



# Streaming Audio-Visual Speech Recognition with Alignment Regularization

Pingchuan Ma<sup>1,2</sup>, Niko Moritz<sup>1</sup>, Stavros Petridis<sup>1,2</sup>, Christian Fuegen<sup>1</sup>, Maja Pantic<sup>1,2</sup>

<sup>1</sup>Meta AI

<sup>2</sup>Imperial College London

pingchuanma@meta.com, nmoritz@meta.com, stavrosp@meta.com

## Abstract

In this work, we propose a streaming AV-ASR system based on a hybrid connectionist temporal classification (CTC)/attention neural network architecture. The audio and the visual encoder neural networks are both based on the conformer architecture, which is made streamable using chunk-wise self-attention (CSA) and causal convolution. Streaming recognition with a decoder neural network is realized by using the triggered attention technique, which performs time-synchronous decoding with joint CTC/attention scoring. Additionally, we propose a novel alignment regularization technique that promotes synchronization of the audio and visual encoder, which in turn results in better word error rates (WERs) at all SNR levels for streaming and offline AV-ASR models. The proposed AV-ASR model achieves WERs of 2.0% and 2.6% on the Lip Reading Sentences 3 (LRS3) dataset in an offline and online setup, respectively, which both present state-of-the-art results when no external training data are used.

**Index Terms:** speech recognition, human-computer interaction, computational paralinguistics

## 1. Introduction

Audio-Visual Automatic Speech Recognition (AV-ASR) can be defined as the task of recognizing speech by using both audio and visual information. It has recently attracted a lot of attention due to the fact that AV-ASR models allow for increased robustness to acoustic noise. However, the vast majority of works have focused on developing AV-ASR models for non-streaming recognition. For example, [1–4] built AV-ASR models by using either a transformer or conformer encoder on top of the acoustic and visual front-ends for speech recognition, followed by a fusion module that jointly learns the audio-visual encoder state sequence. Finally, a decoder network maps the encoder state sequences to output label sequence using the attention mechanism [5, 6]. However, such attention-based encoder-decoder models are limited to offline scenarios since an entire speech utterance is typically required as an input. To the best of our knowledge, Taris [7] is the only AV-ASR model to perform streaming speech recognition but no results are reported on the largest publicly available audio-visual benchmarks.

Various works [8–10] have shown that AV-ASR models do not always outperform audio-only models in clean scenarios [3, 11]. One reason for this is that AV-ASR models tend to ignore the visual cues in clean scenarios, where audio typically outperforms the visual modality. However, because the

visual modality often demonstrates advantages in noisy situations, some works [3, 12] propose to inject noise for the audio-visual training to encourage the model to better take the visual modality into account. Sterpu *et al.* [9, 10] proposed a cross-modal attention mechanism to automatically learn better alignments between audio and visual streams, which also encourages the system to pay attention to both modalities. However, in a streaming system, where access to future frames is limited, techniques such as cross-modal attention, which performs temporal attention between two sequences, may be sub-optimal, because it would introduce an additional delay.

In this work, we propose a novel regularization technique that promotes synchronization of the audio and visual encoder outputs. In an AV-ASR system, audio and visual encoder neural networks are often pre-trained independently and outputs are fused by an audio-visual fusion operation that simply stacks frames of the encoder outputs from both modalities before being further processed by a neural network [3, 13]. However, there is no guarantee that responses of both encoders are aligned and the operation of stacking features fails to compensate for such misalignments. As a result, a frame-level ASR loss, such as CTC or recurrent neural network transducers (RNN-T) [14, 15], would only be able to make decisions based on one of the modalities, because both encoders may respond to inputs at different frames. To mitigate this problem, we propose to use forced alignment information from the AV-ASR model and two auxiliary cross-entropy-based (CE) losses to better synchronize audio and visual encoder outputs. We can show that such regularization indeed leads to better frame-level alignments of separate audio and visual encoder streams, which in turn results in a more robust performance for our AV-ASR models across all noise conditions.

