# FACE LANDMARK-BASED SPEAKER-INDEPENDENT AUDIO-VISUAL SPEECH ENHANCEMENT IN MULTI-TALKER ENVIRONMENTS

Giovanni Morrone*	Luca Pasa†	Vadim Tikhanoff †
Sonia Bergamaschi*	Luciano Fadiga $^{\dagger}$	Leonardo Badino <sup>†</sup>

\*Department of Engineering "Enzo Ferrari", University of Modena and Reggio Emilia, Modena, Italy <sup>†</sup>Istituto Italiano di Tecnologia, Ferrara, Italy

# ABSTRACT

In this paper, we address the problem of enhancing the speech of a speaker of interest in a cocktail party scenario when visual information of the speaker of interest is available.

Contrary to most previous studies, we do not learn visual features on the typically small audio-visual datasets, but use an already available face landmark detector (trained on a separate image dataset).

The landmarks are used by LSTM-based models to generate time-frequency masks which are applied to the acoustic mixed-speech spectrogram. Results show that: (*i*) landmark motion features are very effective features for this task, (*ii*) similarly to previous work, reconstruction of the target speaker's spectrogram mediated by masking is significantly more accurate than direct spectrogram reconstruction, and (*iii*) the best masks depend on both motion landmark features and the input mixed-speech spectrogram.

To the best of our knowledge, our proposed models are the first models trained and evaluated on the limited size GRID and TCD-TIMIT datasets, that achieve speaker-independent speech enhancement in a multi-talker setting.

*Index Terms*— audio-visual speech enhancement, cocktail party problem, time-frequency mask, LSTM, face landmarks

## 1. INTRODUCTION

In the context of speech perception, the *cocktail party effect* [1, 2] is the ability of the brain to recognize speech in complex and adverse listening conditions where the attended speech is mixed with competing sounds/speech.

Speech perception studies have shown that watching speaker's face movements could dramatically improve our ability at recognizing the speech of a target speaker in a multi-talker environment [3, 4].

This work aims at extracting the speech of a target speaker from single channel audio of several people talking simultaneously. This is an ill-posed problem in that many different hypotheses about what the target speaker says are consistent with the mixture signal. Yet, it can be solved by exploiting some additional information associated to the speaker of interest and/or by leveraging some prior knowledge about speech signal properties (e.g., [5]). In this work we use face movements of the target speaker as additional information.

This paper (*i*) proposes the use of face landmark's movements, extracted using Dlib [6, 7] and (*ii*) compares different ways of mapping such visual features into time-frequency (T-F) masks, then applied to clean the acoustic mixed-speech spectrogram.

By using Dlib extracted landmarks we relieve our models from the task of learning useful visual features from raw pixels. That aspect is particularly relevant when the training audio-visual datasets are small.

The analysis of landmark-dependent masking strategies is motivated by the fact that speech enhancement mediated by an explicit masking is often more effective than mask-free enhancement [8].

All our models were trained and evaluated on the GRID [9] and TCD-TIMIT [10] datasets in a speaker-independent setting.

# 1.1. Related work

Speech enhancement aims at extracting the voice of a target speaker, while speech separation refers to the problem of separating each sound source in a mixture. Recently proposed audio-only single-channel methods have achieved very promising results [11, 12, 13]. However the task still remains challenging. Additionally, audio-only systems need separate models in order to associate the estimated separated audio sources to each speaker, while vision easily allow that in a unified model.

Regarding audio-visual speech enhancement and separation methods an extensive review is provided in [14]. Here we focus on the deep-learning methods that are most related to the present work.

Our first architecture (Section 2.1) is inspired by [15], where a pre-trained convolutional neural network (CNN) is used to generate a clean spectrogram from silent video [16]. Rather than directly computing a time-frequency (T-F) mask, the mask is computed by thresholding the estimated clean spectrogram. This approach is not very effective since the pre-trained CNN is designed for a different task (video-tospeech synthesis). In [17] a CNN is trained to directly estimate clean speech from noisy audio and input video. A similar model is used in [18], where the model jointly generates clean speech and input video in a denoising-autoender architecture.

