BERT-base	286	greenhouse_0	greenhouse_2	greenhouse_4	greenhouse_1	emissions_2	greenhouse_3	emissions_5	carbon_5	emissions_6	emissions_0	[environment]
	343	before_7	shortly_5	shortly_8	before_1	before_9	shortly_0	within_2	before_6	prior_1	soon_9	[timings]
	373	hue_1	colors_4	color_0	hue_2	color_3	colors_2	colors_3	color_2	color_4	color_1	[colors]
	470	adhere_0	size_9	soft_0	drop_5	output_0	reserve_2	against_8	band_0	apply_3	flavor_0	[properties]
634	afternoon_1	midday_0	afternoon_6	morning_3	midday_1	morning_8	daily_7	daylight_0	morning_6	evenings_0	[timeframes]	
RoBERTa-base	80	greenhouse_1	carbon_8	emissions_5	greenhouse_2	greenhouse_0	emissions_0	climate_4	climate_2	gases_0	emissions_6	[environment]
	219	photos_2	images_2	images_3	pictures_3	photos_1	image_2	photos_0	photos_3	image_7		[images]
	330	wireless_3	wireless_7	wireless_1	wireless_0	networks_3	broadband_1	broadband_2	wireless_4	connectivity_0	networks_0	[networks]
	350	green_5	hue_0	blue_0	colors_0	color_5	muted_0	colored_1	gray_3	oured_0	color_4	[colors]
755	whole_3	grain_0	light_1	whole_4	guardian_1	grains_2	emb_0	guardian_3	fiber_0	ju_0	[grains]	
GPT-2	28	to_0	attracts_0	in_0	specialized_0	genes_0	ting_0	cause_5	adopt_0	ocytes_0	produced_1	[biology]
	273	arsenic_0	nitrogen_1	phosphorus_0	nitrogen_0	ate_4	dioxide_0	calcium_0	carbon_9	opes_0	fumes_0	[elements]
	349	ultraviolet_4	light_1	_0	light_0	accurate_1	-1	ultraviolet_1	color_4	ultraviolet_3	ultraviolet_0	[light]
	512	domain_1	computer_2	identify_1	bypass_0	mail_6	web_7	internet_8	allows_6	user_2	mail_5	[technology]
619	restaurant_0	proposal_2	project_3	project_5	soldier_2	farmer_0	scheme_6	designer_1	company_7	company_6	[occupations]	
Pythia-160m	262	light_1	ultraviolet_3	light_0	ultraviolet_4	ultraviolet_1	-0	ultraviolet_0	-1	ultraviolet_2	light_0	[light]
	371	aud_6	aud_8	aud_0	aud_2	aud_3	aud_1	aud_7	aud_4	aud_5	jan_4	[subwords]
	485	for_2	and_3	both_0	and_2	and_1	,4	of_2	,5	limited_0	)_0	[function words]
	619	scheme_6	restaurant_0	campaign_3	discovery_2	investor_1	resolution_2	island_0	let_5	designer_1	shoes_4	[business]
660	special_7	special_4	special_3	special_8	special_2	special_6	special_1	extraordinary_2	special_9	special_5	[exceptional]	

(a) Normalized ICA-transformed contextualized embeddings

Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning	
BERT-base	45	150_0	110_0	drilling_0	saddam_0	september_0	10_1	september_7	sunni_8	mist_4	saddam_1	[PC45]
	58	bodyguards_1	danger_0	13_1	threatening_1	angeles_1	castillo_0	tackle_1	bodyguards_0	angeles_7	13_0	[PC58]
	85	deutschland_0	equipment_6	explosions_2	germany_8	germany_5	germany_7	berlin_1	berlin_0	investors_5	munich_2	[PC85]
	91	jones_4	prior_1	h_5	side_9	evidence_4	present_3	burns_1	surgeon_1	news_0	video_6	[PC91]
473	tongue_1	spend_0	matter_6	romney_1	schedule_1	specify_1	weekly_0	finding_5	._7	era_7	[PC473]	
RoBERTa-base	38	examples_0	easy_4	expressed_3	offset_2	peaked_1	extreme_2	sketches_0	rare_4	spike_0	expressed_1	[PC38]
	278	returned_7	\xEF\xBF\xBD_2	worry_3	stories_3	returned_4	\xEF\xBF\xBD_3	_3	j_0	ight_0		[PC278]
	517	culture_4	worms_0	examined_0	machine_8	crew_3	architecture_1	rising_9	behalf_0	style_1	igator_0	[PC517]
	682	il_0	ung_6	fire_0	he_8	side_9	situation_1	dj_2	colleague_0	association_6	free_4	[PC682]
729	pick_0	sn_4	class_4	flo_0	bacon_0	agreement_4	apiece_0	caught_9	defensively_0	mble_0	[PC729]	
GPT-2	128	airlines_0	ids_0	buy_7	bought_3	pleaded_2	club_5	flights_4	values_4	aff_2	pleading_0	[PC128]
	254	exchange_5	21_0	international_6	telephone_0	political_9	mistakes_0	wine_9	partially_2	holder_0	diplomatic_0	[PC254]
	374	nine_3	ib_5	divorced_1	14_8	ultraviolet_1	ib_7	interviewed_1	appealing_0	ane_9	of_1	[PC374]
	560	forcing_7	14_5	results_4	az_0	kn_3	mont_6	really_6	fighting_6	rs_0	tre_0	[PC560]
763	these_3	company_1	ful_9	here_4	ki_0	said_1	ash_9	ec_9	gg_2	aud_0	[PC763]	
Pythia-160m	171	rev_2	rate_6	government_2	intensity_0	vibration_0	regiment_0	ash_3	pet_1	2_4	steady_2	[PC171]
	264	body_1	v_4	radiation_1	victory_3	article_6	article_0	article_5	concrete_1	62_0	yield_1	[PC264]
	270	rain_3	ideology_2	possible_3	bucks_0	ula_2	populated_0	damp_0	nick_3	vis_0	vis_2	[PC270]
	493	expect_8	farms_2	sugar_0	and_0	else_8	slightly_6	mel_1	orb_0	sensations_0	both_8	[PC493]
635	book_2	selling_0	kil_3	undergoing_0	if_5	if_4	if_0	if_7	if_1	if_3	[PC635]	

(b) Normalized PCA-transformed contextualized embeddings

Table 9: The top 10 words for the axes of the top 5 component values of the normalized embeddings of *ultraviolet\_3* in Figs. 6, 13, and 14. We also provide the semantic interpretations of these axes. Note that we removed the prefix ## for BERT and the prefix Ġ for other models. Note that all instances of the corrupted character in the top words of RoBERTa in Table 9b were replaced with the Unicode character \xEF\xBF\xBD.

and their Bonferroni-corrected values is derived from Tables 9a, 10 and 11.

### D.3.3 Comparison of ICA and PCA

We also compare ICA and PCA for contextualized embeddings.

First, as in Fig. 6 in Section 5, Fig. 13 shows bar graphs for the normalized ICA-transformed embeddings of BERT-base and RoBERTa-base. Addition-

ally, Fig. 14 shows bar graphs for the normalized PCA-transformed embeddings of the four models.

