



**HAL**  
open science

# SLOGD: Speaker Location Guided Deflation Approach to Speech Separation

Sunit Sivasankaran, Emmanuel Vincent, Dominique Fohr

► **To cite this version:**

Sunit Sivasankaran, Emmanuel Vincent, Dominique Fohr. SLOGD: Speaker Location Guided Deflation Approach to Speech Separation. ICASSP 2020 - 45th International Conference on Acoustics, Speech, and Signal Processing, May 2020, Barcelona, Spain. hal-02355613v2

**HAL Id: hal-02355613**

**<https://inria.hal.science/hal-02355613v2>**

Submitted on 14 Feb 2020

**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.

# SLOGD: SPEAKER LOCATION GUIDED DEFLATION APPROACH TO SPEECH SEPARATION

*Sunit Sivasankaran, Emmanuel Vincent, Dominique Fohr*

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

## ABSTRACT

Speech separation is the process of separating multiple speakers from an audio recording. In this work we propose to separate the sources using a Speaker Localization Guided Deflation (SLOGD) approach wherein we estimate the sources iteratively. In each iteration we first estimate the location of the speaker and use it to estimate a mask corresponding to the localized speaker. The estimated source is removed from the mixture before estimating the location and mask of the next source. Experiments are conducted on a reverberated, noisy multichannel version of the well-studied WSJ-2MIX dataset using word error rate (WER) as a metric. The proposed method achieves a WER of 44.2 %, a 34% relative improvement over the system without separation and 17% relative improvement over Conv-TasNet.

*Index Terms*— Speech separation, deflation, localization

## 1. INTRODUCTION

The performance of automatic speech recognition (ASR) system degrades in the presence of overlapping speech. To address this issue, the speech signals must be separated before feeding them to the ASR system. A body of work on monaural speech separation has recently been published using time-frequency clustering [1–3] or the raw waveform [4]. These results were however reported in clean conditions without the influence of noise and reverberation which are known to impact the speech separation performance.

The problem of speech separation has also been studied in the multichannel setting. Computational auditory scene analysis [5] based systems cluster time-frequency bins dominated by the same source using cues such as pitch and interaural time and level differences. In [6], a neural network is trained using the phase differences of the multichannel short-time Fourier transform (STFT) as input features to learn such cues. Deep clustering [1] is then used to associate the cues to the right source. This method generalizes well to unseen speakers but clustering faces inherent limitations such as estimating the number of clusters [7] and choosing an appropriate clustering algorithm to optimally model the embedding space [2]. The studies in [8, 9] propose a spatial location-based approach instead, where the direction of arrival (DOA) of each speaker is

assumed to be known and is used to beamform the multichannel signal in that direction. The beamformed signal is input to a neural network that estimates a time-frequency mask corresponding to the speaker. The beamformed signal contains enhanced components of the desired signal which allow the network to estimate the relevant mask. In our recent work [10], we found that the signal-to-interference ratio (SIR) has a big impact on the separation performance even when the speakers are well separated in space: lower SIR makes it harder for the network to estimate the mask even when the true DOA is known. Localization errors also have an impact on the separation performance, albeit smaller [10].

To overcome these limitations, we propose to iteratively estimate the sources using a Speaker Localization Guided Deflation (SLOGD) approach. The concept of deflation was introduced in blind source separation and refers to the iterative estimation and removal of one source at a time [11]. Our intuition is that the dominant sources are easier to estimate in the first few iterations and, once they have been removed from the mixture, it becomes easier to estimate the other sources. In order to implement this approach in a modern deep learning based multichannel signal processing framework, we estimate the location of a first speaker, compute the beamformed signal, and use features of that signal as inputs to a neural network to derive the corresponding mask. The location and the mask of the next source are estimated by removing the first speaker from the mixture’s features using the first estimated mask, and so on. The mask estimation network is trained on beamformed signals computed from potentially erroneous DOA estimates, which results in increased robustness to localization errors.