The proposed alignment regularization technique is applied for building offline and online AV-ASR systems, where the latter must recognize a word shortly after it is spoken. The model architecture is based on [3], where the audio-only and visual-only conformer encoders are made streamable by using chunk-wise restricted self-attention (CSA) and causal convolution [16, 17]. We also use the idea of triggered attention, which leverages the alignment information of the CTC output to trigger an attention-based decoder with a restricted look-ahead onto the encoder sequence [18].

## 2. Model Setup

**Architecture** Our models are built based on the AV-ASR architecture proposed in [3]. The VSR encoder consists of a modified 2D ResNet-18 backbone as the visual front-end [19], followed by a 12-layer conformer encoder, whereas the ASR encoder consists of 1D ResNet-18 as the audio front-end and a 12-

Only non-Meta co-authors downloaded, accessed, and used the datasets. Only non-Meta authors conducted any of the dataset pre-processing (no dataset pre-processing took place on Meta’s servers or facilities).

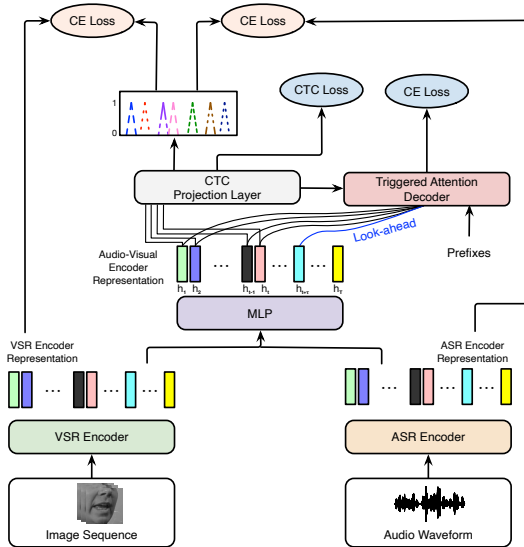


Figure 1: *Triggered attention-based streaming AV-ASR model with alignment regularization, which uses the alignment information from the joint audio-visual CTC output to promote temporally aligned responses of the VSR and ASR encoders through auxiliary CE losses.  $t$  is the time index for a current trigger position and  $\tau$  denotes the look-ahead hyper-parameter of the TA decoder.*

layer conformer encoder. The frame rate at the output of the 1D ResNet-18 is sub-sampled to 25 frames per second (fps), which matches the temporal resolution of the visual stream. The Conformer encoder is made streamable by using chunk-wise self-attention with non-overlapping chunks in order to not increase the computational complexity, which has also demonstrated superior results compared to restricted self-attention [16]. In the streaming setup, causal convolution is applied for the convolutional modules of the conformer architecture [17]. For fusing the audio and visual encoder outputs, a 2-layer multi-layer perceptron (MLP) is used. The joint audio-visual representations, which are computed by the MLP module, are further fed to both CTC and a triggered attention decoder for the joint training. The architecture of the proposed AV-ASR model is illustrated in Fig. 1.

**Triggered attention** A set of solutions [20, 21] have been proposed to enable the application of streaming attention-based encoder-decoder models. Triggered attention (TA) [18, 21, 22], as one of the most popular and successful methods, uses the alignment properties of CTC to determine time positions at which an ASR output is generated and then triggers a attention-based decoder to perform attention over past encoder frames up to the trigger frame position plus some look-ahead encoder frames. During training, such trigger positions can be determined using CTC-based forced alignment to train the decoder with limited encoder inputs [21]. However, in this work we did not apply TA for training, since training with limited context did not help to improve ASR results when performing TA decoding. Thus, for simplicity, the decoder is trained in full-sequence mode and TA is only activated at inference time for streaming recognition. At inference time, the decoding procedure follows the frame-synchronous TA decoding algorithm proposed by [18, 22]. TA decoding leverages frame-synchronous CTC beam search to determine frame/time positions at which a new ASR token is detected to trigger the attention decoder. The TA attention decoder receives as an input all the past encoder

frames up to the trigger frame plus some additional look-ahead frames. This generates additional ASR scores that are used for the joint CTC/attention scoring of ASR hypothesis.

