

How Familiar Does That Sound? Cross-Lingual Representational Similarity Analysis of Acoustic Word Embeddings

Badr M. Abdullah Iuliia Zaitova Tania Avgustinova
Bernd Möbius Dietrich Klakow

Department of Language Science and Technology (LST)
Saarland Informatics Campus, Saarland University, Germany
Corresponding author: babdullah@lsv.uni-saarland.de

Abstract

How do neural networks “perceive” speech sounds from unknown languages? Does the typological similarity between the model’s training language (L1) and an unknown language (L2) have an impact on the model representations of L2 speech signals? To answer these questions, we present a novel experimental design based on representational similarity analysis (RSA) to analyze acoustic word embeddings (AWEs)—vector representations of variable-duration spoken-word segments. First, we train monolingual AWE models on seven Indo-European languages with various degrees of typological similarity. We then employ RSA to quantify the cross-lingual similarity by simulating native and non-native spoken-word processing using AWEs. Our experiments show that typological similarity indeed affects the representational similarity of the models in our study. We further discuss the implications of our work on modeling speech processing and language similarity with neural networks.

1 Introduction

Mastering a foreign language is a process that requires (human) language learners to invest time and effort. If the foreign language (L2) is very distant from our native language (L1), not much of our prior knowledge of language processing would be relevant in the learning process. On the other hand, learning a language that is similar to our native language is much easier since our prior knowledge becomes more useful in establishing the correspondences between L1 and L2 (Ringbom, 2006). In some cases where L1 and L2 are closely related and typologically similar, it is possible for an L1 speaker to comprehend L2 linguistic expressions to a great degree without prior exposure to L2. The term *receptive multilingualism* (Zeevaert, 2007) has been coined in the sociolinguistics literature to describe this ability of a listener to comprehend utter-

ances of an unknown but related language without being able to produce it.

Human speech perception has been an active area of research in the past five decades which has produced a wealth of documented behavioral studies and experimental findings. Recently, there has been a growing scientific interest in the cognitive modeling community to leverage the recent advances in speech representation learning to formalize and test theories of speech perception using computational simulations on the one hand, and to investigate whether neural networks exhibit similar behavior to humans on the other hand (Räsänen et al., 2016; Alishahi et al., 2017; Dupoux, 2018; Scharenborg et al., 2019; Gelderloos et al., 2020; Matushevych et al., 2020b; Magnuson et al., 2020).

In the domain of modeling non-native speech perception, Schatz and Feldman (2018) have shown that a neural speech recognition system (ASR) predicts Japanese speakers’ difficulty with the English phonemic contrast /l/-/ɫ/ as well as English speakers’ difficulty with the Japanese vowel length distinction. Matushevych et al. (2021) have shown that a model of non-native spoken word processing based on neural networks predicts lexical processing difficulty of English-speaking learners of Russian. The latter model is based on word-level representations that are induced from naturalistic speech data known in the speech technology community as acoustic word embeddings (AWEs). AWE models map variable-duration spoken-word segments onto fixed-size representations in a vector space such that instances of the same word type are (ideally) projected onto the same point in space (Levin et al., 2013). In contrast to word embeddings in natural language processing (NLP), an AWE encodes information about the acoustic-phonetic and phonological structure of the word, not its semantic content.

Although the model of Matushevych et al. (2021) has shown similar effects to what has been ob-

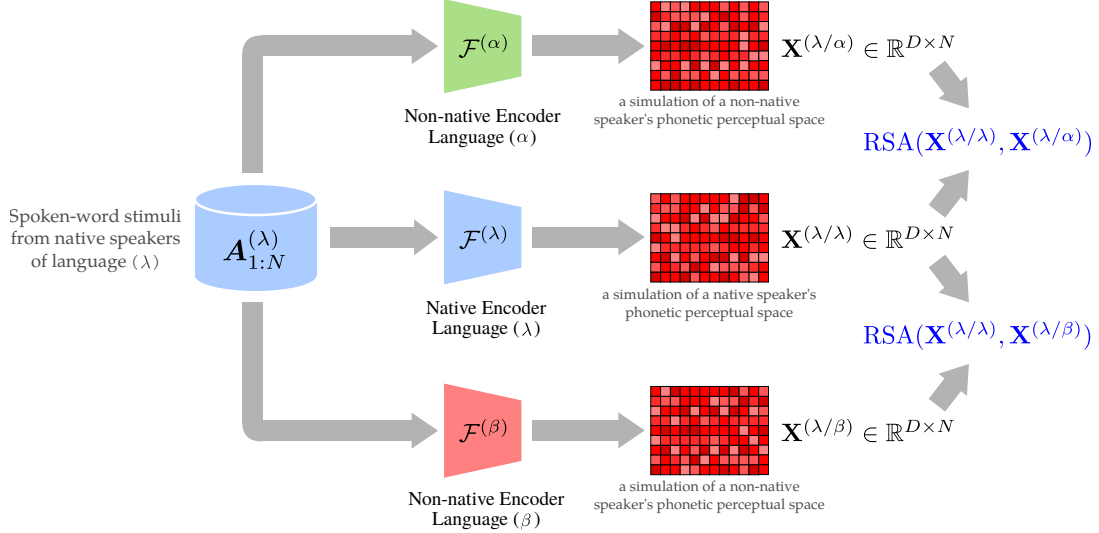


Figure 1: An illustrated example of our experimental design. A set of N spoken-word stimuli from language λ are embedded using the encoder $\mathcal{F}^{(\lambda)}$ which was trained on language λ to obtain a (native) view of the data: $\mathbf{X}^{(\lambda/\lambda)} \in \mathbb{R}^{D \times N}$. Simultaneously, the same stimuli are embedded using encoders trained on other languages, namely $\mathcal{F}^{(\alpha)}$ and $\mathcal{F}^{(\beta)}$, to obtain two different (non-native) views of the data: $\mathbf{X}^{(\lambda/\alpha)}$ and $\mathbf{X}^{(\lambda/\beta)}$. We quantify the cross-lingual similarity between two languages by measuring the association between their embedding spaces using the representational similarity analysis (RSA) framework.

served in behavioral studies (Cook et al., 2016), it remains unclear to what extent AWE models can predict a facilitatory effect of language similarity on cross-language spoken-word processing. In this paper, we present a novel experimental design to probe the *receptive multilingual* knowledge of monolingual AWE models (i.e., trained without L2 exposure) using the representational similarity analysis (RSA) framework. In a controlled experimental setup, we use AWE models to simulate spoken-word processing of native speakers of seven languages with various degrees of typological similarity. We then employ RSA to characterize how language similarity affects the emergent representations of the models when tested on the same spoken-word stimuli (see Figure 1 for an illustrated overview of our approach). Our experiments demonstrate that neural AWE models of different languages exhibit a higher degree of representational similarity if their training languages are typologically similar.

