LOW-LATENCY DEEP CLUSTERING FOR SPEECH SEPARATION

Shanshan Wang, Gaurav Naithani, Tuomas Virtanen

Laboratory of Signal Processing, Tampere University of Technology, Tampere, Finland Email:{ shanshan.2.wang, gaurav.naithani, tuomas.virtanen}@tut.fi

ABSTRACT

This paper proposes a low algorithmic latency adaptation of the deep clustering approach to speaker-independent speech separation. It consists of three parts: a) the usage of longshort-term-memory (LSTM) networks instead of their bidirectional variant used in the original work, b) using a short synthesis window (here 8 ms) required for low-latency operation, and, c) using a buffer in the beginning of audio mixture to estimate cluster centres corresponding to constituent speakers which are then utilized to separate speakers within the rest of the signal. The buffer duration would serve as an initialization phase after which the system is capable of operating with 8 ms algorithmic latency. We evaluate our proposed approach on two-speaker mixtures from Wall Street Journal (WSJ0) corpus. We observe that the use of LSTM yields around one dB lower SDR as compared to the baseline bidirectional LSTM in terms of source to distortion ratio (SDR). Moreover, using an 8 ms synthesis window instead of 32 ms degrades the separation performance by around 2.1 dB as compared to the baseline. Finally, we also report separation performance with different buffer durations noting that separation can be achieved even for buffer duration as low as 300 ms.

Index Terms— Monaural speech separation, Low latency, Deep clustering.

1. INTRODUCTION

Single channel speech separation is the problem of recovering the constituent speech signals from an acoustic mixture signal when information from only a single microphone is available [1]. In recent years, data-driven methods relying on deep neural networks (DNN) [2, 3] have yielded dramatic improvements in performance in comparison to the previously used methods, e.g., model-based approaches [4]) and matrix factorization [5, 6]. In particular, speaker-independent speech separation has been addressed by approaches like deep clustering [7, 8], permutation invariant training [9], and more recently deep attractor networks [10] which is the current state-of-the-art. Improvements to the original deep clustering framework [7] have been proposed in terms of, e.g., better objective functions [11]; and, improved regularization and curriculum training [8]. These studies have considered offline separation scenario where the signal to be separated is available at once.

Low-latency processing is important when these DNNbased methods are applied to applications like hearing aids [12] and cochlear implants [13], In particular, for hearing aids, the latency requirements are quite restrictive as the sound is perceived by the listener via hearing aid as well as the direct path. Several studies have documented the subjective disturbance experienced by the listeners (e.g., [14]). Notably, Agnew *et al.* [15] found the delays above 10 ms to be objectionable while delays as low as 3 to 5 ms to be noticeable by hearing-impaired listeners.

For the above applications, the offline DNN-based methods run into two main problems. Firstly, we do not have access to the future temporal information hence DNN models like bidirectional long-short-term-memory networks (BLSTM), as used in [3, 16, 7, 8], cannot be used. Secondly, for shorttime Fourier transform (STFT) based systems, the algorithmic latency is at least equal to the frame length of synthesis window used for signal reconstruction. This limits us from using window sizes used in conventional speech processing (e.g., 20 -40 ms [17]). Speech separation methods with algorithmic latencies below 10 ms have been reported, e.g., using nonnegative matrix factorization [18], and DNNs [19, 20, 21].

In this paper, we investigate a low-latency adaption of the deep clustering framework first introduced in [7]. The original framework involves using a BLSTM network to estimate high-dimensional embeddings for each time-frequency bin in the mixture STFT which is then partitioned into clusters corresponding to the constituent speakers. Our focus in this work is three-fold: a) investigation of separation performance with LSTM instead of BLSTM to allow online processing, b) investigation of separation performance for short synthesis window (8 ms in this work) instead of longer 32 ms used in original work, and c) investigation of using a certain duration in the beginning of acoustic mixture to estimate the cluster centres corresponding to the constituent sources. We refer to this time duration as the buffer. The estimation of cluster centres here thus serves as an initialization phase and the method is capable of doing online separation after the buffer duration.

The authors wish to thank CSC-IT Centre of Science Ltd., Finland, for providing computational resources used in experiments reported in this paper.