### 3. Alignment regularization

Frame-level sequence-to-sequence ASR losses, such as CTC and RNN-T, learn an alignment correspondence between an input sequence, e.g., a sequence of acoustic or visual features, and a typically shorter output sequence, the ASR labels. This is achieved by allowing the system to output blanks at frames at which no ASR label can be detected. However, there are various ways where to output blanks or labels, which each denote a different alignment. CTC and RNN-T solve this problem by optimizing the sum of all possible alignments, which results in emphasizing one alignment, because alignments are mutually exclusive. However, such learned alignments are not deterministic and also depend on factors such as the neural network architecture or the kind of input features. This fact can be problematic when combining different modalities from separate encoder streams, because responses may not be aligned, even though audio-visual streams are synchronous at the input, which hinders the model to make joint decisions at frame-level. Sanabria *et al.* [23] provide an analysis of the response offset for each modality to demonstrate that their proposed audio-visual ASR systems reduce the response offset between audio and visual encoder state sequences.

In this work, we propose to use the alignments derived from trained CTC outputs and auxiliary CE losses to synchronize the audio and visual streams, where the best alignment path is determined using the Viterbi algorithm [21]. It is worth pointing out that our approach allows alignment information to be extracted on-the-fly without generating alignments in advance from a hybrid ASR system. Thus, the proposed alignment regularization technique aims to temporally synchronize the emission of label information by the audio and visual encoder neural networks, which in turn helps to increase robustness of the AV-ASR system. We train our model using the following loss

$$\mathcal{L} = \alpha \mathcal{L}_{\text{CTC}} + (1 - \alpha) \mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{ALIGN}} \quad (1)$$

where  $\alpha$  denotes the relative weight to control the contribution between the CTC loss  $\mathcal{L}_{\text{CTC}}$  and the attention decoder loss  $\mathcal{L}_{\text{CE}}$ .  $\mathcal{L}_{\text{ALIGN}}$  represents an auxiliary loss for the alignment regularization, which is computed as follows:

$$\mathcal{L}_{\text{ALIGN}} = \beta_{\text{AV2V}} \log p_{\text{CE}}^{\text{AV2V}}(\mathbf{z}^{\text{AV}} | \mathbf{h}^{\text{V}}) + \beta_{\text{AV2A}} \log p_{\text{CE}}^{\text{AV2A}}(\mathbf{z}^{\text{AV}} | \mathbf{h}^{\text{A}}) \quad (2)$$

In Eq. 2,  $\mathbf{h}^{\text{A}} = (h_1^{\text{A}}, \dots, h_T^{\text{A}})$  and  $\mathbf{h}^{\text{V}} = (h_1^{\text{V}}, \dots, h_T^{\text{V}})$  denote the ASR and VSR encoder output sequences, respectively, with  $T$  being the number of encoder frames. The frame-level audio and visual encoder representations are denoted as  $h^{\text{A}}$  (audio) and  $h^{\text{V}}$  (visual).  $\mathbf{z}^{\text{AV}} = (z_1^{\text{AV}}, \dots, z_T^{\text{AV}})$  denotes the best alignment path that is derived by CTC-based forced alignment using the output of the CTC projection layer for the joint audio-visual encoder and the ground truth label sequence, where each  $z_t^{\text{AV}}$  represent a label such as a blank or non-blank symbol with  $t$  indicating the frame position. In particular, the position of output labels in this path corresponds to the frames at which they are most likely emitted. The probabilities  $p_{\text{CE}}^{\text{AV2V}}$  and  $p_{\text{CE}}^{\text{AV2A}}$  are derived from the ASR and VSR encoder state sequences by re-using their CTC projection layers from the pre-training stage (i.e., the stage where we train the audio and visual encoders)

but here a CE loss is used instead of CTC, because the target alignment  $\mathbf{z}^{AV}$  is known.  $\beta_{AV2A}$  and  $\beta_{AV2V}$  do control the relative weights for each CE loss term. These auxiliary CE losses are illustrated in Fig. 1. We should note that the alignments can also be extracted from either the ASR or VSR encoder outputs but instead, we demonstrate that the alignments from joint audio-visual output is optimal (see Section 5).