2 Proposed Methodology

A neural AWE model can be formally described as an encoder function $\mathcal{F} : \mathcal{A} \rightarrow \mathbb{R}^D$, where \mathcal{A} is the (continuous) space of acoustic sequences and D is the dimensionality of the embedding. Given an acoustic word signal represented as a temporal sequence of T acoustic events $\mathbf{A} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_T)$,

where $\mathbf{a}_t \in \mathbb{R}^k$ is a spectral vector of k coefficients, an embedding is computed as

$$\mathbf{x} = \mathcal{F}(\mathbf{A}; \boldsymbol{\theta}) \in \mathbb{R}^D \quad (1)$$

Here, $\boldsymbol{\theta}$ are the parameters of the encoder, which are learned by training the AWE model in a monolingual supervised setting. That is, the training spoken-word segments (i.e., speech intervals corresponding to spoken words) are sampled from utterances of native speakers of a single language where the word identity of each segment is known. The model is trained with an objective that maps different spoken segments of the same word type onto similar embeddings. To encourage the model to abstract away from speaker variability, the training samples are obtained from multiple speakers, while the resulting AWEs are evaluated on a held-out set of speakers.

Our research objective in this paper is to study the discrepancy of the representational geometry of two AWE encoders that are trained on different languages when tested on the same set of (monolingual) spoken stimuli. To this end, we train AWE encoders on different languages where the training data and conditions are comparable across languages with respect to size, domain, and speaker variability. We therefore have access to several encoders $\{\mathcal{F}^{(\alpha)}, \mathcal{F}^{(\beta)}, \dots, \mathcal{F}^{(\omega)}\}$, where the superscripts $\{\alpha, \beta, \dots, \omega\}$ denote the language of the

training samples.

Now consider a set of N held-out spoken-word stimuli produced by native speakers of language λ : $\mathbf{A}_{1:N}^{(\lambda)} = \{\mathbf{A}_1^{(\lambda)}, \dots, \mathbf{A}_N^{(\lambda)}\}$. First, each acoustic word stimulus in this set is mapped onto an embedding using the encoder $\mathcal{F}^{(\lambda)}$, which yields a matrix $\mathbf{X}^{(\lambda/\lambda)} \in \mathbb{R}^{D \times N}$. Since the encoder $\mathcal{F}^{(\lambda)}$ was trained on language λ , we refer to it as the *native encoder* and consider the matrix $\mathbf{X}^{(\lambda/\lambda)}$ as a simulation of a native speaker’s phonetic perceptual space. To simulate the phonetic perceptual space of a non-native speaker, say a speaker of language α , the stimuli $\mathbf{A}_{1:N}^{(\lambda)}$ are embedded using the encoder $\mathcal{F}^{(\alpha)}$, which yields a matrix $\mathbf{X}^{(\lambda/\alpha)} \in \mathbb{R}^{D \times N}$. Here, we read the superscript notation (λ/α) as word stimuli of language λ encoded by a model trained on language α . Thus, the two matrices $\mathbf{X}^{(\lambda/\lambda)}$ and $\mathbf{X}^{(\lambda/\alpha)}$ represent two different views of the same stimuli. Our main hypothesis is that the cross-lingual representational similarity between emergent embedding spaces should reflect the acoustic-phonetic and phonological similarities between the languages λ and α . That is, the more distant languages λ and α are, the more dissimilar their corresponding representation spaces are. To quantify this cross-lingual representational similarity, we use Centered Kernel Alignment (CKA) (Kornblith et al., 2019). CKA is a representation-level similarity measure that emphasizes the distributivity of information and therefore it obviates the need to establish the correspondence mapping between single neurons in the embeddings of two different models. Moreover, CKA has been shown to be invariant to orthogonal transformation and isotropic scaling which makes it suitable for our analysis when comparing different languages and learning objectives. Using CKA, we quantify the similarity between languages λ and α as

$$\text{sim}(\lambda, \alpha) := \text{CKA}(\mathbf{X}^{(\lambda/\lambda)}, \mathbf{X}^{(\lambda/\alpha)}) \quad (2)$$

Here, $\text{sim}(\lambda, \alpha) \in [0, 1]$ is a scalar that measures the correlation between the responses of the two encoders, i.e., native $\mathcal{F}^{(\lambda)}$ and non-native $\mathcal{F}^{(\alpha)}$, when tested with spoken-word stimuli $\mathbf{A}_{1:N}^{(\lambda)}$. $\text{sim}(\lambda, \alpha) = 1$ is interpreted as perfect association between the representational geometry of models trained on languages λ and α while $\text{sim}(\lambda, \alpha) = 0$ indicates that no association can be established. Likewise, we obtain another non-native view $\mathbf{X}^{(\lambda/\beta)}$ of the stimuli $\mathbf{A}_{1:N}^{(\lambda)}$ using another non-native encoder $\mathcal{F}^{(\beta)}$. We then quan-

tify the cross-lingual representational similarity between languages λ and β as

$$\text{sim}(\lambda, \beta) := \text{CKA}(\mathbf{X}^{(\lambda/\lambda)}, \mathbf{X}^{(\lambda/\beta)}) \quad (3)$$

If $\text{sim}(\lambda, \alpha) > \text{sim}(\lambda, \beta)$, then we interpret this as an indication that the native phonetic perceptual space of language λ is more similar to the non-native phonetic perceptual space of language α , compared to that of language β . Note that while $\text{CKA}(\cdot, \cdot)$ is a symmetric metric (i.e., $\text{CKA}(\mathbf{X}, \mathbf{Y}) = \text{CKA}(\mathbf{Y}, \mathbf{X})$), our established similarity metric $\text{sim}(\cdot, \cdot)$ is not symmetric (i.e., $\text{sim}(\lambda, \alpha) \neq \text{sim}(\alpha, \lambda)$). To estimate $\text{sim}(\alpha, \lambda)$, we use word stimuli of language α and collect the matrices $\mathbf{X}^{(\alpha/\alpha)}$ and $\mathbf{X}^{(\alpha/\lambda)}$. Then we compute

$$\text{sim}(\alpha, \lambda) := \text{CKA}(\mathbf{X}^{(\alpha/\alpha)}, \mathbf{X}^{(\alpha/\lambda)}) \quad (4)$$

When we apply the proposed experimental pipeline across M different languages, the effect of language similarity can be characterized by constructing a cross-lingual representational similarity matrix (xRSM) which is an asymmetric $M \times M$ matrix where each cell represents the correlation (or agreement) between two embedding spaces.

