



Analysis of Deep Clustering as Preprocessing for Automatic Speech Recognition of Sparsely Overlapping Speech

Tobias Menne^{1,2}, Ilya Sklyar¹, Ralf Schlüter¹, Hermann Ney^{1,2}

¹Human Language Technology and Pattern Recognition, Computer Science Department, RWTH Aachen University, 52074 Aachen, Germany

²AppTek GmbH, 52062 Aachen, Germany

{menne, schlueter, ney}@cs.rwth-aachen.de
sklyar@i6.informatik.rwth-aachen.de

Abstract

Significant performance degradation of automatic speech recognition (ASR) systems is observed when the audio signal contains cross-talk. One of the recently proposed approaches to solve the problem of multi-speaker ASR is the deep clustering (DPCL) approach. Combining DPCL with a state-of-the-art hybrid acoustic model, we obtain a word error rate (WER) of 16.5% on the commonly used wsj0-2mix dataset, which is the best performance reported thus far to the best of our knowledge. The wsj0-2mix dataset contains simulated cross-talk where the speech of multiple speakers overlaps for almost the entire utterance. In a more realistic ASR scenario the audio signal contains significant portions of single-speaker speech and only part of the signal contains speech of multiple competing speakers. This paper investigates obstacles of applying DPCL as a preprocessing method for ASR in such a scenario of sparsely overlapping speech. To this end we present a data simulation approach, closely related to the wsj0-2mix dataset, generating sparsely overlapping speech datasets of arbitrary overlap ratio. The analysis of applying DPCL to sparsely overlapping speech is an important interim step between the fully overlapping datasets like wsj0-2mix and more realistic ASR datasets, such as CHiME-5 or AMI.

Index Terms: deep clustering, ASR, speaker separation, multi-speaker ASR

1. Introduction

The performance of ASR systems on relatively clean, close-talking recordings has been improved drastically over the recent years. This scenario can e.g. be found in telephony speech or readings of audio books. On standard tasks for this scenario, as switchboard and librispeech [1, 2], typical WERs are below 10%. Nevertheless, ASR on noisy data remains challenging. ASR performance decreases especially when audio is recorded from a larger distance or when multiple speakers are talking simultaneously [3, 4]. There has thus been a growing interest in multi-speaker ASR. A special focus in this area lies on single-channel recordings [5, 6, 7]. This scenario is not only of interest if only one recording channel can be obtained, but also if multi-channel processing steps, like beamforming, can not separate two speakers because they are spatially too close to each other.

In [6] a purely end-to-end system has been proposed which aims at directly recognizing multiple speakers with an attention based sequence-to-sequence model. This system employs no separate source separation stage. Other recently proposed solutions use a separate source separation stage. This preprocessing usually employs masks in the time-frequency domain to separate multiple sources from the mixture signal. One method to obtain those masks, is to directly infer them using an artificial neural network (ANN). This ANN is trained by a permutation-free objective function [8, 5, 9]. The method is often referred to as permutation invariant training (PIT). The integrated application of PIT on the level of the ASR cost function is presented in [10]. A different approach to obtain the masks, is through the utilization of embedding vectors in the time-frequency do-

main. This is done in the DPCL and deep attractor network approaches [8, 11, 12]. The focus of the work presented here lies on the DPCL approach. In DPCL a ANN is trained to map each time-frequency bin to an embedding vector. Those embeddings are then used to allocate the time-frequency bins to different speakers through k-means clustering. The objective function is designed such that embedding vectors of time-frequency bins belonging to the same speaker are close to each other. Embeddings belonging to different speakers have a larger distance.

DPCL has shown good potential when used in a preprocessing step for an ASR system based on a Gaussian mixture model (GMM)-hidden Markov model (HMM) acoustic model [11]. Good performance was also achieved in combination with a sequence-to-sequence ASR system [13]. Results on the combination with a state-of-the-art hybrid deep neural network (DNN)-HMM system have, to the best of our knowledge, not been published so far. Thus DPCL tends to perform worse than the sequence-to-sequence and PIT approaches in recent comparisons on wsj0-2mix [6, 7]. By combining DPCL with a state of the art hybrid DNN-HMM we obtain a WER of 16.5%, which to the best of our knowledge is the best performance reported on wsj0-2mix thus far.