## 4. Experimental Setup

**Datasets** In this work, we use LRS3 [24], which is the largest public benchmarks for audio-visual speech recognition. The dataset is collected from more than 5000 speakers in TED and TEDx talks and contains 151 189 utterances with a total of 438.9 hours. LRS3 is divided into 408 hours of pre-training data, 30 hours of training-validation data, and 0.9 hours of test data.

**Pre-processing** For the visual stream, we follow the pre-processing pipeline from [25] to extract the mouth region of interests (ROIs) with a size of  $96 \times 96$ . Each frame is normalized by subtracting the mean and dividing by the standard deviation of the training set. Audio stream has been normalized by  $z$ -normalization.

**Training details** Similar to [3], the audio and visual encoder neural networks of the proposed AV-ASR system are using 12 conformer blocks each with an attention dimension of 256, a head size of 8 for the audio stream and 4 for the visual stream, feed-forward dimensions of 2048 and a convolution kernel sizes of 31. The number of parameters in the audio front-end, visual front-end and conformer back-end are 3.9 M, 11.2 M and 31.8 M, respectively. Visual streams are augmented with horizontal flipping, random cropping [3], and adaptive time masking [25] while audio streams are augmented with the adaptive time masking. For adaptive time masking, we use one mask per second and for each mask, we randomly mask a duration up to 0.4 second. The audio-only and visual-only models are trained for 75 epochs with AdamW optimizer [26] with a dynamic batching strategy, where the maximum number of frames in a batch is set to 2400. The learning rate increases linearly with the first 25 000 steps, yielding a peak learning rate of 0.0006 and thereafter decreases with a cosine annealing strategy. We use SentencePiece [27] sub-word units with a vocabulary size of 5000.

**Algorithmic latency** The individual encoders of the proposed streaming AV-ASR system are using a total algorithmic delay of 515 ms (audio encoder) and 580 ms (visual encoder), whereby the larger delay is decisive. This delay is composed of the ResNet front-end delay, which amounts to 35 ms (audio) and 100 ms (visual), and the delay caused by the chunk-wise self-attention of the encoder, where we use a total of 12 look-ahead frames which amounts to 480 ms (40 ms frame rate). The triggered attention decoder also uses 12 look-ahead frames which brings to total algorithmic delay to 580 ms (encoder) + 480 ms (decoder) = 1060 ms.

## 5. Results

**Non-streaming and streaming models** Results for our streaming and non-streaming ASR, VSR and AV-ASR models are shown in Table 1. Note that our models are trained using only the LRS3 dataset, with a total of 438 hours. Our non-streaming audio-only and audio-visual models lead to state-of-the-art performance with a WER of 2.2% and 2.0% WER, respectively. Our non-streaming visual-only model achieves

Table 1: Comparison with state-of-the-art audio-only, visual-only and audio-visual models trained only on the LRS3 dataset. ‘CM-seq2seq’ denotes the conformer-based architecture used in [3].

Method	Type	Auxiliary Losses	WER (%)
<i>Non-Streaming Models</i>			
CM-seq2seq [3]	V	-	46.9
AVHuBERT [28]		-	41.6
CM-seq2seq [25]		-	37.9
Ours		-	<b>38.6</b>
CM-seq2seq [3]	A	-	2.3
Ours		-	<b>2.2</b>
CM-seq2seq [3]	A+V	-	2.3
Ours		-	<b>2.2</b>
Ours		✓	<b>2.0</b>
<i>Streaming Models</i>			
Ours	V	-	<b>50.6</b>
	A	-	<b>2.5</b>
	A+V	-	<b>2.6</b>
		✓	<b>2.6</b>

Table 2: WER [%] of our models on the LRS3 dataset.

