

Unspeech: Unsupervised Speech Context Embeddings

Benjamin Milde, Chris Biemann

Language Technology, Universität Hamburg
{milde,biemann}@informatik.uni-hamburg.de

Abstract

We introduce "Unspeech" embeddings, which are based on unsupervised learning of context feature representations for spoken language. The embeddings were trained on up to 9500 hours of crawled English speech data without transcriptions or speaker information, by using a straightforward learning objective based on context and non-context discrimination with negative sampling. We use a Siamese convolutional neural network architecture to train Unspeech embeddings and evaluate them on speaker comparison, utterance clustering and as a context feature in TDNN-HMM acoustic models trained on TED-LIUM, comparing it to i-vector baselines. Particularly decoding out-of-domain speech data from the recently released Common Voice corpus shows consistent WER reductions. We release our source code and pre-trained Unspeech models under a permissive open source license.

Index Terms: Unsupervised learning, speech embeddings, context embeddings, speaker clustering

1. Introduction

Variance and variability in recordings of speech and its representations are a common problem in automatic speech processing tasks. E.g. speaker, environment characteristics and the type of microphone will cause large differences in typical speech representations (e.g. FBANK, MFCC), making direct similarity comparisons difficult. We can describe such factors of variance also as the context of an utterance; speech sounds that occur close in time share similar contexts. Based on this idea, we propose to learn representations of such contexts in an unsupervised way, without needing further speaker IDs, channel information or transcriptions of the data.

Recent acoustic models for automatic speech recognition (ASR) incorporate some form of (deep) neural network that can learn to deal with part of this variance by using supervised training data in combination with the ability to learn representations as part of the model. A growing trend is to incorporate larger context views of the data explicitly into the neural network. In Deep Neural Network Hidden Markov Model (DNN-HMM) hybrids, fixed-length speaker embeddings like i-vectors are made available for the neural network as additional input features [1]. Typically, larger temporal windows than single speech frames are also used as input to the neural network to make context available to local predictions. This can e.g. be achieved by either stacking consecutive speech frames or by using Time-Delayed Neural Networks (TDNNs) [2, 3].

On the other hand, "Unspeech" embedding models embed a window of speech into a fixed-length vector so that corresponding points are close, if they share similar contexts. Unsupervised training of the embedding function is inspired by negative sampling in word2vec [4] – where words that share a similar meaning are embedded in similar regions in a dense vector space. In this work, we demonstrate that the learned Un-

speech context embeddings encode speaker characteristics and also can be used to cluster a speech corpus. As an additional context input feature, they can also improve supervised speech recognition tasks with TDNN-HMM acoustic models, in particular when adaptation to out-of-domain data is needed.

2. Related Work

Speaker embeddings and phonetic embeddings are two major groups of proposed embeddings in speech: While speaker embeddings seek to model utterances from the same speaker so that they share similar regions in a dense vector space, in phonetic embeddings, the same or similar phonetic content is close.

I-vectors [5] are well-known, popular speaker vectors. Recently, supervised neural network-based speaker embeddings also succeeded to show good speaker discriminative properties [6, 7, 8], particularly on short utterances. Bengio and Heigold proposed supervised word embeddings for speech recognition [9], where words are nearby in the vector space if they sound alike. Kamper et al. [10] showed that auto-encoders can also be used in conjunction with top-down information for unsupervised phonetic representation learning in speech. Chung et al. [11] proposed audio word2vec, based on sequence-to-sequence auto-encoders trained on a dictionary of isolated spoken words. By analogy of auto-encoders, Pathak et al. introduced context encoders [12], a class of models that learn context embeddings in images. There is also growing interest in representation learning on non-speech audio by using learning objectives directly related to contexts. Jansen et al. [13] encoded the notion that (non-speech) sounds occurring in context are more related.

Bromley et al. [14] introduced Siamese neural networks: two (time-delayed) neural networks that embed digital signatures and a learning objective based on discriminating between true and false signatures. This idea has recently been revisited in the context of joint phoneme and speaker embeddings learning in a weakly supervised setting, where speaker annotation, same word information and segmentation is available [15, 16]. Gutmann et al. [17] introduced Noise-Contrastive Estimation (NCE), an estimator based on discriminating between observed data and some artificially generated noise. Jati et al. proposed Speaker2vec [18] for speaker segmentation, with unsupervised training using a neural encoder/decoder. Very recently and in parallel to our efforts, Jati et al. also proposed (unsupervised) neural predictive coding to learn speaker characteristics [19].