Then, we compare the ICA bar graphs (Figs. 6 and 13) with the PCA bar graphs (Fig. 14). As in Fig. 2 in Section 1, for the normalized ICA-transformed embeddings, the component values are sparse overall, and the component-wise products for the common semantic components (for exam-

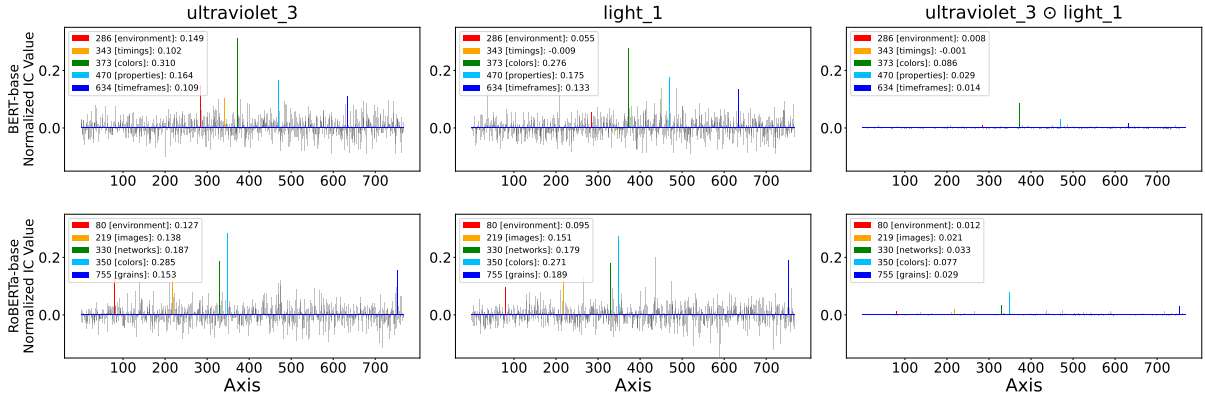


Figure 13: Similar to Fig. 2a, the component values and their component-wise products for the normalized ICA-transformed contextualized embeddings of *ultraviolet\_3* and *light\_1* are shown in bar graphs. The cosine similarity is 0.532 for BERT-base and 0.608 for RoBERTa-base.

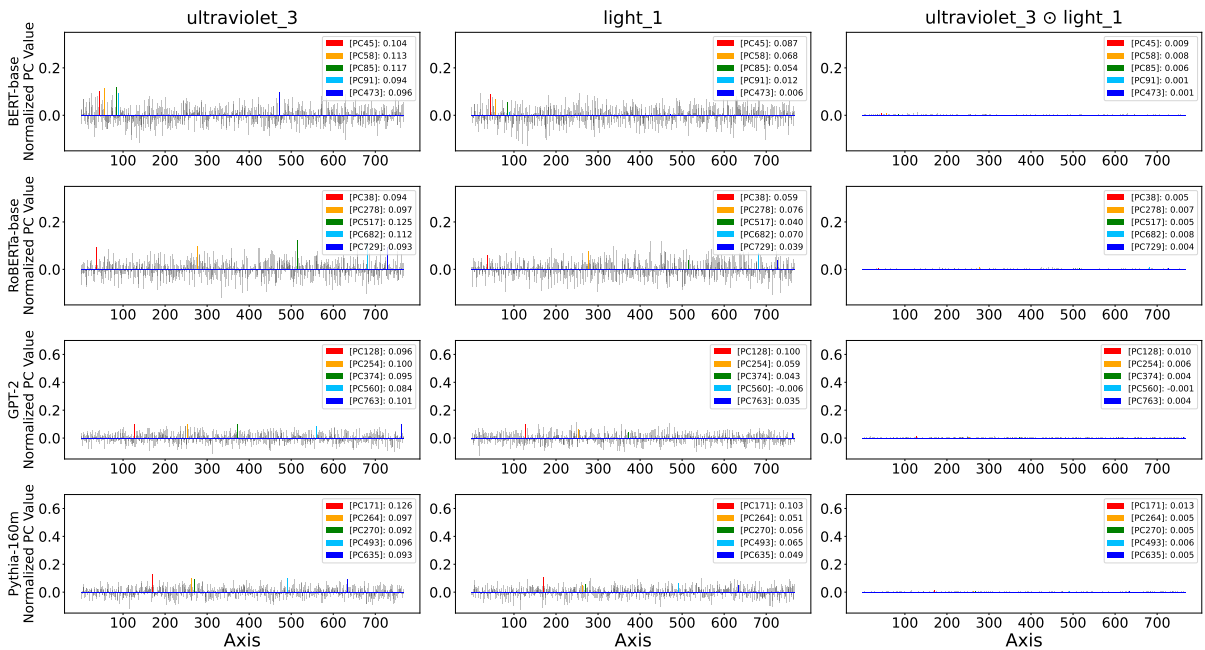


Figure 14: Similar to Fig. 2b, the component values and their component-wise products for the normalized PCA-transformed contextualized embeddings of *ultraviolet\_3* and *light\_1* are shown in bar graphs. The cosine similarity is 0.532 for BERT-base, 0.608 for RoBERTa-base, 0.576 for GPT-2, and 0.699 for Pythia-160m.

ple, [*light*] and [*colors*]) are significantly larger than others. In contrast, these products are not observed in the normalized PCA-transformed embeddings.

#### D.4 Comparison of normalization of embeddings in Appendix B

Based on the normalization experiments described in Appendix B, we also conducted comparative experiments on normalization using the four contextualized embeddings after ICA and PCA transformations.

Focusing on the component values of the 150th axis of each embedding, Fig. 15 shows scatterplots

of the ranks of component values along the embeddings and those along the axes, both before and after normalization. Similar to the norm-derived artifacts observed in Fig. 10 in Appendix B, artifacts are seen in Fig. 15a before normalization. However, these artifacts disappear in Fig. 15b after normalization.

#### D.5 Comparison of component values of ICA and PCA in Section 6.3

We also compare the ICA-transformed and PCA-transformed contextualized embeddings after normalization, based on Section 6.3. These embeddings have a size of 50,000 and a dimensionality

	Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning
$light\_1$	349	ultraviolet_4	light_1	-.0	light_0	accurate_1	-.1	ultraviolet_1	color_4	ultraviolet_3	ultraviolet_0	[light]
	273	arsenic_0	nitrogen_1	phosphorus_0	nitrogen_0	ate_4	dioxide_0	calcium_0	carbon_9	opes_0	fumes_0	[elements]
	619	restaurant_0	proposal_2	project_3	project_5	soldier_2	farmer_0	scheme_6	designer_1	company_7	company_6	[occupations]
	402	rain_8	rain_9	rain_2	rain_1	rain_3	storm_1	rain_7	rain_6	thunder_0	storm_3	[rainstorm]
	284	dangers_2	risks_2	dangers_0	dangers_3	capabilities_1	iciencies_1	behavior_1	difficulties_1	danger_2	inequalities_0	[challenges]
$ultraviolet\_3 \odot light\_1$	349	ultraviolet_4	light_1	-.0	light_0	accurate_1	-.1	ultraviolet_1	color_4	ultraviolet_3	ultraviolet_0	[light]
	273	arsenic_0	nitrogen_1	phosphorus_0	nitrogen_0	ate_4	dioxide_0	calcium_0	carbon_9	opes_0	fumes_0	[elements]
	28	to_0	attracts_0	in_0	specialized_0	genes_0	ting_0	cause_5	adopt_0	ocytes_0	produced_1	[biology]
	619	restaurant_0	proposal_2	project_3	project_5	soldier_2	farmer_0	scheme_6	designer_1	company_7	company_6	[occupations]
	402	rain_8	rain_9	rain_2	rain_1	rain_3	storm_1	rain_7	rain_6	thunder_0	storm_3	[rainstorm]

Table 10: For the normalized ICA-transformed GPT-2 embeddings of  $light\_1$  and the component-wise products  $ultraviolet\_3 \odot light\_1$  in Table 2, the axes of the top 5 values are focused on, and their top 10 words are shown. Axes are sorted by component values. The meanings of the axes are interpreted from these listed words.

	Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning
$light\_1$	262	light_1	ultraviolet_3	light_0	ultraviolet_4	ultraviolet_1	-.0	ultraviolet_0	-.1	ultraviolet_2	light_0	[light]
	619	scheme_6	restaurant_0	campaign_3	discovery_2	investor_1	resolution_2	island_0	let_5	designer_1	shoes_4	[business]
	264	carbon_7	carbon_6	carbon_8	carbon_0	carbon_4	carbon_5	nitrogen_1	nitrogen_0	phosphorus_0	arsenic_0	[elements]
	485	for_2	and_3	both_0	and_2	and_1	.,_4	of_2	.,_5	limited_0	)_0	[function words]
	548	fre_1	fre_4	fre_0	fre_2	fre_3	eas_0	gre_0	gre_1	eas_3	eas_5	[subwords]
$ultraviolet\_3 \odot light\_1$	262	light_1	ultraviolet_3	light_0	ultraviolet_4	ultraviolet_1	-.0	ultraviolet_0	-.1	ultraviolet_2	light_0	[light]
	619	scheme_6	restaurant_0	campaign_3	discovery_2	investor_1	resolution_2	island_0	let_5	designer_1	shoes_4	[business]
	485	for_2	and_3	both_0	and_2	and_1	.,_4	of_2	.,_5	limited_0	)_0	[function words]
	264	carbon_7	carbon_6	carbon_8	carbon_0	carbon_4	carbon_5	nitrogen_1	nitrogen_0	phosphorus_0	arsenic_0	[elements]
	336	made_3	significant_9	ful_7	making_5	make_3	specific_4	making_1	making_2	grave_2	makes_0	[creation]

Table 11: For the normalized ICA-transformed Pythia-160m embeddings of  $light\_1$  and the component-wise products  $ultraviolet\_3 \odot light\_1$  in Table 3, the axes of the top 5 values are focused on, and their top 10 words are shown. Axes are sorted by component values. The meanings of the axes are interpreted from these listed words.

of 768.

