

Scenario-Aware Audio-Visual TF-GridNet for Target Speech Extraction

Pan, Zexu; Wichern, Gordon; Masuyama, Yoshiki; Germain, François G; Khurana, Sameer; Hori, Chiori; Le Roux, Jonathan

TR2023-152 December 20, 2023

Abstract

Target speech extraction aims to extract, based on a given conditioning cue, a target speech signal that is corrupted by interfering sources, such as noise or competing speakers. Building upon the achievements of the state-of-the-art (SOTA) time-frequency speaker separation model TF-GridNet, we propose AV-GridNet, a visual-grounded variant that incorporates the face recording of a target speaker as a conditioning factor during the extraction process. Recognizing the inherent dissimilarities between speech and noise signals as interfering sources, we also propose SAV-GridNet, a scenario-aware model that identifies the type of interfering scenario first and then applies a dedicated expert model trained specifically for that scenario. Our proposed model achieves SOTA results on the second COG-MHEAR Audio-Visual Speech Enhancement Challenge, outperforming other models by a significant margin, objectively and in a listening test. We also perform an extensive analysis of the results under the two scenarios.

IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU) 2023

SCENARIO-AWARE AUDIO-VISUAL TF-GRIDNET FOR TARGET SPEECH EXTRACTION

Zexu Pan¹, Gordon Wichern¹, Yoshiki Masuyama^{1,2}, François G. Germain¹,
Sameer Khurana¹, Chiori Hori¹, Jonathan Le Roux¹

¹Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

²Tokyo Metropolitan University, Tokyo, Japan

ABSTRACT

Target speech extraction aims to extract, based on a given conditioning cue, a target speech signal that is corrupted by interfering sources, such as noise or competing speakers. Building upon the achievements of the state-of-the-art (SOTA) time-frequency speaker separation model TF-GridNet, we propose AV-GridNet, a visual-grounded variant that incorporates the face recording of a target speaker as a conditioning factor during the extraction process. Recognizing the inherent dissimilarities between speech and noise signals as interfering sources, we also propose SAV-GridNet, a scenario-aware model that identifies the type of interfering scenario first and then applies a dedicated expert model trained specifically for that scenario. Our proposed model achieves SOTA results on the second COG-MHEAR Audio-Visual Speech Enhancement Challenge, outperforming other models by a significant margin, objectively and in a listening test. We also perform an extensive analysis of the results under the two scenarios.

Index Terms— Audio-visual, scenario-aware, speaker extraction, TF-GridNet, time-frequency

1. INTRODUCTION

Speech, as the most natural form of human communication, effectively delivers rich information about a speaker’s emotion, identity, location, or spoken content. Many speech processing algorithms have been developed to extract such information [1–5]. These algorithms are often optimized for clean speech signals, while real-world speech signals are typically contaminated by interfering signals such as noise and irrelevant speakers, which is exemplified as the “cocktail party problem” [6]. Therefore, it is often beneficial to incorporate a pre-processing step to extract the speech signal of interest, a task commonly known as target speech extraction [7].

Target speech extraction algorithms are usually conditioned on an auxiliary reference or cue, to distinguish the target speaker in the case of overlapping speakers. A widely studied case is that of a pre-recorded speech utterance being used as the reference, where the network extracts the speech that sounds similar to the speech signal reference [8–14].

A drawback of such an approach is that pre-enrollment of each target speaker is needed, which is cumbersome or even unfeasible in some circumstances.

Human attention is known to be multi-modal [15], involving various sensory stimuli that are processed interactively as described by the reentry theory [16]. Notably, studies have shown that watching a speaker significantly improves speech comprehension in a challenging “cocktail party” scenario [17–19]. Motivated by these studies and the robustness of visual cues against acoustic noise, there have been various attempts to condition the target speech algorithm on visual signals. For example, the FaceFilter model explores the face-voice correspondence using a single face image [20], the reentry model explores the speech-lip synchronization using a lip recording [21], and the SEG model explores the speech-gesture association using an upper-body recording [22].

Among visual cues, face recordings are generally understood to be the most effective and are the most often used, as visemes provide the exact places of articulation [23–32]. Researchers have explored the use of face recordings in various target speech networks, such as frequency-domain bidirectional long short-term memory (BLSTM) networks [33, 34], time-domain temporal convolutional networks (TCN) [21, 26, 35–37], or dual-path recurrent neural networks (DPRNN) [38, 39]. In this work, motivated by the recent success of the time-frequency-domain speaker separation model TF-GridNet, we propose to condition it on face recordings for audio-visual target speech extraction, referring to this model as AV-GridNet.

While a target speaker extraction algorithm can work independently to the type of interfering signals, we contend that, the characteristics of speech and noise being very different, it may be more advantageous to individually optimize a model for each interference scenario, as a model separating speech from speech is likely to more heavily rely on the intrinsic structure of speech signals, while a model separating speech from noise is likely to more heavily rely on the differences between the characteristics of speech and noise.

