

MASK-DEPENDENT PHASE ESTIMATION FOR MONAURAL SPEAKER SEPARATION

Zhaoheng Ni

The Graduate Center, CUNY
New York, NY, USA
zni@gradcenter.cuny.edu

Michael I Mandel

Brooklyn College, CUNY
Brooklyn, NY, USA
mim@sci.brooklyn.cuny.edu

ABSTRACT

Speaker separation refers to isolating speech of interest in a multi-talker environment. Most methods apply real-valued Time-Frequency (T-F) masks to the mixture Short-Time Fourier Transform (STFT) to reconstruct the clean speech. Hence there is an unavoidable mismatch between the phase of the reconstruction and the original phase of the clean speech. In this paper, we propose a simple yet effective phase estimation network that predicts the phase of the clean speech based on a T-F mask predicted by a chimera++ network. To overcome the label-permutation problem for both the T-F mask and the phase, we propose a mask-dependent permutation invariant training (PIT) criterion to select the phase signal based on the loss from the T-F mask prediction. We also propose an Inverse Mask Weighted Loss Function for phase prediction to focus the model on the T-F regions in which the phase is more difficult to predict. Results on the WSJ0-2mix dataset show that the phase estimation network achieves comparable performance to models that use iterative phase reconstruction or end-to-end time-domain loss functions, but in a more straightforward manner.

Index Terms— Speech separation, phase estimation, permutation invariant training, chimera++, deep learning

1. INTRODUCTION

Recently, many deep-learning approaches have been proposed for the monaural multi-talker speech separation problem. [1] proposed a deep clustering method that projects each time-frequency bin to a high-dimensional vector space using a bi-direction long short-term memory (BLSTM) network. T-F masks can be generated by clustering these vectors using the k-means algorithm. While deep clustering can separate mixtures when the number of speakers is unknown without changing the model architecture, it produces binary time-frequency masks, which may hurt the intelligibility of the speech. [2] showed that phase-sensitive masks perform better than ideal binary masks on the speech enhancement task. To overcome the label-permutation problem in multi-talker speech separation, [3] proposed the permutation invariant training (PIT) criterion which lets the deep neural network select the assignment of output sources to ground truth sources using the lowest loss value among all permutations. [4] combined deep clustering with a mask-estimation model in a single hybrid network, called the chimera network. Their results show that adding deep clustering leads to better mask predictions, thus achieving better separation performance. [5] made slight modifications to the original chimera network and developed several alternative loss functions for both deep clustering and mask-inference outputs. The new network (dubbed “chimera++”) boosted separation performance using a much simpler architecture.

While there have also been many recent developments in mask-based separation approaches, there are two main issues remaining to be solved:

- Since most masks are restricted to be between 0 and 1, there will be an unavoidable error if the magnitude of the clean speech is greater than that of the mixture due to phase cancellation.
- Most of these methods utilize the phase of the mixture as that of the separated speech, which causes an unavoidable phase difference error.

The phase-sensitive mask (PSM) was introduced [2] to decrease the impact of such phase differences. The PSM as a mask prediction target is defined as:

$$M_{\text{PSM}} = \frac{|S|}{|X|} \odot \cos(\angle S - \angle X), \quad (1)$$

where S and X are the STFT of the clean speech and the speech mixture respectively at any particular time-frequency point, \odot is the element-wise multiplication. By applying the PSM to the mixture, the real component of the estimated STFT is close to that of the ground truth.

Besides the PSM, [6] introduced the complex ratio mask (cRM) to enhance both the magnitude and the phase spectra by operating in the complex domain. The real component of the cRM is equivalent to the PSM. However, [7] showed that the imaginary component is difficult to predict directly from the noisy magnitude due to the randomness of the phase pattern.

To estimate the phase of each separated source, [5] started with the estimated magnitude and the noisy phase and jointly reconstructed the phase of all sources using the multiple-input spectrogram inversion (MISI) algorithm [8]. Later, [9] unfolded the iterations in the MISI algorithm into separate STFT and iSTFT layers on top of a chimera++ network and trained the network to minimize the error in the time domain. Their results showed that it is possible to predict the phase of separated sources from a magnitude estimate and the noisy phase.