Figure 16 shows the sorted component values of these embeddings. In Fig. 16a, the ICA components are larger than the PCA components for up to approximately 1,000 embeddings. Additionally, in Fig. 16b, the ICA components are larger than the PCA components for up to approximately the 10th axis. These results show a trend similar to that shown by GloVe in Figs. 8a and 8b in Section 6.3.

## E Details of Word Intrusion Task in Section 6.2

**Selection of the intruder word.** Our objective is to evaluate the interpretability of the word embeddings  $\mathbf{Y} \in \mathbb{R}^{n \times d}$ , where each row vector  $\mathbf{y}_i \in \mathbb{R}^d$  corresponds to a word  $w_i$ . In order to select the intruder word for the set of top  $k$  ( $= 5$ ) words of each  $\ell$ -th axis ( $\ell \in \{1, \dots, d\}$ ), denoted as  $\text{top}_k(\ell)$ , we randomly chose a word from a pool of words that satisfy both of the following two criteria simultane-


ously: (i) the word ranks in the lower 50% in terms of the component value on the  $\ell$ -th axis, and (ii) it ranks in the top 10% in terms of the component value on some axis other than the  $\ell$ -th axis. For each axis,  $L$  ( $= 10$ ) intruder words are randomly selected, and  $W_{\text{int}}(\ell)$  denotes the set of these  $L$  intruder words.

**Evaluation metric.** We adopted the metric by Sun et al. (2016), but we mitigated the effect of outliers by using the median instead of the mean in the final score computation described below:

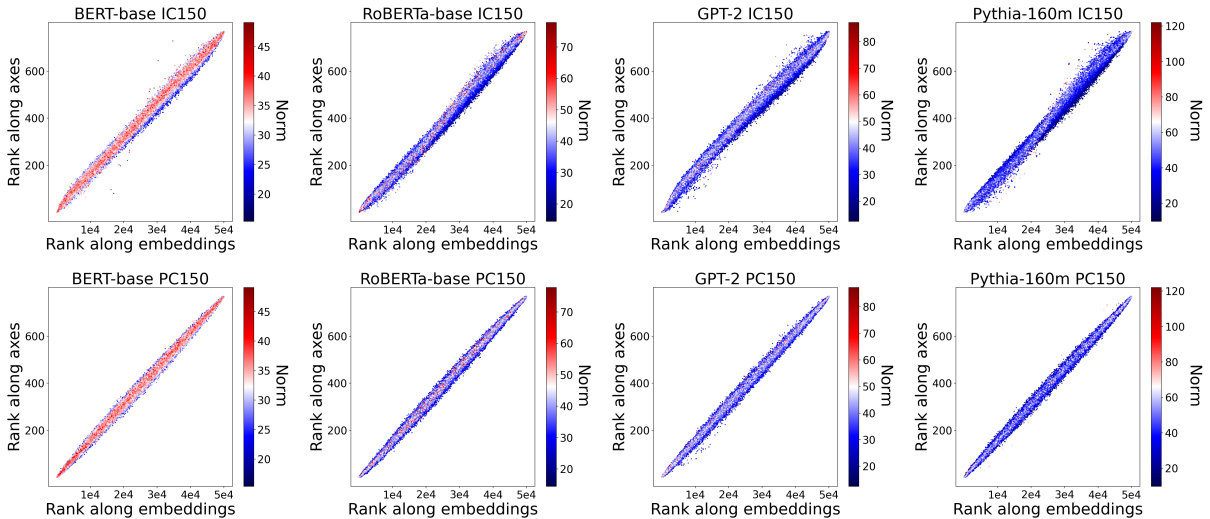
$$\text{IntruderScore} = \text{median}_{\ell=1, \dots, d} \frac{\text{InterDist}(\ell)}{\text{IntraDist}(\ell)},$$

$$\text{IntraDist}(\ell) = \sum_{\substack{w_i, w_j \in \text{top}_k(\ell) \\ w_i \neq w_j}} \frac{\text{dist}(w_i, w_j)}{k(k-1)},$$

$$\text{InterDist}(\ell) = \text{mean}_{w \in W_{\text{int}}(\ell)} \sum_{w_i \in \text{top}_k(\ell)} \frac{\text{dist}(w_i, w)}{k}.$$



(a) Before Normalization



(b) After Normalization

Figure 15: Scatterplots for the 150th axis of the (top) ICA-transformed and (bottom) PCA-transformed contextualized embeddings, (a) before and (b) after normalization, showing the rankings of component values along the embeddings and those along the axes, colored by the norms. The larger the norm of an embedding, the more it is plotted in the foreground.

In this formula, we defined  $\text{dist}(w_i, w_j) = \|y_i - y_j\|$ . Here,  $\text{IntraDist}(\ell)$  denotes the average distance between the top  $k$  words, and  $\text{InterDist}(\ell)$  represents the average distance between the top words and the intruder word. The score is higher when the intruder word is further away from the set  $\text{top}_k(\ell)$ . Therefore, this score serves as a quantitative measure of the ability to identify the intruder word; thus, it is used as a measure of the consistency of the meaning of the top  $k$  words and the interpretability of axes.

## F Details of the Downstream Tasks in Section 6.4

We performed analogy and word similarity tasks following the settings of Yamagiwa et al. (2024), using the Word Embedding Benchmark (Jastrzebski et al., 2017)<sup>8</sup>. We explain the details of the tasks using the ICA-transformed embeddings, and apply the same procedure to the PCA-transformed embeddings. We define  $\llbracket d \rrbracket := \{1, \dots, d\}$  and let  $p (\leq d)$  be the number of non-zero product components.

<sup>8</sup><https://github.com/kudkudak/word-embeddings-benchmarks>

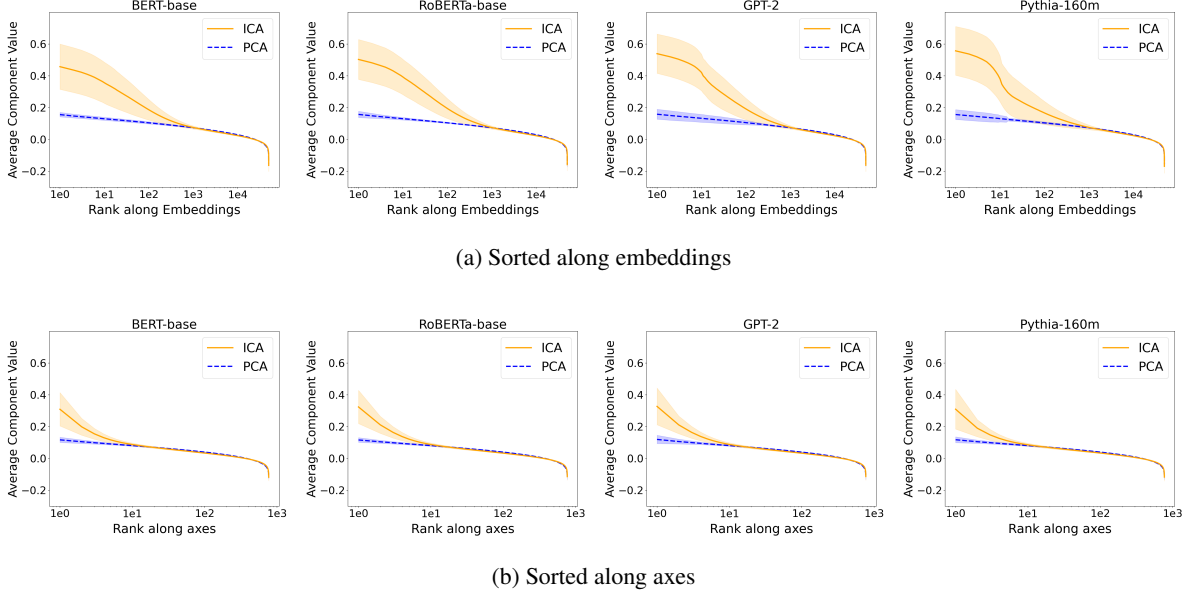


Figure 16: Comparison of component values of normalized ICA-transformed and PCA-transformed contextualized embeddings. The component values are averaged after being sorted in descending order (a) along embeddings for each axis, and (b) along axes for each embedding. The range of  $\pm 1\sigma$  is shown, where  $\sigma$  is the standard deviation of the component values.

