



HAL
open science

Speech Sequence Embeddings using Nearest Neighbors Contrastive Learning

Robin Algayres, Adel Nabli, Benoît Sagot, Emmanuel Dupoux

► **To cite this version:**

Robin Algayres, Adel Nabli, Benoît Sagot, Emmanuel Dupoux. Speech Sequence Embeddings using Nearest Neighbors Contrastive Learning. Interspeech 2022 - 23rd INTERSPEECH Conference, Sep 2022, Incheon, South Korea. hal-03831888

HAL Id: hal-03831888

<https://inria.hal.science/hal-03831888v1>

Submitted on 27 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Speech Sequence Embeddings using Nearest Neighbors Contrastive Learning

Robin Algayres^{1,2}, Adel Nabli^{1,2}, Benoît Sagot², Emmanuel Dupoux^{1,2,3}

¹ENS-PSL, EHESS, CNRS, Paris ²Inria, Paris, ³Facebook AI Research

robin.algayres@inria.fr, adel.nabli@umontreal.ca, benoit.sagot@inria.fr,
emmanuel.dupoux@inria.fr

Abstract

We introduce a simple neural encoder architecture that can be trained using an unsupervised contrastive learning objective which gets its positive samples from data-augmented k-Nearest Neighbors search. We show that when built on top of recent self-supervised audio representations [1, 2, 3], this method can be applied iteratively and yield competitive SSE as evaluated on two tasks: query-by-example of random sequences of speech, and spoken term discovery. On both tasks our method pushes the state-of-the-art by a significant margin across 5 different languages. Finally, we establish a benchmark on a query-by-example task on the LibriSpeech dataset to monitor future improvements in the field.

Index Terms: unsupervised speech sequence embeddings, k -nearest neighbors, data augmentation

1. Introduction

Recently, research into self-supervised representations for speech [3, 2, 1] has made impressive progress, resulting in audio features that are much better than the traditional MFCC features in representing speaker invariant phonetic content. The learned representations typically span a rather short and fixed-length time window (10 or 20ms) and can be fine-tuned for ASR [2] with few labels. However, such representations are still not good enough for other downstream tasks that need to handle larger word-like units that typically span larger time windows (100ms-1000ms) and have varying duration depending on speech rate. These tasks include word segmentation [4, 5] term discovery [6, 7, 8, 9, 10, 11] and unsupervised alignment of text and speech [12, 13] Despite recent research efforts into developing word embeddings also known as speech sequence embeddings (SSE), [14, 15, 16, 17, 18] their performances are low and difficult to compare for lack of a common evaluation benchmark.

Our contribution is twofold. First, we introduce a benchmark for SSE based on a query-by-example task built on LibriSpeech [19]. It uses the Mean Average Precision (MAP) metric computed over random query phoneme-ngrams. Previous work have shown that MAP is a good proxy for word segmentation [20, 21]. We show that MAP based on phoneme-ngrams is preferable to a MAP on words, which is too easy and does not reflect the nature of the downstream tasks in unsegmented continuous speech.

Second, we introduce a new method for learning SSE based on a combination of ideas from representation learning. The model is a convolution and transformer-based embedder trained with the NTXEnt contrastive loss [22]. Building on similar ideas in vision and speech, we select our positive examples through a mix of time-stretching data augmentation [23] and k-Nearest Neighbors search [24, 25]. Figure 1 gives an overview

of our method. To evaluate our method, we test our model on 5 types of acoustic features: MFCCs, CPC [3], HuBERT [1] and Wav2Vec 2.0 (Base and Large) [2]. We pick the best method from our LibriSpeech benchmark and show that when applied without any change to the task of spoken term discovery as defined in the zero resource challenges [26], we beat the state of the art on the NED/COV metric by a large margin in 5 new datasets.

2. Related work

Speech Sequence Embeddings. One straightforward method to obtain an SSE is to perform a max pooling along the time axis of the frame-wise speech features. This method (combined with a PCA) have been used in [12] to represent speech sequences obtained by unsupervised phoneme segmentation. Another method involves downsampling the sequence into a fixed number of frames and concatenate them into a single vector [15]. Other methods use bottleneck auto-encoders [15]. More advanced methods use Siamese loss [16] or contrastive learning [27] and rely on unsupervised ways to obtain positive pairs. Similarly, auto-encoders with a correspondence objective [28] can encode a member of a pair and reconstruct the other member of the pair. Correspondance and Siamese losses can be combined effectively in a model called CAE-Siamese that is on par with the current state-of-the-art in SSE models [14, 21].