In [20] unsupervised speaker clustering was proposed to yield labels for speaker adaptation in acoustic models, based on the idea that consecutive windows of speech are likely from the same speaker. Several forms of context/speaker embeddings have also been used for (speaker) adaptation in state-of-the-art speech recognition acoustic models: i-vector speaker embeddings are by far the most popular [1, 21, 22]. Vesely et al. proposed sequence summary neural networks for speaker adaptation, where utterance context vectors are averaged from the speech feature representation [23]. Gupta et al. [24] showed




Figure 1: The initial sequence with unnormalized FBANK vectors: we choose one target window and two left and right contexts. All windows are of the same size.




Figure 2: Sampling examples for two left contexts and two right contexts from the figure above. Positive example pairs are of class $C = 1$, negative sampled pairs are $C = 0$. In this example, each window has a size of 50 FBANK frames (0.5 seconds).

that visual features, in the form of activations from a pre-trained ConvNet for object detection on videos can also be used as context vectors in the acoustic model.

3. Proposed Models

We construct an artificial binary classification task with logistic regression, where two fixed sized windows are compared. One target window can have multiple context windows, depending on the number of left and right contexts. For every left and right context, a pair with the target is created. Figure 1 illustrates this with two right and left contexts, yielding four positive contexts and four randomly sampled negative contexts. For the target window we denote emb_t as the target embedding transformation, taking a window of FBANK features and producing a fixed sized output vector, emb_c as the embedding of a true context and emb_{neg} as the embedding of a randomly sampled context. The pair of (embedded) speech windows is considered to be of class $C = 1$ if one window is the context of the other, or $C = 0$ if they are not. For $C = 1$, we sample the pairs from consecutive windows, for $C = 0$ we use negative sampling to construct a pair of speech windows that are unrelated with high probability: We uniformly sample a random utterance u and then uniformly a random position in u .

3.1. Objective Function

With the scalar x as the output of the model for a particular data point, σ the sigmoid function and C its true class $\in (0, 1)$, the logistic loss for a binary classification task is:

$$loss(x, C) = C(-\log(\sigma(x))) + (1 - C)(-\log(1 - \sigma(x))) \quad (1)$$

with $x = emb_t^T emb_c$, the dot product over target and context embedding transformations if $C = 1$ and $x = emb_{neg1_i}^T emb_{neg2_i}$, the dot product over two negative sampled embedding transformations if $C = 0$, for $k =$ number of negative samples we can thus obtain:

$$NEG_{loss} = -k \cdot \log(\sigma(emb_t^T emb_c)) - \sum_{i=1}^k \log(1 - \sigma(emb_{neg1_i}^T emb_{neg2_i})) \quad (2)$$

Note that in the similar NCE loss formulation [17], $P(C = 1) = P(C = 0) = \frac{1}{2}$, i.e. the number of data points where $C = 1$ and $C = 0$ are the same, while we could have more negative than positive target/context embedding pairs, depending on k . Instead, for $C = 1$ we multiply with $k =$ the number of negative samples, to penalize errors on positive and negative context pairs equally. Another difference is that we sample two unrelated embedding windows instead of one.

3.2. Model Architecture

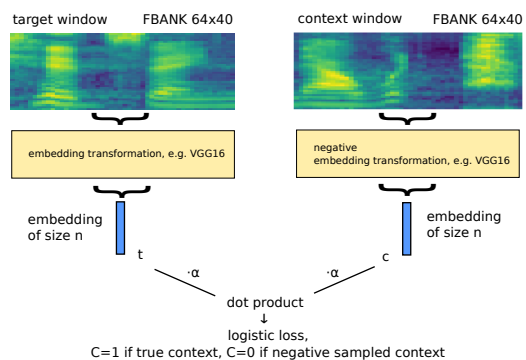


Figure 3: Unspeech embeddings are trained using a Siamese neural network architecture combined with a dot product. We use the VGG16A network as embedding transformation in the yellow boxes (a convolutional neural network with 16 layers).

