



This is a repository copy of *Improving acoustic word embeddings through correspondence training of self-supervised speech representations*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/229425/>

Version: Published Version

Proceedings Paper:

Meghanani, A. orcid.org/0000-0002-0811-274X and Hain, T. orcid.org/0000-0003-0939-3464 (2024) Improving acoustic word embeddings through correspondence training of self-supervised speech representations. In: Graham, Y. and Purver, M., (eds.) Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers). 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2024), 17-22 Mar 2024, St. Julian's, Malta. Association for Computational Linguistics , pp. 1959-1967. ISBN 9798891760882

<https://doi.org/10.18653/v1/2024.eacl-long.118>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Improving Acoustic Word Embeddings through Correspondence Training of Self-supervised Speech Representations

Amit Meghanani, Thomas Hain

Speech and Hearing Research Group

Department of Computer Science, The University of Sheffield, United Kingdom

{ameghanani1,t.hain}@sheffield.ac.uk

Abstract

Acoustic word embeddings (AWEs) are vector representations of spoken words. An effective method for obtaining AWEs is the Correspondence Auto-Encoder (CAE). In the past, the CAE method has been associated with traditional MFCC features. Representations obtained from self-supervised learning (SSL)-based speech models such as HuBERT, Wav2vec2, etc., are outperforming MFCC in many downstream tasks. However, they have not been well studied in the context of learning AWEs. This work explores the effectiveness of CAE with SSL-based speech representations to obtain improved AWEs. Additionally, the capabilities of SSL-based speech models are explored in cross-lingual scenarios for obtaining AWEs. Experiments are conducted on five languages: Polish, Portuguese, Spanish, French, and English. HuBERT-based CAE model achieves the best results for word discrimination in all languages, despite HuBERT being pre-trained on English only. Also, the HuBERT-based CAE model works well in cross-lingual settings. It outperforms MFCC-based CAE models trained on the target languages when trained on one source language and tested on target languages.

1 Introduction

Self-supervised learning (SSL)-based speech representations are becoming popular in speech processing and producing state-of-the-art results in many downstream tasks such as automatic speech recognition, speaker verification, keyword spotting, voice conversion, etc (Yang et al., 2021). These representations are obtained using self-supervised learning on large amounts of unlabelled speech data. Wav2vec2 (Baevski et al., 2020), HuBERT (Hsu et al., 2021), and WavLM (Chen et al., 2022) are a few examples of such SSL-based speech models. However, representations obtained from these models have not been extensively explored

in the context of learning acoustic word embeddings (AWEs). AWEs are fixed-dimensional vector representations of spoken words that find applications in various downstream tasks, such as query-by-example search (Settle et al., 2017; Yuan et al., 2018; Hu et al., 2021), keyword spotting (Barakat et al., 2011), providing clues for human lexical processing (Matusevych et al., 2020), hate speech detection in low resource settings (Jacobs et al., 2023), etc.

Recently, the work (Sanabria et al., 2023) proposed extracting AWEs from SSL-based speech representations using a mean pooling mechanism. The authors suggest that SSL-based speech representations, which are contextualized, can be effectively converted into AWEs using a straightforward pooling mechanism. On the other hand, Correspondence Auto-Encoder (CAE) based training strategies for AWEs (Kamper, 2019) using MFCC (Davis and Mermelstein, 1980) features are shown to be promising in the literature. Correspondence training involves an auto-encoder where a spoken word serves as the input to the encoder, and the target output of the decoder is a different instance of the same spoken word. This approach helps to preserve acoustic-phonetic information while filtering out unnecessary details such as speaker, acoustic environment, and duration, etc. Both encoder and decoder are typically recurrent neural networks (RNNs). More details about the model will be presented in Sec. 2 and Sec. 4.2. Correspondence training has also been explored in the work (Meghanani and Hain, 2024) to improve content representations of SSL-based speech models.

The work (Lin et al., 2023) uses a Correspondence Transformer Encoder (CTE) for obtaining robust AWEs, trained from scratch and a large-scale unlabelled speech corpus. In contrast, in this work, pre-trained SSL speech models are coupled with a simple RNN based auto-encoder for correspondence training to obtain robust AWEs.