At present, there is no common corpus and evaluation metric to evaluate the intrinsic quality of SSE models. For instance, [27, 16, 21] use pre-segmented word tokens to perform a query-by-example task measured by MAP [20]. Others rely on an ABX test on word triplets [15]. All these research papers compute their metrics on different corpora featuring different languages. This plurality of metrics and corpora makes it hard to monitor improvement in the field. We propose to create a common benchmark using the most popular metric so far: the MAP, on a widely known corpus: the LibriSpeech dataset [19].

Spoken term discovery. This task consists in searching audio recordings for speech sequences that are similar in acoustic space. One popular method is to use Segmental DTW [29] on hand-crafted features, mainly MFCC and PLP, to find subparts of two utterances that have a good DTW alignment score. The main problem with this method, as shown by [30], is that for some unknown reason, even though it works well on traditional hand-crafted speech features like MFCC and PLP features, it fails to benefit from the theoretically higher quality frame-level self-supervised embeddings like Wav2vec 2.0, CPC or Hubert. Another method which does not share this limitation, has been proposed by [25]. It works by encoding all speech sequences of a corpus of a given length range with an SSE and look for pairs with low cosine distance using accelerated k-Nearest Neighbors search. We adopt this method here, and the NED/COV metric from the Zerospeech challenges [26] to evaluate spoken term

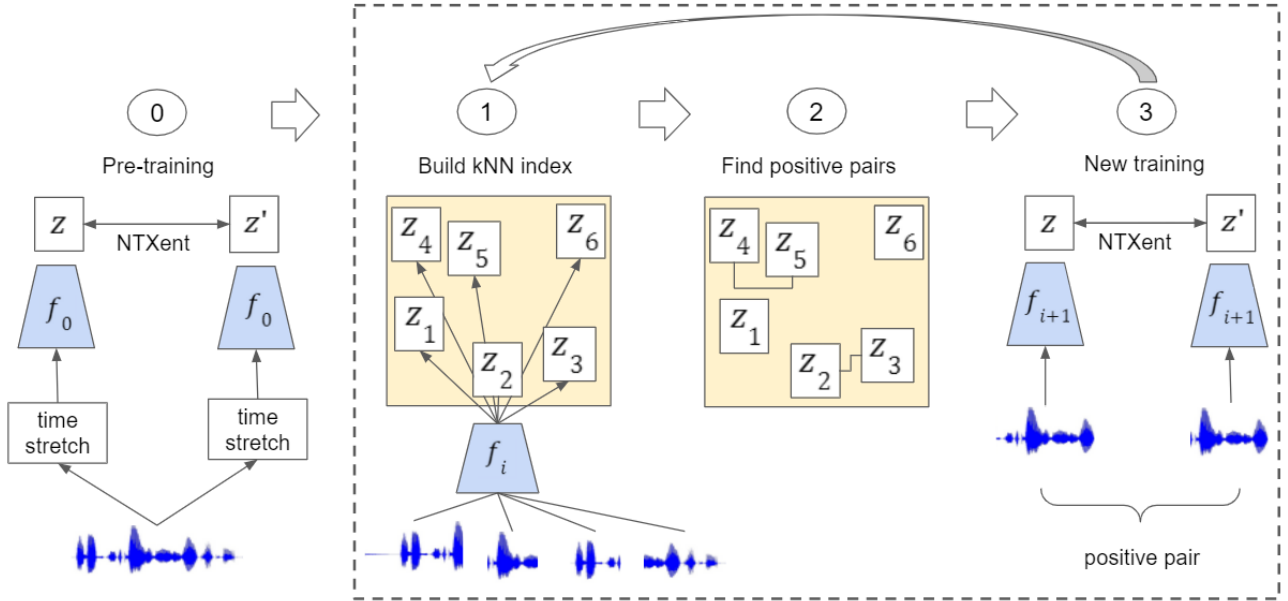


Figure 1: *The main steps of our method. 0) Train a model f_0 with the NTXent loss on time-stretched pairs of sequences of speech. 1) Use f_0 to encode many random sequences of speech from the corpus. 2) List close embeddings in the kNN as positive pairs. 3) Use the speech sequences associated to the positive pairs to train a new model f_1 . Go back to step 1 using f_1 instead of f_0 and iterate.*

discovery.