## F.1 Word similarity tasks

### F.1.1 Settings

We used several datasets, including MEN (Bruni et al., 2014), MTurk (Radinsky et al., 2011), RG65 (Rubenstein and Goodenough, 1965), RW (Luong et al., 2013), SimLex999 (Hill et al., 2015), WS353 (Finkelstein et al., 2002), WS353R (WS353 Relatedness), and WS353S (WS353 Similarity). In these tasks, we compute the cosine similarity  $\cos(w_i, w_j)$  between words  $w_i$  and  $w_j$  and compare it with human-rated similarity scores. Spearman’s rank correlation is used as the evaluation metric.

### F.1.2 Details of $p$ non-zero products

We explain the top  $p$  component-wise products used for the similarity tasks.

As seen in (6), since the cosine similarity between words  $w_i$  and  $w_j$  is expressed as the sum of the component-wise products, we consider the index set of the top  $p$  component-wise products:

$$\text{Top}_p := \operatorname{argmax}_p \hat{s}_i^{(\ell)} \hat{s}_j^{(\ell)} \subset \llbracket d \rrbracket. \quad (12)$$

If  $p = d$ , then  $\text{Top}_p = \llbracket d \rrbracket$ .

Then, based on the cosine similarity expression in (6), we express the top  $p$  component-wise prod-

ucts as

$$\sum_{\ell \in \text{Top}_p} \hat{s}_i^{(\ell)} \hat{s}_j^{(\ell)}. \quad (13)$$

## F.2 Analogy tasks

### F.2.1 Settings

We used the Google Analogy Test Set (Mikolov et al., 2013a), which consists of 14 types of analogy tasks, and the Microsoft Research Syntactic Analogies Dataset (Mikolov et al., 2013d), which consists of 16 types. In these tasks, if  $w_{i_1}$  corresponds to  $w_{i_2}$ , then we predict  $w_{i_4}$  to which  $w_{i_3}$  corresponds. To do this, the vector  $\mathbf{s}_{i_2} - \mathbf{s}_{i_1} + \mathbf{s}_{i_3}$  is computed, and if the index of the closest word embedding to the vector in terms of cosine similarity,

$$\operatorname{argmax}_{i \in \llbracket n \rrbracket} \cos(\mathbf{s}_{i_2} - \mathbf{s}_{i_1} + \mathbf{s}_{i_3}, \mathbf{s}_i), \quad (14)$$

is  $i_4$ , then it is considered correct (top 1 accuracy).

### F.2.2 Details of $p$ non-zero products

We explain the top  $p$  component-wise products used for the analogy tasks.

Based on Appendix F.2.1, we define  $\mathbf{s}_{i_1, i_2, i_3} := \mathbf{s}_{i_2} - \mathbf{s}_{i_1} + \mathbf{s}_{i_3}$  and normalize it as

$$\hat{\mathbf{s}}_{i_1, i_2, i_3} := \frac{\mathbf{s}_{i_1, i_2, i_3}}{\|\mathbf{s}_{i_1, i_2, i_3}\|} = \left( \hat{s}_{i_1, i_2, i_3}^{(\ell)} \right)_{\ell=1}^d \in \mathbb{R}^d. \quad (15)$$

Tasks	$p = 1$		$p = 5$		$p = 10$		$p = 50$		$p = 100$		$p = 300$		
	PCA	ICA	PCA	ICA	PCA	ICA	PCA	ICA	PCA	ICA	PCA	ICA	
Similarity	MEN	0.11	0.45	0.31	0.59	0.45	0.63	0.69	0.72	0.73	0.74	0.75	0.75
	WS353	-0.02	0.19	0.12	0.45	0.25	0.49	0.49	0.55	0.54	0.56	0.57	0.57
	WS353R	0.04	0.16	0.15	0.44	0.22	0.46	0.43	0.49	0.48	0.49	0.51	0.51
	WS353S	0.01	0.31	0.17	0.56	0.34	0.60	0.62	0.67	0.66	0.70	0.69	0.69
	SimLex999	0.00	0.10	0.11	0.21	0.19	0.25	0.35	0.35	0.38	0.37	0.40	0.40
	RW	0.08	0.13	0.15	0.20	0.18	0.22	0.27	0.30	0.30	0.32	0.34	0.34
	RG65	0.42	0.42	0.53	0.57	0.61	0.64	0.71	0.76	0.74	0.79	0.78	0.78
	MTurk	0.22	0.42	0.36	0.58	0.49	0.61	0.64	0.65	0.66	0.66	0.64	0.64
	Average	0.11	<b>0.27</b>	0.24	<b>0.45</b>	0.34	<b>0.49</b>	0.53	<b>0.56</b>	0.56	<b>0.58</b>	0.59	0.59
Analogy	capital-common-countries	0.00	0.22	0.01	0.51	0.24	0.62	0.93	0.94	0.95	0.94	0.95	0.95
	capital-world	0.01	0.05	0.04	0.15	0.21	0.29	0.91	0.90	0.95	0.93	0.95	0.95
	city-in-state	0.00	0.00	0.00	0.13	0.05	0.21	0.49	0.55	0.59	0.60	0.67	0.67
	currency	0.00	0.00	0.00	0.00	0.00	0.03	0.06	0.11	0.11	0.12	0.12	0.12
	family	0.00	0.00	0.01	0.24	0.14	0.36	0.74	0.74	0.84	0.84	0.88	0.88
	gram1-adjective-to-adverb	0.00	0.00	0.00	0.00	0.02	0.02	0.16	0.16	0.19	0.20	0.21	0.21
	gram2-opposite	0.00	0.00	0.00	0.00	0.02	0.02	0.18	0.18	0.24	0.23	0.26	0.26
	gram3-comparative	0.00	0.00	0.02	0.19	0.17	0.42	0.79	0.82	0.84	0.86	0.88	0.88
	gram4-superlative	0.00	0.00	0.04	0.15	0.14	0.28	0.57	0.63	0.66	0.69	0.69	0.69
	gram5-present-participle	0.00	0.00	0.01	0.15	0.07	0.36	0.64	0.68	0.69	0.70	0.69	0.69
	gram6-nationality-adjective	0.01	0.22	0.14	0.37	0.52	0.49	0.91	0.92	0.92	0.92	0.93	0.93
	gram7-past-tense	0.00	0.00	0.03	0.09	0.09	0.21	0.50	0.52	0.56	0.57	0.60	0.60
	gram8-plural	0.00	0.00	0.01	0.15	0.13	0.27	0.70	0.69	0.75	0.74	0.76	0.76
	gram9-plural-verbs	0.00	0.00	0.02	0.20	0.12	0.37	0.48	0.57	0.57	0.60	0.58	0.58
	jj_jjr	0.00	0.00	0.03	0.13	0.14	0.29	0.53	0.56	0.60	0.63	0.66	0.66
	jj_jjs	0.00	0.00	0.03	0.11	0.10	0.22	0.41	0.49	0.49	0.54	0.51	0.51
	jjr_jj	0.00	0.00	0.01	0.03	0.05	0.10	0.45	0.46	0.53	0.54	0.54	0.54
	jjr_jjs	0.00	0.00	0.03	0.08	0.09	0.20	0.38	0.49	0.49	0.56	0.55	0.55
	jj_s_jj	0.00	0.00	0.02	0.05	0.05	0.10	0.36	0.36	0.43	0.44	0.48	0.48
	jj_s_jjr	0.00	0.01	0.01	0.14	0.09	0.29	0.55	0.59	0.60	0.63	0.63	0.63
	nn_nnpos	0.00	0.02	0.00	0.13	0.06	0.22	0.35	0.35	0.39	0.39	0.42	0.42
	nn_nns	0.00	0.03	0.00	0.24	0.10	0.39	0.62	0.63	0.69	0.68	0.74	0.74
	nnpos_nn	0.00	0.06	0.00	0.21	0.05	0.29	0.41	0.41	0.43	0.42	0.45	0.45
	nns_nn	0.00	0.06	0.01	0.24	0.11	0.36	0.56	0.55	0.61	0.59	0.64	0.64
	vb_vbd	0.00	0.01	0.03	0.10	0.13	0.23	0.55	0.53	0.56	0.54	0.58	0.58
	vb_vbz	0.00	0.00	0.02	0.16	0.11	0.34	0.66	0.72	0.74	0.76	0.76	0.76
	vbd_vb	0.00	0.01	0.02	0.09	0.08	0.27	0.60	0.64	0.66	0.67	0.69	0.69
vbd_vbz	0.00	0.00	0.03	0.11	0.10	0.29	0.56	0.60	0.63	0.63	0.63	0.63	
vbz_vb	0.00	0.00	0.02	0.09	0.15	0.27	0.77	0.75	0.80	0.79	0.82	0.82	
vbz_vbd	0.01	0.00	0.04	0.05	0.14	0.14	0.57	0.51	0.59	0.56	0.55	0.55	
Average	0.00	<b>0.02</b>	0.02	<b>0.14</b>	0.12	<b>0.26</b>	0.55	<b>0.57</b>	0.60	<b>0.61</b>	0.63	0.63	