A few earlier works have proposed to iteratively estimate the sources using deep neural networks in a single-channel setting [2, 12]. To the best of our knowledge, this paper is the first such study in a multichannel setting, where we estimate both the DOAs and the masks of all speakers. The estimation of DOA is crucial since location-based speech separation is the only method which works well in the presence of reverberation and noise as shown in [9].

The rest of the paper is organized as follows. Section 2 introduces the proposed method. Section 3 describes the dataset and the experimental settings. The results are discussed in Section 4 and we conclude in Section 5.

## 2. PROPOSED METHOD

### 2.1. Signal model

The multichannel signal  $\mathbf{x}(t) = [x_1(t), \dots, x_I(t)]^T$  captured by an array of  $I$  microphones can be expressed as

$$\mathbf{x}(t) = \sum_{j=1}^J \mathbf{c}_j(t) \quad (1)$$

where  $\mathbf{c}_j(t) = [c_{1j}(t), \dots, c_{Ij}(t)]^T$  is the spatial image of source  $j$ , i.e., the signal emitted by the source and captured at the microphones. Similarly to [13], the microphone index and the time index are denoted by  $i$  and  $t$ , respectively, and  $J$  is the total number of sources. This general formulation is valid for both point sources as well as diffuse noise. For point sources such as human speakers, the spatial image can be expressed as a linear convolution of the room impulse response (RIR)  $\mathbf{a}_j(t, \tau) = [a_{1j}(t, \tau), \dots, a_{Ij}(t, \tau)]^T$  and a single-channel source signal  $s_j(t)$  as

$$\mathbf{c}_j(t) = \sum_{\tau=0}^{\infty} \mathbf{a}_j(t, \tau) s_j(t - \tau). \quad (2)$$

Under the narrowband approximation,  $\mathbf{c}_j$  in the time-frequency domain can be written as

$$\mathbf{c}_j(t, f) = \mathbf{a}_j(t, f) s_j(t, f) \quad (3)$$

If the RIR is time-invariant, that is if the speakers and microphones are not moving,  $\mathbf{a}_j(t, f) = \mathbf{a}_j(f)$ .

### 2.2. Estimating the sources

Our objective is to estimate the spatial image of each of the  $J$  sources. We do it iteratively as shown in Fig. 1 and detailed below.

*Step 1, Estimating the first DOA:* In the first step we estimate the DOA of a first speaker using a neural network. The cosines and sines of the phase differences between all pairs of microphones [6, 14], called cosine-sine interchannel phase difference (CSIPD) features, and the short-term magnitude spectrum of one of the channels (in the following, channel 1) are used as input features (see Section 3.2):

$$p^1(\theta) = \text{DOA\_DNN}_1([\text{CSIPD}, |X_1|]) \quad (4)$$

Similar to [14, 15], the network is trained as a classifier to estimate the DOA within a discrete grid of DOAs in every time frame. A non-speech class is also included at the output so that the network can classify a frame as non-speech if no speaker is active in that particular frame, thereby providing voice-activity detection (VAD) information. VAD labels obtained from ASR alignments are used as targets while training the network. The frame-level DOA probabilities output

by the network are averaged across time, and the DOA corresponding to the highest average probability is used as the first speaker DOA.

*Step 2, Estimating the first mask:* Given the estimated DOA, we enhance the signal by simple delay-and-sum beamforming (DS) as detailed in [10]. The magnitude spectrum of the beamformed signal ( $\text{Mag}_{\text{DS}}^1$ ), its phase difference with respect to the first channel ( $\text{CSIPD}_{\text{DS}}^1$ ) and the output of the DOA network ( $p^1(\theta)$ ) are concatenated and fed to a second neural network which estimates the time-frequency mask  $M_1$  corresponding to the first speaker:

$$M_1 = \text{MASK\_DNN}_1([\text{Mag}_{\text{DS}}^1, \text{CSIPD}_{\text{DS}}^1, p^1(\theta)]). \quad (5)$$

*Step 3, Removing the estimated source:* There are multiple ways to remove the estimated source from the mixture. In [2], a remainder mask was computed after each iteration and appended to the network inputs to estimate the next source. In this work we use a similar idea wherein we compute a remainder mask  $(1 - M_1)$  but instead of appending it to the network inputs we multiply the CSIPD and magnitude spectrum features with the remainder mask before feeding them as input to the following DOA estimation and mask estimation stages. Indeed, mask multiplication was shown to perform better than mask concatenation for speaker localization [14].