Furthermore we present an in depth analysis of DPCL as a preprocessing step for ASR on a more realistic scenario of only sparsely overlapping speech. In this scenario only a small part of the signal contains multi-speaker segments whereas the majority of the signal contains single-speaker segments. This scenario is much closer to the scenarios of meeting room or smart home recordings. To this end we introduce a data simulation approach to obtain sparsely overlapping speech with a fixed but configurable overlap ratio based on the Wall Street Journal (WSJ) data. This data bridges the gap between wsj0-2mix and more realistic datasets as for example Chime-5 [3] and AMI [4]. The use of a simulated dataset offers the advantage of investigating various overlap ratios. Furthermore it allows the utilization of oracle knowledge to analyze separate aspects of the pipeline and analyze their influence on ASR performance in a controlled environment. Our analysis points out obstacles of applying DPCL to sparsely overlapping data which are cloaked in the experiments on fully overlapping signals. We also propose initial solutions.

The paper is organized as follows. An overview of the datasets used in this work is given in Section 2. Section 3 describes the DPCL and ASR system, before the experimental setup and results are presented in Section 4 and discussed in Section 5. Potential future research directions are discussed in Section 6.

2. Data

We report ASR results on the commonly used dataset wsj0-2mix introduced in [8]. This dataset is created by artificial mixing of speakers from the WSJ data. The main drawback of this datasets for ASR experiments is, that all generated utterances contain fully overlapping speech. This means that both

speakers talk for almost the complete length of the utterance. In a realistic ASR scenario for overlapping speech, as they can be found e.g. in the CHIME-5 or AMI datasets [3, 4], speakers are only overlapping for smaller portions of an utterance. This means that each utterance contains significant portions of single-speaker speech.

To study this scenario of more sparsely overlapping speech in the effect of DPCL as a preprocessing step for ASR, we create datasets containing sparsely overlapping speech utterances. Other aspects of the artificial mixing, such as the signal-to-noise ratio (SNR) distribution, are kept as similar as possible to the algorithm used in [8]. The data simulation pipeline is described in the following.

Two separate signal tracks are generated, each containing speech from a single speaker. The ASR system described in Section 3.1 is used to obtain a forced alignment for the source datasets, where the speech segments are sampled from. This alignment is used to cut leading and trailing silences from the utterances. This ensures that only pauses in between words remains as silence in those utterances, which is neglectable in the computation of the overlap ratio. Furthermore, a silence set containing those leading and trailing silence segments is created.

For each speaker three utterances are sampled from the source dataset. One signal track is created for each of the two speakers where the sampled segments are separated by silence gaps. The lengths of the silence gaps are randomly sampled with the constraint that the overlap of speech after adding the two signal tracks has a given overlap ratio and that the ratio of the mixed signal containing no speech does not exceed a certain threshold (here 10%). The silence gaps are then filled with silence signals sampled from the silence set mentioned above, where the energy of the silence segments used to fill the gap is scaled to the leading and trailing silences of the original speech segments. The two signal tracks are then mixed with a given SNR value similar to the data simulation of wsj0-2mix, where the signal energy of the two signal tracks is computed by only considering the speech segments and not the silence segments.

3. System

3.1. ASR system

A state of the art hybrid DNN-HMM acoustic model, trained on the WSJ-SI84 subset (15 h) of the WSJ dataset, is used for the experiments. The input features are unnormalized 80 dimensional log-Mel filterbank features based on a short-time Fourier transform (STFT) employing the Hanning window applied to a 25 ms frame with a frame shift of 10 ms. Since the input features are unnormalized the first layer of the acoustic model is an 80 dimensional linear layer employing batch normalization [14]. The linear layer is followed by 5 bidirectional long short-term memory (BLSTM) layers with 600 units each. The output is a softmax layer with 1501 units. A 3-gram language model is used during recognition. Table 1 shows the ASR performance of the system on the standard 5k and 20k development and evaluation datasets of WSJ, using the respective language model is used.

The system is implemented using RETURNN and RASR [15, 16].

Table 1: WER in % of the ASR system on the standard development and evaluation sets of the WSJ dataset. The data contained in the sets is different for 5k and 20k.

Vocabulary size	dev93	eval92	eval93
5k	4.6	1.8	4.0
20k	10.0	6.9	9.4

3.2. Source separation system

Source separation and ASR are handled in two separate stages. The source separation is done by applying DPCL to the mixed speech signal creating one signal per speaker, which is referred

to as speaker track in the following. The speaker tracks are then fed separately into the ASR system. The following sections give a quick summary of DPCL and how it is applied here.

