Exploring Intra and Inter-language Consistency in Embeddings with ICA

Rongzhi Li[1], **Takeru Matsuda**[1,2], **Hitomi Yanaka**[1]
[1]The University of Tokyo, [2]RIKEN Center for Brain Science
`{iimori-eiji,hyanaka}@is.s.u-tokyo.ac.jp`
`matsuda@mist.i.u-tokyo.ac.jp`

Abstract

Word embeddings represent words as multidimensional real vectors, facilitating data analysis and processing, but are often challenging to interpret. Independent Component Analysis (ICA) creates clearer semantic axes by identifying independent key features. Previous research has shown ICA's potential to reveal universal semantic axes across languages. However, it lacked verification of the consistency of independent components within and across languages. We investigated the consistency of semantic axes in two ways: both within a single language and across multiple languages. We first probed into intra-language consistency, focusing on the reproducibility of axes by running the ICA algorithm multiple times and clustering the outcomes. Then, we statistically examined inter-language consistency by verifying those axes' correspondences using statistical tests. We newly applied statistical methods to establish a robust framework that ensures the reliability and universality of semantic axes.

1 Introduction

Word embedding is a technique that represents words from natural languages as multidimensional real vectors in a space (e.g., Euclidean space), making it easier to handle them as data. These embeddings create a continuous representation of words and sentences, facilitating data analysis and processing. However, word embeddings are challenging to interpret because the values vary greatly depending on the training data and the dimension of the embedding space (Levy and Goldberg, 2014). For example, it is unclear what the embedding exactly means even if we say the embedding of "Argentina" is $[0.0088871, -0.02218, \dots]$.

In order to cope with the interpretability problem, several approaches were suggested, such as Principle Component Analysis (PCA) and Independent Component Analysis (ICA, Hyvärinen et al., 2001). ICA gives a more interpretable representation of semantic axes (i.e., components labeled with high-relation words) over PCA (Musil and Mareček, 2024). For example, if an independent component scores high on the words "apple", "banana", and "peach", the semantic axis can be interpreted as the concept of fruits and labeled as [apple banana peach].

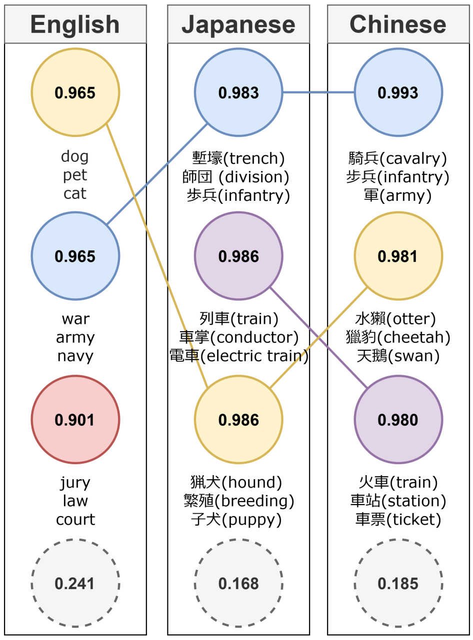

Figure 1: An illustration of clustering of independent components within and between languages. The circles represent the clusters created by Icasso, and the numbers indicate their quality indexes. Clusters with high-quality indexes were given interpretations using words. The circles connected by straight lines show components grouped together by checking consistency among languages.

Yamagiwa et al. (2023) demonstrated that Independent Component Analysis (ICA) can break down multilingual word embeddings into interpretable axes, suggesting that certain semantic axes may be universal across languages. However, their study had two main limitations. First, it focused

Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, pages 19104–19111
November 12-16, 2024 ©2024 Association for Computational Linguistics

solely on calculating the correlation coefficients for the correspondence of semantic axes between languages. Second, it lacked verification of the *consistency* of independent components within and across languages. In this study, consistency refers to the reliability of independent components that appear in different runs within one language **(intra-language consistency)** and the accurate correspondence of semantic axes among multiple languages **(inter-language consistency)**.