In this paper, on top of a universal AV-GridNet that processes the mixture speech signal irrespective of the type of interfering signals, we thus propose a scenario-aware AV-GridNet model, SAV-GridNet, that explicitly integrates the different interfering scenarios. SAV-GridNet is a cascaded

This work was performed while Y. Masuyama was an intern at MERL.

model that first identifies the type of interfering signals with a classifier model, and then applies a dedicated expert AV-GridNet model that is trained specifically for that scenario. To validate the effectiveness of our proposed networks, we participated in the second COG-MHEAR Audio-Visual Speech Enhancement Challenge [40], and achieved state-of-the-art (SOTA) results in terms of objective measures such as perceptual evaluation of speech quality (PESQ), short-time objective intelligibility (STOI), and scale-invariant signal-to-distortion ratio (SI-SDR), in addition to word intelligibility in a listening test, outperforming the other teams and the baseline [41] by a significant margin. We also perform an extensive analysis of the results under the two scenarios.

2. METHODOLOGY

2.1. Related work: TF-GridNet

Our proposed system is built upon the SOTA time-frequency-domain model called TF-GridNet, which has demonstrated promising results in various tasks including speech separation [42] and multi-channel audio-only target speech extraction [43]. TF-GridNet directly estimates the real and imaginary components of the target speech signal from those of the mixture speech. It encodes a speech signal into TF representations of dimension $T \times F \times D$ with a short-time Fourier transform (STFT) followed by a 2-dimensional convolution (Conv2D) and a layer normalization (LN) operation. Then, B repetitions of GridNet blocks are applied to refine the TF representations. Finally, the real and imaginary components of the target speech are obtained by applying 2-dimensional deconvolution to the output of the final GridNet block, followed by an inverse STFT (iSTFT) operation to output the separated speech signals.

Each GridNet block consists of three successive modules:

- 1) An intra-frame spectral module that views the TF representation as T separate sequences of D -dimensional embeddings, each sequence having length F , and applies a BLSTM layer with H units and a 1-dimensional deconvolution layer with kernel size I and stride J ;
- 2) A sub-band temporal module that views the TF representation as F separate sequences of D -dimensional embeddings, each sequence having length T , and performs a similar procedure as in the intra-frame spectral module;
- 3) A full-band self-attention module that first reshapes the TF representation into a single sequence of length T with $F \times D$ channels, and applies a multi-head self-attention operation with L heads.

2.2. Proposed AV-GridNet

Instead of separating all speakers into individual streams, we propose a modified version of TF-GridNet for audio-visual target speaker extraction called AV-GridNet. AV-GridNet is conditioned on the face recording v of the target speaker, and it extracts only the corresponding target speech \hat{s} from the mixture speech signal x , irrespective of the type of interference signals. The architecture of AV-GridNet is illustrated

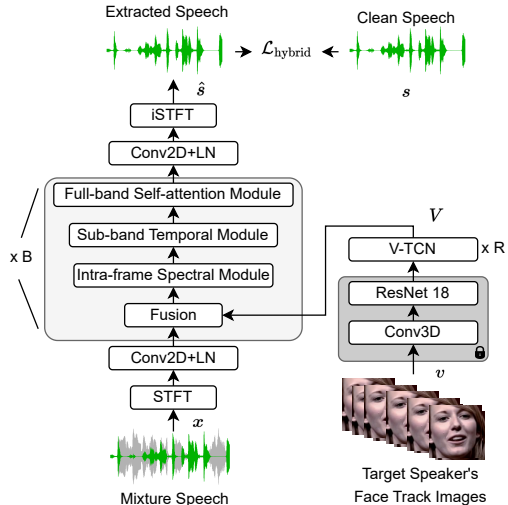


Fig. 1. Our AV-GridNet model extracts the target speech conditioned on the target’s face recording.

in Fig. 1, incorporating an additional visual conditioning network to extract visual features V from the face recording v .

2.2.1. Visual conditioning network

The visual conditioning network comprises a 3-dimensional convolutional layer (Conv3D), a ResNet 18 layer, and R repetitions of visual temporal convolutional network (V-TCN), as depicted in Fig. 1. The Conv3D and ResNet 18 layers are pre-trained on lip-reading tasks and are kept frozen during the training of AV-GridNet¹. This allows the network to retain the ability to encode viseme movements that synchronize with the phoneme sequence of speech. Additionally, we employ V-TCN layers as an adaptation, similar to [21, 39], to adapt the visual embeddings towards speech extraction.

The visual embedding V typically has a lower temporal resolution compared to speech embeddings. To address this mismatch, we linearly interpolate the visual embeddings along the time dimension to match the resolution of the speech embeddings. We fuse the same visual embeddings V using a fusion layer to the start of each GridNet block. Specifically, we concatenate the audio and visual embeddings along the channel dimension, and project them back to the original channel dimension of the audio embeddings before fusion with a linear layer.