3 Acoustic Word Embedding Models

We investigate three different approaches of training AWE models that have been previously introduced in the literature. In this section, we formally describe each one of them.

3.1 Phonologically Guided Encoder

The phonologically guided encoder (PGE) is a sequence-to-sequence model in which the network is trained as a word-level acoustic model (Abdullah et al., 2021). Given an acoustic sequence \mathbf{A} and its corresponding phonological sequence $\varphi = (\varphi_1, \dots, \varphi_\tau)$,¹ the acoustic encoder \mathcal{F} is trained to take \mathbf{A} as input and produce an AWE \mathbf{x} , which is then fed into a phonological decoder \mathcal{G} whose goal is to generate the sequence φ (Fig. 2–a). The objective is to minimize a categorical cross-entropy loss at each timestep in the decoder, which is equivalent to

$$\mathcal{L} = - \sum_{i=1}^{\tau} \log \mathbf{P}_{\mathcal{G}}(\varphi_i | \varphi_{<i}, \mathbf{x}) \quad (5)$$

¹We use the phonetic transcription in IPA symbols as the phonological sequence φ .

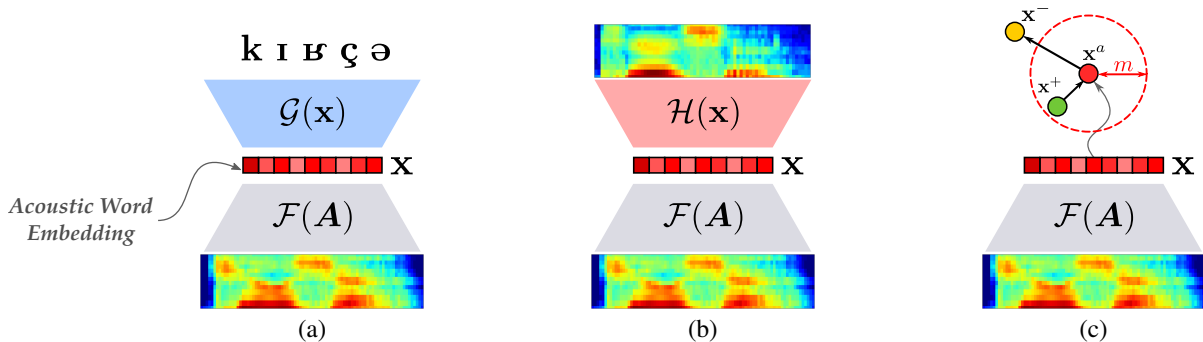


Figure 2: A visual illustration of the different learning objectives for training AWE encoders: (a) phonologically guided encoder (PGE): a sequence-to-sequence network with a phonological decoder, (b) correspondence autoencoder (CAE): a sequence-to-sequence network with an acoustic decoder, and (c) contrastive siamese encoder (CSE): a contrastive network trained via triplet margin loss.

where $\mathbf{P}_{\mathcal{G}}$ is the probability of the phone φ_i at timestep i , conditioned on the previous phone sequence $\varphi_{<i}$ and the embedding \mathbf{x} . The intuition of this objective is that spoken-word segments of the same word type would have identical phonological sequences, thus they are expected to end up nearby in the embedding space.

3.2 Correspondence Autoencoder

In the correspondence autoencoder (CAE) (Kamper, 2019), each training acoustic segment \mathbf{A} is paired with another segment that corresponds to the same word type $\mathbf{A}^+ = (\mathbf{a}_1^+, \mathbf{a}_2^+, \dots, \mathbf{a}_{T^+}^+)$. The acoustic encoder \mathcal{F} then takes \mathbf{A} as input and computes an AWE \mathbf{x} , which is then fed into an acoustic decoder \mathcal{H} that attempts to reconstruct the corresponding acoustic sequence \mathbf{A}^+ (Fig. 2–b). The objective is to minimize a vector-wise mean square error (MSE) loss at each timestep in the decoder, which is equivalent to

$$\mathcal{L} = - \sum_{i=1}^{T^+} \|\mathbf{a}_i^+ - \mathcal{H}_i(\mathbf{x})\|^2 \quad (6)$$

where \mathbf{a}_i^+ is the ground-truth spectral vector at timestep i and $\mathcal{H}_i(\mathbf{x})$ is the reconstructed spectral vector at timestep i as a function of the embedding \mathbf{x} . It has been hypothesized that learning the correspondence between two acoustic realizations of the same word type encourages the model to build up speaker-invariant word representations while preserving phonetic information (Matusevych et al., 2020a).

3.3 Contrastive Siamese Encoder

The contrastive siamese encoder (CSE) is based on a siamese network (Bromley et al., 1994) and

has been explored in the AWEs literature with different underlying architectures (Settle and Livescu, 2016; Kamper et al., 2016). This learning objective differs from the first two as it explicitly minimizes/maximizes relative distances between AWEs of the same/different word types (Fig. 2–c). First, each acoustic word instance is paired with another instance of the same word type, which yields the tuple $(\mathbf{A}, \mathbf{A}^+)$. Given the AWEs of the instances in this tuple $(\mathbf{x}^a, \mathbf{x}^+)$, the objective is then to minimize a triplet margin loss

$$\mathcal{L} = \max \left[0, m + d(\mathbf{x}^a, \mathbf{x}^+) - d(\mathbf{x}^a, \mathbf{x}^-) \right] \quad (7)$$

Here, $d: \mathbb{R}^D \times \mathbb{R}^D \rightarrow [0, 1]$ is the cosine distance and \mathbf{x}^- is an AWE that corresponds to a different word type sampled from the mini-batch such that the term $d(\mathbf{x}^a, \mathbf{x}^-)$ is minimized. This objective aims to map acoustic segments of the same word type closer in the embedding space while pushing away segments of other word types by a distance defined by the margin hyperparameter m .

4 Experiments

4.1 Experimental Data

The data in our study is drawn from the Global-Phone speech (GPS) database (Schultz et al., 2013). GPS is a multilingual read speech resource where the utterances are recorded from native speakers (with self-reported gender labels) in a controlled recording environment with minimal noise. Therefore, the recording conditions across languages are comparable which enables us to conduct a cross-linguistic comparison. We experiment with seven Indo-European languages: Czech (CZE), Polish (POL), Russian (RUS), Bulgarian (BUL), Brazilian Portuguese (POR), French (FRA), and German

(DEU). We acknowledge that our language sample is not typologically diverse. However, one of our objectives in this paper is to investigate whether the cross-lingual similarity of AWEs can predict the degree of mutual intelligibility between related languages. Therefore, we focus on the Slavic languages in this sample (CZE, POL, RUS, and BUL), which are known to be typologically similar and mutually intelligible to various degrees.