While Musil and Mareček (2024) and Yamagiwa et al. (2023) have made significant progress in uncovering semantic axes within individual languages, the extent to which these axes are shared across languages remains unclear. Our study addresses this gap by quantitatively demonstrating the similarity of semantic axes across languages, providing insights that are difficult to achieve with alignment methods alone. To do this, we first test the reliability of each language's independent components using Icasso (Himberg et al., 2004), a method based on running the ICA algorithm multiple times and clusters the results to ensure consistency. After labeling the components with words at the cluster centers, we analyze the statistical correspondence of these semantic axes. We then use the method by Hyvärinen and Ramkumar (2013), originally created for neuroscience, to identify common independent components across different languages.

Our contribution improves the interpretability and consistency of word embeddings. By ensuring robust independent components within languages through Icasso and verifying their correspondences across languages with statistical tests, we provide a more rigorous framework for analyzing semantic structures in multilingual word embeddings.[1]

2 Background

Interpretability in Word Embeddings The interpretability problem of word embeddings has been actively discussed. Various approaches have been proposed to enhance the interpretability of word embeddings, including aligning embedding spaces (Panigrahi et al., 2019; Park et al., 2017; Bommasani et al., 2020) and applying techniques such as loss function design, sparse overcomplete vectors(Yamagiwa et al., 2023). While these methods have shown promise, they primarily focus on ensuring that words with similar meanings are close to each other in the embedding space, without necessarily providing deeper interpretive insights into the semantic structure.

ICA has emerged as a relatively new method for interpreting word embeddings, showing great potential in explaining semantic axes. Musil and Mareček (2022) and Musil and Mareček (2024) applied ICA to word embeddings and presented the semantic axes of the words. Yamagiwa et al. (2023) conducted both PCA and ICA on word embeddings, demonstrating that there are intrinsic semantic axes among them. Their comparison revealed that ICA showed more distinctive axes than PCA.

ICA ICA is a method to extract statistically independent components from multivariate data (Hyvärinen et al., 2001). Let $\mathbf{X} \in \mathbb{R}^{d\times n}$ be a data matrix, where d is the data dimension and n is the number of observations. ICA is based on the assumption that $\mathbf{X}$ is represented as

$$\mathbf{X} = \mathbf{AS},$$

where $\mathbf{A} \in \mathbb{R}^{d\times d}$ is called the mixing matrix and $\mathbf{S} \in \mathbb{R}^{d\times n}$ is the matrix of independent components. Namely, the rows of $\mathbf{S}$ correspond to d latent factors that are statistically independent, and $\mathbf{A}$ indicates how these factors are combined in each of the d observed variables. ICA employs the non-Gaussianity of independent components to compute $(\mathbf{A}, \mathbf{S})$ from $\mathbf{X}$. It has been applied to various data (e.g., audio, neuroimaging data) for signal separation and feature extraction (Hyvärinen et al., 2001).

While ICA is non-deterministic, which could potentially limit its reproducibility, our research addresses this issue by employing Icasso. This method clusters multiple runs of ICA to ensure the reliability of independent components. This approach mitigates the non-deterministic nature of ICA by providing a robust statistical framework, thereby enhancing the interpretability and consistency of the semantic axes derived from ICA. Our study extends the work of Yamagiwa et al. (2023) by applying Icasso and using the cluster correspondence method by Hyvärinen and Ramkumar (2013) to verify these components' intra-language and inter-language consistency. The method ensures that the results of semantic axes are meaningful but reproducible and reliable, making ICA a more exciting and valuable method than PCA for analyzing

[1]The code will be published in the following repository: `https://github.com/des737/ExploreIca`

the semantic structure of word embeddings across languages.

3 Consistency in Embeddings

3.1 Intra-language Consistency

Unlike PCA, the result of ICA can be different between different runs due to the random initialization in the algorithm, such as FastICA (Hyvärinen, 1999), an insufficient number of observations, and the presence of noise in the data. Thus, the reproducibility of the independent components needs to be verified. Himberg et al. (2004) developed a method called Icasso for assessing the algorithmic and statistical reliability of independent components. In this study, we apply Icasso to word embedding vectors to evaluate intrinsic semantic axes' consistency in each language.