*Step 4, Estimating the second DOA:* Similarly to Step 1, we estimate the DOA of the second speaker by pooling over time the frame-level DOA probabilities

$$p^2(\theta) = \text{DOA\_DNN}_2([(1 - M_1) \times \text{CSIPD}, (1 - M_1) \times |X_1|]).$$

*Step 5, Estimating the second mask:* Similarly to Step 3, we apply DS beamforming using the estimated DOA and derive the mask for the second speaker as

$$\bar{M}_2 = \text{MASK\_DNN}_2([(1 - M_1) \times \text{Mag}_{\text{DS}}^2, (1 - M_1) \times \text{CSIPD}_{\text{DS}}^2, p^2(\theta)]). \quad (6)$$

Since this mask applies to  $(1 - M_1) \times \mathbf{x}$ , the equivalent mask to be applied to the original signal  $\mathbf{x}$  is  $M_2 = \bar{M}_2 \times (1 - M_1)$ .

The proposed method can in theory estimate DOAs and masks for any number  $J$  of sources using a source counting method [2]. In the following, we assume  $J = 2$ .

### 2.3. Adaptive beamforming to extract the source signals

The obtained masks, namely  $M_1$  for the first speaker,  $M_2$  for the second speaker, and  $M_3 = 1 - M_1 - M_2$  for the noise are used to estimate the speech signals by adaptive beamforming. For a given speaker  $j \in \{1, 2\}$ , the covariance matrix of that speaker is first estimated as

$$\Sigma_j(t, f) = \alpha \Sigma_j(t - 1, f) + (1 - \alpha) M_j(t, f) \mathbf{x}(t, f) \mathbf{x}^H(t, f) \quad (7)$$

where  $\alpha$  is a forgetting factor and  $^H$  denotes Hermitian transposition. Similarly, the noise covariance matrix  $\Sigma_n$ , which




Fig. 1. Iterative estimation of DOA and mask

includes the statistics corresponding to the other speaker and background noise, can be estimated as

$$\begin{aligned} \Sigma_{\mathbf{n}}(t, f) &= \alpha \Sigma_{\mathbf{n}}(t - 1, f) \\ &+ (1 - \alpha)(1 - M_j(t, f)) \mathbf{x}(t, f) \mathbf{x}^H(t, f). \end{aligned} \quad (8)$$

The first channel of  $\mathbf{c}_j(t, f)$  is then estimated via the rank-1 constrained multichannel Wiener filter (R1-MWF) [16], which imposes a rank-1 constraint on  $\Sigma_j(t, f)$ :

$$W_{\text{R1-MWF}}(t, f) = \frac{\Sigma_{\mathbf{n}}^{-1}(t, f) \Sigma_{R1}(t, f)}{\mu + \lambda(t, f)} \mathbf{u}_1 \quad (9)$$

where  $\Sigma_{R1}(t, f) = \sigma_j(t, f) \mathbf{h}_j(t, f) \mathbf{h}_j^H(t, f)$ ,  $\mathbf{h}_j(t, f)$  is the principal eigenvector of  $\Sigma_{\mathbf{n}}^{-1}(t, f) \Sigma_j(t, f)$ ,  $\sigma_j(t, f) = \text{tr}\{\Sigma_j(t, f)\} / \|\mathbf{h}_j(t, f)\|^2$ ,  $\lambda(t, f) = \text{tr}\{\Sigma_{\mathbf{n}}^{-1}(t, f) \Sigma_{R1}(t, f)\}$ . We chose this beamformer due to its higher performance compared to other popular alternatives [8, 10, 16].