This work attempts to use the correspondence training of auto-encoder to obtain the AWEs by leveraging SSL-based speech representations instead of MFCC features as input features to the CAE model. Further, cross-lingual capabilities are also examined for SSL-based AWEs trained with CAE method. The SSL models (HuBERT, Wav2vec2, and WavLM) used in this work are pre-trained on English data. However, it has been demonstrated that these models work well as feature extractors for the all the languages considered in this study. The performances on the word-discrimination task for all the languages (Polish, Portuguese, Spanish, and French) are as good as on the English language (Sec. 5). A detailed analysis is also conducted to assess the importance of contextual information in spoken words by comparing feature extraction with and without context. For this work, we obtain spoken words for all five languages from the subset of Multilingual LibriSpeech (MLS) (Pratap et al., 2020) dataset. The derived dataset consists of five languages (Polish, Portuguese, Spanish, French, English) with start and end timestamps of spoken words¹. We chose MLS dataset for our experiments as many previous works (Matusevych et al., 2020; Abdullah et al., 2021a, 2022, 2021b; Kamper et al., 2021) on AWEs rely on GobaPhone (Schultz et al., 2013) dataset, which is not freely available.

The main contributions of this work are as follows:

1. Utilizing corresponding training with SSL-based speech representations to obtain highly discriminative AWEs.
2. Showing effectiveness of SSL models, pre-trained only on English, as feature extractors in cross-lingual scenarios for obtaining high-quality AWEs.
3. Quantitatively demonstrating that incorporating the context of the spoken word in SSL-based speech representations leads to the production of more robust AWEs.

The rest of the paper is as follows: Sec. 2 describes the correspondence auto-encoder methodology to obtain AWEs; Sec. 3 describes the data preparation and data statistics; Sec. 4 describes the details of the experiments; Sec. 5 describes the results and analysis; Sec. 6 concludes the work with possible future directions. Sec. 7 describes the limitations of the work.

¹https://github.com/Trikaldarshi/SSL_AWE




Figure 1: CAE-RNN training setup for extracting AWEs (Kamper, 2019).

2 Methodology

Correspondence auto-encoder is trained with input as a spoken word and target output as a different instance of the same spoken word. Typically, Recurrent Neural Network (RNN) based encoder and decoder are used, hence the model is referred to as CAE-RNN. The rationale behind this training method is that CAE-RNN will preserve only the acoustic-phonetic information and filter out the other unnecessary information factors such as speaker, duration, acoustic environment, etc (Kamper, 2019). Fig. 1 shows the CAE-RNN model setup. The input to the ENC is a sequence of acoustic feature vectors ($X = X_1, X_2, \dots, X_m$) of a spoken word. The target output is the sequence of acoustic feature vectors of the different instance of the same spoken word ($X' = X'_1, X'_2, \dots, X'_n$). The encoder produces the AWE (e) of the spoken word X , which is then fed to the decoder to reconstruct X' . The output of the decoder is represented as $Y = Y_1, Y_2, \dots, Y_n$. The mean squared loss function for a single training pair (X, X') can be described as following:

$$L = \sum_{k=1}^n \|X'_k - Y_k\|^2 \quad (1)$$

, where $X' = X'_1, X'_2, \dots, X'_n$ is the target output and $Y = Y_1, Y_2, \dots, Y_n$ is the output of the decoder as shown in Fig. 1.

3 Data Preparation

The Multilingual LibriSpeech (MLS) dataset (Pratap et al., 2020) is utilized to obtain spoken words. Five languages, namely Polish, Portuguese, Spanish, French, and English, are selected from MLS. For each language, approximately 25,000 utterances are selected for the training set, 500 for the development set, and 500 for the test set. These selected utterances are force-aligned to obtain the spoken word boundaries using the Montreal Forced

Data	Polish			Portuguese			Spanish			French			English		
Statistics	Train	Dev	Test	Train	Dev	Test	Train	Dev	Test	Train	Dev	Test	Train	Dev	Test
# Spoken Words	104448	4595	4563	117820	4964	4659	82258	3721	3601	84267	3004	3114	92352	3147	3192
# Unique Spoken Words	9346	3818	3887	9785	3696	3539	7085	2678	2769	7221	2394	2433	7157	2448	2527
# Speaker	11	4	4	42	10	10	79	20	20	120	18	18	182	42	41
Total Duration (hours)	18.5	0.8	0.8	21.9	0.93	0.87	14.7	0.67	0.65	14.9	0.55	0.56	16.6	0.56	0.57