3. Method

3.1. Contrastive learning

The NTXent is a contrastive loss [22] that takes as input a batch P^+ of positive pairs of embedded speech sequences $P^+ = ((z_0, z_0^+), (z_1, z_1^+), \dots, (z_n, z_n^+))$. For each item taken from a positive pair, all remaining $2(n-1)$ items from the other pairs are used to form negative pairs. This loss minimizes the distance within the positive pairs while maximizing the distance within the negative pairs. For each item $z_i \in P^+$, the loss function is written:

$$L_{nce}(z_i, P^+) = -\log\left(\frac{\exp(\text{sim}(z_i, z_i^+)/\tau)}{\sum_{j \leq 2n, j \neq i} \exp(\text{sim}(z_i, z_j)/\tau)}\right) \quad (1)$$

where τ is the softmax temperature parameter and sim is the cosine similarity.

3.2. Stretch-invariant pretraining

At initialisation (step 0 in Figure 1), our SSE model is pre-trained with the NTXent loss on batches of positive pairs obtained by data augmentation. First, we apply Voice Activity Detection (rVAD [31]) to segment LibriSpeech into continuous Vocal Activity (VA) segments without any silence. Next, we use data augmentation on these speech sequences. Preliminary experiments showed that of the various augmentations (pitch, reverberation, additive noise, masking, etc), only time stretch had a large and systematic effect. This makes sense as variations in speech rate is a cause of a major change in the acoustic representation of a given speech sequence: its length. We therefore use time stretch as the sole data augmentation method.

We create a set of positive pairs from a single sampled VA segment, v of duration d as follows. From v , we create two variants by stretching its duration by a random factor in $[0.5, 1.8]$, giving rise to v_1 and v_2 with respective durations d_1 and d_2 .

These variants are converted into sequences of frame-level features, f_1 and f_2 , using an off-the-self acoustic features encoder with frame rate r ($r = 50$ or $r = 100$ frames per second depending on the encoder). We create positive pairs by sampling sub-sequences of frames in f_1 and f_2 that can be realigned to the same original acoustic signal from v . Specifically, we sample in f_1 a starting frame at index s and ending frame at index e so that $r * 0.08 \leq e - s \leq r$. We do not want to sample speech sequences that are shorter than one phoneme (80ms) or longer than most words (1s). Then the corresponding sub-sequence in f_2 is found between the frames at indices $\lfloor s \frac{d_2}{d_1} \rfloor$ and $\lfloor e \frac{d_2}{d_1} \rfloor$. We iterate this process until the number of positive pairs sampled this way reaches the batch size. The negative pairs are formed by recombining items from different positive pairs. Therefore, it is important that the sequences sampled from f_1 are temporally shifted by at least one phoneme length (80ms) so that they correspond to distinct phonetic contents. To ensure that, when sampling a starting frame index s in f_1 , we ensure that $s \% (0.08 * r) = 0$. The same procedure is applied for the end index e .

3.3. Self-labelling with k-NN search

We apply the method described in [25] to find more positive pairs with k-Nearest-Neighbors search (step 1 and 2 in Figure 1). We extract sub-sequences of frames within each VA ensuring they are all shifted by the average phoneme length (80ms). Next, we apply the SSE model pre-trained according to our initialisation method (section 3.2) to turn sampled sub-sequences of frames into embeddings, which are indexed for k-NN search using the FAISS library [32]. For each embedding s , the search returns N closest neighbors according to the cosine distance. We trim down this number to construct pairs that are more likely to have the same phonetic content as follows: First, all neighbors that are temporally overlapping with s are removed. Sec-

Supervision	Models	dev-clean	dev-other	test-clean	test-other
unsup.	Max-pooling	0,07	0,048	0,07	0,047
self-sup.	CAE-Siamese	0,21	0,154	0,212	0,151
self-sup.	Ours (iter. 2)	0,398	0,307	0,399	0,305
weakly-sup.	Topline	0,789	0,647	0,784	0,648