[19] shows that using information about lip positions can help to improve speech enhancement. The video feature vector is obtained computing pair-wise distances between any mouth landmarks. Similarly to our approach their visual features are not learned on the audio-visual dataset but are provided by a system trained on different dataset. Contrary to our approach, [19] uses position-based features while we use motion features (of the whole face) that in our experiments turned out to be much more effective than positional features.

Although the aforementioned audio-visual methods work well, they have only been evaluated in a speaker-dependent setting. Only the availability of new large and heterogeneous audio-visual datasets has allowed the training of deep neural network-based speaker-independent speech enhancement models [20, 21, 22].

The present work shows that huge audio-visual datasets are not a necessary requirement for speaker-independent audio-visual speech enhancement. Although we have only considered datasets with simple visual scenarios (i.e., the target speaker is always facing the camera), we expect our methods to perform well in more complex scenarios thanks to the robust landmark extraction.

# 2. MODEL ARCHITECTURES

We experimented with the four models shown in Fig. 1. All models receive in input the target speaker's landmark motion vectors and the power-law compressed spectrogram of the single-channel mixed-speech signal. All of them perform some kind of masking operation.

# 2.1. VL2M model

At each time frame, the video-landmark to mask (VL2M) model (Fig. 1a) estimates a T-F mask from visual features only (of the target speaker). Formally, given a video sequence  $\mathbf{v} = [\mathbf{v}_1, \dots, \mathbf{v}_T], \mathbf{v}_t \in \mathbb{R}^n$  and a target mask sequence  $\mathbf{m} = [\mathbf{m}_1, \dots, \mathbf{m}_T], \mathbf{m}_t \in \mathbb{R}^d$ , VL2M perform a function  $\mathscr{F}_{vl2m}(\mathbf{v}) = \hat{\mathbf{m}}$ , where  $\hat{\mathbf{m}}$  is the estimated mask.

The training objective for VL2M is a Target Binary Mask (TBM) [23, 24], computed using the spectrogram of the target speaker only. This is motivated by our goal of extracting the speech of a target speaker as much as possible independently of the concurrent speakers, so that, e.g., we do not need to estimate their number. An additional motivations is that the model takes as only input the visual features of the

target speaker, and a target TBM that only depends on the target speaker allows VL2M to learn a function (rather than approximating an ill-posed one-to-many mapping).

Given a clean speech spectrogram of a speaker  $\mathbf{s} = [\mathbf{s}_1, \ldots, \mathbf{s}_T], \mathbf{s}_t \in \mathbb{R}^d$ , the TBM is defined by comparing, at each frequency bin  $f \in [1, \ldots, d]$ , the target speaker value  $\mathbf{s}_t[f]$  vs. a reference threshold  $\tau[f]$ . As in [15], we use a function of long-term average speech spectrum (LTASS) as reference threshold. This threshold indicates if a T-F unit is generated by the speaker or refers to silence or noise. The process to compute the speaker's TBM is as follows:

- 1. The mean  $\pi[f]$  and the standard deviation  $\sigma[f]$  are computed for all frequency bins of all seen spectrograms in speaker's data.
- 2. The threshold  $\tau[f]$  is defined as  $\tau[f] = \pi[f] + 0.6 \cdot \sigma[f]$ where 0.6 is a value selected by manual inspection of several spectrogram-TBM pairs.
- 3. The threshold is applied to every speaker's speech spectrogram s.

$$\mathbf{m}_t[f] = \begin{cases} 1, & \text{if } \mathbf{s}_t[f] \ge \tau[f], \\ 0, & \text{otherwise.} \end{cases}$$

The mapping  $\mathscr{F}_{vl2m}(\cdot)$  is carried out by a stacked bidirectional Long Short-Term Memory (BLSTM) network [25]. The BLSTM outputs are then forced to lay within the [0,1] range. Finally the computed TBM  $\hat{\mathbf{m}}$  and the noisy spectrogram  $\mathbf{y}$  are element-wise multiplied to obtain the estimated clean spectrogram  $\hat{\mathbf{s}}^{\mathbf{m}} = \hat{\mathbf{m}} \circ \mathbf{y}$ , where  $\mathbf{y} = [\mathbf{y}_1, \dots \mathbf{y}_T], \mathbf{y}_t \in \mathbb{R}^d$ .