Here, we briefly explain the procedure of Icasso. See Appendix A and Himberg et al. (2004) for technical details. First, we run ICA on the data matrix $\mathbf{X} \in \mathbb{R}^{d\times n}$ m times and obtain m sets of d independent components. Then, we compute a similarity measure for each pair of two independent components from different sets (i.e., $m(m-1)/2 \cdot d^2$ pairs). By using this similarity, we perform agglomerative hierarchical clustering of independent components. Namely, starting from md clusters of size one containing each independent component, we iteratively merge two clusters with the maximum similarity. The reliability of each cluster is quantified by the quality index introduced in Himberg et al. (2004), which takes a value from 0 to 1. Clusters with a quality index close to 1 represent highly reproducible independent components, which correspond to consistent semantic axes in the case of word embeddings.

3.2 Inter-language Consistency

While ICA can extract semantic axes for each language, it has not been quantitatively examined whether there is correspondence between the semantic axes of several languages. Thus, we investigate the consistency of semantic axes across languages with statistical significance evaluation. We utilize the method by Hyvärinen and Ramkumar (2013) for clustering independent components from several data. This method was initially developed to find common independent components across subjects in neuroimaging data analysis.

Here, we explain the method of Hyvärinen and Ramkumar (2013) for the case of studying consistency across English and Japanese. Suppose that we have n pairs of English and Japanese words with the same meanings (e.g., "word" and "単語" (word)). Let $\mathbf{X}_{\mathrm{E}} \in \mathbb{R}^{d\times n}$ and $\mathbf{X}_{\mathrm{J}} \in \mathbb{R}^{d\times n}$ be the matrices composed of their d-dimensional English and Japanese embedding vectors, respectively. We apply ICA and obtain $\mathbf{X}_{\mathrm{E}} = \mathbf{A}_{\mathrm{E}}\mathbf{S}_{\mathrm{E}}$ and $\mathbf{X}_{\mathrm{J}} = \mathbf{A}_{\mathrm{J}}\mathbf{S}_{\mathrm{J}}$. Recall that each row of $\mathbf{S}_{\mathrm{E}}$ and $\mathbf{S}_{\mathrm{J}}$ represents the activation pattern of each independent component. Then, for $i, j = 1, \ldots, d$, we compute the p-value (with multiplicity correction) of the null hypothesis that the i-th row of $\mathbf{S}_{\mathrm{E}}$ and the j-th row of $\mathbf{S}_{\mathrm{J}}$ are independent. If this p-value is small, the i-th independent component of English and the j-th independent component of Japanese are significantly similar.

In the above way, we compute the p-values for each pair of languages. Then, we utilize them as a similarity measure for agglomerative hierarchical clustering of the independent components from multiple languages. The obtained clusters indicate the consistency of semantic axes across languages.

3.3 Interpretation of Independent Components

We use three representative words selected as follows to interpret independent components obtained from word embedding vectors. Recall that ICA is given by $\mathbf{X} = \mathbf{AS}$. Thus, the embedding vector of the j-th word (the j-th column vector of $\mathbf{X}$) is represented as

$$\mathbf{x}_j = s_{1j}\mathbf{a}_1 + \cdots + s_{dj}\mathbf{a}_d,$$

where $\mathbf{a}_i$ is the i-th column vector of $\mathbf{A}$ and s_{ij} is the (i, j)-th entry of $\mathbf{S}$. Therefore, s_{ij} quantifies how much the j-th word is related to the i-th independent component. Based on this observation, we sort the i-th row of $\mathbf{S}$ to $s_{ij_1} > s_{ij_2} > s_{ij_3} > \ldots$ and take the j_1, j_2, j_3-th words as the representatives of the i-th independent component. These words provide an intuitive understanding of the independent components as semantic axes.

4 Experimental Settings

We conducted the intra-language experiment focusing on the consistency of each language and then conducted the inter-language experiment focusing on the consistency among the languages. To align with Yamagiwa et al. (2023), we used the same FastText (Joulin et al., 2016) embeddings obtained by training on 157 different languages.

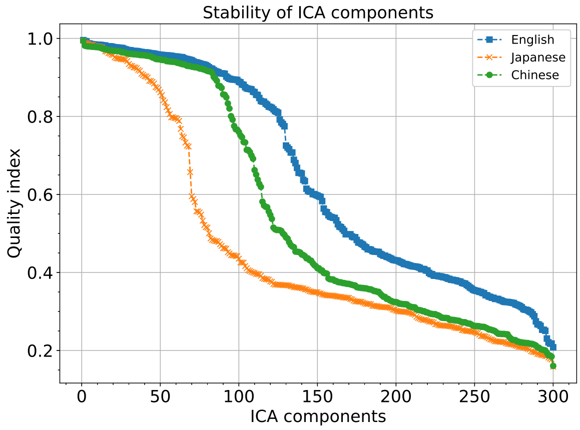

Figure 2: Quality index for FastText embeddings.