2.2.2. Loss function to train the AV-GridNet

In time-domain end-to-end speaker extraction network training, the negative SI-SDR loss function [45] has been widely used in most methods. It is formulated as follows:

$$\mathcal{L}_{\text{SI-SDR}}(s, \hat{s}) = -20 \log_{10} \frac{\|\langle \hat{s}, s \rangle_s\|}{\|\hat{s} - \frac{\langle \hat{s}, s \rangle_s}{\|s\|^2} s\|}. \quad (1)$$

¹The pre-trained visual network can be found at https://github.com/smeetr/s/deep_avsr [44]

In this work, we adopt the hybrid loss proposed in [41], which reduces the over-suppression error and leads to improved perceptual quality and intelligibility for the extracted speech. The hybrid loss consists of the time-domain SI-SDR loss as shown in Eq. (1), together with a frequency-domain multi-resolution delta spectrum loss²:

$$\mathcal{L}_{\text{hybrid}}(s, \hat{s}) = \mathcal{L}_{\text{SI-SDR}}(s, \hat{s}) + \gamma \frac{1}{M} \sum_{m=1}^M \mathcal{L}_{\text{freq-}\Delta}^m(s, \hat{s}), \quad (2)$$

where the delta spectrum loss $\mathcal{L}_{\text{freq-}\Delta}^m$ is calculated at $M = 3$ different resolutions, using the following triplets of parameters for {FFT size, hop size, window length} in samples: {512, 50, 240}, {1024, 120, 600}, and {2048, 240, 1200}. γ is a balancing weight that is set to 1 in this paper.

2.3. Proposed scenario-aware SAV-GridNet

2.3.1. Motivation

Speech and noise exhibit very distinct characteristics, and scenarios where one or the other acts as interfering signal of a target speaker may thus require different strategies. Indeed, a model trained to separate speech from speech is likely to heavily rely on the structure of speech, and thus to be different from a model trained to separate speech from noise, which has the opportunity to rely on the intrinsic differences between the characteristics of the two signals to be separated. Therefore, we advocate that a dedicated expert AV-GridNet model that is trained specifically for noise or speech interference may better handle each scenario.

To this end, we propose a model, referred to as SAV-GridNet, that is aware of the different interference scenarios as depicted in Fig. 2. SAV-GridNet first identifies the type of interfering scenario with a classifier network, and then applies a dedicated expert model that is trained specifically for that scenario. The classifier model and the expert models AV-GridNet_n (for noise interference) and AV-GridNet_s (for speech interference) are trained independently.

Note that we consider here scenarios involving either noise or a single speaker as interference because of the particular setting of the COG-MHEAR challenge. A generalized and arguably more realistic setting for practical applications would be to consider noise-only interference on one hand, and one or more speakers with or without background noise on the other. While our proposed classifier-based approach can be readily extended to this setting, with the corresponding expert models, we leave a thorough investigation of the performance of such a system to future work.

2.3.2. Classifier network

The classifier network, detailed in the left panel of Fig. 2, accepts both x and v as inputs. Although the task could be performed with only x , it may be beneficial to include the

²Code for the hybrid loss function can be found at https://github.com/zexupan/avse_hybrid_loss [41]

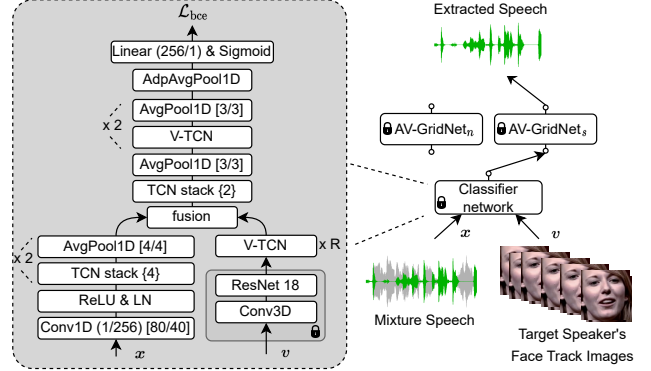


Fig. 2. Our scenario-aware SAV-GridNet model is a cascaded model that first classifies the type of interfering signals with a classifier network, and then applies dedicated expert models AV-GridNet_n for noise interference or AV-GridNet_s for speech interference. In the classifier network, the values inside “(/)” represent the input and output feature sizes, the values inside “[/]” represent kernel size and stride, and the value inside “{ }” represents the upper-bound of the dilation value in a TCN stack [35]. AvgPool1D and AdpAvgPool1D denote 1D average pooling and 1D adaptive average pooling operations, respectively.

visual signals here as it could serve as an anchor point for the target speech. The classifier network design is motivated by the audio-visual SLSyn network in [21], which consists of a visual front-end, a speech front-end, and an audio-visual back-end. It is worth noting that the visual front-end here also consists of Conv3D, ResNet 18, and V-TCN layers. As with the visual conditioning network in Fig. 1, the Conv3D and ResNet 18 layers are pre-trained on lip-reading tasks and are kept frozen during the training of the classifier model.