Table 1: *Phoneme-ngrams MAP computed on LibriSpeech held-out sets for different SSE models. All models take as input features the Wav2vec2.0 Base at layer 8*

ond, when two neighbors are temporally overlapping with each other, we keep only the one that has the smallest cosine distance with s (Non-Maximal-Suppression [25]). Finally, we apply a distance threshold above which all pairs are discarded. We set it such that half of speech sequences indexed in the k-NN graph have at least one selected pair. This was estimated on a sample of LibriSpeech (which functions as our dev set) where we found out that the proportion of speech sequence (excluding silences) with a unique transcription was around 50%. All selected pairs were used to form batches of positive pairs to retrain a new SSE model from scratch (step 3 in Figure 1). We iterate the procedure by re-embedding the speech sequences into a new k-NN index, re-sampling positive pairs and retraining another SSE model.

3.4. Evaluation metrics

We evaluated the SSE models on the basis of two tasks: a query-by-example task measured by MAP and a word discovery task measured by NED/COV.

Query-by-example. Given a query (a speech sequence), the task is to retrieve all sequences with the same transcription in an audio file. Previous works have used lists of words extracted from an audio file using forced alignment[20]. This task is a little too easy, since the words are disjoint intervals in the speech signal, while a retrieval task on an unsegmented audio file would have to deal with potentially overlapping speech sequences that differ from the query by only one phoneme. To push the task closer to an unsegmented use case, here, we compute MAPs over randomly selected phoneme-ngrams, i.e non-silent speech sequences that start and end at phoneme boundaries. To avoid long computation of MAP, we reduce the number of such sequences by picking only those that are shorter than 1 second. Most of them will be phoneme-ngrams with unique transcription in force-alignment. We also remove all phoneme-ngrams with unique transcriptions. Prior experiments showed that it does not change the MAP scores much but drastically reduces computation time. Given a query, we can compute precision and recall of the retrieved sequences as a function of the distance threshold r . The MAP score is obtained by varying r between 0 and infinity and integrating the precision/recall curve. It results in a score between 0 and 1 (higher is better). We compute it using the pytorch-metric-learning library [35].

Term Discovery metric. In its simplest form, spoken term discovery consist in finding clusters or pairs of speech segments that have the same transcription from unsegmented audio. It is evaluated by the NED/COV as introduced in the zero speech challenges [26]. For each discovered pair, the Normalised Edit Distance (NED) is computed as the edit distance normalised by the length of the longest item in the pair. The coverage (COV) is given by the proportion of phonemes in the corpus that is covered by least one speech sequence from the sets of pairs. There is typically a tradeoff between these two measures: low NED systems have typically low COV, and high COV have high NED.

Models	buckeye	xitsonga	mandarin	french	english	average
Max-pooling	0,05	0,053	0,075	0,039	0,052	0,054
CAE-Siamese	0,16	0,23	0,26	0,2	0,19	0,208
Ours (iter. 2)	0,235	0,362	0,277	0,283	0,346	0,301
Topline	0,751	0,948	0,822	0,71	0,857	0,818

Table 2: *Phoneme-ngrams MAP computed on Zerospeech corpora for different SSE models. All models take as input features the Wav2vec2.0 Base at layer 8*

3.5. Datasets

For training and evaluation of our models we used six different speech datasets that have been force-aligned with a phonetic transcription. Models are trained on the speech only and phonetic force-alignments are used for evaluation metrics. First, we use the LibriSpeech dataset [19] composed of a 960 hours training set and four 5 hours held-out datasets for development and test called: dev-clean, base at test-clean and test-other. Then, to study generalization, we used five speech corpora from the Zerospeech challenges [26]: English (45hours), Xitsonga (2h30), Mandarin (2h30), French (24h) and the Buckeye (5 hours of conversational English). As these corpora do not have a held-out test set, evaluation is performed using the training set, a common practice in unsupervised speech processing [26].

To compute a MAP score on the LibriSpeech dataset, we sample phoneme-ngrams from the four held-out datasets as described in section 3.4. It results in around 700,000 phoneme-ngrams in each held-out dataset. On the Zerospeech corpora, we hypothesise a test-set by sampling from the training set 150,000 non unique phoneme-ngrams.