Instead of using an iterative STFT and iSTFT to reconstruct the phase, the current paper proposes an approach to reconstruct the phase directly from the magnitude estimate and the noisy phase. The problem can thus be formulated as estimating the magnitude and phase of each source from the noisy STFT, X

$$S_c = |X| \odot M_c \odot \cos \theta_c + j|X| \odot M_c \odot \sin \theta_c, \quad (2)$$

where $|X|$ is the noisy magnitude, M_c and θ_c are the mask and phase estimates for source c , respectively.

Some studies have tried to predict the mask and phase directly without iterative algorithms. In [10] and [11], the authors utilized

visual feature to drive the model to predict the corresponding mask and phase. However, some datasets may not contain visual data, or the video may not capture the speaker of interest. Though the PIT criterion solves the label-permutation problem for mask estimates, integrating PIT with both mask and phase estimates is still an open problem. Thus, it is still essential to figure out how to apply PIT criterion in phase estimation with only the audio data.

Another problem in phase estimation is that its difficulty varies across regions. Since the noisy STFT $X = \sum_c S_c$, where S_c represents the clean STFT for source c , if a T-F bin is only dominated by one speaker (i.e., the mask value is close to one), the phase difference between noisy and clean is close to zero. In regions where the mask is not close to one, the phase difference is influenced by the magnitude and phase of all significant sources, and is thus more difficult to estimate.

Hence this study proposes a joint-training algorithm that estimates the ideal ratio mask (IRM) by using a chimera++ network and then estimates the clean phase based on the mask estimate and the noisy phase for each source. The label permutation problem is solved by using a PIT criterion based only on the estimated masks, with the phase estimated after this matching. Three different weighed loss functions are proposed to compare the influence of different T-F regions over the training of the phase estimation model.

2. CHIMERA++ NETWORK

This section describes the architecture of the chimera++ network and its loss function, which is also incorporated into our algorithm. The left half of Figure 1 shows the architecture of the chimera++ network. Each time-frequency bin is projected to a D -dimensional vector $v_i \in \mathbb{R}^{1 \times D}$ via a deep clustering layer [1], where i corresponds to a particular pair of time and frequency indices. Each T-F bin has a one-hot label vector $y_i \in \mathbb{R}^{1 \times C}$ indicating the speaker among C speakers who dominates this bin. Stacking all embedding vectors and labels produces the embedding matrix $V \in \mathbb{R}^{T \times F \times D}$ and the label matrix $Y \in \mathbb{R}^{T \times F \times C}$. The objective of deep clustering is to group together the embedding vectors of T-F bins from the same speaker and make those from different speakers orthogonal. The loss function is defined as:

$$L_{DC, \text{classic}} = \left\| VV^T - YY^T \right\|_F^2 \quad (3)$$

[5] showed reducing the influence of silence T-F bins benefits training the deep clustering network. They introduced binary voice activity weights W_{VA} to the loss function:

$$L_{DC, \text{classic}, W} = \left\| W_{VA}^{\frac{1}{2}} (VV^T - YY^T) W_{VA}^{\frac{1}{2}} \right\|_F^2 \quad (4)$$

In terms of the loss function for the mask inference layer, [5] recommends using a loss function based on the truncated phase-sensitive spectrum approximation (tPSA), defined as:

$$L_{MI, \text{tPSA}} = \min_{\pi \in P} \sum_c \left\| \hat{M}_c \odot |X| - T_0^{|X|} (|S_{\pi(c)}| \odot \cos(\theta_X - \theta_{\pi(c)})) \right\|_F^1, \quad (5)$$

where P is the set of permutations on $\{1, \dots, C\}$, X and S are the STFT of the noisy and clean speech respectively, T is the truncation function defined as $T_a^b(x) = \min(\max(x, a), b)$, θ_X and θ_c are the noisy phase and the true phase of source c , respectively.

Model	Loss Function	Mask Type	SDR
Chimera++ [5]	DC(W, W_{VA})	tPSA	10.9
Our implementation	DC(classic, W_{VA})	tPSA	11.0

Table 1. Comparison between the chimera++ in [5] and our implementation

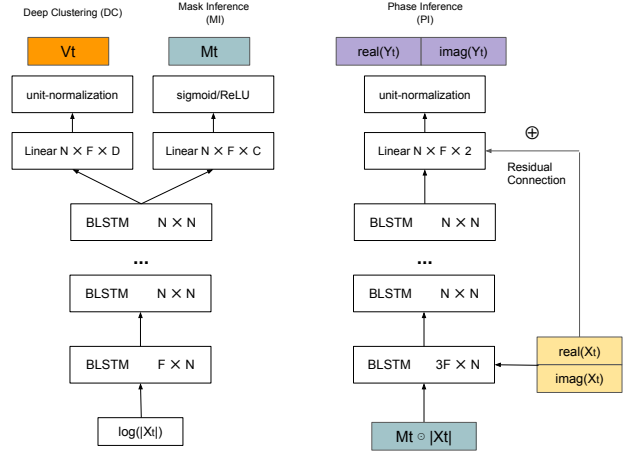


Fig. 1. Model architecture of the proposed phase estimation network. Note that the left tower is the original chimera++ network.