We obtained 300-dimensional embedding vectors of 50000 words for English, Japanese, and Chinese, respectively, with matrices $\mathbf{X_0}, \mathbf{X_1}, \mathbf{X_2} \in \mathbb{R}^{300\times 50000}$. The 50000 words consist of 6903 common words among the three languages selected from the multilingual word dictionary (Conneau et al., 2017) and 43097 words selected in order of their frequency of occurrence in each language by Wordfreq (Speer, 2022). We applied Icasso implemented by Captier et al. (2022) to FastText's word embeddings of English, Japanese, and Chinese with 10 runs, designated 300 as the number of clusters. We then tested the consistency among the components by the method proposed by Hyvärinen and Ramkumar (2013) by setting the false discovery rate and the false positive rate at 1%. Detailed explanations are in Appendix B.

5 Results and Discussion

5.1 Overall Results

Figure 2 shows the results of Icasso. There is a clear drop after the quality index reaches 0.8. The number of clusters with a cluster quality index exceeding 0.8 was 118 for English, 64 for Japanese, and 104 for Chinese.

As a result of the inter-language analysis, 47 clusters, 120 out of 354 (118×3) vectors, were found, which means the average number of vectors per cluster is 2.55. The language pairs of clusters of English-Japanese, Japanese-Chinese, Chinese-English, and all languages were 7, 10, 4, and 26, respectively. These results suggest that Japanese and Chinese share more semantic similarities, while English and Chinese have the least overlap. The high number of clusters shared across all three languages (26) indicates the presence of universal semantic concepts.

Based on the results of Icasso applied to static word embeddings, we identified a maximum of 118 consistent components for each language with the clusters' quality index above 0.8. We selected the independent component located at the center of each cluster in order of the highest quality index. Consequently, we constructed independent component matrices $\mathbf{S_0}, \mathbf{S_1}, \mathbf{S_2} \in \mathbb{R}^{118\times 50000}$.

Table 1 presents part of the results of the semantic axes after interpretation. Each axis is related to themes such as "words", "fishery", "religion", "film", "mathematical terms" and "army" demonstrating clear alignment of the axes among different languages.

5.2 Quantitative Evaluations

We performed human judgment experiments to evaluate the aligned components. Five participants proficient in all three languages took part in the experiments, conducting binary classification to determine if the semantic axes were sufficiently similar. We tested Fleiss' kappa, which is defined in Appendix A. κ was 0.364, with $\bar{P}$, $\bar{P}_e$ being 0.702 and 0.531, respectively. This suggests that our semantic axes extracted agree fairly with human valuation since κ is between 0.2 and 0.4 (Landis and Koch, 1977).

5.3 Discussion

The variation in the number of stable clusters across languages (Figure 2) provides meaningful insights into the structure of multilingual embeddings. English, as the source language in the multilingual dictionary (Conneau et al., 2017), exhibits the highest number of stable clusters (118). This result is expected, as it likely reflects the central role of English in the embedding space and its influence on anchoring more stable semantic clusters.

In comparison, the higher number of stable clusters in Chinese (104) compared to Japanese (64) is particularly intriguing. Based on the dictionary data[2], we observed that Chinese words in the dataset tend to correspond to a larger number of English words on average than Japanese words. This suggests that Chinese words may have broader semantic coverage or exhibit greater polysemy. By corresponding to more English words, Chinese words may be capturing a wider range of semantic concepts, resulting in the emergence of

[2]Refer to Appendix D.