Figure 3 shows our architecture for FBANK input features. We project two dimensional input windows into two fixed sized embeddings, which are combined with a dot product. Since this is sensitive to the scaling of output embeddings, we multiply them with a single scalar parameter α , which is trained with the rest of the network. Direct normalization to unit length (making the dot product equivalent to a cosine distance) did hamper convergence of the loss and was discarded early on.

Many different architectures are possible for converting the input representation into a fixed sized embedding, but we mainly evaluated with a VGG-style ConvNet architecture (Model A in [25]), as it is well established and it can exploit the two dimensional structure in the FBANK signal. We share the weights of the convolutional layers of both embedding transformations, but keep the fully connected layers separate. Dropout is used for fully connected layers (0.1) and L2 regularization is used on the weights (0.0001), for all experiments we optimize with ADAM [26]. We make use of leaky ReLUs [27].

4. Evaluation

Table 1: Comparison of English speech data sets.

dataset	hours			speakers		
	train	dev	test	train	dev	test
TED-LIUM V2	211	2	1	1273+3	14+4	13+2
Common Voice V1	242	5	5	16677	2728	2768
TEDx (crawled)	9505			41520 talks		

Table 1 characterizes the datasets we used in our evaluation. TED-LIUM V2 [28] has a comparatively small number of speakers, especially in the development and test set of the corpus. TED-LIUM and Common Voice [29] are segmented at the utterance level, both are similar in the number of hours. In Common Voice, volunteers from all over the world recorded predefined prompts, in TED-LIUM utterances are segmented from and aligned to TED talks. In order to explore large-scale training, we also downloaded all TEDx talks from 01-01-2016 until 03-01-2018 from YouTube, giving us 41520 talks (0.5 TB compressed audio data) with a total of 9505 hours of unannotated audio. While the majority of TEDx talks are in English, a very small number of them are in other languages or contain only music. We did not segment or clean the TEDx data.

4.1. Same/different Speaker Experiment

In the same/different speaker experiment, we evaluate a binary classification task: given two utterances, are they uttered from the same or different speakers? Our hypothesis is that Unspeech embeddings can be used for this task, because one strategy to discriminate samples of true contexts from negative sampled ones is modelling speaker traits. In Table 2 we show equal error rates (EER)¹ on same/different speaker comparisons of all utterance pairs, limiting the number of speakers to 100 in the train sets of TED-LIUM and Common Voice corpus in this experiment. The embedding dimension is 100 in all experiments, we train Unspeech models with different target window widths (32, 64, 128) and i-vectors are trained/extracted with Kaldi [31]. For all experiments, we use two left and two right context windows.

Table 2: Equal error rates (EER) on TED-LIUM V2 – Unspeech embeddings correlate with speaker embeddings.

Embedding	EER		
	train	dev	test
TED-LIUM:			
(1) i-vector	7.59%	0.46%	1.09%
(2) i-vector-sp	7.57%	0.47%	0.93%
(3) unspeech-32-sp	13.84%	5.56%	3.73%
(4) unspeech-64	15.42%	5.35%	2.40%
(5) unspeech-64-sp	13.92%	3.4%	3.31%
(6) unspeech-64-tedx	19.56%	7.96%	4.96%
(7) unspeech-128-tedx	20.32%	5.56%	5.45%

The distance function $d_1(a, b) = \sigma(\text{emb}_t(a)^T \text{emb}_c(b))$ to compare two segments a and b correspond to the distance function in the Unspeech training process. The cosine distance, or equally after normalization to unit length, the Euclidean distance on vectors produced by emb_t also produces good comparison results, so that $d_2(a, b) = \|\text{emb}_t(a) - \text{emb}_t(b)\|$ can be used for comparing two Unspeech segments. Sequences that are longer than the trained target window can be windowed and averaged to obtain a single vector for the whole sequence, since vectors that are close in time share contexts and correlate highly. However, EER on i-vectors trained with supervised speaker labels compared with the cosine distance (results with $d_2(a, b)$ are identical after normalization) are lower than on Unspeech embeddings with $d_2(a, b)$ (1,2 vs 3,4,5). Training Unspeech on TEDx talks instead of TED-LIUM does also produce higher EER as a speaker embedding (6,7). ”-sp” denotes training on speed-perturbed data: adding copies of the raw training data at 0.9 and 1.1 playing speed, as recommended in [32].