Table 12: The performance of ICA-transformed and PCA-transformed embeddings when reducing the number of non-zero normalized component-wise products to  $p$  and then computing cosine similarity. The values in the table correspond to Top 1 accuracy for analogy tasks and Spearman’s rank correlation for word similarity tasks.

Then, following the cosine similarity expression in (6), (14) can be rewritten as

$$\operatorname{argmax}_{i \in \llbracket n \rrbracket} \sum_{\ell=1}^d \hat{s}_{i_1, i_2, i_3}^{(\ell)} \hat{s}_i^{(\ell)}. \quad (16)$$

Based on (16), we consider the index set of top  $p$  component-wise products for each word  $w_i$ :

$$\operatorname{Top}_p^i := \operatorname{argmax}_p \sum_{\ell \in \llbracket d \rrbracket} \hat{s}_{i_1, i_2, i_3}^{(\ell)} \hat{s}_i^{(\ell)} \subset \llbracket d \rrbracket. \quad (17)$$

If  $p = d$ , then  $\operatorname{Top}_p^i = \llbracket d \rrbracket$ .

Then, following the expression in (16), we use the top  $p$  component-wise products and predict the index of the answer word as follows:

$$\operatorname{argmax}_{i \in \llbracket n \rrbracket} \sum_{\ell \in \operatorname{Top}_p^i} \hat{s}_{i_1, i_2, i_3}^{(\ell)} \hat{s}_i^{(\ell)}. \quad (18)$$

### F.3 Results

Table 12 shows the results for the tasks with  $p = 1, 5, 10, 50, 100$ , and 300. The performance of the ICA-transformed embeddings is better than that of the PCA-transformed embeddings. As seen in




Figure 17: Comparison of the scatterplots for the sum of the top  $p$  component-wise products and the cosine similarity for the normalized ICA-transformed and PCA-transformed embeddings at (a)  $p = 1$ , (b) 10, and (c) 100. We used word pairs from the word similarity tasks. For  $p = 1, 10,$  and  $100$ , the correlation coefficients for ICA were 0.619, 0.876, and 0.979, respectively, while for PCA they were 0.432, 0.761, and 0.982. See Fig. 18 for the correlation coefficients for other values of  $p$ .




Figure 18: Comparison of correlation coefficients between the sum of the top  $p$  component-wise products and the cosine similarity for normalized ICA-transformed and PCA-transformed embeddings. We used word pairs from the word similarity tasks.

Fig. 9, the difference in performance between ICA and PCA is larger for the word similarity tasks than for the analogy tasks. This is probably because in the word similarity tasks, the semantic components are used directly, as shown in (13). On the other hand, in the analogy tasks, the component values of the normalized vector  $\hat{s}_{i_1, i_2, i_3}$  are used, as shown in (18).

#### F.4 Relation between the sum of the top $p$ products and cosine similarity

As seen in Appendix F.1, in the word similarity tasks, we compared performance using the sum of the top  $p$  component-wise products for the normalized ICA-transformed and PCA-transformed GloVe embeddings. In this section, we examine the relation between this sum and the original cosine similarity.

Figure 17 shows the comparison of the scatterplots for the sum and the cosine similarity for the normalized ICA-transformed and PCA-transformed embeddings. For smaller values of  $p$ , such as  $p = 1$  or  $p = 10$ , the correlation of the scatterplots is stronger for ICA than for PCA. At  $p = 100$ , the scatterplot shows strong correlations for both ICA and PCA.

Figure 18 compares the correlation coefficients at different values of  $p$  for these normalized embeddings. As seen in Fig. 17, the correlation coefficients for ICA are stronger than those for PCA at smaller values of  $p$ . As  $p$  increases, the correlation coefficients for both ICA and PCA increase and the difference between them decreases. These results indicate the favorable sparsity of the component-wise products for the normalized ICA-transformed embeddings.

## G Retrieval of Embeddings with Selected Semantic Components

In this section, we define ideal embeddings only with specific semantic components and investigate words with high cosine similarity to these embeddings. This analysis aims to provide a deeper understanding of the sparsity properties of ICA-transformed embeddings and the interpretation of cosine similarity presented in this study.

### G.1 Settings

The ideal embedding  $\hat{\mathbf{q}} = (\hat{q}^{(\ell)})_{\ell=1}^d \in \mathbb{R}^d$  with the semantic components at indices  $\ell_1, \dots, \ell_m$  is






Figure 19: Search results for seven ideal embeddings derived from ICA-transformed GloVe embeddings, containing only the semantic components of *[food]*, *[animals]*, and *[plants]*. For each embedding, the top 10 words and their cosine similarities are displayed. The search is performed sequentially from top to bottom for each combination of semantic components. Each word has its rank and is assigned a color corresponding to the combination of semantic components where it first appears.

defined as follows<sup>9</sup>:

$$\hat{q}^{(\ell)} := \begin{cases} 1/\sqrt{m} & \text{if } \ell = \ell_1, \dots, \ell_m, \\ 0 & \text{otherwise,} \end{cases} \quad (19)$$

where  $\|\hat{q}\| = 1$ . Note that searching using an ideal embedding with only one semantic component is the same as examining the top words of the axis of the normalized ICA-transformed embeddings. Using the ideal embedding  $\hat{q}$ , we search for the top  $k$  words by the inner product, i.e., cosine similarity:

$$\operatorname{argmax}_{i \in [n]} k \hat{q}^\top \hat{s}_i, \quad (20)$$

where  $[n] := \{1, \dots, n\}$ .

## G.2 Specific examples

As an example, we use ideal embeddings derived from ICA-transformed GloVe embeddings. We focus on the semantic components of *[food]*, *[animals]*, and *[plants]*, which yield insightful results. Figure 19 shows these results. The results for one semantic component show that the meanings of the selected axes represent *[food]*, *[animals]*, and *[plants]*. For the combinations of two semantic components, top words include *chicken*, *meat*, and *fish* for *[food]* & *[animals]*; *vegetables*, *tomato*, and *potatoes* for *[food]* & *[plants]*; and *species*,

*insects*, and *wild* for *[animals]* & *[plants]*. In addition, for *[food]* & *[animals]* & *[plants]*, top words include words such as *mushrooms*, *fruits*, and *figs*, which are plant-based food favored by animals. As the number of semantic components increases, the ambiguity increases, resulting in lower cosine similarity values.