English	Japanese	Chinese
verb noun word	流暢 発音 方言 (fluency) (pronunciation) (dialect)	話 流利 諺語 (speech) (fluency) (proverb)
boat sail buoy	漁業 漁師 捕鯨 (fishing industry) (fisherman) (whaling)	漁民 舢 捕鯨 (fisherman) (sampan) (whaling)
nun pope monk	教義 礼拝 会衆 (doctrine) (worship) (congregation)	恩典 基督 禱告 (grace) (christ) (prayer)
film gore cinema	演技 俳優 演劇 (acting) (actor) (drama)	放映 喜劇 戲服 (screening) (comedy) (costume)
sum cosine ray	乗法 整数 写像 (multiplication) (integer) (mapping)	方程 向量 切線 (equation) (vector) (tangent)
war army navy	塹壕 師団 歩兵 (trench) (division) (infantry)	騎兵 步兵 軍 (cavalry) (infantry) (army)

Table 1: Interpretation of clusters.

more distinct and stable semantic axes during the analysis.

Despite the variation in the number of stable clusters across languages, identifying about 30% similar semantic axes among languages (47 out of 118 maximum possible clusters) supports universal semantic concepts across these diverse languages. Regarding our methodology for representing semantic axes, we used three representative words for each axis, an extension of the approach used by Yamagiwa et al. (2023), who used a single word. This decision made the axes easier for humans to interpret by providing a more nuanced representation. The choice of three words is solely for improving human interpretability and does not impact our approach's core experimental results or statistical validity.

In future work, we plan to explore these applications and investigate whether the observed patterns hold across a broader range of languages and embedding types. Additionally, we intend to experiment directly with contextual representations to determine if similar insights can be gained as those obtained from static embeddings.

6 Conclusion

Our study statistically confirmed the consistency of semantic axes within and across languages using ICA components. Recognizing the inherent instability of ICA, we employed Icasso to ensure robustness by running multiple iterations and clustering the results. This process resulted in high-quality, reproducible semantic axes for English, Japanese, and Chinese. We then statistically verified inter-language consistency by identifying common semantic axes shared among these languages, supported by rigorous statistical tests. Our primary contribution is the innovative use of statistical methods to ensure the reliability and universality of semantic axes. The validation underscores the effectiveness of our approach in achieving consistent and interpretable word embeddings and highlights the potential for improved multilingual natural language processing applications.

Limitations

We also conducted experiments with BERT because it would have been beneficial to include an analysis of contextualized word embeddings compared to static word embeddings like FastText. It would be necessary to have a parallel corpus across English, Japanese, and Chinese to gain word embeddings from the same context. However, there currently needs to be more data on the multilingual parallel corpus, especially in English, Japanese, and Chinese. For preliminary experiments, we used TED Multilingual Parallel Corpus[3]. However, ICA did not converge, which was mainly attributed to the low amount of data, so we did not include the results in this paper.

Also, the number of independent components was limited to the dimensionality because the linear trait of ICA. Non-linear ICA proposed by Hyvärinen et al. (2019) were not implemented due to time constraints but can be applied to word embeddings in the future.

Acknowledgements

We thank the three anonymous reviewers for their helpful comments and feedback. This work was partially supported by PRESTO, JST Grant Number JPMJPR21C8, Japan, and JSPS KAKENHI Grant Number 22K17865.

[3]https://github.com/ajinkyakulkarni14/TED-Multilingual-Parallel-Corpus

References

Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.

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*, pages 4758–4781, Online. Association for Computational Linguistics.

Nicolas Captier, Jane Merlevede, Askhat Molkenov, Ainur Ashenova, Altynbek Zhubanchaliyev, Petr V Nazarov, Emmanuel Barillot, Ulykbek Kairov, and Andrei Zinovyev. 2022. Biodica: a computational environment for independent component analysis of omics data. *Bioinformatics*, 38(10):2963–2964.

Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.

Johan Himberg, Aapo Hyvärinen, and Fabrizio Esposito. 2004. Validating the independent components of neuroimaging time series via clustering and visualization. *NeuroImage*, 22(3):1214–1222.

Aapo Hyvärinen. 1999. Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks*, 10(3):626–634.

Aapo Hyvärinen. 2011. Testing the ICA mixing matrix based on inter-subject or inter-session consistency. *NeuroImage*, 58(1):122–136.

Aapo Hyvärinen, Juha Karhunen, and Erkki Oja. 2001. Independent component analysis, adaptive and learning systems for signal processing, communications, and control. *John Wiley & Sons, Inc*, 1:11–14.