We minimize the following binary cross-entropy loss for the scenario classifier network training:

$$\mathcal{L}_{\text{bce}} = -y \log(\hat{y}) - (1 - y) \log(1 - \hat{y}), \quad (3)$$

where $y \in \{0, 1\}$ indicates whether the interfering signal is speech or noise, while \hat{y} is the predicted probability. We arbitrarily set speech interference to be the negative class and noise interference to be the positive class.

2.3.3. Classifier post-processing

If the classifier makes a mistake and the wrong expert model is used, the results may be detrimental as there is a mismatch between training and inference for the AV-GridNet. We empirically find that a model trained only on speech interference generalizes remarkably well on noise interference, but not vice versa. Therefore, we propose two post-processing strategies to mitigate the false-positive cases (i.e., predicting noise while the ground-truth label is speech).

For the first post-processing strategy (post-proc₁), if the classifier prediction is noise and the criteria in Eq. (4) is met, which indicates that the universal model is more in agreement

with the noise expert than with the speech expert, we classify the interference as noise, otherwise as speech:

$$\mathcal{L}_{\text{SI-SDR}}(\hat{s}, \hat{s}_n) < \mathcal{L}_{\text{SI-SDR}}(\hat{s}, \hat{s}_s), \quad (4)$$

where

$$\hat{s}_n = \text{AV-GridNet}_n(x, v), \quad (5)$$

$$\hat{s}_s = \text{AV-GridNet}_s(x, v), \quad (6)$$

$$\hat{s} = \text{AV-GridNet}(x, v). \quad (7)$$

For the second post-processing strategy (post-proc₂), if the classifier prediction is noise, and either of the criteria in Eq. (4) or Eq. (8) is met, with this latter criterion indicating that the original mixture is further to the output of the noise expert than to that of the speech expert, we classify the interference as noise, otherwise as speech:

$$\mathcal{L}_{\text{SI-SDR}}(x, \hat{s}_n) > \mathcal{L}_{\text{SI-SDR}}(x, \hat{s}_s) \quad (8)$$

For the samples that are classified as noise interference initially by the classifier but are changed to speech interference by the post-processing, we use the extracted speech from the universal AV-GridNet model. For all other samples, the model indicated by the classifier is used.

3. EXPERIMENTAL SETUP

3.1. Dataset

We participated in the second COG-MHEAR Audio-Visual Speech Enhancement Challenge³ and evaluated our proposed method on its benchmark dataset. The speech dataset is from the Lip Reading Sentences 3 (LRS3) [46], which consists of thousands of spoken sentences from TED videos. The noise datasets are from the Clarity challenge [47], which comprises around 7 hours of domestic noises, the DEMAND dataset [48], which includes recordings of 18 soundscapes that represent over 1 hour of data, and the Deep Noise Suppression (DNS) challenge [49], for which only the noise signals from Freesound [50] are considered. The challenge has two tracks: systems in track 1 can only use the above-mentioned provided datasets and unimodal pre-trained models, while systems in track 2 have no limitations in the datasets and pre-trained models used.

The training, development, and evaluation sets consist of 34519, 3300, and 2792 scenes respectively. There are two scenarios in total, a target speaker mixed with a competing speaker at random signal-to-noise ratio (SNR) levels that range from -15 dB to $+5$ dB, or a target speaker mixed with a noise signal at random SNR levels that range from -10 dB to $+10$ dB. The clean speech signals and scenario labels are only available for the training and development set. The audio signals are sampled at 16 kHz, while the video has a frame rate of 25 per second. The target face tracks are provided for all the samples.

³<https://challenge.cogmhear.org/>

3.2. Baselines

We use the AV-DPRNN network [39, 41] as our main baseline, as it is currently one of the best-performing audio-visual speech extraction networks. There are three main differences between AV-DPRNN and AV-GridNet: 1) The speech encoder and decoders used by AV-DPRNN are in the time domain, while those of AV-GridNet are in the time-frequency domain; 2) AV-DPRNN is a mask-based method that uses dual-path BLSTM as the extractor, while AV-GridNet directly maps the signals using the GridNet blocks as the extractor; 3) The visual embeddings are only fused at the first repeat of the extractor for AV-DPRNN, while the visual embeddings are fused at every repeat of the extractor for AV-GridNet. We also report results by the official baseline released by the challenge organization. It has a similar architecture to the AV-DPRNN network [39, 41], but with no visual pre-training involved.

3.3. Model and training settings

For the baseline AV-DPRNN, the hyperparameter setting follows [39, 41]. For the classifier network in SAV-GridNet, the TCN stack hyperparameter follows [35]. For AV-GridNet, the V-TCN hyperparameter follows [39, 41]. We set $D = 48$, $B = 6$, and $R = 5$. The STFT window size is 256, the hop size is 128, and the square root Hann window is used. A 256-point discrete Fourier transform is applied to extract 129-dimensional complex spectra at each frame. For other hyperparameters in the GridNet block, we set $I = 4$, $J = 1$, $H = 192$, $E = 4$, and $L = 4$ [42].