### 3. EXPERIMENTAL SETTINGS

#### 3.1. Dataset

Experiments are conducted on the multichannel, reverberated, noisy version of the WSJ-2MIX dataset [1] introduced in [10]. Each mixture was generated by convolving two clean Wall Street Journal (WSJ) utterances with room impulse responses (RIRs) simulated for a 4-channel Kinect-like microphone array and adding real CHiME-5 [17] noise. Strong reverberation and real noise make this data quite realistic and challenging. The noise segments were randomly drawn using the VAD labels in [18]. The RIRs were simulated using RIR-generator [19]. For each mixture, we generated a random room with dimensions in the range of [3 – 9] m and a random RT60 in the range of [0.3 – 1.0] s. Two speakers were placed randomly in the room at a distance of [0.5 – 5.5] m from the microphone array. The signal-to-interference ratio between the speakers was kept at the same level as the original WSJ-2MIX dataset. Noise was added at a signal to noise ratio (SNR) of [0 – 10] dB with respect to the first speaker. The ‘maximum’ version of the WSJ-2MIX dataset was used

wherein the length of the mixed signal is equal to the maximum of the length of the two speakers. This version addresses some of the issues raised in [20] regarding the ‘minimum’ version of the dataset which does not contain mixture signals with single-speaker segments. The simulated dataset contains 30, 10, 5 hours for train, dev and test respectively. The speech and noise signals used in train, dev and test come from different speakers and different rooms.

#### 3.2. Features and network architectures

The STFT was computed using 50 ms windows with 25 ms shift, resulting in 801 frequency bins. The 4 microphone channels result in 6 microphone pairs. Since the CSIPD features consist of the sines and cosines of the phase differences between all pairs of microphones, 12 CSIPD feature vectors, each of dimension 801 were obtained for every time frame.

*DOA estimation networks:* The DOA estimation network (for both sources) contain a 2D convolutional neural network (CNN) which takes the 12 CSIPD channels as inputs and throws out a single-channel output using a  $5 \times 5$  filter. This is followed by a rectified linear unit (ReLU) nonlinearity, a dropout layer, and a max pooling layer of kernel size  $2 \times 1$  along the frequency dimension to obtain a phase projection. Similarly, another 2D CNN is used to obtain a magnitude projection with  $|X_1|$  as input. The magnitude and phase projections are concatenated and fed to a bidirectional long short-term memory (BLSTM) layer followed by a linear layer and a sigmoid nonlinearity. The DOA space is divided into 181 discrete angles as in [14]. Together with the non-speech class, this results in  $181 + 1$  output classes. Since multiple classes are active in every time frame due to the presence of several speakers, binary cross entropy (BCE) was used as the training cost for  $\text{DOA\_DNN}_1$ . After source removal, only a single active speaker remains, hence we used cross entropy as the cost function for  $\text{DOA\_DNN}_2$ .

*Mask estimation networks:* The mask estimation networks are 2-layer BLSTMs. The true reverberated mask corresponding to the localized speaker was used as the training target, with mean square error (MSE) as the cost function. We trained the four networks one after another in the follow-

**Table 1.** Baseline WER(%) results.

| Input | Single Speaker | Single speaker + noise | 2 speakers + noise |
|-------|----------------|------------------------|--------------------|
| WER   | 12.5           | 25.5                   | 66.5               |

ing order:  $\text{DOA\_DNN}_1 \Rightarrow \text{MASK\_DNN}_1 \Rightarrow \text{DOA\_DNN}_2 \Rightarrow \text{MASK\_DNN}_2$ . We also tried to train the DOA and mask networks jointly which yielded poor results. All networks were trained using Adam as the optimizer.

In order to evaluate the impact of localization errors, we also trained a 2-layer BLSTM mask estimation network using the true speaker DOAs. Only the  $\text{CSIPD}_{DS}^1$  and  $\text{Mag}_{DS}^1$  features obtained after DS beamforming were used, since  $p^j(\theta)$  is irrelevant in this case. We also evaluate the performance of this network (without retraining) using the DOAs estimated via the classical generalized cross-correlation with phase transform (GCC-PHAT) method [21]. Finally, we trained the single-channel separation method Conv-TasNet [4]<sup>1</sup> on our dataset for comparison.