Table 1: A summary of the data statistics for all five languages across the train, dev, and test splits

Aligner toolkit (McAuliffe et al., 2017). Only spoken words with a duration of 0.5 seconds or longer are included in the derived dataset, following the standard practice in the literature (He et al., 2017). Spoken words with a frequency greater than 50 or less than 5 are excluded from the derived dataset. Table 1 presents a summary of the statistics for the final extracted dataset, encompassing all five languages. The speakers across different sets are non-overlapping, which is a desirable characteristic for evaluating AWEs as they should be robust to speaker variations. Polish language had limited available data and consequently has the fewest number of speakers, while English has the highest number. The duration in Table 1 represents the total time duration of spoken words across the different sets.

4 Experimental Setup

Experiments are conducted on all the five languages with SSL-based speech representations as input features extracted from Wav2vec2, HuBERT, and WavLM. Experiments are also conducted with MFCC as input features. First, the feature extraction methods for various SSL-based speech representations and MFCCs are described. Then, the configuration of the CAE-RNN model is explained, along with the mean pooling baseline (Sanabria et al., 2023) and the AE-RNN method (without correspondence training), for comparison. Next, the word discrimination task is described, which is used for evaluating the quality of the extracted AWEs. Finally, the training details of the CAE-RNN and other models are provided.

4.1 Feature Extraction

4.1.1 SSL-based Speech Representations

SSL models are pre-trained on large amount of unlabelled speech data. The task defined for the pre-training is known as the pretext task. Each model differs based on how the pretext task was defined, the data used for pre-training, and the model architecture. In this work, the “BASE” architectures of Wav2vec2, HuBERT, and WavLM model

² (all with $\approx 95\text{M}$ parameters) are used for feature extraction. All these models are pre-trained on 960 hours of LibriSpeech data (Panayotov et al., 2015). A “BASE” architecture typically has a multi-layer CNN-based feature encoder followed by 12 Transformer layers. In this work, representations from each model are extracted from the final (i.e. 12th) Transformer layer. For all the above mentioned SSL models, 768-dimensional feature vectors are obtained for each spoken word at a framerate of 20 ms.

SSL-based speech representations are extracted in two different ways: the first one is extracted using the context around the spoken word, and the other one is extracted without the context, as described here:

1. **With context:** In this case, first the SSL-based speech representations of the entire spoken utterance are computed and then the time boundaries of the spoken word is used to get the representations of the segment belonging to the spoken word. This ensures that the extracted representations capture the context around the spoken word as the entire utterance is processed by the SSL model. Let us assume U represents an utterance and X represents a spoken word instance present in the utterance U with start and end timestamps denoted as t_1 and t_2 . If f represents the SSL model, then the SSL-based speech representation for the entire utterance is computed $Z = f(U)$. Then the speech representations for the spoken word X will be $Z_{t_1:t_2}$.
2. **Without context:** In this case, no context is considered and SSL-based speech representations are extracted by inputting only speech segments belonging to the spoken words to the SSL models. Hence, in this case, the SSL-based speech representation for the spoken word X will be $Z = f(U_{t_1:t_2})$.

²<https://github.com/pytorch/fairseq>

4.1.2 MFCC Features

For each spoken word, 20-dimensional MFCC features are extracted with 30 ms window size and 20 ms shift along with delta and delta-delta features, which leads to 60-dimensional MFCC feature vectors.

4.2 Model Details

A 4-layer Bidirectional GRU with a hidden dimension of 256 is used for both the encoder and decoder in the CAE-RNN model. Dropout rate is set to 0.2. The final hidden state of the encoder-GRU is fed to a fully connected layer to obtain 128-dimensional AWE (e) as shown in Fig. 1. This embedding is then fed to the decoder at each time step as input to the decoder (Kamper, 2019). The output of the decoder is then fed to a fully connected layer to produce the target output.