To train our AWE models, we obtain time-aligned spoken-word segments using the Montreal Forced Aligner (McAuliffe et al., 2017). Then, we sample 42 speakers of balanced gender from each language. For each language, we sample ~32k spoken-word segments that are longer than 3 phonemes in length and shorter than 1.1 seconds in duration (see Table 2 in Appendix A for word-level summary statistics of the data). For each word type, we obtain an IPA transcription using the grapheme-to-phoneme (G2P) module of the automatic speech synthesizer, *eSpeak*. Each acoustic word segment is parametrized as a sequence of 39-dimensional Mel-frequency spectral coefficients where frames are extracted over intervals of 25ms with 10ms overlap.

4.2 Architecture and Hyperparameters

Acoustic Encoder We employ a 2-layer recurrent neural network with a bidirectional Gated Recurrent Unit (BGRU) of hidden state dimension of 512, which yields a 1024-dimensional AWE.

Training Details All models in this study are trained for 100 epochs with a batch size of 256 using the ADAM optimizer (Kingma and Ba, 2015) and an initial learning rate (LR) of 0.001. The LR is reduced by a factor of 0.5 if the mean average precision (mAP) for word discrimination on the validation set does not improve for 10 epochs. The epoch with the best validation performance during training is used for evaluation on the test set.

Implementation We build our models using PyTorch (Paszke et al., 2019) and use FAISS (Johnson et al., 2017) for efficient similarity search during evaluation. Our code is publicly available on GitHub.²

4.3 Quantitative Evaluation

We evaluate the models using the standard intrinsic evaluation of AWEs: the same-different word discrimination task (Carlin et al., 2011). This task

²<https://github.com/uds-lsv/xRSA-AWEs>

Language	Encoder type		
	PGE	CAE	CSE
CZE	78.3	76.1	82.9
POL	67.6	63.5	73.8
RUS	64.3	57.7	71.0
BUL	72.1	68.9	78.4
POR	74.5	72.2	80.4
FRA	65.6	64.5	68.5
DEU	67.9	70.3	75.8

Table 1: mAP performance on evaluation sets.

aims to assess the ability of a model to determine whether or not two given speech segments correspond to the same word type, which is quantified using a retrieval metric (mAP) reported in Table 1.

5 Representational Similarity Analysis

Figure 3 shows the cross-lingual representational similarity matrices (xRSMs) across the three different models using the linear CKA similarity metric.³ Warmer colors indicate a higher similarity between two representational spaces. One can observe striking differences between the PGE-CAE models on the one hand, and the CSE model on the other hand. The PGE-CAE models yield representations that are cross-lingually more similar to each other compared to those obtained from the contrastive CSE model. For example, the highest similarity score from the PGE model is $\text{sim}(\text{RUS}, \text{BUL}) = 0.748$, which means that the representations of the Russian word stimuli from the Bulgarian model $\mathcal{F}^{(\text{BUL})}$ exhibit the highest representational similarity to the representations of the native Russian model $\mathcal{F}^{(\text{RUS})}$. On the other hand, the lowest similarity score from the PGE model is $\text{sim}(\text{POL}, \text{DEU}) = 0.622$, which shows that the German model’s view of the Polish word stimuli is the view that differs the most from the view of the native Polish model. Likewise, the highest similarity score we observe from the CAE model is $\text{sim}(\text{BUL}, \text{RUS}) = 0.763$, while the lowest similarity is $\text{sim}(\text{DEU}, \text{BUL}) = 0.649$. If we compare these scores to those of the contrastive CSE model, we observe significant cross-language differences as the highest similarity scores are $\text{sim}(\text{POL}, \text{RUS}) = \text{sim}(\text{POR}, \text{RUS}) = 0.269$, while the lowest is $\text{sim}(\text{BUL}, \text{DEU}) = 0.170$. This discrepancy between the contrastive model and the

³The xRSMs obtained using non-linear CKA with an RBF kernel are shown in Figure 6 in Appendix B, which show very similar trends to those observed in Figure 3.

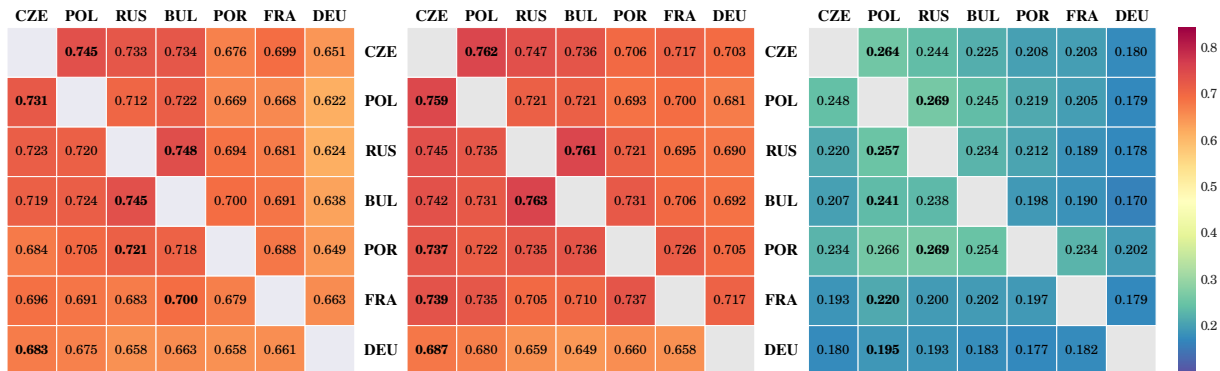


Figure 3: The cross-lingual representational similarity matrix (xRSM) for each model: PGE (Left), CAE (Middle), and CSE (Right). Each row corresponds to the language of the spoken-word stimuli while each column corresponds to the language of the encoder. Note that the matrices are not symmetric. For example in the PGE model, the cell at row CZE and column POL holds the value of $\text{sim}(\text{CZE}, \text{POL}) = \text{CKA}(\mathbf{X}^{(\text{CZE}/\text{CZE})}, \mathbf{X}^{(\text{CZE}/\text{POL})}) = 0.745$, while the cell at row POL and column CZE holds the value of $\text{sim}(\text{POL}, \text{CZE}) = \text{CKA}(\mathbf{X}^{(\text{POL}/\text{POL})}, \mathbf{X}^{(\text{POL}/\text{CZE})}) = 0.731$.

other models suggests that training AWE models with a contrastive objective hinders the ability of the encoder to learn high-level phonological abstractions and therefore contrastive encoders are more sensitive to the cross-lingual variation during inference compared to their sequence-to-sequence counterparts.