The model parameters are estimated to minimize the loss:

$$J_{vl2m} = \sum_{t=1}^{T} \sum_{f=1}^{d} -\mathbf{m}_t[f] \cdot \log(\hat{\mathbf{m}}_t[f]) - (1 - \mathbf{m}_t[f]) \cdot \log(1 - \hat{\mathbf{m}}_t[f])$$

# 2.2. VL2M\_ref model

VL2M generates T-F masks that are independent of the acoustic context. We may want to refine the masking by including such context. This is what the novel VL2M\_ref does (Fig. 1b). The computed TBM  $\hat{\mathbf{m}}$  and the input spectrogram  $\mathbf{y}$  are the input to a function that outputs an Ideal Amplitude Mask (IAM)  $\mathbf{p}$  (known as FFT-MASK in [8]). Given the target clean spectrogram  $\mathbf{s}$  and the noisy spectrogram  $\mathbf{y}$ , the IAM is defined as:

$$\mathbf{p}_t[f] = \frac{\mathbf{s}_t[f]}{\mathbf{y}_t[f]}$$

Note that although IAM generation requires the mixed-speech spectrogram, separate spectrograms for each concurrent speakers are not required.

The target speaker's spectrogram s is reconstructed by multiplying the input spectrogram with the estimated IAM. Values greater than 10 in the IAM are clipped to 10 in order to obtain better numerical stability as suggested in [8].



Fig. 1. Model architectures.

The model performs a function  $\mathscr{F}_{mr}(\mathbf{v}, \mathbf{y}) = \hat{\mathbf{p}}$  that consists of a VL2M component plus three different BLSTMs  $\mathscr{G}_m$ ,  $\mathscr{G}_y$  and  $\mathscr{H}$ .

 $\mathscr{G}_m(\mathscr{F}_{vl2m}(\mathbf{v})) = \mathbf{r}_m$  receives the VL2M mask  $\hat{\mathbf{m}}$  as input, and  $\mathscr{G}_y(\mathbf{y}) = \mathbf{r}_y$  is fed with the noisy spectrogram. Their output  $\mathbf{r}_m, \mathbf{r}_y \in \mathbb{R}^z$  are fused in a joint audio-visual representation  $\mathbf{h} = [\mathbf{h}_1, \dots, \mathbf{h}_T]$ , where  $\mathbf{h}_t$  is a linear combination of  $\mathbf{r}_{m_t}$  and  $\mathbf{r}_{y_t}$ :  $\mathbf{h}_t = \mathbf{W}_{hm} \cdot \mathbf{r}_{m_t} + \mathbf{W}_{hy} \cdot \mathbf{r}_{y_t} + \mathbf{b}_h$ .  $\mathbf{h}$  is the input of the third BLSTM  $\mathscr{H}(\mathbf{h}) = \hat{\mathbf{p}}$ , where  $\hat{\mathbf{p}}$  lays in the [0,10] range.

## 2.3. Audio-Visual concat model

The third model (Fig. 1c) performs early fusion of audiovisual features. This model consists of a single stacked BLSTM that computes the IAM mask  $\hat{\mathbf{p}}$  from the concatenated  $[\mathbf{v}, \mathbf{y}]$ . The training loss is the same  $J_{mr}$  used to train VL2M\_ref. This model can be regarded as a simplification of VL2M\_ref, where the VL2M operation is not performed.

#### 2.4. Audio-Visual concat-ref model

The fourth model (Fig. 1d) is an improved version of the model described in section 2.3. The only difference is the input of the stacked BLSTM that is replaced by  $[\hat{s}^{m}, y]$  where  $\hat{s}^{m}$  is the denoised spectrogram returned by VL2M operation.

#### 3. EXPERIMENTAL SETUP

### 3.1. Dataset

All experiments were carried out using the GRID [9] and TCD-TIMIT [10] audio-visual datasets. For each of them, we created a mixed-speech version.

Regarding the GRID corpus, for each of the 33 speakers (one had to be discarded) we first randomly selected 200 utterances (out of 1000). Then, for each utterance, we created 3 different audio-mixed samples. Each audio-mixed sample was created by mixing the chosen utterance with one utterance from a different speaker.

That resulted in 600 audio-mixed samples per speaker.