A regular auto-encoder RNN (AE-RNN) model is also used as one of the baselines with similar configurations. AE-RNN model is an auto-encoder model where input and target output is exactly the same spoken word, i.e. input-output training pair is (X, X) . A mean pooling model is also used as baseline (Sanabria et al., 2023), which does not require any training. This method computes the mean of the SSL-based speech representations to get the 768-dimensional AWE of a spoken word.

4.3 Word Discrimination Task

To evaluate the AWEs, the same-different word-discrimination task is used (Kamper et al., 2015; Carlin et al., 2011). First, all possible spoken word pairs are generated. For example, if there are total N spoken words, then the total generated spoken word pairs for comparison will be $\binom{N}{2} = \frac{N(N-1)}{2}$. After that, the cosine distance between the AWEs of these pairs are computed and compared with a threshold to decide whether the spoken words are same or different. The average precision (AP) is calculated by varying all the possible threshold values, which is the area under the precision-recall curve. AP is reported for the same-different word discrimination task. Word-discrimination task is applied on the test set, which has unseen speakers during training. Also, a subset of the test set is created for all five languages in such a way that none of the words in the subset are encountered during training. This particular subset is referred to as **test'**. The word-discrimination task is also conducted on **test'**. The total number of generated

Spoken Word Pairs (in million)	Polish	Portuguese	Spanish	French	English
test	10.4 M	10.8 M	6.4 M	4.8 M	5.1 M
test'	4.7	2.6 M	1.8 M	1.3 M	1.4 M

Table 2: The total number of spoken word pairs generated for the **test** and **test'** sets.

spoken word pairs for both the test and **test'** sets is described in Table 2 for all languages.

4.4 Training Details

The total number of generated correspondence training pairs (X, X') for each language is as follows: 9,55,106 for Polish, 11,63,468 for Portuguese, 7,80,197 for Spanish, 7,95,613 for French, and 9,72,532 for English. The remaining training details are as follows for various inputs:

- **SSL-based Speech Representations as Input:** CAE-RNN models are trained for 30 epochs, using a learning rate of 0.0001 with Adam optimizer and a batch size of 512. In each run, the model with the best performance on the development set in terms of word-discrimination is selected as the final model for evaluation on the test set. AE-RNN models are trained for 50 epochs, keeping all other parameters same as mentioned above for the CAE-RNN models.
- **MFCC as Input:** Both AE-RNN and CAE-RNN models with MFCC as inputs are trained for 100 epochs, using a learning rate of 0.0001 with Adam optimizer. The batch size for the AE-RNN model was chosen as 64, while for the CAE-RNN model it was set to 256 based on preliminary experiments for better convergence. Similarly to the previous case, the model with the best performance on the development set in terms of word-discrimination is selected as the final model for evaluation on the test set.

Model	Polish	Portuguese	Spanish	French	English
AE-RNN	0.20	0.10	0.17	0.01	0.01
CAE-RNN	0.56	0.41	0.57	0.43	0.24

Table 3: AP on the **test** set for word-discrimination task using MFCCs as input features for AE-RNN and CAE-RNN models in different languages.

Model	Polish	Portuguese	Spanish	French	English
AE-RNN	0.21	0.10	0.24	0.03	0.01
CAE-RNN	0.54	0.47	0.63	0.57	0.33

Table 4: AP on the **test'** set for word-discrimination task using MFCCs as input features for AE-RNN and CAE-RNN models in different languages.

5 Results and Analysis

Table 3 shows the baseline results with MFCC features as input for the AE-RNN and CAE-RNN models. This demonstrates the effectiveness of the CAE-RNN model over the AE-RNN model for the word-discrimination task, as the CAE-RNN consistently outperforms the AE-RNN for all languages. Table 4 presents the results for the derived subset of the test set (**test'**) with similar trends. It is worth noting that the AP on the **test'** set is relatively better than that of the original test set in most cases. This is likely due to the fact that the number of spoken word pairs generated for the evaluation on the **test'** is significantly fewer compared to the original test set, as mentioned in Table 2.