¹The error rate where the number of false positive and false negatives is the same, calculated using pyannotate.metrics [30]

4.2. Clustering Utterances

We can also use the generated vectors to cluster a corpus of utterances to gain insight into what kind of utterances get clustered together. We use HDBSCAN [33], a modern hierarchical density-based clustering algorithm for our experiments since it scales very well to a large number of utterances and the number of clusters does not need to be known apriori. It uses an approximate nearest neighbor search if the comparison metric is Euclidean, making it significantly faster on a large number of utterances as compared to other speaker clustering methods that require distance computations of all utterance pairs (including greedy hierarchical clustering with BIC [34]). We use Adjusted Rand Index (ARI) [35] and Normalized Mutual Information (NMI) [36] to compare the clusters to the speaker IDs provided by the TED-LIUM corpus in Table 3. We found that Unspeech embeddings and i-vectors give a sensible number of clusters, without much tweaking of HDBSCAN’s two parameters (min. cluster size, min. samples)². On the train set, Unspeech embeddings will provide slightly higher cluster scores, while i-vectors provide better scores for dev and test (that have a significantly smaller number of speakers). Unspeech-64 is slightly better than Unspeech-32 on the dev and test set. ARI is sensitive to the absolute number of outliers – we found NMI to be a better metric to compare the results on the train set.

Taking a closer look at the clustered Unspeech embeddings, we observed that different speakers in the same talk tend to get clustered into distinct clusters (making the clustered output very often more accurate than the train speaker IDs provided in TED-LIUM), while the same speaker across different talks and also the same speaker in one talk with significantly different background noises tends to be clustered into distinct clusters. This implies that Unspeech embeds more than just speaker traits.

4.3. Acoustic Models With Unspeech Cluster IDs

We can also train acoustic models with the cluster IDs provided and use them in lieu of speaker IDs for HMM-GMM speaker adaptation and online i-vector training for the TDNN-HMM model. We use the TED-LIUM TDNN-HMM chain recipe (s5_r2) in Kaldi [37] and show WER before (plain) and after rescoring with the standard 4-gram Cantab TED-LIUM LM (resc.). Table 4 shows WER on different speaker separation strategies on the train set, with one speaker per talk being the default in the s5_r2 recipe. All models pre-trained online i-vectors based on the given speaker IDs and use those as additional input features. The standard recipe computes a fixed affine transform on the combined input features (40 dim hi-res MFCC + 100 dim i-vector), c.f. Appendix C.6 of [38]. GMM-HMM bootstrap models will perform about 15% worse and TDNN-HMM trained on them will perform about 10% worse if no speaker information is available. Using the cluster IDs from clustering Unspeech embeddings of all utterances, the baseline WER can not only be recovered, but even slightly improved upon. For all TDNN-HMM models we set the width of a layer to 1024.

4.4. Unspeech Context Vectors in TDNN-HMM models

We can also replace the i-vector representation used in training the TDNN-HMM with the Unspeech context vector. In Table 5, we selected the strongest baseline from Table 4 according to the dev set (Unspeech 64-sp clusters) and show WERs on the

²We use 5/3 for all experiments shown in Table 3, but other parameters in the range 3-10 will give similar results.

Table 3: Comparing clustered utterances from TED-LIUM using *i*-vectors and (normalized) Unspeech embeddings with speaker labels from the corpus. "-sp" denotes embeddings trained with speed-perturbed training data.