5.1 Cross-Lingual Comparison

To get further insights into how language similarity affects the model representations of non-native spoken-word segments, we apply hierarchical clustering on the xRSMs in Figure 3 using the Ward algorithm (Ward, 1963) with Euclidean distance. The result of the clustering analysis is shown in Figure 4. Surprisingly, the generated trees from the xRSMs of the PGE and CAE models are identical, which could indicate that these two models induce similar representations when trained on the same data. Diving a level deeper into the cluster structure of these two models, we observe that the Slavic languages form a pure cluster. This could be explained by the observation that some of the highest pair-wise similarity scores among the PGE models are observed between Russian and Bulgarian, i.e., $\text{sim}(\text{RUS}, \text{BUL}) = 0.748$ and $\text{sim}(\text{BUL}, \text{RUS}) = 0.745$, and Czech and Polish, i.e., $\text{sim}(\text{CZE}, \text{POL}) = 0.745$. The same trend can be observed in the CAE models: i.e., $\text{sim}(\text{RUS}, \text{BUL}) = 0.761$, $\text{sim}(\text{BUL}, \text{RUS}) = 0.763$, and $\text{sim}(\text{CZE}, \text{POL}) = 0.762$. Within the Slavic cluster, the West Slavic languages Czech and Polish are grouped together, while the Russian (East Slavic) is first grouped with Bulgarian (South Slavic) before joining the West Slavic group

to make the Slavic cluster. Although Russian and Bulgarian belong to two different Slavic branches, we observe that this pair forms the first sub-cluster in both trees, at a distance smaller than that of the West Slavic cluster (Czech and Polish). At first, this might seem surprising as we would expect the West Slavic cluster to be formed at a lower distance given the similarities among the West Slavic languages which facilitates cross-language speech comprehension, as documented by sociolinguistic studies (Golubovic, 2016). However, Russian and Bulgarian share typological features at the acoustic-phonetic and phonological levels which distinguish them from West Slavic languages.⁴ We further elaborate on these typological features in §6. Even though Portuguese was grouped with French in the cluster analysis, which one might expect given that both are Romance languages descended from Latin, it is worth pointing out that the representations of the Portuguese word stimuli from the Slavic models show a higher similarity to the representations of the native Portuguese model compared to these obtained from the French model (with only two exceptions, the Czech PGE model and Polish CAE model). We also provide an explanation of why this might be the case in §6.

The generated language cluster from the CSE model does not show any clear internal structure with respect to language groups since all cluster pairs are grouped at a much higher distance compared to the PGE and CAE models. Furthermore, the Slavic languages in the generated tree do not

⁴Note that our models do not have access to word orthography. Thus, the similarity cannot be due to the fact that Russian and Bulgarian use Cyrillic script.




Figure 4: Hierarchical clustering analysis on the cross-lingual representational similarity matrices (using linear CKA) of the three models: PGE (Left), CAE (Middle), and CSE (Right).

form a pure cluster since Portuguese was placed inside the Slavic cluster. We also do not observe the West Slavic group as in the other two models since Polish was grouped first with Russian, and not Czech. We believe that this unexpected behavior of the CSE models can be related to the previously attributed to the poor performance of the contrastive AWEs in capturing word-form similarity (Abdullah et al., 2021).

Moreover, it is interesting to observe that German seems to be the most distant language to the other languages in our study. This observation holds across all three encoders since the representations of the German word stimuli by non-native models are the most dissimilar compared to the representations of the native German model.

5.2 Cross-Model Comparison

To verify our hypothesis that the models trained with decoding objectives (PGE and CAE) learn similar representations when trained on the same data, we conduct a similarity analysis between the models in a setting where we obtain views of the spoken-word stimuli from native models, then compare these views across different learning objectives while keeping the language fixed using CKA as we do in our cross-lingual comparison. The results of this analysis is shown in Figure 5. We observe that across different languages the pairwise representational similarity of the PGE-CAE models (linear CKA = 0.721) is very high in comparison to that of PGE-CSE models (linear CKA = 0.239) and CAE-CSE models (linear CKA = 0.230).⁵ Although the non-linear CKA similarity scores tend to be higher, the general trend remains identical in both measures. This finding validates our hypothesis that the PGE and CAE models yield similar

representation spaces when trained on the same data even though their decoding objectives operate over different modalities. That is, the decoder of the PGE model aims to generate the word’s phonological structure in the form of a sequence of discrete phonological units, while the decoder of the CAE model aims to generate an instance of the same word represented as a sequence of (continuous) spectral vectors. Moreover, the sequences that these decoders aim to generate vary in length (the mean phonological sequence length is ~6 phonemes while mean spectral sequence length is ~50 vectors). Although the CAE model has no access to abstract phonological information of the word-forms it is trained on, it seems that this model learns non-trivial knowledge about word phonological structure as demonstrated by the representational similarity of its embeddings to those of a word embedding model that has access to word phonology (i.e., PGE) across different languages.

6 Discussion

6.1 Relevance to Related Work

Although the idea of representational similarity analysis (RSA) has originated in the neuroscience literature (Kriegeskorte et al., 2008), researchers in NLP and speech technology have employed a similar set of techniques to analyze emergent representations of multi-layer neural networks. For example, RSA has previously been employed to analyze the correlation between neural and symbolic representations (Chrupała and Alishahi, 2019), contextualized word representations (Abnar et al., 2019; Abdou et al., 2019; Lepori and McCoy, 2020; Wu et al., 2020), representations of self-supervised speech models (Chung et al., 2021) and visually-grounded speech models (Chrupała et al., 2020). We take the inspiration from previous work and

⁵CKA scores are averaged over languages.