3.2.1. Deep clustering network

The network architecture for DPCL described in [11] was reimplemented using RETURNN [15]. The architecture consists of an ANN, which computes a 40 dimensional embedding vector for each time-frequency bin of the input signal. As input features the STFT of the input signal is computed with a window size of 32 ms, a frame shift of 8 ms and a discrete Fourier transform (DFT) of dimension 512 is used as input features. The embedding vectors are used to cluster the time-frequency bins into multiple classes (one for each speaker) using soft clustering. A binary mask is generated from the classification of the time-frequency bins. Those masks are applied to the input signal obtaining a separate speaker track for each speaker. The resulting signals are the input to an enhancement network as described in [11]. The architecture of the embedding network consists of 4 BLSTM layers with 600 units each. Curriculum learning is applied as described in [11] with an input size of 100 frames for 100 epochs and 400 frames for 100 epochs. The architecture of the enhancement network consists of 2 BLSTM layers with 300 units each. The signal to distortion ratio (SDR) improvements obtained by this systems on the wsj0-2mix dataset are shown in Table 3 and are in line with the results described in [11].

3.2.2. Application of DPCL to sparsely overlapping speech

Applying DPCL to sparsely overlapping speech signals as described in Section 2 can be done in the same manner as for the fully overlapping speech. This approach is referred to as full-sequence mode hereafter. This approach can potentially suffer from signal quality degradation of the single speech segments, due to erroneous masking. An alternative approach is to apply DPCL to the multi-speaker segments only. This second approach requires to deal with the segmentation problem, meaning how to separate the input signal into multi-speaker and single-speaker segments. We experimented with various ways to solve the segmentation problem, but those results will be presented in future work, since they go beyond the scope of this work and do not serve to further the conclusions presented here. The results reported here use the oracle knowledge for segmentation.

The embedding vectors are computed for the complete signal. The single-speaker segments remain unprocessed, while one output signal per speaker is generated for the multi-speaker segments by computing the masks based on the embeddings of only that segment. This creates a segment permutation problem, where the resulting output segments need to be allocated to an output speaker track. For the experiments presented here a fixed number of 2 speaker tracks is used.

Three approaches to handle the permutation problem are used in this work. First the oracle knowledge is used. This is done by computing the correlations to the respective segment of the source signals tracks for each of the outputs per segment. The output segment is allocated to the source signal track with the higher correlation and thus to a speaker track.

The second approach is hereafter referred to as affinity approach. In this approach the mean of the embedding vectors for each speaker in the multi-speaker segments is calculated. For each possible permutation of multi-speaker segments the average distance of the resulting group of mean vectors is computed and the permutation with the lowest average distance is selected. Then a mean embedding vector for each speaker track based on the selected permutation is computed. The single-speaker segments are then allocated to the speaker track which mean embedding vector is closest to the mean embedding vector of the single-speaker segment.

The third approach is hereafter referred to as speaker-Id approach. In this approach the DPCL network is trained in a multi-task approach similar to [17]. The network is extended by a second output which is utilized for speaker identification. Different embedding vectors are computed for the speaker identi-

fication part of the network. Those speaker-Id embedding vectors are then used to handle the permutation problem in the same manner as is done in the affinity approach. Details about the network architecture and the cost function are described in [17]. The main difference of our network architecture to the one presented in [17] is that the deep attractor network from [17] is replaced by the DPCL network described above and the dimension of the speaker-Id embedding vectors which is 40. Furthermore both cost functions are weighted equally in the multi-task training.

4. Experimental setup and results

Table 2 shows the WER of the system on the fully overlapping dataset wsj0-2mix described in Section 2. The table also shows the more recently published performance of sequence-to-sequence systems. In past publications those systems have only been compared to a system employing DPCL in combination with a GMM-HMM acoustic model [11] and have been shown to yield better WERs [6, 7]. But the results in Table 2 show DPCL to be superior to the integrated approaches, when combined with a state of the art acoustic model.

Table 2: WER in % of the described system compared to recently published results on wsj0-2mix. The acoustic models from [6] and [7] differ in details.

System				Eval
Separation	Acoustic model	Language model	Reference	
DPCL	GMM-HMM	n/a	[11]	30.8
	DNN-HMM	3-gram	proposed here	16.5
integrated	joint CTC/attention	word- & char-level RNNLMs	[6]	28.2
		word-level RNNLMs	[7]	25.4

The results in Table 3 show, that DPCL is much more reliable, when the competing speaker is of different gender to the dominant speaker. This effect can also be seen in the signal quality metric SDR as presented in previous publications [11] and confirmed by our experiments. In our experiments the effect seems to be slightly stronger for WER than for SDR.