Aapo Hyvärinen and Pavan Ramkumar. 2013. Testing independent component patterns by inter-subject or inter-session consistency. *Frontiers in Human Neuroscience*, 7:94.

Aapo Hyvärinen, Hiroaki Sasaki, and Richard Turner. 2019. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 859–868. PMLR.

Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hérve Jégou, and Tomas Mikolov. 2016. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.

J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.

Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 302–308, Baltimore, Maryland. Association for Computational Linguistics.

Tomáš Musil and David Mareček. 2022. Independent components of word embeddings represent semantic features. *ArXiv*, abs/2212.09580.

Tomáš 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)*, pages 6922–6928, Torino, Italia. ELRA and ICCL.

Abhishek Panigrahi, Harsha Vardhan Simhadri, and Chiranjib Bhattacharyya. 2019. Word2Sense: Sparse interpretable word embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5692–5705, Florence, Italy. 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*, pages 401–411, Copenhagen, Denmark. Association for Computational Linguistics.

Robyn Speer. 2022. rspeer/wordfreq: v3.0.

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*, pages 4647–4675, Singapore. Association for Computational Linguistics.

A Definitions

A.1 Similarity

The similarity σ_{ij} between s_i and s_j is defined as follows:

$$\sigma_{ij} = \left| \frac{\frac{1}{d}\sum_k s_{ik}s_{jk}}{\sqrt{\frac{1}{d}\sum_k s_{ik}^2}\sqrt{\frac{1}{d}\sum_k s_{jk}^2}} \right|$$

In other words, σ is the absolute value of the correlation coefficient, and the degree of difference is given by

$$d_{ij} = 1 - \sigma_{ij}.$$

A.2 Quality Index

The quality index I_q is defined as follows:

$$I_q(C_m) = \frac{1}{|C_m|^2} \sum_{i,j \in C_m} \sigma_{ij} - \frac{1}{|C_m||C_{-m}|} \sum_{i \in C_m} \sum_{j \in C_{-m}} \sigma_{ij},$$

where C_m refers to cluster m and C_{-m} refers to all independent components except cluster m. $|C_m|$ is the number of components in a cluster.

A.3 Fleiss' Kappa

Fleiss' kappa is defined as below:

$$\kappa = \frac{\bar{P} - \bar{P}_e}{1 - \bar{P}_e}.$$

B FDR and FPR

In multiple inter-language tests, the null hypothesis may be rejected by chance as the number of tests increases. For example, if we consider the test at 5% significance level, when all the null hypotheses are true, as many as 50 null hypotheses are rejected by chance in 1000 tests. Therefore, as discussed below, a correction is often made to account for this. The false discovery rate (FDR) is defined as

$$\text{FDR} = \frac{\text{False rejections when } H_0 \text{ is true}}{\text{Total rejections}},$$

where H_0 is the null hypothesis. To keep the FDR below a specified value of α_{FD} in overall tests, the corrected FDR α_{FD}^{corr} is calculated in each test by the method proposed by Benjamini and Hochberg (1995). In addition to FDR, we also consider the false positive rate (FPR). The FPR is defined as follows:

$$\text{FPR} = \frac{\text{False rejections when } H_0 \text{ is true}}{\text{Cases where } H_0 \text{ is false}}.$$

To control FDR below α_{FP}, corrected value α_{FP}^{corr} is also calculated by Bonferroni correction (Hyvärinen, 2011). In the experiment, FPR was used to confirm the existence of clusters among languages, and FDR was used to decide which components should be clustered into existing clusters.

C Detailed Results

Table 2 shows detailed results of ICA. The distribution of similarities is illustrated in Figure 3, Figure 4, and Figure 5. The red lines in the figures represent the top 5% line of similarities.

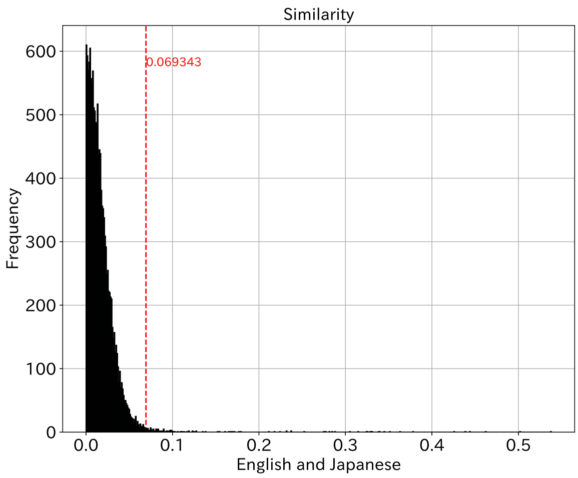

Figure 3: Similarity of Independent Components - English and Japanese.