Type	Auxiliary Losses	SNR Levels [dB]				
		12.5	7.5	2.5	-2.5	-7.5
<i>Non-Streaming Models</i>						
A	-	3.1	5.1	12.1	38.6	89.9
A+V	-	2.9	3.9	7.5	18.8	52.9
	✓	2.7	3.7	7.8	18.4	43.9
<i>Streaming Models</i>						
A	-	4.2	8.4	23.6	64.9	97.9
A+V	-	4.0	7.5	16.0	38.0	74.0
	✓	3.8	6.8	14.9	32.4	57.5

slightly lower performance than the existing state-of-the-art model, which uses additional auxiliary losses during training [25]. Our streaming audio-only model leads to a small increase in the WER of 0.3% compared to the non-streaming model (2.3% WER). On the other hand, the streaming visual-only model results in a large increase in the WER of 12%. Presumably, the visual-only model is relying on more future context, which is available in the non-streaming case, to resolve visual ambiguities. Finally, the streaming audio-visual model achieves a WER of 2.6% which is 0.6% higher than the non-streaming model. This demonstrates the effectiveness of the proposed model. We should also note that the streaming audio-visual model is slightly worse than the streaming audio-only model mainly due to degraded performance of the streaming visual-only model.

**Noise experiments** In order to test the robustness of the AV-ASR models, we add pink noise from the Speech Commands dataset [29] to the test set at different signal-to-noise ratio (SNR) levels. For training, babble noise from NOISEX [30] is added to the clean speech with a random SNR level selected from the set [-5 dB, 0 dB, 5 dB, 10 dB, 15 dB, 20 dB,  $\infty$  dB] with a uniform distribution. The results shown in Table 2 demonstrate that overall the performance gap between audio-

Table 3: WER [%] of our AV-ASR models as a function of noise levels on the LRS3 dataset. We use the alignment information from the **target modality (T. M.)** and force the response of **predictive modality (P. M.)** to be closer to it. “A Frozen” and “V Frozen” denote that the weights of ASR encoder and VSR encoder loaded from the pre-training stage and are kept frozen during fine-tuning the AV-ASR system. “A Fine-Tuned” and “V Fine-Tuned” denote that the weights of the ASR encoder and VSR encoder are loaded from the pre-training stage and are fine-tuned when training the AV-ASR system.

T. M. → P. M.	Auxiliary Losses	SNR Levels [dB]				
		12.5	7.5	2.5	-2.5	-7.5
<i>A Frozen, V Frozen</i>						
-	-	4.4	8.5	18.2	38.2	59.8
<i>A Frozen, V Fine-Tuned</i>						
-	-	4.5	8.3	19.0	43.5	80.5
A → V	✓	4.6	8.3	17.8	35.8	56.7
<i>A Fine-Tuned, V Frozen</i>						
-	-	4.8	8.1	17.3	39.1	69.1
V → A	✓	4.4	7.7	16.6	34.8	56.1
<i>A Fine-Tuned, V Fine-Tuned</i>						
-	-	4.0	7.5	16.0	38.0	74.0
AV → {A, V}	✓	3.8	6.8	14.9	32.4	57.5

only and audio-visual models is increasing as the level of noise increases. For example, at the SNR level of -7.5 dB the streaming audio-visual model results in an absolute improvement of 23.9 % WER compared with its audio-only counterpart. When including alignment regularization by using the auxiliary CE losses, a further reduction of 16.5 % WER is achieved. A similar improvement can also be observed for the non-streaming models. A reduction of 9.0 % in WER is achieved at the SNR level of -7.5 dB for audio-visual non-streaming models when alignment regularization is applied.

**Using different modalities for alignment** In this section, the impact of using different target alignments from different modalities is studied for alignment regularization, i.e., target alignments generated by the audio or visual encoder are compared against using target alignments from the joint audio-visual responses (default case). Results presented in Table 3 show that independent of the modality for the target alignment, the overall performance is improved when alignment regularization is applied. For example, when using the alignments from the visual encoder (targeted modality) to supervise the audio counterpart (predictive modality), i.e., V → A in Table 3, the performance of the AV-ASR model (the fifth row in Table 3) outperforms the model when using the alignments from the audio encoder to supervise the visual counterpart, i.e., A → V (the third row in Table 3). It is likely that for A → V, i.e., forcing the visual stream to temporally align with the audio stream, the output latency is reduced, because the audio encoder tends to respond earlier than the visual encoder (see Table 4). However, the best WER are obtained when using target alignments generated from the joint AV output.