The resulting dataset was split into disjoint sets of 25/4/4 speakers for training/validation/testing respectively.

The TCD-TIMIT corpus consists of 59 speakers (we excluded 3 professionally-trained lipspeakers) and 98 utterances per speaker. The mixed-speech version was created following the same procedure as for GRID, with one difference. Contrary to GRID, TCD-TIMIT utterances have different duration. Thus 2 utterances were mixed only if their duration difference did not exceed 2 seconds. For each utterance pair, we forced the non-target speaker's utterance to match the duration of the target speaker utterance. If it was longer, the utterance was cut at its end, whereas if it was shorter, silence samples were equally added at its start and end.

The resulting dataset was split into disjoint sets of 51/4/4 speakers for training/validation/testing respectively.

### 3.2. LSTM training

In all experiments, the models were trained using the Adam optimizer [26]. Early stopping was applied when the error on the validation set did not decrease over 5 consecutive epochs.

VL2M, AV concat and AV concat-ref had 5, 3 and 3 stacked BLSTM layers respectively. All BLSTMs had 250 units. Hyper-parameters selection was performed by using random search with a limited number of samples, therefore all the reported results may improve through a deeper hyper-parameters validation phase.

VL2M\_ref and AV concat-ref training was performed in 2 steps. We first pre-trained the models using the oracle TBM **m**. Then we substituted the oracle masks with the VL2M component and retrained the models while freezing the parameters of the VL2M component.

#### 3.3. Audio pre- and post-processing

The original waveforms were resampled to 16 kHz. Short-Time Fourier Transform (STFT) x was computed using FFT size of 512, Hann window of length 25 ms (400 samples), and hop length of 10 ms (160 samples). The input spectrogram was obtained taking the STFT magnitude and performing power-law compression  $|\mathbf{x}|^p$  with p = 0.3. Finally we applied per-speaker 0-mean 1-std normalization.

In the post-processing stage, the enhanced waveform generated by the speech enhancement models was reconstructed by applying the inverse STFT to the estimated clean spectrogram and using the phase of the noisy input signal.

	SDR	PESQ	ViSQOL
Noisy	-1.06	1.81	2.11
VL2M	3.17	1.51	1.16
VL2M_ref	6.50	2.58	2.99
AV concat	6.31	2.49	2.83
AV c-ref	6.17	2.58	2.96

 Table 1. GRID results - speaker-dependent. The "Noisy" row refers to the metric values of the input mixed-speech signal.

	2 Speakers		3 Speakers			
	SDR 2	PESQ	ViSQOL	SDR	PESQ	ViSQOL
Noisy	0.21	1.94	2.58	-5.34	1.43	1.62
VL2M	3.02	1.81	1.70	-2.03	1.43	1.25
VL2M_ref	6.52	2.53	3.02	2.83	2.19	2.53
AV concat	7.37	2.65	3.03	3.02	2.24	2.49
AV c-ref	8.05	2.70	3.07	4.02	2.33	2.64

Table 2. GRID results - speaker-independent.

# 3.4. Video pre-processing

Face landmarks were extracted from video using the Dlib [7] implementation of the face landmark estimator described in [6]. It returns 68 x-y points, for an overall 136 values. We upsampled from 25/29.97 fps (GRID/TCD-TIMIT) to 100 fps to match the frame rate of the audio spectrogram. Upsampling was carried out through linear interpolation over time.

The final video feature vector  $\mathbf{v}$  was obtained by computing the per-speaker normalized motion vector of the face landmarks by simply subtracting every frame with the previous one. The motion vector of the first frame was set to zero.

## 4. RESULTS

In order to compare our models to previous works in both speech enhancement and separation, we evaluated the performance of the proposed models using both speech separation and enhancement metrics. Specifically, we measured the capability of separating the target utterance from the concurrent utterance with the source-to-distortion ratio (SDR) [27, 28].

	2 Speakers		3 Speakers			
	SDR	PESQ	ViSQOL	SDR	PESQ	ViSQOL
Noisy	0.21	2.22	2.74	-3.42	1.92	2.04
VL2M	2.88	2.25	2.62	-0.51	1.99	1.98
VL2M_ref	9.24	2.81	3.09	5.27	2.44	2.54
AV concat	9.56	2.80	3.09	5.15	2.41	2.52
AV c-ref	10.55	3.03	3.21	5.37	<b>2.45</b>	2.58

Table 3. TCD-TIMIT results - speaker-independent.