Table 3: WER and SDR of the proposed system differentiated by gender of competing speaker and speaker dominance.

Speaker	Gender	Eval WER (%)	Eval SDR
dominant	same	17.6	10.7
	diff	12.1	12.7
non dominant	same	22.2	8.0
	diff	14.0	10.1

The drawback of the wsj0-2mix datasets is, that it does not cover scenarios in which a significant portion of the speech signal contains single-speaker segments and only part of the signal contains multi-speaker segments. This is what one would e.g. expect in meeting recordings or the smart home scenario. The following experiments investigate the additional obstacles that those scenarios pose for the application of DPCL as a preprocessing for ASR.

The data used in the following experiments has been created as described in Section 2. The data for the evaluation set is sampled from `si_dt_05` and `si_et_05`, which are both not used for training, cross validation or hyper parameter tuning. We chose those source datasets to stay as close as possible to the original wsj0-2mix dataset presented in Table 2 and 3 to make the following results most comparable to the fully overlapping

scenarios. As before decoding was done using a 3-gram language model with a vocabulary size of 20k.

The WER on the separate signal tracks before mixing can be considered a lower boundary for the WERs and is referred to as clean in the following. The WER on the mixed signal is dominated by insertion errors induced by the single-speaker speech segments of the competing speaker. Therefore this WER is less useful to investigate degradation which stem from the multi-speaker segments. An alternative sensible upper reference for the WER can be described by a perfect speaker identification system, which allocates multi-speaker segments to both speakers. In our experiments this is the same as not applying any source separation to the multi-speaker segments and using oracle knowledge for segmentation and permutation. Figure 1 shows the WERs of the various processing approaches over the evaluation sets with various overlap ratios.

5. Discussion

As expected the ASR performance strongly improves with decreasing overlap ratio when no separation is applied to the multi-speaker segments and oracle knowledge for segmentation and permutation is used. On the other hand the ASR performance when applying DPCL in full-sequence mode, decreases for decreasing overlap ratios. The results shows that even though DPCL works extremely well on fully overlapping speech a simple direct application to sparsely overlapping speech could potentially even hurt ASR performance compared to no processing of the mixed signal.

Figure 1d shows, that if oracle knowledge is used for the segmentation and permutation problem and DPCL is only applied to the multi-speaker segments, the WER improves for decreasing overlap ratios. The gap between the application of DPCL in full-sequence mode and its application to only the multi-speaker segments with utilization of oracle knowledge for segmentation and permutation shows the maximum potential performance gain that can be obtained if the segmentation and permutation problems are solved optimally. Especially for low overlap ratios the potential for improving ASR performance is large. Closing this gap is crucial for the applicability of DPCL in real world ASR.

The intuitive expectation when applying DPCL to only the multi-speaker segments with utilization of oracle knowledge for segmentation and permutation is an increase in WER for increasing overlap ratios, since larger portions of the signal will suffer from quality degradation due to the required source separation. But Figure 1d, shows only a minor decrease in ASR performance over increasing overlap ratios. A differentiation of scenarios in which the competing speaker has the same or different gender as the dominant speaker as done in Figures 1e and 1f reveal that the decrease in ASR performance for larger overlap ratios stems almost exclusively from the same gender scenario.

Furthermore the performance of DPCL applied in full-sequence mode is significantly worse for the same gender scenario throughout all overlap ratios. One explanation for the higher difficulty of the same gender scenario is that mask based separation approaches rely on sparsity of the acoustic features along the frequency domain, which is more dominant in the different gender scenario. This can explain why, in the same gender scenario, the performance increases more strongly with decreasing overlap ratio, when applying DPCL to only the multi-speaker segments with utilization of oracle knowledge for segmentation and permutation. If the sparsity is also the main obstacle for the application of DPCL on the full-sequence mode, the performance gap between the different gender scenario and the same gender scenario should shrink for decreasing overlap ratios. This is not the case as can be seen in Figures 1e and 1f. This indicates that the main problem is not the masking of the multi-speaker segments, but the handling of single-speaker segments.