**Response offset analysis** In order to better understand how the proposed auxiliary tasks contribute to learning better alignments between audio and visual streams, we analyze the response offset between the joint audio-visual representations and either the audio or visual encoder outputs for the proposed

Table 4: The mean response offset over the test set of LRS3 using a streaming AVSR model. “V → AV” with a positive N denotes the visual encoder outputs are relatively slower to the the joint audio-visual representations with N frames, and vice versa.

Auxiliary Losses	Relative Offset	SNR Levels [dB]				
		12.5	7.5	2.5	-2.5	-7.5
<i>A Frozen, V Frozen</i>						
-	V → AV	1.11	1.02	0.86	0.57	0.20
-	A → AV	0.01	0.02	0.02	0.22	5.68
<i>A Fine-Tuned, V Fine-Tuned</i>						
-	V → AV	0.22	0.18	0.11	0.07	0.31
-	A → AV	0.00	0.01	0.04	0.23	-3.34
<i>A Fine-Tuned, V Fine-Tuned</i>						
✓	V → AV	0.14	0.10	0.04	0.02	0.03
✓	A → AV	0.00	0.01	0.04	0.20	1.60

streaming AV-ASR models. In particular, we calculate the average frame distance between the forced alignment sequences  $z^{AV}$  and either  $z^A$  or  $z^V$  for the ASR labels over the entire test set of LRS3. Results are shown in Table 4. It can be observed that without alignment regularization, the joint AV response tends to be better aligned with the audio modality at high SNR levels and only for low SNR levels with the visual modality. This indicates that the modality which is better aligned with the joint audio-visual representations is more likely responsible for the final predictions. In other words, the modality with larger delay relative to the joint audio-visual representations is more likely to be ignored. This property is consistent with the results reported in Table 2, i.e., the performance of streaming audio-visual models outperforms the streaming audio-only counterpart especially at high SNR levels. Table 4 also demonstrates that fine-tuning the audio and visual encoders leads to better alignments between both modalities and better ASR results as shown in Table 3. However, when using alignment regularization in training, we observe that the maximum response offset between the VSR encoder and the joint audio-visual representations is less than 0.14 frames (5.6 ms) across all SNR levels. This indicates that a better synchronization between ASR and VSR encoder state sequences leads to a more robust performance (the last row in Table 3), especially when the audio stream is heavily corrupted by noise.

## 6. Conclusion

In this work, we presented an end-to-end streaming AV-ASR system based on the conformer architecture. We showed that visual-only models are performing substantially worse as compared to the audio-only counterpart when comparing their streaming and non-streaming results, probably due to the assumption that visual streams relies more on the future context. Furthermore, we showed that by adding alignment regularization, the performance of streaming and non-streaming AV-ASR models is much improved with an increasing impact for higher noise levels. The response time between the audio and visual modalities is analyzed and we demonstrated that: 1) The modality with smaller latency relative to the joint audio-visual representations is more responsible to the final predictions. 2) The proposed alignment regularization technique results in a better response synchronization between the audio and visual encoder streams leading to improved joint audio-visual predictions and an increased robustness.