3.6. Models architectures

Our SSE model (f in Figure 1) is a neural network which starts with a layer-normalisation and a 1-dimensional convolution layer with 512 channels (kernel size 4, stride 1) followed by Gated-linear-units and drop out. Sinusoidal Positional embeddings are added after the convolution layer. After going through a transformer layer with 4 attention heads and 512 units, the resulting features are max-pooled along the time axis to form a fixed-size embedding of the input speech sequence. During training, we add a projection head composed of two 512 neurons linear layers and a ReLU activation in between the layers. The model is trained by back-propagation with Adam optimiser and 0.0001 initial learning rate. The temperature parameter of the NTXEnt contrastive loss is set to 0.15.

We feed our SSE model on 4 frame-level pre-trained features: CPC (last layer as in [3]), Wav2vec2.0 (Base at layer 8 and Large at layer 15 as in [12]), and Hubert Large (last layer as in [1]). We add MFCC (40 coefficients) for comparison. All hyper-parameters were chosen by gridsearch on the LibriSpeech to maximizes the MAP scores on the dev-clean set.

We compare this model to two other methods: Max-pooling and CAE-Siamese. Max-pooling is a common baseline SSE model. It consists in max-pooling input frame-level features across time to get a fixed-size vector. The CAE-Siamese referenced in [14, 21] is on par with the current state-of-the-art in SSE models. We took the implementation from [21] which is a three-layers LSTM encoder-decoder trained with a weighted average of the Siamese loss as in [16] and the Correspondance loss from [28]. Positive pairs for training are sampled using PLP features as in [21]. We know from [36] that the Correspondance loss benefits from recent frame-level speech representation. Therefore, we train the CAE-Siamese with Wav2vec2.0 features (Base model at layer 8) to give a general idea of the current best practices in SSE modelling.

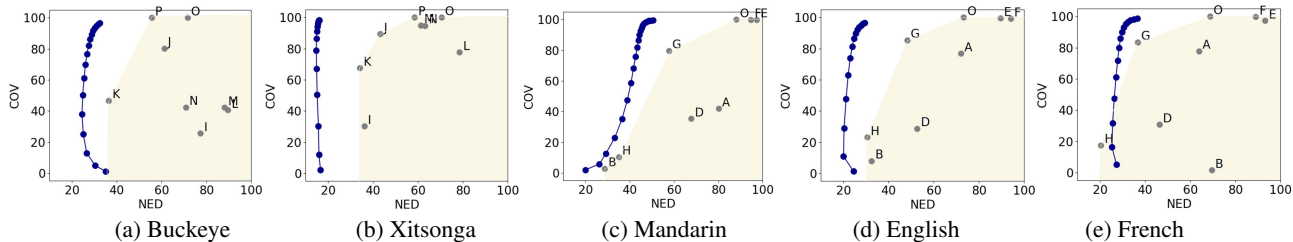


Figure 2: *NED/COV* curves on *Zerospeech* corpora. Our method is represented with a line of blue dots and each other competitors as grey points: Garcia-Granada (A [33]), Jansen (B,[6]), Räsänen (L,M,N [7] and D [8] and G,H [9]), Kamper (O [4] and P [34]), Lysinski (I,L,J [10]), Bhati (E,F[11]).

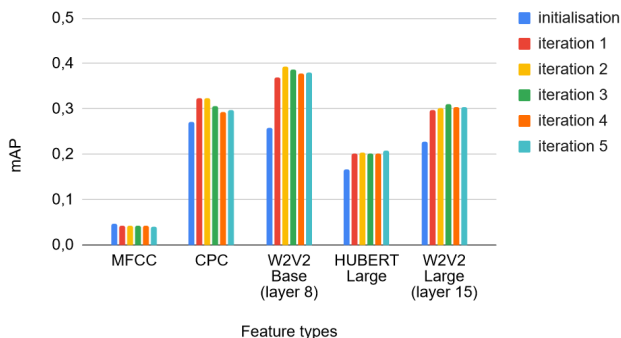


Figure 3: *Phoneme-ngrams* MAP computed on *LibriSpeech dev-clean* for different iteration of our model and different input feature types

4. Results and discussion

4.1. Mean Average Precision

Figure 3 shows the MAP scores evaluated on LibriSpeech dev-clean datasets for different frame-level input features and different training steps (initialisation and k-NN iterations). Overall, our method reaches a peak of performance after two iterations of k-NN self-labelling. Across input features, the Wa2vec2.0 features from the Base model at layer 8 lead to the best MAP scores. From now on, we use this input features for all following experiments.