## H Retrieval of Embeddings after Ablating a Semantic Component

In this section, we ablate one semantic component from a normalized ICA-transformed embedding, effectively setting the semantic similarity on the axis to zero. We then examine how the words with high similarity change before and after this ablation. This analysis aims to explore the impact of changes in a single semantic similarity and to provide an illustrative example for interpreting cosine similarity as the sum of semantic similarities.

### H.1 Settings

By setting the  $\ell_*$ -th semantic component to zero for the normalized ICA-transformed embeddings, we define the ideal embedding of the  $i_*$ -th word as

$$\hat{s}_{i_* \ominus \ell_*} := (\hat{s}_{i_*}^{(1)}, \dots, \hat{s}_{i_*}^{(\ell_*-1)}, 0, \hat{s}_{i_*}^{(\ell_*+1)}, \dots, \hat{s}_{i_*}^{(d)}). \quad (21)$$

Since  $\|\hat{s}_{i_*}\| = 1$ , it follows that  $\|\hat{s}_{i_* \ominus \ell_*}\| \leq 1$ . For  $\hat{s}_{i_*}$  and  $\hat{s}_{i_* \ominus \ell_*}$ , based on (6), we search for the top

<sup>9</sup>This study considers only uniform weights for simplicity.

$w_i$	<b>M</b>	<b>F</b>	$\hat{\mathbf{s}}_{i_{\text{woman}}}^\top \hat{\mathbf{s}}_i$	$\hat{\mathbf{s}}_{i_{\text{woman} \ominus \ell_{[\text{female}]}}}^\top \hat{\mathbf{s}}_i$
<i>girl</i>		✓	0.691	0.595 (↓ 0.096)
<i>man</i>	✓		0.679	0.720 (↑ 0.041)
<i>mother</i>		✓	0.632	0.467 (↓ 0.165)
<i>person</i>	✓	✓	0.579	0.576 (↓ 0.003)
<i>female</i>		✓	0.575	0.518 (↓ 0.057)
<i>she</i>		✓	0.568	0.422 (↓ 0.146)
<i>herself</i>		✓	0.567	0.429 (↓ 0.138)
<i>wife</i>		✓	0.553	0.392 (↓ 0.161)
<i>women</i>		✓	0.544	0.468 (↓ 0.076)
<i>daughter</i>		✓	0.535	0.382 (↓ 0.153)

(a) *woman* top 10 words

$w_i$	<b>M</b>	<b>F</b>	$\hat{\mathbf{s}}_{i_{\text{woman}}}^\top \hat{\mathbf{s}}_i$	$\hat{\mathbf{s}}_{i_{\text{woman} \ominus \ell_{[\text{female}]}}}^\top \hat{\mathbf{s}}_i$
<i>man</i>	✓		0.679	0.720 (↑ 0.041)
<i>girl</i>		✓	0.691	0.595 (↓ 0.096)
<i>person</i>	✓	✓	0.579	0.576 (↓ 0.003)
<i>female</i>		✓	0.575	0.518 (↓ 0.057)
<i>boy</i>	✓		0.497	0.509 (↑ 0.012)
<i>someone</i>	✓	✓	0.490	0.482 (↓ 0.008)
<i>teenager</i>	✓	✓	0.524	0.481 (↓ 0.044)
<i>men</i>	✓		0.432	0.471 (↑ 0.039)
<i>victim</i>	✓	✓	0.507	0.471 (↓ 0.036)
<i>women</i>		✓	0.544	0.468 (↓ 0.076)

(b) *woman*⊖*[female]* top 10 words

Table 13: The top 10 words for (a) *woman* and (b) *woman*⊖*[female]* are shown with their inner product values against both *woman* and *woman*⊖*[female]*. Check marks indicate whether words are masculine (M) or feminine (F). The differences between their inner product values for *woman* and *woman*⊖*[female]*, computed for each word  $w_i$  as  $\hat{\mathbf{s}}_{i_{\text{woman} \ominus \ell_{[\text{female}]}}}^\top \hat{\mathbf{s}}_i - \hat{\mathbf{s}}_{i_{\text{woman}}}^\top \hat{\mathbf{s}}_i$ , are indicated with ↑ for increases and ↓ for decreases, with the magnitude shown as the absolute value.




Figure 20: Bar graphs of the normalized ICA-transformed GloVe embeddings of *woman*, *girl*, and *man*, showing each component value of *[female]*.

$k$  words by their inner products as follows:

$$\operatorname{argmax}_{i \in [n] \setminus \{i_*\}} k \hat{\mathbf{s}}_{i_*}^\top \hat{\mathbf{s}}_i, \quad (22)$$

$$\operatorname{argmax}_{i \in [n] \setminus \{i_*\}} k \hat{\mathbf{s}}_{i_* \ominus \ell_*}^\top \hat{\mathbf{s}}_i. \quad (23)$$

Note that the original  $i_*$ -th word is excluded from the candidates. In particular, as evident from the definition of  $\hat{\mathbf{s}}_{i_* \ominus \ell_*}$  in (21), (23) considers the situation where the  $\ell_*$ -th semantic similarity is set to zero in the cosine similarity in (1).

## H.2 Specific examples

As an example, consider the ablation of a semantic component in the normalized ICA-transformed GloVe embedding of *woman*. We focused on one of the axes with large component values in the embedding, where the top 5 words were *her*, *wife*, *mother*, *daughter*, and *actress*. Therefore, we interpret the meaning of this axis as *[female]*. We then set the semantic component of *[female]* in the embedding to zero and define the ideal embedding. We call the corresponding ideal word for this ideal embedding *woman*⊖*[female]*.

### H.2.1 Top words and their inner products

Based on (22), we first searched for the top 10 words of *woman*, and Table 13a shows their inner product values with both *woman* and *woman*⊖*[female]*. Words such as *mother*, *wife*, and *daughter*, which have high inner product values with *woman*, have lower values with *woman*⊖*[female]*. Conversely, *man* has higher values with *woman*⊖*[female]* than with *woman*.

Next, based on (23), for the top 10 words of *woman*⊖*[female]*, Table 13b shows their inner product values with both *woman* and *woman*⊖*[female]*. Compared to the top 10 words of *woman*, that of *woman*⊖*[female]* includes masculine words such as *boy* and *men*, as well as words like *someone* and *teenager* that are both masculine and feminine. Furthermore, for the masculine words, the inner product values with *woman*⊖*[female]* are larger than those with *woman*. On the other hand, for feminine words (including those that are both masculine and feminine), the opposite trend is observed.

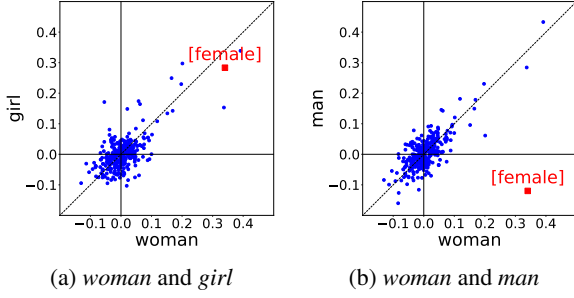


Figure 21: Scatterplots of component values of normalized ICA-transformed GloVe embeddings for (a) the pair *woman* and *girl* and (b) the pair *woman* and *man*. The semantic component of *[female]* is marked with ■ and others with •. Figure 20 in Appendix H shows bar graphs for each embedding.




Figure 22: Scatterplot of  $\hat{s}_{i_{woman \ominus \ell_{[female]}}} \top \hat{s}_i$  and  $\hat{s}_i^{(\ell_{[female]})}$  for each  $i \in \llbracket n \rrbracket$ . The position of *woman* is marked with ★. Words from Table 13a (*woman* top 10 words) and Table 13b (*woman*  $\ominus$  *[female]* top 10 words) are categorized as follows: Words that appear in the top 10 for both *woman* and *woman*  $\ominus$  *[female]* are marked with ◆, those that appear only in the top 10 for *woman* with ▲, those that appear only in the top 10 for *woman*  $\ominus$  *[female]* with ▼. Other words are marked with •. The value of  $\hat{s}_{i_{woman}} \top \hat{s}_i$ , which represents the cosine similarity  $\cos(woman, w_i)$  between *woman* and the word  $w_i$ , is plotted as contours. The contour for the 10th closest word *daughter* to *woman* is shown as a dash-dotted line. For the 10th closest word *women* to *woman*  $\ominus$  *[female]*, their inner product  $\hat{s}_{i_{woman \ominus \ell_{[female]}}} \top \hat{s}_{i_{women}}$  is shown as a dashed line.