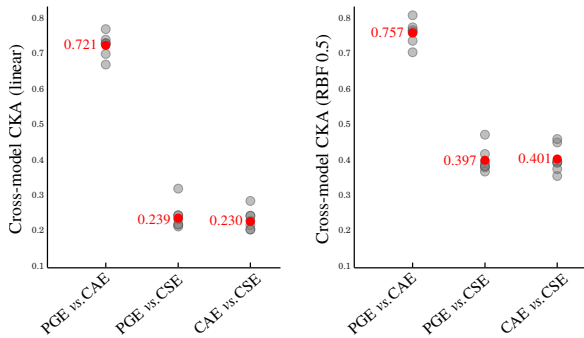


Figure 5: Cross-model CKA scores: linear CKA (left) and non-linear CKA (right). Each point in this plot is the within-language representational similarity of two models trained with different objectives when tested on the same (monolingual) word stimuli. Each red point is the average CKA score per model pair.

apply RSA in a unique setting: to analyze the impact of typological similarity between languages on the representational geometry. Our analysis focuses on neural models of spoken-word processing that are trained on naturalistic speech corpora. Our goal is two-fold: (1) to investigate whether or not neural networks exhibit predictable behavior when tested on speech from a different language (L2), and (2) to examine the extent to which the training strategy affects the emergent representations of L2 spoken-word stimuli. To the best of our knowledge, our study is the first to analyze the similarity of emergent representations in neural acoustic models from a cross-linguistic perspective.

6.2 A Typological Discussion

Given the cross-linguistic nature of our study, we choose to discuss our findings on the representational similarity analysis from a language typology point of view. Language typology is a sub-field within linguistics that is concerned with the study and categorization of the world’s languages based on their linguistic structural properties (Comrie, 1988; Croft, 2002). Since acoustic word embeddings are word-level representations induced from actual acoustic realizations of word-forms, we focus on the phonetic and phonological properties of the languages in our study. We emphasize that our goal is not to use the method we propose in this paper as a computational approach for phylogenetic reconstruction, that is, discovering the historical relations between languages. However, typological similarity usually correlates with phylogenetic distance since languages that have diverged from a common historical ancestor inherited most of their

typological features from the ancestor language (see Bjerva et al. (2019) for a discussion on how typological similarity and genetic relationships interact when analyzing neural embeddings).

Within the language sample we study in this paper, four of these languages belong to the Slavic branch of Indo-European languages. Compared to the Romance and Germanic branches of Indo-European, Slavic languages are known to be remarkably more similar to each other not only at lexical and syntactic levels, but also at the pre-lexical level (acoustic-phonetic and phonological features). These cross-linguistically shared features between Slavic languages include rich consonant inventories, phonemic iotation and complex consonant clusters. The similarities at different linguistic levels facilitate spoken intercomprehension, i.e., the ability of a listener to comprehend an utterance (to some degree) in a language that is unknown, but related to their mother tongue. Several sociolinguistic studies of mutual intelligibility have reported a higher degree of intercomprehension among speakers of Slavic languages compared to other language groups in Europe (Golubovic, 2016; Gooskens et al., 2018). On the other hand, and even though French and Portuguese are both Romance languages, they are less mutually intelligible compared to Slavic language pairs as they have diverged in their lexicons and phonological structures (Gooskens et al., 2018). This shows that cross-language speech intelligibility is not mainly driven by shared historical relationships, but by contemporary typological similarities.

Therefore, and given the documented phonological similarities between Slavic languages, it is not surprising that Slavic languages form a pure cluster in our clustering analysis over the xRSMs of two of the models we investigate. Furthermore, the grouping of Czech-Polish and Russian-Bulgarian in the Slavic cluster can be explained if we consider word-specific suprasegmental features. Besides the fact that Czech and Polish are both West Slavic languages that form a spatial continuum of language variation, both languages have fixed stress. The word stress in Czech is always on the initial syllable, while Polish has penultimate stress, that is, stress falls on the syllable preceding the last syllable of the word (Sussex and Cubberley, 2006). On the other hand, Russian and Bulgarian languages belong to different Slavic sub-groups—Russian is East Slavic while Bulgarian is South

Slavic. The representational similarity between Russian and Bulgarian can be attributed to the typological similarities between Bulgarian and East Slavic languages. In contrast to West Slavic languages, the stress in Russian and Bulgarian is free (it can occur on any syllable in the word) and movable (it moves between syllables within morphological paradigms). Slavic languages with free stress tend to have a stronger contrast between stressed and unstressed syllables, and vowels in unstressed syllables undergo a process that is known as vowel-quality alternation, or lexicalized vowel reduction (Barry and Andreeva, 2001). For example, consider the Russian word *ruka* (‘hand’). In the singular nominative case the word-form is phonetically realized as *ruká* [rʊˈka] while in the singular accusative case *ríku* [ˈrukʊ] (Sussex and Cubberley, 2006). Here, the unstressed high back vowel /u/ is realized as [ʊ] in both word-forms. Although vowel-quality alternations in Bulgarian follow different patterns, Russian and Bulgarian vowels in unstressed syllables are reduced in temporal duration and quality. Therefore, the high representational similarity between Russian and Bulgarian could be explained if we consider their typological similarities in word-specific phonetics and prosody.

The Portuguese language presents us with an interesting case study of language variation. From a phylogenetic point of view, Portuguese and French have diverged from Latin and therefore they are both categorized as Romance languages. The clustering of the xRSMs of the PGE and CAE models groups Portuguese and French together at a higher distance compared to the Slavic group, while Portuguese was grouped with Slavic languages when analyzing the representations of the contrastive CSE model. Similar to Russian and Polish, sibilant consonants (e.g., /ʃ/, /ʒ/) and the velarized (dark) *L*-sound (i.e., /ɫ/) are frequent in Portuguese. We hypothesize that contrastive training encourages the model to pay more attention to the segmental information (i.e., individual phones) in the speech signal at the expense of phonotactics (i.e., phone sequences). Given that contrastive learning is a predominant paradigm in speech representation learning (van den Oord et al., 2018; Schneider et al., 2019), we encourage further research to analyze whether or not speech processing models trained with contrastive objectives exhibit a similar behavior to that observed in human listeners and closely examine their plausibility for cognitive modeling.

6.3 Implications on Cross-Lingual Transfer Learning

A recent study on zero-resource AWEs has shown that cross-lingual transfer is more successful when the source (L1) and target (L2) languages are more related (Jacobs and Kamper, 2021). We conducted a preliminary experiment on the cross-lingual word discrimination performance of the models in our study and observed a similar effect. However, we also observed that cross-lingual transfer using the contrastive model is less effective compared to the models trained with decoding objectives. From our reported RSA analysis, we showed that the contrastive objective yields models that are cross-lingually dissimilar compared to the other objectives. Therefore, future work could investigate the degree to which our proposed RSA approach predicts the effectiveness of different models in a cross-lingual zero-shot scenario.