Table 1 is our benchmark across different SSE models on LibriSpeech held-out datasets. As expected, the CAE-Siamese is consistently better than the Max-pooling. Yet it is far from reaching the performance of our method after two iterations of k-NN self-labelling.

As a topline system, we train our neural network architecture on positive pairs of speech sequences that have each the same phonetic transcription. These positive pairs are sampled using the time-aligned transcriptions of our training sets. Unsurprisingly, the MAP scores of this weakly-supervised topline are way above all others.

To show generalisation, we chose the configuration of our SSE model that works best on our LibriSpeech benchmark and applied it directly on the five corpora of the Zerospeech challenges. In table 2 again, Max-pooling and CAE-Siamese get much lower scores than our method as well as the weakly-supervised topline. Interestingly, the Wav2vec2.0 were pre-trained on the LibriSpeech corpus but transfer well as frozen features across corpora and languages.

4.2. NED/COV

Figure 2 shows the NED/COV scores of our model on the Zerospeech’s dataset compared to other models submitted for previous editions of the challenge. Some of these models performed additional tasks that what we do not do here such as clustering and finding word boundaries. Therefore we only run the evaluation metrics relevant to the matching task. Competitors’ submissions to the Zerospeech challenges are represented with grey dots on figure 2 as they were asked to submit only one set of clusters. We do not impose this constraint on us. For each corpus, using our trained SSE model, by varying the threshold mentioned in section 3.4, we create a range of twenty submissions per corpora which appear as a blue dots on figure 2. Strong NED/COV scores are in the upper left corners. Regarding the NED/COV results, whatever the threshold that we choose for our method, we get better NED/COV than all other systems on all datasets (except one submission on the French dataset but we beat that submission on the other datasets).

5. Conclusion

We introduced a new Speech Sequence Embedding model that leverages time-stretching data augmentation as well as nearest neighbors search. Across several corpora, our model shows much higher results on our ngram-based MAP LibriSpeech metric than a previous competitive model (CAE-Siamese), although there is still room for improvements based on our weak supervision experiment. We open source our model and our MAP metric in the hope it will encourage researchers to use the latter to evaluate SSE models. On the NED/COV metric, our model push the state-of-the-art by a large margin compared to other submissions from the Zerospeech challenges, opening up the possibility of further improvement in speech segmentation tasks.

6. Acknowledgements

This work was funded in part, to the authors in their academic capacities, by the Agence Nationale pour la Recherche (ANR-17-EURE-0017 Frontcog, ANR-10-IDEX-0001-02 PSL*, ANR-19-P3IA-0001 PRAIRIE 3IA Institute), CIFAR (Learning in Machines and Brains) and Facebook AI Research (Research Grant). This work was performed using HPC resources from GENCI-IDRIS (Grant 2021-[AD011011217]).