While the quality of estimated target speech was measured with the perceptual PESQ [29] and ViSQOL [30] metrics. For PESQ we used the narrow band mode while for ViSQOL we used the wide band mode.

As a very first experiment we compared landmark position vs. landmark motion vectors. It turned out that landmark positions performed poorly, thus all results reported here refer to landmark motion vectors only.

We then carried out some speaker-dependent experiments to compare our models with previous studies as, to the best of our knowledge, there are no reported results of speakerindependent systems trained and tested on GRID and TCD-TIMIT to compare with. Table 1 reports the test-set evaluation of speaker-dependent models on the GRID corpus with landmark motion vectors. Results are comparable with previous state-of-the-art studies in an almost identical setting [15, 17].

Table 2 and 3 show speaker-independent test-set results on the GRID and TCD-TIMIT datasets respectively. V2ML performs significantly worse than the other three models indicating that a successful mask generation has to depend on the acoustic context. The performance of the three models in the speaker-independent setting is comparable to that in the speaker-dependent setting.

AV concat-ref outperforms V2ML\_ref and AV concat for both datasets. This supports the utility of a refinement strategy and suggests that the refinement is more effective when it directly refines the estimated clean spectrogram, rather than refining the estimated mask.

Finally, we evaluated the systems in a more challenging testing condition where the target utterance was mixed with 2 utterances from 2 competing speakers. Despite the model was trained with mixtures of two speakers, the decrease of performance was not dramatic.

Code and some testing examples of our models are available at https://goo.gl/3h1NgE.

# 5. CONCLUSION

This paper proposes the use of face landmark motion vectors for audio-visual speech enhancement in a single-channel multi-talker scenario. Different models are tested where landmark motion vectors are used to generate time-frequency (T-F) masks that extract the target speaker's spectrogram from the acoustic mixed-speech spectrogram.

To the best of our knowledge, some of the proposed models are the first models trained and evaluated on the limited size GRID and TCD-TIMIT datasets that accomplish speakerindependent speech enhancement in the multi-talker setting, with a quality of enhancement comparable to that achieved in a speaker-dependent setting.

#### 6. REFERENCES

- E. Colin Cherry, "Some experiments on the recognition of speech, with one and with two ears," *The Journal of the Acoustical Society of America*, vol. 25, no. 5, pp. 975–979, 1953.
- [2] Josh H McDermott, "The cocktail party problem," *Current Biology*, vol. 19, no. 22, pp. R1024–R1027, 2009.
- [3] Elana Zion Golumbic, Gregory B. Cogan, Charles E. Schroeder, and David Poeppel, "Visual input enhances selective speech envelope tracking in auditory cortex at a "cocktail party"," *Journal of Neuroscience*, vol. 33, no. 4, pp. 1417–1426, 2013.
- [4] Wei Ji Ma, Xiang Zhou, Lars A. Ross, John J. Foxe, and Lucas C. Parra, "Lip-reading aids word recognition most in moderate noise: A bayesian explanation using high-dimensional feature space," *PLOS ONE*, vol. 4, no. 3, pp. 1–14, 03 2009.
- [5] Albert S Bregman, Auditory scene analysis: The perceptual organization of sound, MIT press, 1994.
- [6] Vahid Kazemi and Josephine Sullivan, "One millisecond face alignment with an ensemble of regression trees," in *The IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), June 2014.
- [7] Davis E. King, "Dlib-ml: A machine learning toolkit," Journal of Machine Learning Research, vol. 10, pp. 1755–1758, 2009.
- [8] Yuxuan Wang, Arun Narayanan, and DeLiang Wang, "On Training Targets for Supervised Speech Separation," *IEEE/ACM Transactions* on Audio, Speech, and Language Processing, vol. 22, no. 12, pp. 1849–1858, Dec. 2014.
- [9] Martin Cooke, Jon Barker, Stuart Cunningham, and Xu Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *The Journal of the Acoustical Society of America*, vol. 120, no. 5, pp. 2421–2424, Nov. 2006.
- [10] Naomi Harte and Eoin Gillen, "TCD-TIMIT: An Audio-Visual Corpus of Continuous Speech," *IEEE Transactions on Multimedia*, vol. 17, no. 5, pp. 603–615, May 2015.
- [11] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), March 2017, pp. 246–250.
- [12] Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R. Hershey, "Single-channel multi-speaker separation using deep clustering," in *Interspeech*, 2016.
- [13] Morten Kolbaek, Dong Yu, Zheng-Hua Tan, Jesper Jensen, 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 Trans. Audio, Speech and Lang. Proc.*, vol. 25, no. 10, pp. 1901–1913, Oct. 2017.
- [14] Bertrand Rivet, Wenwu Wang, Syed Mohsen Naqvi, and Jonathon Chambers, "Audiovisual Speech Source Separation: An overview of key methodologies," *IEEE Signal Processing Magazine*, vol. 31, no. 3, pp. 125–134, May 2014.
- [15] Aviv Gabbay, Ariel Ephrat, Tavi Halperin, and Shmuel Peleg, "Seeing through noise: Visually driven speaker separation and enhancement," in *ICASSP*. 2018, pp. 3051–3055, IEEE.
- [16] Ariel Ephrat, Tavi Halperin, and Shmuel Peleg, "Improved speech reconstruction from silent video," *ICCV 2017 Workshop on Computer Vision for Audio-Visual Media*, 2017.