We evaluate separately the effect of the above three modifications. We observe one dB lower SDR while using an LSTM instead of the BLSTM network as was used in the original work [7]. The separation performance degrades by around 2.1 dB while shortening the window length from 32 ms to 8 ms as compared to the baseline. Moreover, we show that it is possible to estimate reasonably well cluster centres using just the beginning of the signal yielding good separation for the rest of the signal.

The paper is structured as follows: Section 2 describes the baseline deep clustering approach proposed in [7]. Section 3 describes its low-latency adaptation. Section 4 describes the evaluation procedure, experimental set up, and obtained results. Finally, Section 5 concludes the paper.

2. DEEP CLUSTERING FOR SPEAKER SEPARATION

In this section, we summarize the deep clustering method proposed in [7, 8]. Deep clustering can be thought of as a combination of supervised learning and unsupervised learning. Unlike the traditional DNN-based speech separation methods that predict a time-frequency mask or separate speech spectrum for the mixture input in a supervised manner [2, 3], it generates an embedding vector for each time-frequency bin and then uses the unsupervised learning approach such as kmeans to cluster the embedding vectors in order to get the time-frequency masks.

Given a mixture audio signal in the time domain x(n), firstly, features are extracted by calculating its log magnitude short-time Fourier transform (STFT). The features are then inputted to a neural network that will output an embedding vector for each of the time-frequency points. In the original deep clustering framework, BLSTM network was used [7], and therefore we choose it as the baseline here. The output of the neural networks is an embedding matrix $\mathbf{V} \in \mathbb{R}^{TF \times D}$, where T denotes the number of frames, F the number of frequency bins, and D the embedding dimension. Finally, kmeans clustering is employed to partition the embedding vectors into clusters corresponding to different constituent speakers. Binary time-frequency masks for each speaker is then obtained using these cluster assignments by assigning 1 to all the time-frequency bins within the cluster of the speaker, and 0 to the rest of the bins.

The neural network is trained to minimize the difference between the estimated affinity matrix $\mathbf{VV}^{\mathcal{T}}$ derived from the embeddings \mathbf{V} predicted by the neural network and the target affinity matrix $\mathbf{YY}^{\mathcal{T}}$, where $\mathbf{Y} \in \mathbb{R}^{TF \times C}$ is the ideal binary mask. *C* indicates the number of speakers in the mixture. The training loss function \mathcal{L} is computed as,

$$\mathcal{L} = \left\| \mathbf{V}\mathbf{V}^{\mathcal{T}} - \mathbf{Y}\mathbf{Y}^{\mathcal{T}} \right\|_{F}$$
$$= \left\| \mathbf{V}^{\mathcal{T}}\mathbf{V} \right\|_{F}^{2} - 2\left\| \mathbf{V}^{\mathcal{T}}\mathbf{Y} \right\|_{F}^{2} + \left\| \mathbf{Y}^{\mathcal{T}}\mathbf{Y} \right\|_{F}^{2}, \quad (1)$$



Fig. 1. The block diagram of the proposed low-latency deep clustering method.

where F denotes the Frobenius norm of the matrix. In order to remove the contribution of noisy/silent regions in the network training, a voice active detection (VAD) threshold is employed. Only the embeddings corresponding to timefrequency bins with magnitude greater than the VAD threshold (-40 dB below the maximum amplitude as in [7]) contribute to the above loss calculation.

It should be noted that at the test stage the k-means algorithm is employed to cluster the embeddings using the entire test signal, thus making the method unsuitable for low-latency processing. In the test stage, the estimated binary masks are applied to the complex spectrogram of the mixture hence mixture phase is utilized. Inverse STFT and overlap-add processing is applied to obtain separated signals in the time domain.

3. LOW-LATENCY DEEP CLUSTERING

In order to make the deep clustering based separation operate with low latency, there are three parts that need to be adapted: *a*) The topology of the neural network is changed from BLSTM as was used in [7] to LSTM in order to produce embedding vectors in an online manner for each frame as they are inputted to the network; *b*) In the baseline method [7], 32 ms synthesis window length is used. The resulting latency may be prohibitive for certain applications, e.g., hearing aids [15] as was explained in the introduction. Hence we shorten the window length to 8 ms; and; c) Instead of using the whole signal, we propose using only a certain length in the beginning of the mixture, which we refer as the buffer, to get the cluster centres. These cluster centres can be then used to estimate the masks for the rest of the mixture. Please note that since the first few seconds of the signal are used to estimate the cluster centres, the method is not able to separate sources during this initialization stage. However, after the buffer, the rest of the signal will be processed in an online manner. The process of the low-latency version of deep clustering method is depicted in Fig. 1 with a buffer duration of 1.5 s.