However, estimating tPSA contradicts with estimating the phase vectors in our problem formulation since both consider the cosine of the clean phase. To adapt the estimate into our problem formulation, we change the tPSA loss function to the magnitude spectrum approximation (MSA) for magnitude reconstruction.

$$L_{MI, \text{MSA}} = \min_{\pi \in P} \sum_c \left\| \hat{M}_c \odot |X| - |S_{\pi(c)}| \right\|_F^1. \quad (6)$$

We compare our chimera++ implementation with that in [5] in Table 1, which shows that we have successfully reproduced the reported result. We find that our implementation achieves slightly better performance by using the classic loss function based on tPSA, compared that using a whitened k-means loss function (W) [5] based on tPSA. This indicates we can fairly compare our phase estimation network with the other methods which are also based on the chimera++ network.

3. PHASE ESTIMATION NETWORK

Figure 1 shows the architecture of the proposed phase network. First the magnitudes of each of the sources are estimated by applying the masks estimated by the chimera++ network to the noisy magnitude. Each magnitude estimate is concatenated with the real and imaginary components of the noisy STFT as the input feature for the phase inference (PI) sub-network. The output is a $T \times F \times 2$ matrix representing the cosine and sine of the phase estimate for each T-F bin. Each phase estimate is generated based on one of the masks. The output of the final linear layer is added to the original real and imaginary components of the noisy phase as a residual connection [12] before unit-normalization. The loss function for this prediction is

$$L_{PI} = - \sum_{c, t, f} \langle \hat{p}_{t, f}^c, p_{t, f}^c \rangle \quad (7)$$

where $p_{t,f}^c$ is the vector $(\cos \theta_c, \sin \theta_c)$ representing the true phase of the t, f bin in source c and $\hat{p}_{t,f}^c$ is the corresponding phase estimate.

3.1. Mask-dependent PIT Criterion

The PIT criterion [3] has been successful in solving the label-permutation problem. To overcome the permutation problem for both the mask and the phase, several combinations of mask and phases losses could be used for selecting the best permutation. Since the phase estimate is dependent on the mask, the simplest combined PIT is based on a combination of both the losses of the MI and PI layers together. We use L_{PIT} to represent this combined loss

$$L_{\text{PIT, MP}} = \min_{\pi \in P} \sum_c \left(\left\| \hat{M}_c \odot |X| - |S_{\pi(c)}| \right\|_F^1 - \sum_{t,f} \langle \hat{p}_{t,f}^c, p_{t,f}^{\pi(c)} \rangle \right) \quad (8)$$

But since the phase is more difficult to learn to predict than the mask, perhaps because of a larger number of local minima, the phase tends to be misleading to this permutation matching in early epochs. Instead we only utilize the mask for matching, and continue this throughout training. Hence we propose a Mask-dependent PIT criterion defined as

$$L_{\text{PIT, MD}} = \sum_c \left\| \hat{M}_c \odot |X| - |S_{\pi(c)}| \right\|_F^1 - \sum_{c,t,f} \langle \hat{p}_{t,f}^c, p_{t,f}^{\pi(c)} \rangle, \quad (9)$$

$$\pi = \arg \min_{\pi} \sum_c \left\| \hat{M}_c \odot |X| - |S_{\pi(c)}| \right\|_F^1.$$

3.2. Weighted Loss Functions for Phase Estimation

Phase is harder to predict in certain T-F regions. Figure 2 shows for one source in a mixture the log magnitude of the clean speech, IRM, cos, and sin of the phase difference $\angle \theta_S - \angle \theta_X$. If the mask value is close to one, meaning the phase of the mixture mostly comes from this source, the cos of the phase difference is close to 1 and the sin of the phase difference is close to 0. Hence the phase difference between the clean and the noisy STFT is very small. When the mask value is small, the phase differences are larger, making it more challenging to predict the phase.