When the oracle permutation is replaced by the affinity approach described in Section 3.2 it can be observed that the WERs are very similar to the WERs of applying DPCL in full-sequence

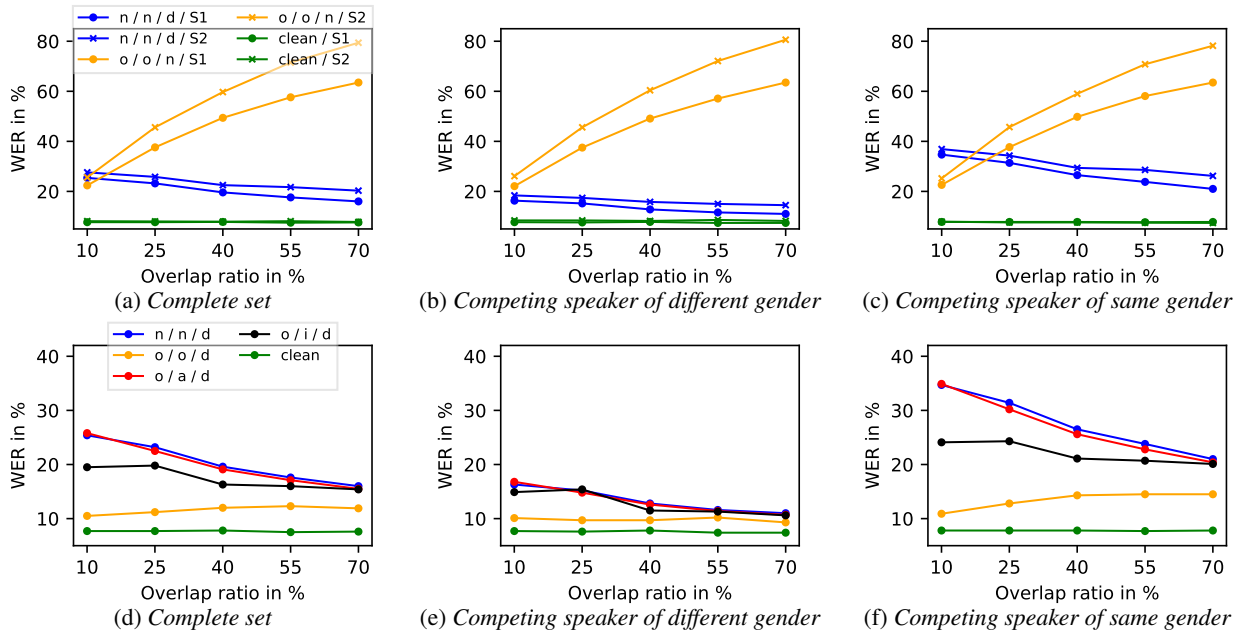


Figure 1: Analysis of DPCL for various overlap ratios. a)-c) compare DPCL to an upper and a lower performance bound. d)-f) compare, for the dominant speaker, the application of DPCL in full-sequence mode and the application of DPCL on the multi-speaker segments only. The legend syntax is as follows: segmentation / permutation / source separation [/ speaker dominance] n = none, o = oracle, a = affinity permutation, i = speaker-Id permutation, d = DPCL, $S1$ = dominant speaker, $S2$ = non dominant speaker

mode. Combined with the findings described above, this indicates that the main obstacle for the application of DPCL to sparsely overlapping speech is not the potential signal degradation from the masking of the single- or multi-speaker segments, but the collective allocation of the time-frequency bins of the single-speaker segments to a speaker in the multi-speaker segments.

A straight forward solution to this problem could be to provide the DPCL network with sparsely overlapping data during training. To this end we trained multiple DPCL networks with training data containing various average overlap ratios. The data was simulated as described in Section 2 and the network was trained as described in Section 3.2. The total amount of source data used for the simulation has been kept constant for a fair comparison with the network trained on fully overlapping signals. Without adjustment of the training cost function we were not able to gain an improvement by only changing the training data in that manner. Investigations into this approach will be future work.

A different solution is to improve the handling of the permutation problem when handling segmentation, permutation and source separation separately. For this we utilize the speaker-Id approach described in Section 3.2.2. With this approach we were able to get improvements of about 20% relative for low overlap ratios. It can also be seen that this improvement is mainly due to the same gender scenario and that the effect vanishes for higher overlap ratios.

6. Conclusions

In this work we have shown that the use of DPCL as a separate speaker separation step for multi-speaker ASR works well on the standard wsj0-2mix dataset if it is combined with a state of the art DNN-HMM acoustic model. To the best of our knowledge the WER of 16.5% obtained by this system is currently the lowest WER reported on wsj0-2mix. Furthermore we presented in depth investigations on the effects of DPCL as a preprocessing step for ASR on sparsely overlapping speech.