4. EVALUATION

4.1. Acoustic material

The evaluation is done using synthetic two-speaker mixtures generated from Wall Street Journal (WSJ0) corpus. The duration of mixtures is on average around 6 s. The training data set consists of 20,000 two-speaker mixtures created by randomly selecting utterances from 101 different speakers from WSJ0 *si_tr_s* that amounts to around 30 hours of training material. Similarly, for the validation data set, we create 5000 two-speaker mixtures that last for around 8 hours having the same speakers as in training data set. The test data is generated from WSJ0 *si_dt_05* and *si_et_05* and consists of 3000 mixtures and lasts around 5 hours having 18 different speakers. The test data has different speakers from training data and validation data for the purpose of evaluating the separation performance in open conditions as described in [7].

We downsample the speech samples from 16 kHz to 8 kHz for reducing the computational requirements and to make the evaluation setup similar to [7]. As the proposed approach (factor c) relies on both the speakers being active during the buffer duration, for a fair investigation of the effect of buffer duration and comparison of offline deep clustering to its online counterpart, the same data should be used for evaluation. Hence in the test set for (c) we firstly remove the silence from the beginning of both speech signals and sum them to form the mixture thus ensuring that both speakers are active during the buffer duration (\geq 100 ms in this work, i.e., all mixtures have both speakers active within at least 100 ms in the beginning). The longer speech signal is trimmed to the length of the shorter utterance before adding to form the mixture. It should also be noted that such test mixtures have a larger degree of overlapped speech and are thus harder to separate.

4.2. Metrics

We use BSS-EVAL toolbox [22] for evaluating the system performance. It consists of three metrics: signal-to-distortionratio (SDR), signal-to-interference-ratio (SIR), and signal-toartifacts-ratio (SAR). The average SDR of test mixtures without any separation is 0.1 dB.

4.3. Experiment setup

In order to analyze the effect of the following different factors, namely, a) BLSTM vs LSTM networks, b) 32 ms vs 8 ms window length, and, c) different buffer duration for low-latency process, we conduct separate experiment for each of these.

The baseline framework is taken to be the one used in [7]. It consists of a BLSTM network with four layers and 600 units in each layer followed by a time-distributed dense layer. The number of units in the time distributed dense layer is the product of the number of embedding dimensions and the number of effective FFT points. Hyperbolic tangent (tanh) is used as the activation function in this layer. After the dense layer, L2 normalization is used to bound the embedding vectors to unit norm. The same parameters have been used for the LSTM network in order to analyze the effect of factor a). To compare the effect of different window length, the same LSTM network and a shorter window length of 8 ms for STFT feature extraction are used. For a fair comparison with the baseline, the network must be trained with the sequences having the same time context. The baseline BLSTM was trained on 100 frame sequences (800 ms). Here we reduce the hop length to 4 ms hence the sequence length is increased to 200 (800 ms). We first train the network for 100 frame sequences and then after convergence continue training with 200 frame sequences, known as curriculum learning used in [8] and first introduced in [23]. The idea is to pre-train a network on an easier task first improves learning and generalization. Finally, c) is studied by varying the buffer duration using the network with the same LSTM network (four layers and 600 units in each layer) with 8 ms window length. The same FFT size (256) is used for both offline (32 and 8 ms frame length) and online deep clustering (8 ms frame length) frameworks, i.e.,

Table 1. Feature and system parameters for offline and online deep clustering experiments.

	offline DC	low-latency DC		
Window length	32 ms	8 ms		
Hop length	8 ms	4 ms		
Sequence length	100	200		
Network	BLSTM	LSTM		
Window	Hanning			
Sampling frequency	8 kHz			
FFT size	256			
Number of layers	4			
Number of LSTM units	600			
Embedding dimension	40			



Fig. 2. Evaluation metrics (dB) for LSTM network with different buffer durations (factor c in the experimental set up.)

zero padding is used wherever required.