### 3.3. ASR

Separation performance was evaluated in terms of the word error rate (WER) on the R1-MWF output. No dereverberation was performed. The acoustic model (AM) for the ASR system was trained in matched conditions: every separation method was applied to the training set and the ASR system was trained on the estimated source signals using accurate senone alignments obtained from the underlying clean single-speaker utterances. The AM was a sub-sampled 15-layered time-delayed neural network trained using lattice-free maximum mutual information [22]. 40-dimensional Mel frequency cepstral coefficients along with 100-dimensional i-vectors were used as features. A 3-gram language model was used while decoding.

## 4. RESULTS

The baseline results are shown in Table 1. A WER of 12.5% was obtained for a single speaker without noise. The WER degrades to 25.5% with noise, showing the difficulty of the real noise for ASR. With an additional overlapping speaker, we observe a substantial WER degradation as expected.

Table 2 shows the ASR performance after speech separation. We obtain a WER of 35.0% using the true speaker DOAs, a 47 % relative improvement with respect to the system without speech separation, showing the positive impact of localization information on speech separation. This is the best WER that can possibly be achieved using location-guided speech separation on this dataset. The WER drops to 63.725 % when using the DOAs estimated by GCC-PHAT, indicating that the localization errors have an adverse impact

**Table 2.** WER(%) results after speech separation

| Method                       | WER  |
|------------------------------|------|
| <b>True DOA</b>              | 35.0 |
| <b>Est DOA with GCC-PHAT</b> | 63.7 |
| <b>SLOGD</b>                 | 44.2 |
| <b>Conv-TasNet</b>           | 53.2 |

on the performance. We obtain a WER of 44.2 % using our proposed method, that is a 34% relative improvement over the baseline without separation and a 19% relative improvement over GCC-PHAT based DOA estimation. In comparison, Conv-TasNet [4] gave a WER of 53.2 % on our dataset.

On analyzing the errors made by the proposed SLOGD approach, we found that, for around 7% of the test dataset, the masks estimated in the two iterations were both closer (in terms of MSE with respect to the true masks) to the same speaker rather than two distinct speakers as expected. This suggests potential for further improvement either in the first mask estimation step or in the source removal stage.

## 5. CONCLUSION

In this work we approached the problem of distant speech separation in challenging noise and reverberation conditions using a deflation based strategy. For each iteration, we train a network to estimate the location of the speaker and use it to estimate a time-frequency mask corresponding to the speaker. The estimated mask is used along with a rank-1 constrained MWF to extract the signal. The estimated source is then removed by masking the signal features before extracting the next source. Using this approach, we obtain a WER of 44.2 % compared to the WER of 53.2 % obtained by Conv-TasNet. Although the proposed method gave large improvements, the problem remains very difficult. Better localization systems and improved strategies for source removal will be studied to improve the source separation performance.

## 6. ACKNOWLEDGEMENT

This work was made with the support of the French National Research Agency, in the framework of the project VOCADOM “Robust voice command adapted to the user and to the context for AAL” (ANR-16-CE33-0006). Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>). The research reported here was conducted at the 2019 Frederick Jelinek Memorial Summer Workshop on Speech and Language Technologies, hosted at L’École de Technologie Supérieure (Montreal, Canada) and sponsored by Johns Hopkins University with unrestricted gifts from Amazon, Facebook, Google, and Microsoft.