### H.3 Bar graphs and scatterplots of the embeddings for *woman*, *girl*, and *man*

In Table 13, *girl* has a high inner product value with *woman* but a low inner product value with *woman*  $\ominus$  *[female]*. Conversely, *man* exhibits the opposite trend. Therefore, in this section, we take a closer look at the normalized ICA-transformed GloVe embeddings of *woman*, *girl*, and *man*.

Figure 20 shows bar graphs of these three embeddings. The graphs for *woman* and *girl* are similar in shape. On the other hand, the graphs for *woman* and *man* show differences in the semantic component of *[female]*, although other parts are broadly similar.

In addition, Fig. 21 shows scatterplots of the normalized ICA-transformed GloVe embeddings for the pair *woman* and *girl*, and the pair *woman* and *man*. For the semantic component of *[female]*, both the embeddings of *woman* and *girl* have positive values, while the embedding of *man* has negative values. Similar to Fig. 20, these results also show that other semantic components for both pairs are roughly correlated.

### H.4 Analysis of changes in top words

The changes in the top words before and after ablating the semantic component of *[female]* are shown in Table 13. In this section, we examine these changes in detail.

Figure 22 shows a scatterplot of the inner product with  $woman \ominus [female]$  and the semantic component of *[female]* for every word. Words in the upper right are close to *woman*, while those with higher values on the horizontal axis are closer to  $woman \ominus [female]$ . The transition from dash-dotted to dashed lines indicates that the ablation for *[female]* replaces the top words for *woman* such as *mother*, *she*, *herself*, *wife*, and *daughter* with the top words for  $woman \ominus [female]$  such as *boy*, *someone*, *teenager*, *men*, and *victim*. These results show that setting one semantic similarity to zero can significantly alter the cosine similarity in (1).

As related work, Ishibashi et al. (2020) proposed a method to invert the meaning of a word by mirroring its embedding across a hyperplane. In contrast, we invert the meaning by simply ablating the component from a particular axis.

## I Case Study of ICA-transformed embeddings

In this section, we introduce interesting case studies of ICA-transformed embeddings using the embeddings published by Yamagiwa et al. (2023).

### I.1 Comparison of the noun and verb senses of *shore* in ICA-transformed BERT Embeddings

Yamagiwa et al. (2023) observed that the noun *shore* has a large component value of *[sea]*, while

that of the verb *shore* is small. We use their embeddings to examine the semantic components across all axes. Figure 23 shows the normalized ICA-transformed embeddings of these three *shore* instances as bar graphs. For the embeddings of *shore\_1* and *shore\_2*, the semantic components of *[sea]* and *[location]* are large. For the embedding of *shore\_0*, these are small, but those of *[control]* and *[causative verbs]* are large. These results explain the large and small relations in the cosine similarity:  $\cos(\textit{shore}_1, \textit{shore}_2) = 0.299$ , while  $\cos(\textit{shore}_0, \textit{shore}_1) = 0.054$  and  $\cos(\textit{shore}_0, \textit{shore}_2) = 0.128$ . Table 14 shows the top 10 words of the 10 axes selected by choosing the axes of the top 4 component values for each *shore*, excluding duplicates.

Additionally, the sentences containing these tokens are shown in Table 15. While *shore\_0* is a verb, both *shore\_1* and *shore\_2* are nouns. Comparing the nouns *shore\_1* and *shore\_2*, we find that since *shore\_1* is specifically part of *Sydney's North Shore*, the normalized ICA-transformed BERT embedding of *shore\_1* in Fig. 23 has a large semantic component of *[australia]*. This result also shows that the BERT embeddings are well contextualized.

## I.2 Comparison of ICA-transformed fastText embeddings in multiple languages

Yamagiwa et al. (2023) showed that common semantic axes exist between ICA-transformed embeddings of different languages, and matched these axes by their correlations. They used the fastText (Bojanowski et al., 2017) embeddings for their experiments, and their ICA-transformed and PCA-transformed fastText embeddings are published<sup>10</sup>.

As an example, we analyzed the embedding of *boat* and compared it with those of its translations in Spanish, Russian, Arabic, Hindi, Chinese, and Japanese<sup>11</sup>. Figures 24 and 25 show the bar graphs of the normalized ICA-transformed and PCA-transformed embeddings, respectively, for these languages. Table 16 shows the top 10 words of the axes of the top 5 component values in each embedding of *boat*<sup>12</sup>. These results show that while the semantic component of *[ship-and-sea]*

is the largest for all normalized ICA-transformed embeddings in Fig. 24, there is no such semantic component for the normalized PCA-transformed embeddings in Fig. 25.

<sup>10</sup><https://github.com/shimo-lab/Universal-Geometry-with-ICA>

<sup>11</sup>We chose *boat* as the example word, which is the top word of the second most correlated axis. Note that while the meaning of the most correlated axis is *[first name]*, first names such as *mike* are the same across languages such as Spanish.

<sup>12</sup>“海” in the top words of the second axis of the normalized ICA-transformed embeddings is the Chinese character for *sea*.




Figure 23: Bar graphs of the normalized ICA-transformed BERT embeddings for *shore\_0*, *shore\_1*, and *shore\_2*. The ten axes whose component values are large for these *shore* are interpreted and colored. While *shore\_0* is a verb, *shore\_1* and *shore\_2* are nouns. See Table 14 for the top words of the axes and Table 15 for the sentences containing each *shore*.

Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning
61	naval_2	sea_1	marine_2	vessels_2	sea_2	sea_6	vessel_0	naval_1	ships_1	ship_2	[sea]
207	australian_1	australian_2	new_1	wales_0	australia_0	sydney_1	south_0	queensland_0	sydney_0	australia_3	[australia]
250	neighborhood_3	street_12	drag_0	of_209	central_7	district_12	street_29	street_8	heart_5	in_680	[district]
303	curb_2	prevent_0	avoid_0	reduce_12	reducing_1	##d_39	reduce_3	control_2	prevent_7	prevent_10	[control]
338	makes_2	allowed_7	forcing_1	bring_11	triggered_1	lead_2	illustrated_0	make_65	triggered_2	force_8	[causative verbs]
344	the_1179	the_1976	of_1927	were_41	were_89	the_1593	also_64	the_1180	been_154	been_81	[numbers]
354	agency_17	agency_1	strategy_2	company_59	country_7	group_41	law_7	charity_1	company_13	pact_0	[organization]
521	bail_2	bail_8	pay_23	walk_0	##avi_1	roll_0	walk_5	cop_1	go_28	turn_5	[action]
558	ground_10	side_2	front_10	side_18	corner_2	scene_3	sides_6	trail_0	hand_8	hand_19	[location]
572	more_167	a_244	##some_0	with_434	and_1493	good_38	and_246	with_318	more_148	the_2200	[linking words]

Table 14: For the normalized ICA-transformed BERT embeddings of *shore\_0*, *shore\_1*, and *shore\_2* in Fig. 23, the axes of the top 4 component values for each embedding are focused. The number of axes is 10, excluding duplicates, and the top 10 words of the axes are shown.

<i>shore_0</i>	Last month , the two companies sliced their dividends and sold billions of dollars of special stock to raise capital and <i>shore</i> up their finances .
<i>shore_1</i>	Working for a Sydney newspaper , my daughter covered a dreadful 1994 fire where , on one of the suburban streets of Sydney 's North <i>Shore</i> , the fire jumped the road and , for some terrifying seconds , took all the oxygen with it .
<i>shore_2</i>	Coastguards from Clevedon and Weston searched the <i>shore</i> while two lifeboats and two helicopters were also involved .

Table 15: Sentences for *shore\_0*, *shore\_1*, and *shore\_2* in Fig. 23. Note that *shore\_0* is a verb, while *shore\_1* and *shore\_2* are nouns.