For all model training, we use the Adam optimizer with an initial learning rate of 0.001, the learning rate is halved if the best development loss (BDL) does not improve for 6 consecutive epochs, and the training stops when the BDL does not improve for 20 consecutive epochs. We train the model on 8 GPUs with 48 GB RAM each. To fit the data in the GPU memory during training, the audio clips are truncated to 3 seconds for AV-GridNet, 12 seconds for AV-DPRNN, and 25 seconds for the classifier network.

4. RESULTS

We evaluate the speech signals extracted by our proposed networks and the baselines using objective measures PESQ, STOI, and SI-SDR. PESQ measures the perceptual quality of the extracted speech signal and is in the range of -0.5 to 4.5 ; STOI measures the intelligibility of the extracted speech signal and is in the range of 0 to 1; and SI-SDR measures the signal quality of the extracted speech signal in dB and is unbounded. The higher the better for all three metrics. We use PESQ as our main measure when describing the results, as other measures show similar trends.

4.1. Comparison with baseline and ablation study

In Table 1, we present the results of our baseline and proposed models on the development set. For our baseline AV-DPRNN, it is seen that using hybrid loss (Sys. 2) outperforms (Sys. 1)

Table 1. Development set results on the 2nd COG-MHEAR Audio-Visual Speech Enhancement Challenge benchmark. Experiments are done for track 1, in which only the provided dataset is used in training. We use the system number (Sys.) to identify different systems. Init. indicates which (if any) other system is used to initialize that system’s parameters (note that the optimizer parameters are always re-initialized). DM stands for dynamic mixing, which involves simulating the mixture utterances on the fly during training using the protocol provided by the challenge organizers. We report performance on the speech+speech scenario, the speech+noise scenario, and overall. The symbol * indicates the model used the oracle scenario labels.

Sys.	Model	Init.	DM	Loss	Speech+Speech			Speech+Noise			Overall		
					PESQ	STOI	SI-SDR	PESQ	STOI	SI-SDR	PESQ	STOI	SI-SDR
-	Noisy	-	-	-	1.17	0.60	-5.0	1.15	0.68	-4.4	1.16	0.64	-4.7
1	AV-DPRNN	-	✗	$\mathcal{L}_{\text{SI-SDR}}$	1.68	0.84	8.8	1.71	0.82	9.8	1.69	0.83	9.3
2		1	✗	$\mathcal{L}_{\text{hybrid}}$	2.23	0.90	12.6	2.02	0.86	11.4	2.12	0.88	12.0
3	AV-GridNet	-	✗	$\mathcal{L}_{\text{hybrid}}$	3.07	0.95	17.3	2.36	0.87	12.8	2.78	0.92	14.6
4		3	✓		3.10	0.95	16.7	2.62	0.91	13.9	2.86	0.93	15.3
5	AV-GridNet _s	3	✓	$\mathcal{L}_{\text{hybrid}}$	3.23	0.95	17.5	2.56	0.90	13.4	2.89	0.93	15.5
6	AV-GridNet _n				1.27	0.61	-4.7	2.68	0.91	14.2	1.98	0.76	4.8
7	SAV-GridNet	-	✓	$\mathcal{L}_{\text{hybrid}}$	3.22	0.95	17.4	2.68	0.91	14.2	2.95	0.93	15.8
8	+ post-proc ₁				3.23	0.95	17.5	2.68	0.91	14.2	2.95	0.93	15.8
9	+ post-proc ₂				3.23	0.95	17.5	2.68	0.91	14.2	2.95	0.93	15.8
10	SAV-GridNet*	-	✓	$\mathcal{L}_{\text{hybrid}}$	3.23	0.95	17.5	2.68	0.91	14.2	2.95	0.93	15.8




Fig. 3. Confusion matrix of the scenario classification on the development set.

by 0.43 on the overall PESQ, showing the effectiveness of the frequency-domain loss on the speech perceptual quality.

Our first proposed AV-GridNet (Sys. 3) outperforms AV-DPRNN (Sys. 2) by 0.66 for PESQ. With additional dynamic mixing (Sys. 4), the PESQ further improves by 0.08. Our expert models, AV-GridNet_s and AV-GridNet_n, improved PESQ compared to AV-GridNet by 0.13 in the speech+speech scenario and by 0.06 in the speech+noise scenario, respectively. Thanks to adaptively choosing the expert models, our proposed SAV-GridNet (Sys. 7) outperforms the AV-GridNet (Sys. 4) regardless of the scenarios. While the post-processing techniques (Sys. 8 and 9) do not appear to show performance improvements on the averaged metrics, they do reduce outliers, as will be investigated in the next subsection. The performance of our best system is nearly identical to that of SAV-GridNet with the oracle scenario labels (Sys. 10).

4.2. Analysis for speech and noise interfering signals

Fig. 3 illustrates the confusion matrix of the scenario classification with and without post-processing. Our scenario classification network achieved accuracy over 99% without post-

processing. Since the false-positive cases (predicting noise when the ground-truth label is speech) severely deteriorate the subsequent target speech extraction performance according to Sys. 6 in Table 1, it is important to reduce the number of false-positive cases, which post-proc₁ does successfully. This however increased the false-negative cases to 58, but post-proc₂ mitigated this increase and performed best overall in terms of target speech extraction.