During the training process, the 'Adam' optimizer is utilized [24]. The Keras [25] and Tensorflow [26] libraries are used for network training, and Librosa library [27] is used for feature extraction and signal reconstruction in this paper. In order to reduce overfitting, early stopping method is used [28] by monitoring the loss on validation data and stopping the training when no decrease in it is observed for 30 epochs. The embedding dimension is set to 40 and the VAD threshold is 40 dB, similar to the original study [7]. The detailed description of the parameters for the network can be found from Table 1.

It should be noted that Fig. 1 depicts the real world use case where a buffer duration in the beginning of an utterance is used for estimating clusters and the separation starts after that. This however makes acoustic material used for evaluation dependent upon the buffer duration if the same utterance is used for both cluster estimation and evaluation. To deal with this mismatch, for each test utterance, we randomly select another utterance (cluster utterance) belonging to the same speaker pair and use it to estimate clusters. Different buffer lengths can thus be sampled from the beginning of this cluster utterance for the same test utterance in order to study the effect of factor c. Moreover, the VAD threshold is used during cluster estimation to exclude the effect of noisy time-frequency bins.

4.4. Results and discussion

We calculate the mean of evaluation metrics, SDR, SIR, and SAR over all the test mixtures. All the test mixtures are formed such that both constituent speech signals are active within the lowest buffer duration used in experiments (100 ms here). Hence they have a higher overlap between the constituent speech. As described in the previous section, we

Table 2. Evaluation metrics (dB) of different variants of the offline method and the online method with 1.5s buffer. Here online refers to factor c in the experimental setup

	Window length	SDR	SIR	SAR
BLSTM	32 ms	7.9	15.6	9.2
LSTM	32 ms	6.9	14.5	8.4
LSTM	8 ms	5.8	13.6	7.2
Online LSTM (1.5s buffer)	8 ms	5.1	12.6	6.7

adopt the strategy of estimating clusters on different utterances than test utterances. Hence the same dataset can be used for evaluation for both offline and online deep clustering methods. The results with baseline offline case is shown in Table 2. Similarly, the evaluation metrics corresponding to the LSTM network with 32 and 8 ms window lengths is shown as well. The online LSTM in Table 2 refers to lowlatency LSTM with 8 ms window length and 1.5 s buffer time. It can be observed that the separation performance is one dB lower in terms of SDR while replacing BLSTM to LSTM while keeping the same window length. Moreover, by decreasing the window length to 8 ms, SDR degrades by about 2.1 dB as compared to the baseline.

The effect of varying buffer duration on performance metrics is shown in Fig. 2. Two interesting observations can be made from it: firstly, even with a short buffer duration, e.g., 100 ms, relatively reasonable separation performance can still be achieved (4.5 dB); and secondly after a certain buffer length more information does not lead to a drastic improvement in separation performance. This means a small buffer duration, even as low as 300 ms, can yield good separation provided both the constituent speakers are active during it.

5. CONCLUSION

The paper proposes a low-latency adaptation of deep clustering based speech separation. In particular, a buffer signal duration in the beginning of audio mixture is used for estimating cluster centres corresponding to the speakers present in the mixture. This duration serves as an 'initialization' period after which the rest of the speech mixture is processed in online manner. Moreover, separation performance of the method using an LSTM network and a short synthesis window length of 8 ms, as required by real-time operation, has been studied. A degradation in SDR of about one dB is observed for the former and 2.1 dB for the latter as compared to the baseline. Finally, we investigate how the buffer duration affects the separation result and observe that even very short buffer duration, e.g. 300 ms, is sufficient to estimate clusters for reasonable separation.