In the MI loss of the chimera++ network, an MSA or tPSA loss function is introduced to increase the weight of the losses where the energy of the noisy magnitude is higher. However, the same weight may not be suitable for estimating the phase. We propose three different weighted loss function for the PI loss to differently weight the contributions to the phase estimates.

We define the Magnitude Weighted Loss Function (MWL) as

$$L_{\text{PI, MWL}} = - \sum_{c,t,f} (\gamma + M_{c,t,f}) \langle \hat{p}_{t,f}^c, p_{t,f}^{\pi(c)} \rangle, \quad (10)$$

where $M_{c,t,f}$ represents the magnitude of source c on the t, f bin, π is identified by the PIT loss, and γ is a tunable parameter that avoids applying 0 loss to points with 0 magnitude. We set it to 0.2 for our experiments. The motivation for this is that higher magnitude values should be emphasized because they are more reliable in training the phase network.

We define the Inverse Magnitude Weighted Loss Function (I-MWL) as

$$L_{\text{PI, I-MWL}} = - \sum_{c,t,f} (\gamma + M_{-c,t,f}) \langle \hat{p}_{t,f}^c, p_{t,f}^{\pi(c)} \rangle, \quad (11)$$

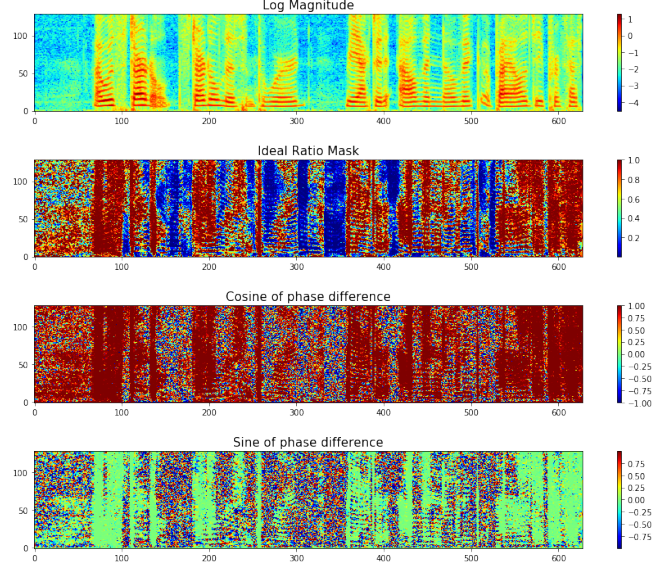


Fig. 2. The clean spectrogram, the IRM, cosine, and sine of phase difference between clean and noisy STFT (sample utterance: cv/s2 (mix) / 011a010g_0.16366_40pc0204_-0.16366.wav)

where $M_{-c,t,f} = \sum_{i \neq c} M_{i,t,f}$ and π is identified by the PIT loss. The motivation for this is that lower mask values are more difficult to predict, so should be emphasized.

We define the Joint Weighted Loss Function (Joint) as

$$L_{\text{PI, Joint}} = - \sum_{c,t,f} \langle \hat{p}_{t,f}^c, p_{t,f}^{\pi(c)} \rangle \sum_i M_{i,t,f}. \quad (12)$$

The motivation for this weighting is to emphasize all active regions in the spectrogram. Only silent regions are ignored, similarly to the loss function of the deep clustering layer. Note that if the noisy magnitude is small while there are other sources whose magnitudes are large, the bin is emphasized more.

The loss function for the whole network thus becomes

$$L_{\text{Phase Network}} = \alpha L_{\text{DC}} + (1 - \alpha) L_{\text{PIT}}, \quad (13)$$

where α is a weight parameter that balances the loss of the DC layers and the MI and PI layers. We use 0.975 for α , following [5].

4. EXPERIMENTS

To evaluate our proposed method, the models are trained and validated on the WSJ0-2mix dataset [1], which is a publicly available dataset for multi-talker speech separation. There are 20,000 mixtures (~ 30 hours) in the training data, 5,000 mixtures (~ 10 hours) in the validation data, and 3,000 mixtures (~ 5 hours) in the test data. The training data and the validation data are generated by randomly mixing two utterances of two random speakers in the WSJ0 training data (si_tr_s). Some speakers in the validation data are included in the training data with different utterances, thus the validation data is considered to be the closed speaker condition (CSC). The test data is generated by mixing two random utterances of two random speakers from the WSJ0 test data (si_et_05). None of the speakers in the test data are included in either the training or the validation data, thus it

Method	SDR	SI-SDR
Chimera++, MSA	10.5	-
+ tPSA [5]	11.5	11.2
+ MISI-5 [5]	11.8	11.5
+ WA-MISI-5 [9]	12.9	12.6
Phasebook, MISI-0 [16]	-	12.6
+ MISI-5 [16]	-	12.8
Chimera++(Encoder-BLSTM-Decoder) [17]	-	11.9
Sign prediction network [17]	15.6	15.3

Table 2. Published SDR/SI-SDR improvements of different phase estimation methods on the open speaker condition (OSC) of the WSJ0-2mix dataset.