This tendency is also confirmed from the distributions of PESQ shown in Figs. 4–6. SAV-GridNet (Sys. 7, Fig. 5) improved the overall performance from AV-GridNet (Sys. 4, Fig. 4), but it had more outliers in the speech+speech scenario. This is likely because AV-GridNet_n was applied to some speech+speech samples due to misclassification. As the post-processing techniques successfully reduced the false-positive cases, Sys. 9 reduced the number of outliers with low PESQ in the left panel of Fig. 6, while substantially preserving the distribution in the speech+noise scenario. Overall, the number of samples with a PESQ value smaller than 1.5 went down from 62 for Sys. 7 to 54 and 53 for Sys. 8 and 9, respectively, post-proc₁ thus reducing the number of such failing samples by 13%, and post-proc₂ by 14.5%.

Analyzing low performing samples: When informally listening to samples with the lowest objective metrics, we noticed that, while a few samples did have mid-utterance switching between target and interfering speakers in the speech+speech scenario, the main issue was that many of the target speech signals were not very clean, e.g., they contained impulsive disturbances from microphone contact or crowd noises such as cheering and clapping. To quantify objectively the quality of the target speech signals, we used the P808




Fig. 4. PESQ of extracted speech signal from system 4 for the speech (left) and noise (right) interfering signals.




Fig. 5. PESQ of extracted speech signal from system 7 for the speech (left) and noise (right) interfering signals.




Fig. 6. PESQ of extracted speech signal from system 9 for the speech (left) and noise (right) interfering signals.

DNSMOS [51] score, a reference-free measure for evaluating overall audio quality. Figure 7 displays the distribution of SI-SDR vs. DNSMOS. We chose SI-SDR over PESQ as the reference metric for this figure because the larger dynamic range makes outliers more visible. We note that all of the lowest SI-SDR samples output by our model have DNSMOS values below 3.0 in Fig. 7, and we found by informal listening that these samples contained noisy target speech. As datasets become larger, using a reference-free speech quality metric could help remove noisy target speech samples.

4.3. Performance on leaderboard

In Table 2, we present the performance of our models on the hidden evaluation set, for which the numbers are obtained from the submissions to the leaderboard⁴. We can see that our AV-DPRNN baseline (Sys. 2) outperforms the challenge baseline by 0.53 in terms of PESQ, thanks to the combined use of the presented pre-trained models, hybrid loss, and training settings. We also report the results of the top 3 other teams on track 1 in terms of PESQ. Results for track 2 are not reported as they did not improve upon those of track 1.

AV-GridNet outperforms AV-DPRNN by 0.65 in terms of

⁴<https://challenge.cogmhear.org/#/results>




Fig. 7. Development set SI-SDR improvement [dB] vs. P808 DNSMOS for system 9 for speech (left) and noise (right) interfering signals. Many of the lowest performing samples appear to have noisy ground-truth signals based on their low score under the reference-free DNSMOS quality measure.

Table 2. Evaluation set results on the second COG-MHEAR Audio-Visual Speech Enhancement Challenge benchmark. We only present the top-3 teams on track 1 out of 8 other teams based on the PESQ ranking.

Sys.	Model	Track	PESQ	STOI	SI-SDR
-	Noisy Baseline	-	1.14	0.44	-5.1
-	Team 1	1	1.61	0.68	8.8
-	Team 2	1	1.66	0.68	6.7
-	Team 3	1	1.76	0.71	7.7
2	AV-DPRNN		1.94	0.73	10.4
4	AV-GridNet		2.59	0.83	13.8
7	SAV-GridNet	1	2.71	0.84	14.5
8	+ post-proc ₁		2.71	0.84	14.5
9	+ post-proc ₂		2.71	0.84	14.5

PESQ, while SAV-GridNet further outperforms AV-GridNet by 0.12. Similarly to the results obtained on the development set, we cannot see a difference on the averaged metrics from using post-processing with SAV-GridNet (Sys. 8 and 9), but we hope that the number of outliers will again be reduced. This will need to be confirmed when/if the evaluation set is released. Furthermore, in a listening test, we achieved an overall word intelligibility score of 84.54%, compared to 57.56% for the baseline and 80.41% for the next best team. The Fisher’s least significant difference (LSD) was 2.14%, indicating that our model offered statistically significant intelligibility improvements compared to all other systems.

5. CONCLUSION

In this work, we explored visually-grounded target speaker extraction based on the TF-GridNet separation architecture. Considering the different characteristics of noise and speech as interfering signals raised by the 2nd COG-MHEAR Audio-Visual Speech Enhancement Challenge, we proposed a scenario-aware model named SAV-GridNet that is capable to apply an expert model to individual scenarios independently. Experimental results show that the scenario-aware model generally improves the quality of the extracted speech, while reducing the number of samples with very low quality.