7 Conclusion

We presented an experimental design based on representational similarity analysis (RSA) whereby we analyzed the impact of language similarity on representational similarity of acoustic word embeddings (AWEs). Our experiments have shown that AWE models trained using decoding objectives exhibit a higher degree of representational similarity if their training languages are typologically similar languages. We discussed our findings from a typological perspective and highlighted pre-lexical features that could have an impact on the models’ representational geometry. Our findings provide evidence that AWE models can predict the facilitatory effect of language similarity on cross-language speech perception and complement ongoing efforts in the community to assess their utility in cognitive modeling. Our work can be further extended by considering speech segments below the word-level (e.g., syllables, phonemes), incorporating semantic representations into the learning procedure, and investigating other neural architectures.

Acknowledgements

We thank the anonymous reviewers for their constructive comments and insightful feedback. We further extend our gratitude to Miriam Schulz and Marius Mosbach for proofreading the paper. This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Project ID 232722074 – SFB 1102.

References

- Mostafa Abdou, Artur Kulmizev, Felix Hill, Daniel M. Low, and Anders Søgaard. 2019. [Higher-order comparisons of sentence encoder representations](#). 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)*, pages 5838–5845, Hong Kong, China. Association for Computational Linguistics.
- Badr M. Abdullah, Marius Mosbach, Iuliia Zaitova, Bernd Möbius, and Dietrich Klakow. 2021. [Do Acoustic Word Embeddings Capture Phonological Similarity? An Empirical Study](#). In *Proc. Interspeech 2021*, pages 4194–4198.
- Samira Abnar, Lisa Beinborn, Rochelle Choenni, and Willem Zuidema. 2019. [Blackbox meets blackbox: Representational similarity & stability analysis of neural language models and brains](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 191–203, Florence, Italy. Association for Computational Linguistics.
- Afra Alishahi, Marie Barking, and Grzegorz Chrupała. 2017. [Encoding of phonology in a recurrent neural model of grounded speech](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 368–378, Vancouver, Canada. Association for Computational Linguistics.
- William Barry and Bistra Andreeva. 2001. Cross-language similarities and differences in spontaneous speech patterns. *Journal of the International Phonetic Association*, 31(1):51–66.
- Johannes Bjerva, Robert Östling, Maria Han Veiga, Jörg Tiedemann, and Isabelle Augenstein. 2019. What do language representations really represent? *Computational Linguistics*, 45(2):381–389.
- Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1994. Signature verification using a "siamese" time delay neural network. In *Proc. NIPS*.
- Michael A Carlin, Samuel Thomas, Aren Jansen, and Hynek Hermansky. 2011. Rapid evaluation of speech representations for spoken term discovery. In *Proc. Interspeech*.
- Grzegorz Chrupała and Afra Alishahi. 2019. [Correlating neural and symbolic representations of language](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2952–2962, Florence, Italy. Association for Computational Linguistics.
- Grzegorz Chrupała, Bertrand Higy, and Afra Alishahi. 2020. [Analyzing analytical methods: The case of phonology in neural models of spoken language](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4146–4156, Online. Association for Computational Linguistics.
- Yu-An Chung, Yonatan Belinkov, and James Glass. 2021. Similarity analysis of self-supervised speech representations. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3040–3044. IEEE.
- Bernard Comrie. 1988. Linguistic typology. *Annual Review of Anthropology*, 17:145–159.
- Svetlana V Cook, Nick B Pandža, Alia K Lancaster, and Kira Gor. 2016. Fuzzy nonnative phonolexical representations lead to fuzzy form-to-meaning mappings. *Frontiers in Psychology*, 7:1345.
- William Croft. 2002. *Typology and Universals*. Cambridge University Press.
- Emmanuel Dupoux. 2018. [Cognitive science in the era of artificial intelligence: A roadmap for reverse-engineering the infant language-learner](#). *Cognition*, 173:43–59.
- Lieke Gelderloos, Grzegorz Chrupała, and Afra Alishahi. 2020. [Learning to understand child-directed and adult-directed speech](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1–6, Online. Association for Computational Linguistics.
- Jelena Golubovic. 2016. Mutual intelligibility in the Slavic language area. *Groningen: Center for Language and Cognition*.
- Charlotte Gooskens, Vincent J van Heuven, Jelena Golubović, Anja Schüppert, Femke Swarte, and Stefanie Voigt. 2018. Mutual intelligibility between closely related languages in Europe. *International Journal of Multilingualism*, 15(2):169–193.
- Christiaan Jacobs and Herman Kamper. 2021. [Multilingual Transfer of Acoustic Word Embeddings Improves When Training on Languages Related to the Target Zero-Resource Language](#). In *Proc. Interspeech*, pages 1549–1553.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*.
- H. Kamper, W. Wang, and Karen Livescu. 2016. Deep convolutional acoustic word embeddings using word-pair side information. In *Proc. ICASSP*.
- Herman Kamper. 2019. Truly unsupervised acoustic word embeddings using weak top-down constraints in encoder-decoder models. In *Proc. ICASSP*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc. ICLR*.

- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pages 3519–3529. PMLR.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bاندettini. 2008. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4.
- Michael Lepori and R. Thomas McCoy. 2020. [Picking BERT’s brain: Probing for linguistic dependencies in contextualized embeddings using representational similarity analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3637–3651, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Keith Levin, Katharine Henry, Aren Jansen, and Karen Livescu. 2013. Fixed-dimensional acoustic embeddings of variable-length segments in low-resource settings. In *Proc. ASRU*.
- James S Magnuson, Heejo You, Sahil Luthra, Monica Li, Hosung Nam, Monty Escabi, Kevin Brown, Paul D Allopenna, Rachel M Theodore, Nicholas Monto, et al. 2020. Earshot: A minimal neural network model of incremental human speech recognition. *Cognitive science*, 44(4):e12823.
- Yevgen Matuselych, Herman Kamper, and Sharon Goldwater. 2020a. Analyzing autoencoder-based acoustic word embeddings. In *Bridging AI and Cognitive Science Workshop@ ICLR 2020*.
- Yevgen Matuselych, Herman Kamper, Thomas Schatz, Naomi Feldman, and Sharon Goldwater. 2021. [A phonetic model of non-native spoken word processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1480–1490, Online. Association for Computational Linguistics.
- Yevgen Matuselych, Thomas Schatz, Herman Kamper, Naomi Feldman, and Sharon Goldwater. 2020b. Evaluating computational models of infant phonetic learning across languages. In *Proc. CogSci*.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal Forced Aligner: Trainable text-speech alignment using Kaldi. In *Interspeech*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Proc. NeuRIPS*.
- Okko Räsänen, Tasha Nagamine, and Nima Mesgarani. 2016. Analyzing distributional learning of phonemic categories in unsupervised deep neural networks. *Annual Conference of the Cognitive Science Society (CogSci)*. *Cognitive Science Society (U.S.) Conference*, 2016:1757–1762.
- Håkan Ringbom. 2006. *Cross-linguistic similarity in foreign language learning*. Multilingual Matters.
- Odette Scharenborg, Nikki van der Gouw, Martha Larson, and Elena Marchiori. 2019. The representation of speech in deep neural networks. In *International Conference on Multimedia Modeling*, pages 194–205. Springer.
- Thomas Schatz and Naomi H Feldman. 2018. Neural network vs. hmm speech recognition systems as models of human cross-linguistic phonetic perception. In *Proceedings of the conference on cognitive computational neuroscience*.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019. [wav2vec: Unsupervised Pre-Training for Speech Recognition](#). In *Proc. Interspeech 2019*, pages 3465–3469.
- Tanja Schultz, Ngoc Thang Vu, and Tim Schlippe. 2013. GlobalPhone: A multilingual text and speech database in 20 languages. In *Proc. ICASSP*.
- Shane Settle and Karen Livescu. 2016. Discriminative acoustic word embeddings: Recurrent neural network-based approaches. In *Proc. IEEE Spoken Language Technology Workshop (SLT)*.
- Roland Sussex and Paul Cubberley. 2006. *The slavic languages*. Cambridge University Press.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *ArXiv*, abs/1807.03748.
- Joe H Ward. 1963. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244.
- John Wu, Yonatan Belinkov, Hassan Sajjad, Nadir Durani, Fahim Dalvi, and James Glass. 2020. Similarity analysis of contextual word representation models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4638–4655.
- Ludger Zeevaert. 2007. Receptive multilingualism. *JD Ten Thijs & L. Zeevaert (Eds.)*, 6:103–137.

Appendices

A Experimental Data Statistics

Table 2 shows a word-level summary statistics of our experimental data extracted from the Global-Phone speech daabase (GPS).

B Representational Similarity with Non-Linear CKA

We provide the cross-lingual representational similarity matrices (xRSMs) constructed by applying the non-linear CKA measure with a radial basis function (RBF) in Figure 6 below. We observe similar trends to those presented in Figure 3.

Applying hierarchical clustering on the xRSMs in Figure 6 yields the language clusters shown in Figure 7. We observe that the clusters are identical to those shown in Figure 4, except for the CSE model where Portuguese is no longer in the Slavic group. However, Portuguese remains the non-Slavic language that is the most similar to Slavic languages.

Lang.	Language group	#Train spkrs	#Train samples	#Eval samples	#Phones/word (mean \pm SD)	Word dur. (sec) (mean \pm SD)	Token-Type ratio
CZE	West Slavic	42	32,063	9,228	6.77 \pm 2.28	0.492 \pm 0.176	0.175
POL	West Slavic	42	31,507	9,709	6.77 \pm 2.27	0.488 \pm 0.175	0.192
RUS	East Slavic	42	31,892	9,005	7.56 \pm 2.77	0.496 \pm 0.163	0.223
BUL	South Slavic	42	31,866	9,063	7.24 \pm 2.60	0.510 \pm 0.171	0.179
POR	Romance	42	32,164	9,393	6.95 \pm 2.36	0.526 \pm 0.190	0.134
FRA	Romance	42	32,497	9,656	6.24 \pm 1.95	0.496 \pm 0.163	0.167
DEU	Germanic	42	32,162	9,865	6.57 \pm 2.47	0.435 \pm 0.178	0.150

Table 2: Summary statistics of our experimental data.

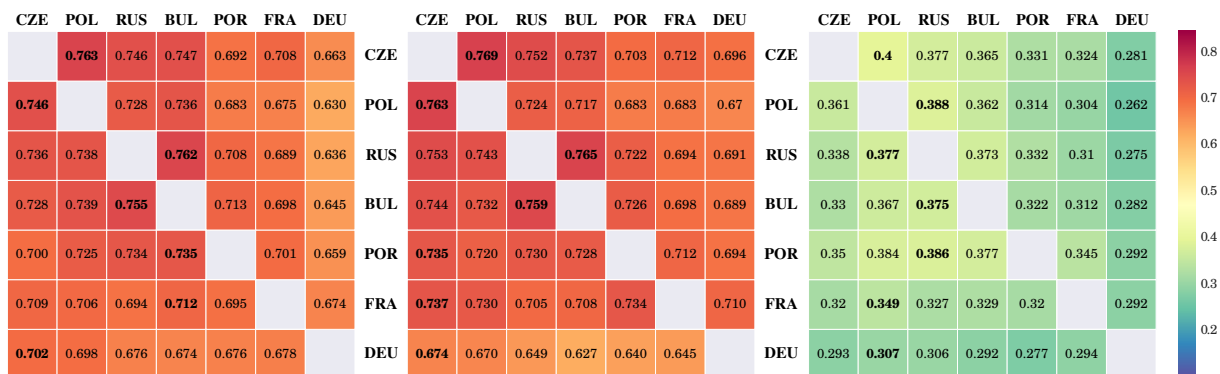


Figure 6: The cross-lingual representational similarity matrix (xRSM) for each model: PGE (Left), CAE (Middle), and CSE (Right), obtained using the non-linear CKA measure. Each row corresponds to the language of the spoken-word stimuli while each column corresponds to the language of the encoder.

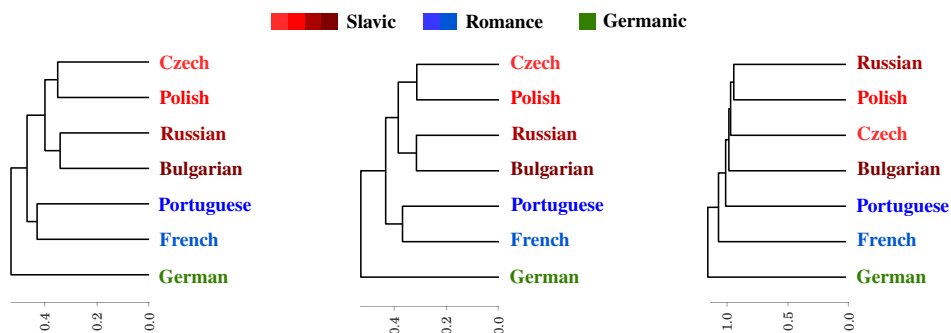


Figure 7: Hierarchical clustering analysis on the cross-lingual representational similarity matrices (using non-linear CKA) of the three models: PGE (Left), CAE (Middle), and CSE (Right).