6. REFERENCES

- [1] E. Vincent, T. Virtanen, and S. Gannot, *Audio source separation and speech enhancement.* John Wiley & Sons, 2018.
- [2] P. Huang, M. Kim, M. Hasegawa-Johnson, and P. Smaragdis, "Deep learning for monaural speech separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 1562–1566.
- [3] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing* (ICASSP), 2015, pp. 708–712.
- [4] S. T. Roweis, "One microphone source separation," in Proc. Advances in Neural Information Processing Systems (NIPS), 2001, pp. 793–799.
- [5] T. Virtanen, "Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 3, pp. 1066–1074, 2007.
- [6] M. N. Schmidt and R. K. Olsson, "Single-channel speech separation using sparse non-negative matrix factorization," in *Proc. International Conference on Spoken Language Process*ing, 2006.
- [7] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 31–35.
- [8] Y. Isik, J. L. Roux, Z. Chen, S. Watanabe, and J. R. Hershey, "Single-channel multi-speaker separation using deep clustering," in *Proc. Interspeech*, 2016, pp. 545–549. [Online]. Available: http://dx.doi.org/10.21437/Interspeech.2016-1176
- [9] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multitalker speech separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 241–245.
- [10] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 246–250.
- [11] Z.-Q. Wang, J. Le Roux, and J. R. Hershey, "Alternative objective functions for deep clustering," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*), 2018, pp. 686–690.
- [12] L. Bramsløw, "Preferred signal path delay and high-pass cut-off in open fittings," *International Journal of Audiology*, vol. 49, no. 9, pp. 634–644, 2010.
- [13] J. Hidalgo, "Low latency audio source separation for speech enhancement in cochlear implants," Master's thesis, Universitat Pompeu Fabra, 2012.
- [14] M. A. Stone, B. C. Moore, K. Meisenbacher, and R. P. Derleth, "Tolerable hearing aid delays. V. estimation of limits for open canal fittings," *Ear and Hearing*, vol. 29, no. 4, pp. 601–617, 2008.

- [15] J. Agnew and J. M. Thornton, "Just noticeable and objectionable group delays in digital hearing aids," *Journal of the American Academy of Audiology*, vol. 11, no. 6, pp. 330–336, 2000.
- [16] Z. Chen, S. Watanabe, H. Erdogan, and J. R. Hershey, "Speech enhancement and recognition using multi-task learning of long short-term memory recurrent neural networks," in *Proc. Inter*speech, 2015.
- [17] K. K. Paliwal, J. Lyons, and K. Wojcicki, "Preference for 20-40 ms window duration in speech analysis," in *Proc. International Conference on Signal Processing and Communication Systems* (ICSPCS), 2011, pp. 1 – 4.
- [18] T. Barker, T. Virtanen, and N. H. Pontoppidan, "Low-latency sound-source-separation using non-negative matrix factorisation with coupled analysis and synthesis dictionaries," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 241–245.
- [19] G. Naithani, T. Barker, G. Parascandolo, L. Bramsløw, N. H. Pontoppidan, and T. Virtanen, "Low latency sound source separation using convolutional recurrent neural networks," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2017, pp. 71–75.
- [20] G. Naithani, G. Parascandolo, T. Barker, N. H. Pontoppidan, and T. Virtanen, "Low-latency sound source separation using deep neural networks," in *Proc. IEEE Global Conference* on Signal and Information Processing (GlobalSIP), 2016, pp. 272–276.
- [21] Y. Luo and N. Mesgarani, "TasNet: time-domain audio separation network for real-time single-channel speech separation," in *Proc. IEEE International Conference on Acoustics, Speech* and Signal Processing (ICASSP), 2018, pp. 696–700.
- [22] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Transactions* on Audio, Speech, and Language Processing, vol. 14, no. 4, pp. 1462–1469, 2006.
- [23] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proc. International Conference on Machine Learning (ICML)*, 2009, pp. 41–48.
- [24] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. International Conference on Learning Representations*, 2014.
- [25] F. Chollet *et al.*, "Keras," 2015. [Online]. Available: https://github.com/fchollet/keras
- [26] M. Abadi *et al.*, "Tensorflow: A system for large-scale machine learning," in *Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2016, pp. 265–283.
- [27] B. McFee *et al.*, "librosa 0.5.0," Feb. 2017. [Online]. Available: https://doi.org/10.5281/zenodo.293021
- [28] R. Caruana, S. Lawrence, and L. Giles, "Overfitting in neural nets: Backpropagation, conjugate gradient, and early stopping," in *Proc. Advances in Neural Information Processing Systems (NIPS)*, vol. 13, 2001, p. 402.