7. References

- [1] w. Hsu, B. Bolte, Y. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *CoRR*, 2021.
- [2] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *CoRR*, 2020.
- [3] M. Rivière, A. Joulin, P. Mazaré, and E. Dupoux, "Unsupervised pretraining transfers well across languages," in *ICASSP*, 2020.
- [4] H. Kamper, K. Livescu, and S. Goldwater, "An embedded segmental k-means model for unsupervised segmentation and clustering of speech," *CoRR*, 2017.
- [5] S. Bhati, J. Villalba, P. Želasko, L. Moro-Velazquez, and N. Dehak, "Segmental contrastive predictive coding for unsupervised word segmentation," *ArXiv*, 2021.
- [6] K. Levin, A. Jansen, and B. V. Durme, "Segmental acoustic indexing for zero resource keyword search," in *ICASSP*, 2015.
- [7] O. Räsänen, G. Doyle, and M. Frank, "Unsupervised word discovery from speech using automatic segmentation into syllable-like units," 2015.
- [8] S. Seshadri, R. Ulpu, and O. Räsänen, "Comparison of non-parametric bayesian mixture models for syllable clustering and zero-resource speech processing," in *INTERSPEECH*, 2017.
- [9] O. Räsänen and M. Blandón, "Unsupervised discovery of recurring speech patterns using probabilistic adaptive metrics," *ArXiv*, 2020.
- [10] V. Lyzinski, G. Sell, and A. Jansen, "An evaluation of graph clustering methods for unsupervised term discovery," in *INTERSPEECH*, 2015.
- [11] S. Bhati, J. Villalba, P. Želasko, and N. Dehak, "Self-expressing autoencoders for unsupervised spoken term discovery," *ArXiv*, 2020.
- [12] A. Baevski, W. Hsu, A. Conneau, and M. Auli, "Unsupervised speech recognition," *CoRR*, 2021.
- [13] Y. Chung, W. Weng, S. Tong, and J. R. Glass, "Unsupervised cross-modal alignment of speech and text embedding spaces," *CoRR*, vol. abs/1805.07467, 2018. [Online]. Available: <http://arxiv.org/abs/1805.07467>
- [14] P. Last, H. A. Engelbrecht, and H. Kamper, "Unsupervised feature learning for speech using correspondence and siamese networks," *IEEE Signal Processing Letters*, 2020.
- [15] N. Holzenberger, D. Mingxing, J. Karadayi, R. Riad, and E. Dupoux, "Learning word embeddings: Unsupervised methods for fixed-size representations of variable-length speech segments," in *INTERSPEECH*, 2018.
- [16] S. Settle and K. Livescu, "Discriminative acoustic word embeddings: Recurrent neural network-based approaches," *CoRR*, 2016.
- [17] Y. Chung, C. Wu, C. Shen, H. Lee, and L. Lee, "Audio word2vec: Unsupervised learning of audio segment representations using sequence-to-sequence autoencoder," *CoRR*, 2016.
- [18] L. Badino, C. Canevari, L. Fadiga, and G. Metta, "An auto-encoder based approach to unsupervised learning of subword units," in *ICASSP*, 2014.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *ICASSP*, 2015.
- [20] M. Carlin, S. Thomas, A. Jansen, and H. Hermansky, "Rapid evaluation of speech representations for spoken term discovery," in *INTERSPEECH*, 2011.
- [21] R. Algayres, S. Zaiem, B. Sagot, and E. Dupoux, "Evaluating the reliability of acoustic speech embeddings," in *INTERSPEECH*, 2020.
- [22] K. Sohn, "Improved deep metric learning with multi-class n-pair loss objective," in *NIPS*, 2016.
- [23] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton, "A simple framework for contrastive learning of visual representations," *CoRR*, 2020.
- [24] D. Dwivedi, Y. Aytar, S. J. Tompson, and A. Zisserman, "With a little help from my friends: Nearest-neighbor contrastive learning of visual representations," *CoRR*, 2021.
- [25] A. Thual, C. Dancette, J. Karadayi, J. Benjumea, and E. Dupoux, "A k-nearest neighbours approach to unsupervised spoken term discovery," in *SLT*, 2018.
- [26] E. Dunbar, J. Karadayi, M. Bernard, X. Cao, R. Algayres, L. Ondel, L. Besacier, S. Sakti, and E. Dupoux, "The zero resource speech challenge 2020: Discovering discrete subword and word units," *CoRR*, 2020.
- [27] C. Jacobs, Y. Matuselych, and H. Kamper, "Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation," *CoRR*, 2021.
- [28] H. Kamper, M. Elsner, A. Jansen, and S. Goldwater, "Unsupervised neural network based feature extraction using weak top-down constraints," in *ICASSP*, 2015.
- [29] A. Park and J. Glass, "Towards unsupervised pattern discovery in speech," in *ASRU*, 2005.
- [30] E. Hermann, H. Kamper, and S. Goldwater, "Multilingual and unsupervised subword modeling for zero-resource languages," *Computer Speech & Language*, 2021.
- [31] Z. Tan, A. Sarkar, and N. Dehak, "rvad: An unsupervised segment-based robust voice activity detection method," *CoRR*, 2019.
- [32] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with gpus," *CoRR*, 2017.
- [33] Garcia-Granada, E. Sanchis, M. Castro-Bleda, J. Gonzalez, and L. Hurtad, "Zerospeech2017 elirf-upv system," 2017.
- [34] H. Kamper, A. Jansen, and S. Goldwater, "Fully unsupervised small-vocabulary speech recognition using a segmental bayesian model," in *INTERSPEECH*, 2015.
- [35] K. Musgrave, S. J. Belongie, and S. Lim, "Pytorch metric learning," *CoRR*, 2020.
- [36] L. van Staden and H. Kamper, "A comparison of self-supervised speech representations as input features for unsupervised acoustic word embeddings," *CoRR*, 2020.