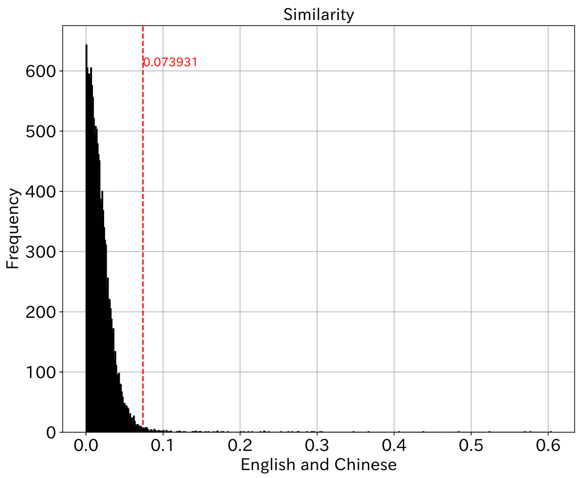

Figure 4: Similarity of Independent Components - English and Chinese.

Number of Clusters Found Number of Clustered Vectors Average Number of Vectors per Cluster	47 120 (33.90% of all vectors) 2.55
Internal Parameters	
α_{FD}^{corr} Minimum Similarity Considered Significant by FDR α_{FP}^{corr} Minimum Similarity Considered Significant by FPR	1.000000×10^{-2} 0.1110 2.754821×10^{-5} 0.1468

Table 2: Detailed results of ICA.

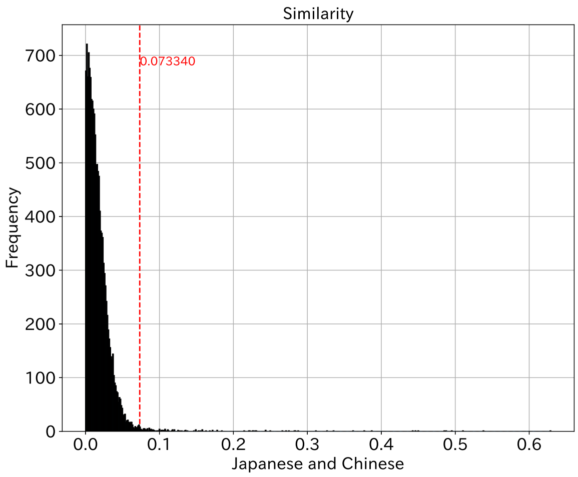

Figure 5: Similarity of Independent Components - Japanese and Chinese.

D Dictionary Statistics

The number of unique Japanese words in the English-Japanese dictionary was 21003, and the number of English words was 22531. The number of unique Chinese words in the English-Chinese dictionary was 13768, and the number of English words was 25969.

E Questionnaire Form

The following questionnaire form E, initially in Japanese, was used to conduct quantitative evaluations of semantic axes. The English translations of Japanese and Chinese words are only for explanation here and were not annotated in the actual questionnaire form.

Questionnaire Form

Below is a list of words in several languages. If you think that the English, Japanese, and Chinese words all belong to the same meaning category, check the box. For example,
en:[‘eyes’ ‘see’ ‘rib’] ja:[‘視界’(‘vision’) ‘網膜’(‘retina’) ‘凝視’(‘stare’)] zh:[‘觀看’(‘look’) ‘凝視’(‘stare’) ‘眼’(‘eye’)]
In this case, the three languages have a meaning associated with the eye, so check the box.
en:[‘deco’ ‘arts’ ‘murals’] ja:[‘礼儀’(‘courtesy’) ‘ひも’(‘string’) ‘冗長’(‘redundancy’)] zh:[‘民俗’(‘folk’) ‘漆器’(‘lacquerware’) ‘壁畫’(‘wall art’)]
In this case, because the list of Japanese words does not make sense or does not match the meaning of the other languages, do not check the box.