PIT Criterion	Mask Activation	Phase Loss	SDR
MP	ReLU	MWL	11.0
MP	Sigmoid	MWL	11.5
MD	Sigmoid	I-MWL	12.0
MD	Sigmoid	MWL	12.6
MD	Sigmoid	Joint	13.0
MD	Sigmoid	Joint, $\alpha = 0.5$	13.6

Table 3. SDR improvements of the proposed method with different settings on the OSC of the WSJ0-2mix dataset. PIT criteria Mask+Phase (MP) and Mask-dependent (MD). Phase losses magnitude-weighted loss (MWL), inverse magnitude-weighted loss (I-MWL), and joint weighted loss (Joint).

is considered to be the open speaker condition (OSC). The sampling rate is 8 kHz and the Signal to Noise Ratio (SNR) of utterances is in a range between 0 and 10 dB.

The chimera++ network and the proposed phase network are implemented in PyTorch 1.0¹ [13]. Both the chimera++ network and the phase network have 4 BLSTM layers, each with 600 neural units in each direction. A 0.3 dropout rate is applied to the outputs of the BLSTM layers. Random chunks of 400 consecutive frames from each utterance are extracted to train the model with a batch size of 16 such chunks. The STFT is calculated by the librosa library [14] utilizing a Hann window. The FFT window size is 256 samples (32 ms) and the hop length is 64 samples (8 ms). The Adam optimizer with a 10^{-3} initial learning rate is used. Training is stopped when the validation loss does not decrease for 10 epochs. The model is trained for a total of 100 epochs if no early stopping mechanism is activated.

To evaluate the separation performance, the Signal to Distortion Ratio (SDR) computed with the mir_eval library [15] as the major evaluation metric. The baseline method is the noisy phase with the MSA magnitude estimated by the chimera++ network.

5. RESULTS

Tables 2 and 3 show the SDR or scale-invariant SDR (SI-SDR) [18] or both improvements of recently published methods and those of the proposed phase estimation networks, respectively. The cells are left blank if they are not reported in the publications. The chimera++ network with MSA loss achieves 10.5 dB SDR. Adding our phase estimation network with PIT_{MP} criterion improves the performances by 1.0 dB, which is the same as that of the chimera++ network with

¹<https://github.com/speechLabBcCuny/onssen>

the tPSA loss. By comparing the activation functions of the MI layer, ReLU does not improve the performance over Sigmoid.

Comparing rows 1 and 3 of Table 3, using the mask-dependent PIT_{MD} criterion achieves 12.6 dB SDR, which is 1.1 dB higher than using the PIT_{MP} criterion. This shows that the loss for the mask estimate is more reliable and stable than combining it with the phase loss for the PIT criterion. In terms of different weighted loss functions, the weight of all magnitudes (Joint) achieves the best result of 13.0 dB SDR, followed by T-F regions of the clean speech (MWL) and then noisy speech (I-MWL). Inspired by [19], we also apply a curriculum training strategy [20] that trains the network using $\alpha = 0.975$ first, then retrains the model using $\alpha = 0.5$, increasing the performance to 13.6 dB.

Comparing our results with published phase-based methods in Table 2, our result is better than the chimera++ with 5 iterations of the MISI algorithm (MISI-5). We get competitive result to the chimera++ with 5 MISI iterations and an end-to-end waveform approximate loss function (WA-MISI-5) without applying the curriculum training strategy. This indicates the proposed network can predict phase directly from the mask generated by chimera++. We do not currently use the waveform approximate loss function in our model, but our mask-dependent PIT criterion is applicable to the such losses. Instead of choosing the minimum waveform difference, we choose the permutation that gives the minimum mask loss. The model still backpropagates to minimize the waveform difference, while the mask-dependent PIT criterion helps train the model in a more reliable way. Our future research will analyze how the mask-dependent PIT criterion can contribute to such end-to-end approaches.