- [17] Aviv Gabbay, Asaph Shamir, and Shmuel Peleg, "Visual speech enhancement," in *Interspeech*. 2018, pp. 1170–1174, ISCA.
- [18] Jen-Cheng Hou, Syu-Siang Wang, Ying-Hui Lai, Yu Tsao, Hsiu-Wen Chang, and Hsin-Min Wang, "Audio-Visual Speech Enhancement Using Multimodal Deep Convolutional Neural Networks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, pp. 117–128, Apr. 2018.
- [19] Jen-Cheng Hou, Syu-Siang Wang, Ying-Hui Lai, Jen-Chun Lin, Yu Tsao, Hsiu-Wen Chang, and Hsin-Min Wang, "Audio-visual speech enhancement using deep neural networks," in 2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA), Jeju, South Korea, Dec. 2016, pp. 1–6, IEEE.
- [20] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T. Freeman, and Michael Rubinstein, "Looking to Listen at the Cocktail Party: A Speaker-Independent Audio-Visual Model for Speech Separation," ACM Transactions on Graphics, vol. 37, no. 4, pp. 1–11, July 2018, arXiv: 1804.03619.
- [21] T. Afouras, J. S. Chung, and A. Zisserman, "The conversation: Deep audio-visual speech enhancement," in *Interspeech*, 2018.
- [22] Andrew Owens and Alexei A Efros, "Audio-visual scene analysis with self-supervised multisensory features," *European Conference on Computer Vision (ECCV)*, 2018.
- [23] Michael C. Anzalone, Lauren Calandruccio, Karen A. Doherty, and Laurel H. Carney, "Determination of the potential benefit of timefrequency gain manipulation," *Ear Hear*, vol. 27, no. 5, pp. 480–492, Oct 2006, 16957499[pmid].
- [24] Ulrik Kjems, Jesper B. Boldt, Michael S. Pedersen, Thomas Lunner, and DeLiang Wang, "Role of mask pattern in intelligibility of ideal binary-masked noisy speech," *The Journal of the Acoustical Society* of America, vol. 126, no. 3, pp. 1415–1426, 2009.
- [25] A. Graves, A. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, May 2013, pp. 6645–6649.
- [26] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.
- [27] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 14, no. 4, pp. 1462–1469, July 2006.
- [28] 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.
- [29] A.W. Rix, J.G. Beerends, M.P. Hollier, and A.P. Hekstra, "Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs," in 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.01CH37221), Salt Lake City, UT, USA, 2001, vol. 2, pp. 749–752, IEEE.
- [30] A. Hines, J. Skoglund, A. Kokaram, and N. Harte, "ViSQOL: The Virtual Speech Quality Objective Listener," in *IWAENC 2012; International Workshop on Acoustic Signal Enhancement*, Sept. 2012, pp. 1–4.