## 7. References

- [1] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, pp. 1–1, 2018.
- [2] S. Petridis, T. Stafylakis, P. Ma, G. Tzimiropoulos, and M. Pantic, "Audio-visual speech recognition with a hybrid ctc/attention architecture," in *Proc. SLT*, 2018, pp. 513–520.
- [3] P. Ma, S. Petridis, and M. Pantic, "End-to-end audio-visual speech recognition with conformers," in *Proc. ICASSP*, 2021, pp. 7613–7617.
- [4] D. Serdyuk, O. Braga, and O. Siohan, "Audio-visual speech recognition is worth 32x32x8 voxels," in *Proc. ASRU*, 2021, pp. 796–802.
- [5] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [6] A. Vaswani *et al.*, "Attention is all you need," in *Proc. NIPS*, 2017, pp. 5998–6008.
- [7] G. Sterpu and N. Harte, "Taris: An online speech recognition framework with sequence to sequence neural networks for both audio-only and audio-visual speech," *Computer Speech & Language*, vol. 74, p. 101349, 2022.
- [8] W. Yu, S. Zeiler, and D. Kolossa, "Fusing information streams in end-to-end audio-visual speech recognition," in *Proc. ICASSP*, 2021, pp. 3430–3434.
- [9] G. Sterpu, C. Saam, and N. Harte, "Attention-based audio-visual fusion for robust automatic speech recognition," in *Proc. ICMI*, 2018, pp. 111–115.
- [10] G. Sterpu, C. Saam, and N. Harte, "How to teach dnns to pay attention to the visual modality in speech recognition," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1052–1064, 2020.
- [11] D. Serdyuk, O. Braga, and O. Siohan, "Transformer-based video front-ends for audio-visual speech recognition for single and multi-person video," in *Proc. Interspeech*, 2022, pp. 2833–2837.
- [12] B. Shi, W. Hsu, and A. Mohamed, "Robust self-supervised audio-visual speech recognition," in *Proc. Interspeech*, 2022, pp. 2118–2122.
- [13] T. Makino *et al.*, "Recurrent neural network transducer for audio-visual speech recognition," in *Proc. ASRU*, 2019, pp. 905–912.
- [14] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, vol. 148, Jun. 2006, pp. 369–376.
- [15] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv:1211.3711*, 2012.
- [16] X. Chen, Y. Wu, Z. Wang, S. Liu, and J. Li, "Developing real-time streaming transformer transducer for speech recognition on large-scale dataset," in *Proc. ICASSP*, 2021, pp. 5904–5908.
- [17] N. Moritz, T. Hori, and J. Le Roux, "Dual causal/non-causal self-attention for streaming end-to-end speech recognition," in *Proc. ICASSP*, 2021, pp. 1822–1826.
- [18] N. Moritz, T. Hori, and J. L. Roux, "Streaming automatic speech recognition with the transformer model," in *Proc. ICASSP*, 2020, pp. 6074–6078.
- [19] S. Petridis *et al.*, "End-to-end audiovisual speech recognition," in *Proc. ICASSP*, 2018, pp. 6548–6552.
- [20] C. Chiu and C. Raffel, "Monotonic chunkwise attention," in *Proc. ICLR*, 2018.
- [21] N. Moritz, T. Hori, and J. L. Roux, "Triggered attention for end-to-end speech recognition," in *Proc. ICASSP*, 2019, pp. 5666–5670.
- [22] N. Moritz, T. Hori, and J. Le Roux, "Streaming end-to-end speech recognition with joint ctc-attention based models," in *Proc. ASRU*, 2019, pp. 936–943.
- [23] R. Sanabria, F. Metze, and F. De La Torre, "Robust end-to-end deep audiovisual speech recognition," *arXiv preprint arXiv:1611.06986*, 2016.
- [24] T. Afouras, J. S. Chung, and A. Zisserman, "LRS3-TED: a large-scale dataset for visual speech recognition," *CoRR*, vol. abs/1809.00496, 2018.
- [25] P. Ma, S. Petridis, and M. Pantic, "Visual speech recognition for multiple languages in the wild," *Nature Machine Intelligence*, pp. 1–10, 2022.
- [26] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. ICLR*, 2019.
- [27] T. Kudo, "Subword regularization: Improving neural network translation models with multiple subword candidates," in *Proc. ACL*, 2018, pp. 66–75.
- [28] B. Shi, W. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," in *Proc. ICLR*, 2022.
- [29] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *CoRR*, vol. abs/1804.03209, 2018.
- [30] A. Varga and H. J. M. Steeneken, "Assessment for automatic speech recognition: II. NOISEX-92: A database and an experiment to study the effect of additive noise on speech recognition systems," *Speech Communication*, vol. 12, no. 3, pp. 247–251, 1993.