Embedding	Num. clusters			Outliers			ARI			NMI		
	train	dev	test	train	dev	test	train	dev	test	train	dev	test
TED-LIUM IDs	1273 (1492)	14	13	3	4	2	1.0	1.0	1.0	1.0	1.0	1.0
i-vector	1630	12	10	8699	1	2	0.8713	0.9717	0.9792	0.9605	0.9804	0.9598
i-vector-sp	1623	12	10	9068	1	2	0.8641	0.9717	0.9792	0.9592	0.9804	0.9598
unspeech-32-sp	1686	16	12	3235	22	32	0.9313	0.9456	0.9178	0.9780	0.9536	0.9146
unspeech-64	1690	16	11	5690	14	21	0.8130	0.9537	0.9458	0.9636	0.9636	0.9493
unspeech-64-sp	1702	15	11	3705	23	25	0.9205	0.9517	0.9340	0.9730	0.9633	0.9366

Table 4: Comparing the effect of two speaker division baselines (One speaker per talk, one speaker per utterance) and clustering with Unspeech on WER with GMM-HMM and TDNN-HMM chain acoustic models trained on TED-LIUM.

Acoustic model	Spk. div.	Dev WER		Test WER	
		plain	resc.	plain	resc.
GMM-HMM	per talk	19.2	18.2	17.6	16.7
TDNN-HMM	per talk	8.6	7.8	8.8	8.2
GMM-HMM	per utt.	19.6	18.7	20.1	19.2
TDNN-HMM	per utt.	8.5	7.9	9.3	9.0
GMM-HMM	Unspeech	18.4	17.4	17.5	16.5
TDNN-HMM	64	8.6	7.8	8.5	8.1
GMM-HMM	Unspeech	18.4	17.5	17.2	16.4
TDNN-HMM	64-sp	8.3	7.5	8.6	8.2

Table 5: WER for TDNN-HMM chain models trained with Unspeech embeddings on TED-LIUM.

Context vector	Dev WER		Test WER	
	plain	resc.	plain	resc.
(1) none	9.1	8.5	9.5	9.1
(2) i-vector-sp-ted	8.3	7.5	8.6	8.2
(3) unspeech-64-sp-ted	9.1	8.3	9.6	9.0
(4) unspeech-64-sp-cv	9.1	8.3	9.5	9.1
(5) unspeech-64-sp-cv + (2)	8.4	7.6	8.5	8.1
(6) unspeech-64-tedx	9.0	8.2	9.4	8.7
(7) unspeech-128-tedx	8.9	8.2	9.4	8.9

TED-LIUM dev and test for different Unspeech context embeddings. We trained Unspeech models with different window sizes (64,128) on TED-LIUM (ted) and Common Voice V1 (cv) and computed them for every 10 frames, like the online *i*-vector baseline. While Unspeech embeddings can slightly improve a baseline model trained without any context vectors, with best results obtained when training on the 9500 hours of TEDx data (6,7), using *i*-vectors (2) yields better WERs compared to Unspeech embeddings. Combining Unspeech embeddings trained on Common Voice and *i*-vectors in the input representation can yield slightly lower WERs than *i*-vectors alone (5).

In Table 6 we show WER on decoding utterances from the Common Voice V1 dev and test sets with TDNN-HMM acoustic models trained on TED-LIUM. Utterances from Common Voice are much harder to recognize, since a lot more noise and variability is present in the recordings and the recording have perceivably a much lower signal-to-noise ratio. Since they also contain over 2700 speakers each using an egregious range of microphones, they provide an excellent dev/test to test how robust the TDNN-HMM models are on out-of-domain data. Unsurprisingly, WERs are fairly high compared to the TED-LIUM test set with mostly clean and well pronounced speech. With

Table 6: Decoding Common Voice V1 utterances. Mozilla’s open source dataset provides a challenging test set, which is out-of-domain for an acoustic model trained on TED-LIUM.

Context vector	Dev WER		Test WER	
	plain	resc.	plain	resc.
(1) none	31.2	29.6	29.9	28.5
(2) i-vector-sp-ted	30.3	29.0	29.9	28.2
(3) unspeech-64-sp-cv	29.5	27.9	28.3	26.9
(4) unspeech-64-sp-cv + (2)	29.6	28.2	28.9	27.4
(5) unspeech-64-tedx	30.2	28.8	29.2	27.5
(6) unspeech-128-tedx	30.1	28.7	29.5	28.0

Common Voice we observed that acoustic models trained with Unspeech embeddings consistently resulted in better WERs compared to the baselines, helping the model to adapt. Particularly pre-training Unspeech models on the Common Voice train data help a TDNN-HMM model trained on TED-LIUM to adapt to the style of Common Voice recordings. Embeddings from Unspeech models trained on TEDx will also perform better than the no context and *i*-vector baseline models. In contrast to the results in Table 5, in this decoding task, *i*-vectors in the acoustic model do not provide much of an improvement over the TDNN-HMM baseline model without context vectors.