<sup>1</sup><https://github.com/kaituoxu/Conv-TasNet>

## 7. REFERENCES

- [1] John R. Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *ICASSP*. 2016, pp. 31–35, IEEE.
- [2] Keisuke Kinoshita, Lukas Drude, Marc Delcroix, and Tomohiro Nakatani, "Listening to each speaker one by one with recurrent selective hearing networks," in *ICASSP*. Apr. 2018, pp. 5064–5068, IEEE.
- [3] Morten Kolbaek, Dong Yu, Zheng-Hua Tan, and Jesper 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.
- [4] Yi Luo and Nima Mesgarani, "Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, Aug. 2019.
- [5] DeLiang Wang and Guy J. Brown, *Computational Auditory Scene Analysis: Principles, Algorithms, and Applications*, Wiley-IEEE Press, 2006.
- [6] Z. Wang and D. Wang, "Combining spectral and spatial features for deep learning based blind speaker separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 2, pp. 457–468, Feb. 2019.
- [7] Takuya Higuchi, Keisuke Kinoshita, Marc Delcroix, Kateřina Žmolíková, and Tomohiro Nakatani, "Deep clustering-based beamforming for separation with unknown number of sources," in *Interspeech*. Aug. 2017, pp. 1183–1187, ISCA.
- [8] L. Perotin, R. Serizel, E. Vincent, and A. Guérin, "Multi-channel speech separation with recurrent neural networks from high-order ambisonics recordings," in *ICASSP*, Apr. 2018, pp. 36–40.
- [9] Z. Chen, X. Xiao, T. Yoshioka, H. Erdogan, J. Li, and Y. Gong, "Multi-Channel overlapped speech recognition with location guided speech extraction network," in *IEEE Spoken Language Technology Workshop (SLT)*, Dec. 2018, pp. 558–565.
- [10] Sunit Sivasankaran, Emmanuel Vincent, and Dominique Fohr, "Analyzing the impact of speaker localization errors on speech separation for automatic speech recognition," in *EUSIPCO (Submitted)*, Amsterdam, 2020.
- [11] Nathalie Delfosse and Philippe Loubaton, "Adaptive blind separation of independent sources: A deflation approach," *Signal Processing*, vol. 45, no. 1, pp. 59–83, July 1995.
- [12] Naoya Takahashi, Sudarsanam Parthasaarathy, Nabarun Goswami, and Yuki Mitsufuji, "Recursive speech separation for unknown number of speakers," in *Interspeech*, Graz, Austria, 2019-09-15-.
- [13] S. Gannot, E. Vincent, S. Markovich-Golan, and A. Ozerov, "A consolidated perspective on multimicrophone speech enhancement and source separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 4, pp. 692–730, Apr. 2017.
- [14] Sunit Sivasankaran, Emmanuel Vincent, and Dominique Fohr, "Keyword-based speaker localization: Localizing a target speaker in a multi-speaker environment," in *Interspeech*, Sept. 2018.
- [15] Sharath Adavanne, Archontis Politis, and Tuomas Virtanen, "Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network," in *EUSIPCO*, Sept. 2018, pp. 1462–1466.
- [16] Ziteng Wang, Emmanuel Vincent, Romain Serizel, and Yonghong Yan, "Rank-1 constrained multichannel Wiener filter for speech recognition in noisy environments," *Computer Speech & Language*, vol. 49, pp. 37–51, May 2018.
- [17] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, "The fifth 'CHiME' speech separation and recognition challenge: Dataset, task and baselines," in *Interspeech*, Sept. 2018, pp. 1561–1565.
- [18] Neville Ryant, Kenneth Church, Christopher Cieri, Alejandrina Cristia, Jun Du, Sriram Ganapathy, and Mark Liberman, "The second DIHARD diarization challenge: Dataset, task, and baselines," in *Interspeech*, Graz, Austria, June 2019.
- [19] Emanuël A. P. Habets, "RIR-Generator: Room impulse response generator," Feb. 2018.
- [20] Tobias Menne, Ilya Sklyar, Ralf Schlüter, and Hermann Ney, "Analysis of deep clustering as preprocessing for automatic speech recognition of sparsely overlapping speech," in *ICASSP*, May 2019.
- [21] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 4, pp. 320–327, Aug. 1976.
- [22] Daniel Povey, Vijayaditya Peddinti, Daniel Galvez, Pegah Ghahremani, Vimal Manohar, Xingyu Na, Yiming Wang, and Sanjeev Khudanpur, "Purely sequence-trained neural networks for ASR based on lattice-free MMI," in *Interspeech*, 2016, pp. 2751–2755.