Table 5 displays the results obtained from using various SSL-based speech representations (Wav2vec2, WavLM, and HuBERT) as input features, combined with different AWE extraction methods (mean pooling, CAE-RNN, and AE-RNN). The results presented in Table 5 represent the AP for the word-discrimination task on the test set, employing different SSL-based speech representation feature extraction setups ('with context' and 'without context'). From Table 5, it is evident that the AWEs derived 'with context' exhibit greater robustness. The AP on the test set for all languages is significantly better when utilizing SSL-based speech representations 'with context' compared to the feature extraction 'without context'. As shown in Table 5, the CAE-RNN model demonstrates superior performance when using SSL-based speech representations as input features compared to the MFCC-based baseline model (Table 3) across all languages. Furthermore, Table 5 provides a comparison of the CAE-RNN model with other baseline models (mean pooling and AE-RNN) when utilizing SSL-based speech representations as input features. CAE-RNN consistently outperforms both the AE-RNN and mean pooling methods for all languages and SSL models. Another advantage of the CAE-RNN model over mean pooling is that the AWEs obtained from CAE-RNN have a dimension of 128, while mean pooling-

based AWEs are 768-dimensional. Based on the results presented in Table 5, it is evident that the HuBERT features consistently achieve the best performance across all configurations and languages. Specifically, when using the CAE-RNN method for AWE extraction and SSL-based speech representations extracted 'with context', the HuBERT achieves the highest AP on the test set: 0.90 for Polish, 0.88 for Portuguese, 0.95 for Spanish, 0.74 for French, and 0.86 for English. The performance order can be sorted as HuBERT > Wav2vec2 > WavLM > MFCCs when using the CAE-RNN-based AWE model and SSL-based speech representations extracted 'with context'.

Table 6 presents the results for the **test'** set, which includes unseen words and speakers. The models exhibit similar trends in performance in this case as well. This provides evidence that the proposed methodology performs equally well on unseen words. One interesting finding is that the SSL-based speech representations considered in this work were pre-trained solely on English language. Despite this, they are capable of generating meaningful features for other languages, resulting in good performance as demonstrated in Table 5 and 6 for the word-discrimination task.

5.1 Cross-lingual Analysis

To assess the effectiveness of SSL speech representation-based CAE-RNN models in cross-lingual settings, a CAE-RNN model trained on one source language (English in this case) is evaluated on four different target languages. This evaluation can be considered a 'zero-shot' evaluation, as no training data from the target languages is required. Table 7 displays the results in terms of AP for the word-discrimination task on the test set and **test'** set for the four target languages (Polish, Portuguese, Spanish, and French). In this scenario as well, the HuBERT-based CAE-RNN model achieves the best performance across all languages, except for French where Wav2vec2 performs the best. It is worth noting that the CAE-RNN model in the 'zero-shot' setting outperforms the mean pooling method (Table 5 and 6) (Sanabria et al., 2023) and the CAE-RNN model trained on the target language with MFCC features (Table 3 and 4). The mean pooling method (Sanabria et al., 2023) can be considered a 'zero-shot' AWE extraction method, as it does not involve additional training on top of the pre-trained SSL models. In a 'zero-shot' setup for target languages, using a CAE-RNN trained on a

AWE Extraction Method	Input Features	Polish		Portuguese		Spanish		French		English	
		with context	without context	with context	without context	with context	without context	with context	without context	with context	without context
Mean Pooling	Wav2vec2	0.01	0.00	0.00	0.00	0.02	0.00	0.03	0.01	0.02	0.00
	WavLM	0.07	0.00	0.01	0.00	0.03	0.00	0.05	0.02	0.07	0.03
	HuBERT	0.17	0.33	0.10	0.15	0.14	0.32	0.22	0.31	0.15	0.24
AE-RNN	Wav2vec2	0.10	0.00	0.08	0.00	0.15	0.01	0.13	0.01	0.07	0.00
	WavLM	0.34	0.14	0.21	0.11	0.43	0.20	0.34	0.23	0.26	0.17
	HuBERT	0.44	0.40	0.36	0.27	0.58	0.52	0.45	0.40	0.36	0.34
CAE-RNN	Wav2vec2	0.86	0.71	0.86	0.63	0.93	0.79	0.71	0.61	0.82	0.52
	WavLM	0.86	0.72	0.76	0.63	0.92	0.85	0.70	0.61	0.66	0.51
	HuBERT	0.90	0.82	0.88	0.71	0.95	0.89	0.74	0.65	0.86	0.65

Table 5: AP scores for the word-discrimination task on the **test** set using SSL-based speech representations as input features for all five languages. AWE extraction methods include mean pooling (Sanabria et al., 2023), AE-RNN, and CAE-RNN.