6. REFERENCES

- [1] Z. Pan, Z. Luo, J. Yang, and H. Li, “Multi-modal attention for speech emotion recognition,” in *Proc. Interspeech*, 2020.
- [2] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust DNN embeddings for speaker recognition,” in *Proc. ICASSP*, 2018.
- [3] X. Qian, M. Madhavi, Z. Pan, J. Wang, and H. Li, “Multi-target DoA estimation with an audio-visual fusion mechanism,” in *Proc. ICASSP*, 2021.
- [4] J. Wang, X. Qian, and H. Li, “Predict-and-Update network: Audio-visual speech recognition inspired by human speech perception,” *arXiv preprint arXiv:2209.01768*, 2022.
- [5] R. Tao, Z. Pan, R. K. Das, X. Qian, M. Z. Shou, and H. Li, “Is someone speaking? Exploring long-term temporal features for audio-visual active speaker detection,” in *Proc. ACM Multimedia*, 2021.
- [6] A. W. Bronkhorst, “The cocktail party phenomenon: A review of research on speech intelligibility in multiple-talker conditions,” *Acta Acust. United Acust.*, vol. 86, no. 1, pp. 117–128, 2000.
- [7] K. Zmolikova, M. Delcroix, T. Ochiai, K. Kinoshita, J. Černocký, and D. Yu, “Neural target speech extraction: An overview,” *IEEE Signal Process. Mag.*, vol. 40, no. 3, pp. 8–29, 2023.
- [8] C. Xu, W. Rao, E. S. Chng, and H. Li, “SpEx: Multi-scale time domain speaker extraction network,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1370–1384, 2020.
- [9] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. R. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, “VoiceFilter: Targeted voice separation by speaker-conditioned spectrogram masking,” in *Proc. Interspeech*, 2019.
- [10] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, “SpEx+: A complete time domain speaker extraction network,” in *Proc. Interspeech*, 2020.
- [11] S. He, H. Li, and X. Zhang, “Speakerfilter: Deep learning-based target speaker extraction using anchor speech,” in *Proc. ICASSP*, 2020.
- [12] X. Xiao, Z. Chen, T. Yoshioka, H. Erdogan, C. Liu, D. Dimitriadis, J. Droppo, and Y. Gong, “Single-channel speech extraction using speaker inventory and attention network,” in *Proc. ICASSP*, 2019.
- [13] J. Shi, J. Xu, Y. Fujita, S. Watanabe, and B. Xu, “Speaker-conditional chain model for speech separation and extraction,” in *Proc. Interspeech*, 2020.
- [14] M. Delcroix, T. Ochiai, K. Zmolikova, K. Kinoshita, N. Tawara, T. Nakatani, and S. Araki, “Improving speaker discrimination of target speech extraction with time-domain SpeakerBeam,” in *Proc. ICASSP*, 2020.
- [15] L. Smith and M. Gasser, “The development of embodied cognition: Six lessons from babies,” *Artif. Life*, vol. 11, no. 1–2, pp. 13–29, 2005.
- [16] G. M. Edelman, *Neural Darwinism: The theory of neuronal group selection.*, Basic Books, 1987.
- [17] W. J. Ma, X. Zhou, L. A. Ross, J. J. Foxe, and L. 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. e4638, 2009.
- [18] E. Z. Golumbic, G. B. Cogan, C. E. Schroeder, and D. Poeppel, “Visual input enhances selective speech envelope tracking in auditory cortex at a “cocktail party”,” *J. Neurosci.*, vol. 33, no. 4, pp. 1417–1426, 2013.
- [19] M. J. Crosse, G. M. Di Liberto, and E. C. Lalor, “Eye can hear clearly now: inverse effectiveness in natural audiovisual speech processing relies on long-term cross-modal temporal integration,” *J. Neurosci.*, vol. 36, no. 38, pp. 9888–9895, 2016.
- [20] S.-W. Chung, S. Choe, J. S. Chung, and H.-G. Kang, “FaceFilter: Audio-visual speech separation using still images,” in *Proc. Interspeech*, 2020.
- [21] Z. Pan, R. Tao, C. Xu, and H. Li, “Selective listening by synchronizing speech with lips,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 1650–1664, 2022.
- [22] Z. Pan, X. Qian, and H. Li, “Speaker extraction with co-speech gestures cue,” *IEEE Signal Process. Lett.*, vol. 29, pp. 1467–1471, 2022.
- [23] D. Michelsanti, Z.-H. Tan, S.-X. Zhang, Y. Xu, M. Yu, D. Yu, and J. Jensen, “An overview of deep-learning-based audio-visual speech enhancement and separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 2021.
- [24] T. Afouras, J. S. Chung, and A. Zisserman, “The conversation: Deep audio-visual speech enhancement,” in *Proc. Interspeech*, 2018.
- [25] J. Wu, Y. Xu, S. Zhang, L. Chen, M. Yu, L. Xie, and D. Yu, “Time domain audio visual speech separation,” in *Proc. ASRU*, 2019.