Though we do not achieve better performance than [17], it is good to mention the chimera++ network applied in [17] is different. It adds several convolutional encoder and decoder layers before and after the chimera++ network to achieve 11.9 dB SI-SDR (~ 12.2 to 12.4 dB SDR). Also, since the model applies the waveform approximate loss function to reconstruct the time-domain signal, similarly to Phasebook, we can adapt our proposed method to the model to improve the phase estimation accuracy further. [21] show that frame-level PIT criterion can find a better local minimum hence significantly improves the performance comparing than the utterance-level PIT. Our future plan is to integrate the frame-level PIT into our mask-dependent criterion and conduct the evaluation.

6. CONCLUSION

This paper proposed a phase estimation method using the mask estimate from the chimera++ network. A mask-dependent PIT criterion is applied to solve the label-permutation problem for both the mask and the phase estimates. The mask-dependent PIT criterion significantly improved the separation performance compared with the PIT over the whole loss. Future study will focus on applying the proposed PIT criterion to end-to-end phase estimation methods (e.g., phasebook [16] and the sign prediction network [17]).

7. ACKNOWLEDGMENTS

This material is based upon work supported by a Google Faculty Research Award and the National Science Foundation (NSF) under Grant IIS-1409431. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

8. REFERENCES

- [1] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [2] Hakan Erdogan, John R Hershey, Shinji Watanabe, and Jonathan Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 708–712.
- [3] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 241–245.
- [4] Yi Luo, Zhuo Chen, John R Hershey, Jonathan Le Roux, and Nima Mesgarani, "Deep clustering and conventional networks for music separation: Stronger together," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 61–65.
- [5] Zhong-Qiu Wang, Jonathan Le Roux, and John R Hershey, "Alternative objective functions for deep clustering," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 686–690.
- [6] Donald S Williamson, Yuxuan Wang, and DeLiang Wang, "Complex ratio masking for monaural speech separation," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 24, no. 3, pp. 483–492, 2016.
- [7] Donald S Williamson and DeLiang Wang, "Speech dereverberation and denoising using complex ratio masks," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5590–5594.
- [8] David Gunawan and Deep Sen, "Iterative phase estimation for the synthesis of separated sources from single-channel mixtures," *IEEE Signal Processing Letters*, vol. 17, no. 5, pp. 421–424, 2010.
- [9] Zhong-Qiu Wang, Jonathan Le Roux, DeLiang Wang, and John R Hershey, "End-to-end speech separation with unfolded iterative phase reconstruction," *arXiv preprint arXiv:1804.10204*, 2018.
- [10] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinandan Hassidim, William T Freeman, and Michael Rubinstein, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," *arXiv preprint arXiv:1804.03619*, 2018.
- [11] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman, "The conversation: Deep audio-visual speech enhancement," *arXiv preprint arXiv:1804.04121*, 2018.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [13] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.
- [14] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto, "librosa: Audio and music signal analysis in python," in *Proceedings of the 14th python in science conference*, 2015, pp. 18–25.
- [15] Colin Raffel, Brian McFee, Eric J Humphrey, Justin Salamon, Oriol Nieto, Dawen Liang, Daniel PW Ellis, and C Colin Raffel, "mir_eval: A transparent implementation of common mir metrics," in *In Proceedings of the 15th International Society for Music Information Retrieval Conference, ISMIR*. Citeseer, 2014.
- [16] Jonathan Le Roux, Gordon Wichern, Shinji Watanabe, Andy Sarroff, and John R Hershey, "Phasebook and friends: Leveraging discrete representations for source separation," *IEEE Journal of Selected Topics in Signal Processing*, 2019.
- [17] Zhong-Qiu Wang, Ke Tan, and DeLiang Wang, "Deep learning based phase reconstruction for speaker separation: A trigonometric perspective," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 71–75.
- [18] Jonathan Le Roux, JR Hershey, A Liutkus, F Stöter, ST Wisdom, and H Erdogan, "Sdr-half-baked or well done?," *Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA, Tech. Rep*, 2018.
- [19] Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R Hershey, "Single-channel multi-speaker separation using deep clustering," *arXiv preprint arXiv:1607.02173*, 2016.
- [20] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 41–48.
- [21] Yuzhou Liu and DeLiang Wang, "Divide and conquer: A deep casa approach to talker-independent monaural speaker separation," *arXiv preprint arXiv:1904.11148*, 2019.