5. Conclusion

Unspeech context embeddings contain and embed speaker characteristics, but supervised speaker embeddings like *i*-vectors would be better suited for tasks like speaker recognition or authentication. However, clustering utterances according to Unspeech contexts and using the cluster IDs for speaker adaptation in HMM-GMM/TDNN-HMM models is a viable alternative if no speaker information is available. While using Unspeech context embeddings as additional input features did not yield significant WER improvements compared to an *i*-vector baseline on TED-LIUM dev and test, we observed consistent WER reductions with out-of-domain data from the Common Voice corpus when we add Unspeech embeddings. This is a compelling use case of Unspeech context embedding for the adaptation of TDNN-HMM models. Better scores on the same/different speaker similarity task was not indicative of WER reduction – our TEDx Unspeech models scored higher EERs, but were at the same time better context vectors in the acoustic models. We are currently also working on modifying the training objective to see if phonetic Unspeech embeddings can be trained using a similar unsupervised training procedure. Furthermore, we are releasing our source code and offer pre-trained models.³

Acknowledgments. We thank Michael Henretty from Mozilla for giving us access to Common Voice V1 speaker information.

³See <http://unspeech.net>, license: Apache 2.0

6. References

- [1] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors." in *ASRU*, Olomouc, Czech Republic, 2013, pp. 55–59.
- [2] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," in *Readings in speech recognition*, 1990, pp. 393–404.
- [3] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proc. Interspeech*, Dresden, Germany, 2015, pp. 3214–3218.
- [4] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed Representations of Words and Phrases and their Compositionality," in *Proc. NIPS*, Lake Tahoe, NV, USA, 2013, pp. 3111–3119.
- [5] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [6] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, Y. Carmiel, and S. Khudanpur, "Deep neural network-based speaker embeddings for end-to-end speaker verification," in *Proc. Spoken Language Technology Workshop (SLT)*, San Diego, CA, USA, 2016, pp. 165–170.
- [7] D. Snyder, D. Garcia-Romero, D. Povey, and S. Khudanpur, "Deep neural network embeddings for text-independent speaker verification," in *Proc. Interspeech 2017*, Stockholm, Sweden, 2017, pp. 999–1003.
- [8] L. Li, Y. Chen, Y. Shi, Z. Tang, and D. Wang, "Deep speaker feature learning for text-independent speaker verification," in *Proc. Interspeech 2017*, Stockholm, Sweden, 2017, pp. 1542–1545.
- [9] S. Bengio and G. Heigold, "Word embeddings for speech recognition," in *Proc. Interspeech*, Singapore, 2014, pp. 1053–1057.
- [10] H. Kamper, M. Elsner, A. Jansen, and S. Goldwater, "Unsupervised neural network based feature extraction using weak top-down constraints," in *Proc. Acoustics, Speech and Signal Processing (ICASSP)*. South Brisbane, Queensland, Australia: IEEE, 2015, pp. 5818–5822.
- [11] Y.-A. Chung, C.-C. Wu, C.-H. Shen, H.-Y. Lee, and L.-S. Lee, "Audio word2vec: Unsupervised learning of audio segment representations using sequence-to-sequence autoencoder," in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 765–769.
- [12] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context encoders: Feature learning by inpainting," in *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, Caesars Palace, Nevada, United States, 2016, pp. 2536–2544.
- [13] A. Jansen, M. Plakal, R. Pandya, D. P. Ellis, S. Hershey, J. Liu, R. C. Moore, and R. A. Saurous, "Unsupervised learning of semantic audio representations," *arXiv preprint arXiv:1711.02209*, 2017.
- [14] J. Bromley, I. Guyon, Y. LeCun, E. Säcker, and R. Shah, "Signature verification using a "Siamese" time delay neural network," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, Denver, CO, USA, 1994, pp. 737–744.