- [26] Y. Wu, C. Li, J. Bai, Z. Wu, and Y. Qian, “Time-domain audio-visual speech separation on low quality videos,” in *Proc. ICASSP. IEEE*, 2022.
- [27] C. Li and Y. Qian, “Deep audio-visual speech separation with attention mechanism,” in *Proc. ICASSP*, 2020.
- [28] J. Li, M. Ge, Z. Pan, R. Cao, L. Wang, J. Dang, and S. Zhang, “Rethinking the visual cues in audio-visual speaker extraction,” in *Proc. Interspeech*, 2023.
- [29] J. Li, M. Ge, Z. Pan, L. Wang, and J. Dang, “VCSE: Time-domain visual-contextual speaker extraction network,” in *Proc. Interspeech*, 2022.
- [30] K. Tan, Y. Xu, S.-X. Zhang, M. Yu, and D. Yu, “Audio-visual speech separation and dereverberation with a two-stage multimodal network,” *IEEE J. Sel. Topics in Signal Process.*, vol. 14, no. 3, pp. 542–553, 2020.
- [31] R. Lu, Z. Duan, and C. Zhang, “Audio-visual deep clustering for speech separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 11, pp. 1697–1712, 2019.
- [32] G. Morrone, S. Bergamaschi, L. Pasa, L. Fadiga, V. Tikhonoff, and L. Badino, “Face landmark-based speaker-independent audio-visual speech enhancement in multi-talker environments,” in *Proc. ICASSP*, 2019.
- [33] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, “Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation,” *ACM Trans. Graph.*, vol. 37, no. 4, pp. 1–11, 2018.
- [34] T. Ochiai, M. Delcroix, K. Kinoshita, A. Ogawa, and T. Nakatani, “Multimodal SpeakerBeam: Single channel target speech extraction with audio-visual speaker clues,” in *Proc. Interspeech*, 2019.
- [35] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [36] Z. Pan, R. Tao, C. Xu, and H. Li, “MuSE: Multi-modal target speaker extraction with visual cues,” in *Proc. ICASSP*, 2021.
- [37] Z. Pan, W. Wang, M. Borsdorf, and H. Li, “ImagineNet: Target speaker extraction with intermittent visual cue through embedding inpainting,” in *Proc. ICASSP*, 2023.
- [38] Y. Luo, Z. Chen, and T. Yoshioka, “Dual-path RNN: Efficient long sequence modeling for time-domain single-channel speech separation,” in *Proc. ICASSP*, 2020.
- [39] Z. Pan, M. Ge, and H. Li, “USEV: Universal speaker extraction with visual cue,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 3032–3045, 2022.
- [40] A. L. A. Blanco, C. Valentini-Botinhao, O. Klejch, M. Gogate, K. Dashtipour, A. Hussain, and P. Bell, “AVSE challenge: Audio-visual speech enhancement challenge,” in *Proc. SLT*, 2023.
- [41] Z. Pan, M. Ge, and H. Li, “A hybrid continuity loss to reduce over-suppression for time-domain target speaker extraction,” in *Proc. Interspeech*, 2022.
- [42] Z.-Q. Wang, S. Cornell, S. Choi, Y. Lee, B.-Y. Kim, and S. Watanabe, “TF-GridNet: Making time-frequency domain models great again for monaural speaker separation,” in *Proc. ICASSP*, 2023.
- [43] S. Cornell, Z.-Q. Wang, Y. Masuyama, S. Watanabe, M. Pariente, and N. Ono, “Multi-channel target speaker extraction with refinement: The WAVLab submission to the second clarity enhancement challenge,” in *Proc. Clarity*, 2022.
- [44] T. Afouras, J. S. Chung, and A. Zisserman, “Deep lip reading: A comparison of models and an online application,” in *Proc. Interspeech*, 2018.
- [45] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR–half-baked or well done?,” in *Proc. ICASSP*, 2019.
- [46] T. Afouras, J. S. Chung, and A. Zisserman, “LRS3-TED: a large-scale dataset for visual speech recognition,” *arXiv preprint arXiv:1809.00496*, 2018.
- [47] S. Graetzer, J. Barker, T. J. Cox, M. Akeroyd, J. F. Culling, G. Naylor, E. Porter, and R. V. Munoz, “Clarity-2021 challenges: Machine learning challenges for advancing hearing aid processing,” in *Proc. Interspeech*, 2021, vol. 2.
- [48] J. Thiemann, N. Ito, and E. Vincent, “DEMAND: a collection of multi-channel recordings of acoustic noise in diverse environments,” in *Proc. Meetings Acoust.*, 2013.
- [49] H. Dubey, A. Aazami, V. Gopal, B. Naderi, S. Braun, R. Cutler, H. Gamper, M. Golestaneh, and R. Aichner, “Deep speech enhancement challenge at ICASSP 2023,” in *Proc. ICASSP*, 2023.
- [50] E. Fonseca, X. Favory, J. Pons, F. Font, and X. Serra, “FSD50K: An open dataset of human-labeled sound events,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 829–852, 2021.
- [51] C. Reddy, V. Gopal, and R. Cutler, “DNSMOS: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors,” in *Proc. ICASSP*, 2021.