AWE Extraction Method	Input Features	Polish		Portuguese		Spanish		French		English	
		with context	without context	with context	without context	with context	without context	with context	without context	with context	without context
Mean Pooling	Wav2vec2	0.01	0.01	0.01	0.00	0.04	0.01	0.06	0.01	0.03	0.00
	WavLM	0.05	0.00	0.02	0.01	0.05	0.01	0.10	0.05	0.09	0.04
	HuBERT	0.18	0.31	0.11	0.17	0.19	0.40	0.30	0.40	0.18	0.29
AE-RNN	Wav2vec2	0.09	0.01	0.09	0.01	0.24	0.02	0.21	0.03	0.08	0.01
	WavLM	0.33	0.12	0.25	0.12	0.55	0.28	0.44	0.33	0.29	0.23
	HuBERT	0.44	0.41	0.41	0.31	0.69	0.64	0.55	0.50	0.43	0.41
CAE-RNN	Wav2vec2	0.87	0.72	0.90	0.65	0.96	0.85	0.84	0.73	0.89	0.67
	WavLM	0.85	0.72	0.81	0.68	0.95	0.90	0.83	0.75	0.75	0.62
	HuBERT	0.90	0.83	0.91	0.75	0.97	0.93	0.86	0.81	0.93	0.75

Table 6: AP scores for the word-discrimination task on the **test'** set using SSL-based speech representations as input features for all five languages. AWE extraction methods include mean pooling (Sanabria et al., 2023), AE-RNN, and CAE-RNN..

well-resourced source language can offer an advantage over the mean pooling method. In conclusion, SSL-based CAE-RNN models have fairly good performance when used crosslingually. There have been earlier studies (Kamper et al., 2021) on acoustic word embeddings for zero-resource languages using multilingual transfer with MFCC features, which worked well. Also, intuitively, some generalisation was expected as the aim of modelling is to compress a small segment of speech into a fixed dimensional vector. There might be a language effect on pre-trained SSL speech models but the basic speech properties are still invariant to changes in the language. The cross-lingual ability of SSL-based CAE-RNN models to obtain AWEs can support many applications such as speech search, indexing and discovery systems for languages with low-resources (Kamper et al., 2021).

5.2 Analysis of Anagram Pairs

Anagrams are words that can be formed by rearranging the letters of another word. Analysing anagram pairs provides insights into the impact of

Input Features	Polish		Portuguese		Spanish		French	
	test	test'	test	test'	test	test'	test	test'
Wav2vec2	0.57	0.57	0.48	0.54	0.60	0.68	0.52	0.68
WavLM	0.48	0.47	0.36	0.40	0.54	0.63	0.50	0.64
HuBERT	0.59	0.60	0.50	0.56	0.62	0.69	0.48	0.65
MFCC	0.18	0.20	0.11	0.15	0.22	0.29	0.22	0.35

Table 7: AP for the word-discrimination task with CAE-RNN model trained on English language with various input features and tested on other four languages.

letter order on AWE representation. Robust AWEs should capture the letter order in spoken words. For this analysis, same spoken word pairs and anagram pairs are chosen from different speakers. Ideally, the cosine distance between the same word pairs should be close to 0, while anagram word pairs should be close to 1. In Table 8, HuBERT-based CAE-RNN AWEs demonstrate cosine distances of approximately 0.01, 0.11, and 0.02 for the same spoken word pairs ‘aside’, ‘this’, and ‘no’, respectively. The anagram pairs of the words ‘aside’, ‘this’, and ‘no’ (i.e., ‘ideas’, ‘hits’, and ‘on’) have distances of 0.99, 0.50, and 0.69, respectively,




Figure 2: t-SNE visualisation of the AWEs derived from HuBERT-based CAE-RNN model for all five languages. From each language, all spoken instances of the top 7 words with the highest frequency count from the test set are chosen.