- [15] G. Synnaeve and E. Dupoux, "Weakly supervised multi-embeddings learning of acoustic models," *arXiv preprint arXiv:1412.6645*, 2014.
- [16] N. Zeghidour, G. Synnaeve, N. Usunier, and E. Dupoux, "Joint learning of speaker and phonetic similarities with siamese networks," in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 1295–1299.
- [17] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *Proc. International Conference on Artificial Intelligence and Statistics*, Sardinia, Italy, 2010, pp. 297–304.
- [18] A. Jati and P. Georgiou, "Speaker2vec: Unsupervised learning and adaptation of a speaker manifold using deep neural networks with an evaluation on speaker segmentation," *Proc. Interspeech 2017*, pp. 3567–3571, 2017.
- [19] —, "Neural predictive coding using convolutional neural networks towards unsupervised learning of speaker characteristics," *arXiv preprint arXiv:1802.07860*, 2018.
- [20] H. Jin, F. Kubala, and R. Schwartz, "Automatic speaker clustering," in *Proceedings of the DARPA speech recognition workshop*, 1997, pp. 108–111.
- [21] A. Senior and I. Lopez-Moreno, "Improving DNN speaker independence with i-vector inputs," in *Proc. Acoustics, Speech and Signal Processing (ICASSP)*, Florence, Italy, 2014, pp. 225–229.
- [22] Y. Miao, H. Zhang, and F. Metze, "Speaker adaptive training of deep neural network acoustic models using i-vectors," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 23, no. 11, pp. 1938–1949, 2015.
- [23] K. Veselý, S. Watanabe, K. Žmolíková, M. Karafiát, L. Burget, and J. H. Černocký, "Sequence summarizing neural network for speaker adaptation," in *Acoustics, Speech and Signal Processing (ICASSP)*, Lujiazui, Shanghai, China, 2016, pp. 5315–5319.
- [24] A. Gupta, Y. Miao, L. Neves, and F. Metze, "Visual features for context-aware speech recognition," in *Proc. Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5020–5024.
- [25] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [26] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *CoRR*, vol. abs/1412.6, 2014.
- [27] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *Proc. ICML*, vol. 30, no. 1, Atlanta, GA, USA, 2013, p. 3.
- [28] A. Rousseau, P. Deléglise, and Y. Estève, "Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks," in *Proc. LREC*, Reykjavik, Iceland, 2014, pp. 3935–3939.
- [29] Mozilla, *Common Voice Corpus VI*, 2018. [Online]. Available: <https://voice.mozilla.org/en/data>
- [30] H. Bredin, "pyannote.metrics: a toolkit for reproducible evaluation, diagnostic, and error analysis of speaker diarization systems," in *Proc. Interspeech*, Stockholm, Sweden, 2017. [Online]. Available: <http://pyannote.github.io/pyannote-metrics>
- [31] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The Kaldi speech recognition toolkit," in *Proc. ASRU*, Atlanta, GA, USA, 2011.
- [32] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Proc. Interspeech*, Dresden, Germany, 2015, pp. 3586–3589.
- [33] L. McInnes, J. Healy, and S. Astels, "hdbscan: Hierarchical density based clustering," *The Journal of Open Source Software*, vol. 2, no. 11, p. 205, 2017.
- [34] B. Zhou and J. H. Hansen, "Unsupervised audio stream segmentation and clustering via the bayesian information criterion," in *International Conference on Spoken Language Processing (ICSLP)*, Beijing, China, 2000, pp. 714–717.
- [35] L. Hubert and P. Arabie, "Comparing partitions," *Journal of classification*, vol. 2, no. 1, pp. 193–218, 1985.
- [36] A. Strehl, "Relationship-based clustering and cluster ensembles for high-dimensional data mining," Ph.D. dissertation, University Of Texas at Austin, 2002.
- [37] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for ASR based on lattice-free MMI," in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 2751–2755.
- [38] D. Povey, X. Zhang, and S. Khudanpur, "Parallel training of DNNs with natural gradient and parameter averaging," *arXiv preprint arXiv:1410.7455*, 2014.