To this end we simulated data with varying overlap ratio of the competing speakers. Those experiments on simulated data aim to further the utilization of DPCL to improve ASR performance on real data as e.g. the AMI or CHiME-5 datasets.

The results presented here indicate the main obstacles to obtain similar ASR gains on real data as can be seen for the wsj0-2mix data. More specifically it has been shown that a major drawback of the basic DPCL approach is the handling of single-speaker segments. The results indicate that the main reason is not a degradation of signal quality of the single-speaker segments by erroneous masking, but rather the problem of allocating the time-frequency bins of a single-speaker segment to a speaker of the multi-speaker segments. The results show that a promising approach is to separate the source separation into three different steps. The first step is the segmentation of the signal into single-speaker and multi-speaker segments. In a second step DPCL is applied to the multi-speaker segments only. Finally the problem of allocating the resulting segments to a speaker track needs to be solved. For this problem we have presented an approach based on speaker identification, which improved the results more than 20% relative for low overlap ratios.

Future work will focus on the modification of DPCL to obtain a better handling of single-speaker segments in the full-sequence approach. This could be attempted by introducing a regularizing term in the cost function during training on sparsely overlapping training data. Furthermore we will explore the use of feedback from the acoustic model to solve the segmentation and permutation problems in the separated approach.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement No. 694537 and under the Marie Skłodowska-Curie grant agreement No. 644283 and from a Google Focused Award. The work reflects only the authors' views and none of the funding agencies is responsible for any use that may be made of the information it contains.

References

- [1] W. Xiong, L. Wu, F. Allewa, J. Droppo, X. Huang, and A. Stolcke, "The Microsoft 2017 conversational speech recognition system," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, Apr 2018, pp. 5934–5938.
- [2] K. J. Han, A. Chandrasekaran, J. Kim, and I. Lane, "The CAPIO 2017 conversational speech recognition system," *arXiv preprint arXiv:1801.00059*, 2017.
- [3] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth 'CHiME' speech separation and recognition challenge: Dataset, task and baselines," in *Proc. Interspeech*, Hyderabad, India, Sep 2018, pp. 1561–1565.
- [4] J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal, B. Lathoud, M. Lincoln, A. Lisowska, I. McCowan, W. Post, D. Reidsma, and P. Wellner, "The AMI meeting corpus: A pre-announcement," in *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 2005, pp. 28–39.
- [5] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, USA, Mar 2017, pp. 241–245.
- [6] H. Seki, T. Hori, S. Watanabe, J. Le Roux, and J. R. Hershey, "A purely end-to-end system for multi-speaker speech recognition," in *Proc. of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, Melbourne, Australia, 2018, pp. 2620–2630.
- [7] X. Chang, Y. Qian, K. Yu, and S. Watanabe, "End-to-end monaural multi-speaker ASR system without pretraining," in *accepted for publication in Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brighton, UK, May 2019.
- [8] J. 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)*, Shanghai, China, Mar 2016, pp. 31–35.
- [9] M. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, Oct 2017.
- [10] D. Yu, X. Chang, and Y. Wian, "Recognizing multi-talker speech with permutation invariant training," in *Proc. Interspeech*, Stockholm, Sweden, Aug 2017, pp. 2456–2460.
- [11] Y. Isik, J. L. Roux, Z. Chen, S. Watanabe, and J. R. Hershey, "Single-channel multi-speaker separation using deep clustering," in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 545–549.
- [12] 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)*, New Orleans, LA, USA, Mar 2017, pp. 246–250.
- [13] S. Settle, J. Le Roux, T. Hori, S. Watanabe, and J. R. Hershey, "End-to-end multi-speaker speech recognition," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, 2018, pp. 4819–4823.
- [14] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, Mar 2015.
- [15] P. Doetsch, A. Zeyer, P. Voigtlaender, I. Kulikov, R. Schlüter, and H. Ney, "RETURNN: The RWTH extensible training framework for universal recurrent neural networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, LA, USA, Mar 2017, pp. 5345–5349.
- [16] S. Wiesler, A. Richard, P. Golik, R. Schlüter, and H. Ney, "RASR/NN: The RWTH neural network toolkit for speech recognition," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Florence, Italy, May 2014, pp. 3281–3285.
- [17] L. Drude, T. von Neumann, and R. Haeb-Umbach, "Deep attractor networks for speaker re-identification and blind source separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Calgary, Canada, Apr 2018, pp. 11–15.