Word1	Word2	Cosine Distance (Mean Pooling)	Cosine Distance (CAE-RNN)	Description
aside	aside	0.23	0.01	Same word
aside	ideas	0.56	0.99	Anagram pair
this	this	0.33	0.11	Same word
this	hits	0.38	0.50	Anagram pair
no	no	0.30	0.02	Same word
no	on	0.63	0.69	Anagram pair

Table 8: Comparison of cosine distances between AWEs of same spoken word pairs and anagram pairs. HuBERT features are used for both CAE-RNN and mean pooling method.

for the HuBERT-based CAE-RNN model. These values are significantly better for both the same word pairs and anagram word pairs when compared to the HuBERT-based mean pooling method (Sanabria et al., 2023). This indicates that the HuBERT-based CAE-RNN model accurately captures the letter order in a word compared to the mean pooling baseline (Sanabria et al., 2023).

5.3 AWE Visualisation

t-SNE visualization is used to plot the 2-dimensional representations of the derived AWEs for all five languages. From each language, all spoken instances of the top 7 words with the highest frequency count from the test set are chosen. The plots demonstrate distinct and well-separated clusters for each spoken word across all languages. One interesting pattern can be observed for the Polish language, where the clusters of the spoken words ‘Owadów’ and ‘Owady’ share the boundary and are closely related in the AWE space. This is likely due to the fact that the first four letters of both the words (o, w, a, d) are shared and these words only differ in their endings.

6 Conclusions and Future Work

It has been demonstrated that SSL-based speech representations with CAE-RNN models outperform mean pooling and AE-RNN models across all languages. They also outperform MFCC-based models. Among all the SSL models, HuBERT performs the best when used as input for the CAE-RNN model, outperforming models such as Wav2vec2 and WavLM. Notably, despite being pre-trained on English data, the SSL models exhibit excellent performance on other languages, showcasing their cross-lingual generalization capability for AWE extraction.

Furthermore, quantitative analysis reveals that incorporating context information of the spoken word

leads to more robust AWEs. The HuBERT-based CAE-RNN model trained on English language and tested on other target languages outperforms the mean pooling method and the CAE-RNN model trained on the target language using MFCC features. This ‘zero-shot’ method to obtain robust AWEs for the target language can be useful in applications for low-resource languages (Jacobs et al., 2023). An analysis was also conducted to show that the CAE-RNN model effectively captures the order of letters in a word.

In future work, experiments will be conducted with the “LARGE” variation of SSL models, as well as multilingual pre-trained SSL models such as Wav2vec2-XLSR (Conneau et al., 2020). Additionally, an interesting experiment would involve training a single universal AWE model on all languages and comparing its performance with language-specific AWE models. Further research will focus on measuring the performance gains of SSL-based CAE-RNN models on downstream tasks such as query-by-example search (Settle et al., 2017; Yuan et al., 2018; Hu et al., 2021) and keyword spotting (Shin et al., 2022).

7 Limitations

This work is focused on the extraction of AWEs and measuring their quality solely based on the word discrimination task. No downstream applications such as query-by-example search and keyword spotting, have been discussed using the improved AWEs. In this work, only the “BASE” versions of the SSL-based speech models are explored for experiments and analysis. There are other variations, such as “LARGE” version, for which this study can be extended. All the languages considered in this work belong to the Indo-European language family. This work does not contain the analysis of languages that belong to another language family, such as Dravidian or Afroasiatic language families. This work does not deal with layer-wise analysis, which can provide better insights for further improving the AWEs.

8 Acknowledgments

This work was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1]. This work was also funded in part by LivePerson, Inc.