	Axis	Top1	Top2	Top3	Top4	Top5	Top6	Top7	Top8	Top9	Top10	Meaning
Normalized ICA	2	boat	sailing	sail	ship	boats	sea	ships	海	open-sea	ocean	[ship-and-sea]
	17	car.	car	bmw	4-door	car-	v-6	2-dr	u.s.-market	car-and	2-door	[cars]
	36	water	rivers	reservoir	water-the	river-water	water.	water-	black-water	basin	de-water	[water]
	129	12-man	five-man	five-person	six-member	seven-man	14-member	12-person	seven-person	12-member	three-person	[multiple people]
	131	race	races	racing	race.	racer	race.-	.race	rider	laps	race-like	[races]
Normalized PCA	69	bit-field	torn	out-of-round	unused	final-	3-space	bad.2.	cup	too-large	.language	[Aligned Axis69]
	85	2-the	name-called	accusations	t-head	down.1.	1-the	relata	attacked	flew	two-place	[Aligned Axis85]
	96	government-run	trading	state-run	military-run	kids-only	trade	floating	e-a	sirven	trade.	[Aligned Axis96]
	99	white-red-white	-green	white-blue	red-green	white-red	.hair	color	'k	voz	poles	[Aligned Axis99]
	142	business-process	source-to-pay	time-to-market	menuado	s-u	5-6-11	pulse	cost	news-	time-to-value	[Aligned Axis142]

Table 16: For the normalized ICA and PCA transformed fastText embeddings of *boat* in Figs 24 and 25, the axes of the top 5 component values are focused and their top 10 words are shown. For PCA, since it is difficult to interpret the meanings of the axes, they are simply labeled such as [Aligned Axis69]. Similar to GloVe, the meanings of the axes of the ICA-transformed fastText embeddings are interpretable.

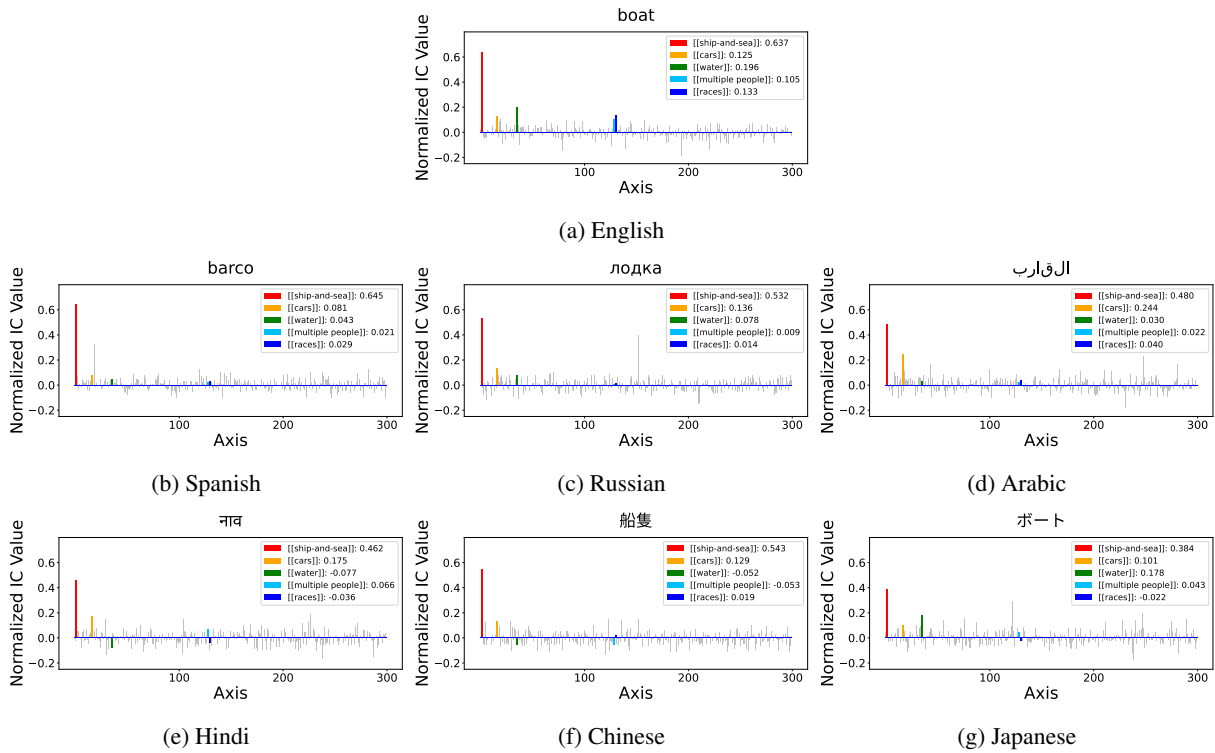


Figure 24: For *boat* and its translations, the normalized ICA-transformed fastText embeddings are shown as bar graphs. These axes are aligned by the correlation coefficients between English and the other languages. The axes of the top 5 component values in English are highlighted with their meanings. The component values of these axes are shown in the bar graphs for each language. See Table 16 for the top 10 words of these axes.

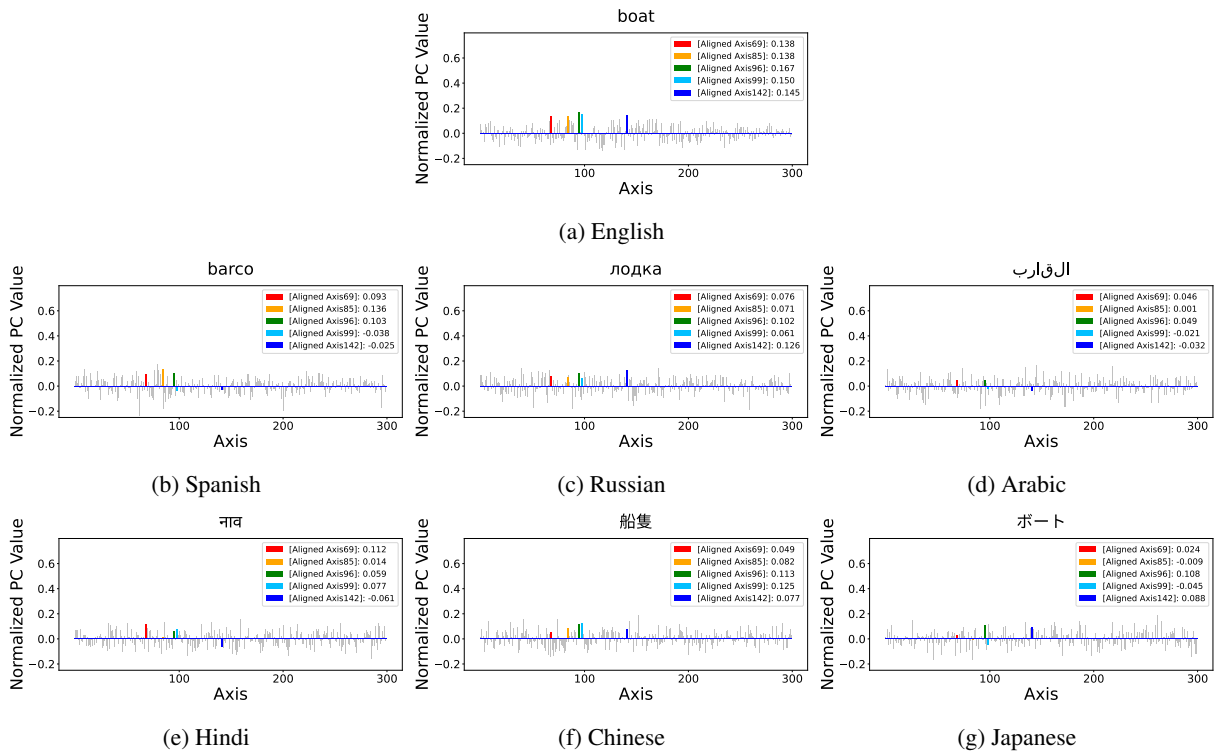


Figure 25: For *boat* and its translations, the normalized PCA-transformed fastText embeddings are shown as bar graphs. These axes are aligned by the correlation coefficients between English and the other languages. The axes of the top 5 component values in English are highlighted with their meanings. The component values of these axes are shown in the bar graphs for each language. See Table 16 for the top 10 words of these axes.