References

- Badr Abdullah, Iuliia Zaitova, Tania Avgustinova, Bernd Möbius, and Dietrich Klakow. 2021a. [How familiar does that sound? cross-lingual representational similarity analysis of acoustic word embeddings](#). In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 407–419, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Badr M. Abdullah, Marius Mosbach, Iuliia Zaitova, Bernd Möbius, and Dietrich Klakow. 2021b. [Do Acoustic Word Embeddings Capture Phonological Similarity? An Empirical Study](#). In *Proc. Interspeech 2021*, pages 4194–4198.
- Badr M. Abdullah, Bernd Möbius, and Dietrich Klakow. 2022. [Integrating Form and Meaning: A Multi-Task Learning Model for Acoustic Word Embeddings](#). In *Proc. Interspeech 2022*, pages 1876–1880.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework for self-supervised learning of speech representations](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12449–12460. Curran Associates, Inc.
- M. S. Barakat, C. H. Ritz, and D. A. Stirling. 2011. [Keyword spotting based on the analysis of template matching distances](#). In *2011 5th International Conference on Signal Processing and Communication Systems (ICSPCS)*, pages 1–6.
- Michael A. Carlin, Samuel Thomas, Aren Jansen, and Hynek Hermansky. 2011. [Rapid evaluation of speech representations for spoken term discovery](#). In *Proc. Interspeech 2011*, pages 821–824.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei. 2022. [Wavlm: Large-scale self-supervised pre-training for full stack speech processing](#). *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. 2020. Unsupervised cross-lingual representation learning for speech recognition. In *Interspeech*.
- S Davis and P Mermelstein. 1980. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *IEEE Trans. Acoust., Speech, Signal Process.*, 28(4):357–366.
- Wanjia He, Weiran Wang, and Karen Livescu. 2017. Multi-view recurrent neural acoustic word embeddings. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.

- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Yushi Hu, Shane Settle, and Karen Livescu. 2021. [Acoustic span embeddings for multilingual query-by-example search](#). In *2021 IEEE Spoken Language Technology Workshop (SLT)*, pages 935–942.
- Christiaan Jacobs, Nathanaël Carraz Rakotonirina, Evelyn Asiko Chimoto, Bruce A. Bassett, and Herman Kamper. 2023. Towards hate speech detection in low-resource languages: Comparing asr to acoustic word embeddings on wolof and swahili. *ArXiv*, abs/2306.00410.
- Herman Kamper. 2019. Truly unsupervised acoustic word embeddings using weak top-down constraints in encoder-decoder models. *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6535–3539.
- Herman Kamper, Micha Elsner, Aren Jansen, and Sharon Goldwater. 2015. [Unsupervised neural network based feature extraction using weak top-down constraints](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5818–5822.
- Herman Kamper, Yevgen Matushevych, and Sharon Goldwater. 2021. [Improved acoustic word embeddings for zero-resource languages using multilingual transfer](#). *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 29:1107–1118.
- Jingru Lin, Xianghu Yue, Junyi Ao, and Haizhou Li. 2023. [Self-Supervised Acoustic Word Embedding Learning via Correspondence Transformer Encoder](#). In *Proc. INTERSPEECH 2023*, pages 2988–2992.
- Yevgen Matushevych, Herman Kamper, and Sharon Goldwater. 2020. Analyzing autoencoder-based acoustic word embeddings. *CoRR*, abs/2004.01647.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal forced aligner: Trainable text-speech alignment using kaldi. In *Interspeech*.
- Amit Meghanani and Thomas Hain. 2024. Score: Self-supervised correspondence fine-tuning for improved content representations. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An asr corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. [Mls: A large-scale multilingual dataset for speech research](#). *Interspeech 2020*.
- Ramon Sanabria, Hao Tang, and Sharon Goldwater. 2023. [Analyzing acoustic word embeddings from pre-trained self-supervised speech models](#). In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Tanja Schultz, Ngoc Thang Vu, and Tim Schlippe. 2013. [Globalphone: A multilingual text & speech database in 20 languages](#). In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8126–8130.
- Shane Settle, Keith Levin, Herman Kamper, and Karen Livescu. 2017. Query-by-example search with discriminative neural acoustic word embeddings. In *INTERSPEECH*.
- Hyeon-Kyeong Shin, Hyewon Han, DoYeon Kim, Soo-Whan Chung, and Hong-Goo Kang. 2022. Learning audio-text agreement for open-vocabulary keyword spotting. In *Interspeech*.
- Shu-Wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Lai, Kushal Lakhotia, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu hsien Huang, Wei-Cheng Tseng, Ko tik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdel rahman Mohamed, and Hung yi Lee. 2021. Superb: Speech processing universal performance benchmark. In *Interspeech*.
- Yougen Yuan, Cheung-Chi Leung, Lei Xie, Hongjie Chen, Bin Ma, and Haizhou Li. 2018. [Learning Acoustic Word Embeddings with Temporal Context for Query-by-Example Speech Search](#). In *Proc. Interspeech 2018*, pages 97–101.