

Leveraging Modality-Specific Representations for Audio-Visual Speech Recognition via Reinforcement Learning

Chen Chen¹, Yuchen Hu¹, Qiang Zhang^{2,3}, Heqing Zou¹, Beier Zhu¹, and Eng Siong Chng¹

¹School of Computer Science and Engineering, Nanyang Technological University

²ZJU-Hangzhou Global Scientific and Technological Innovation Center

³College of Computer Science and Technology, Zhejiang University

Abstract

Audio-visual speech recognition (AVSR) has gained remarkable success for ameliorating the noise-robustness of speech recognition. Mainstream methods focus on fusing audio and visual inputs to obtain modality-invariant representations. However, such representations are prone to over-reliance on audio modality as it is much easier to recognize than video modality in clean conditions. As a result, the AVSR model underestimates the importance of visual stream in face of noise corruption. To this end, we leverage visual modality-specific representations to provide stable complementary information for the AVSR task. Specifically, we propose a reinforcement learning (RL) based framework called MSRL, where the agent dynamically harmonizes modality-invariant and modality-specific representations in the auto-regressive decoding process. We customize a reward function directly related to task-specific metrics (*i.e.*, word error rate), which encourages the MSRL to effectively explore the optimal integration strategy. Experimental results on the LRS3 dataset show that the proposed method achieves state-of-the-art in both clean and various noisy conditions. Furthermore, we demonstrate the better generality of MSRL system than other baselines when test set contains unseen noises.

Introduction

Background noise is inevitable in real world that can dramatically degrade the speech quality and intelligibility, thereby increasing the difficulty of speech recognition task (Hu et al. 2022a,b). In noisy scenarios, human will unconsciously observe the mouth region of speakers, as such noise-invariant visual cues can provide useful information for the corrupted speech understanding (Ma, Petridis, and Pantic 2021).

Similar to this, the audio-visual speech recognition (AVSR) technique couples the audio and visual modalities, which has attracted increasing research interest for several years (Noda et al. 2015). Recent machine learning based AVSR methods successfully demonstrate that deep neural network (DNN) can process and fuse raw acoustic and visual inputs to improve the noise-robustness for recognition through a supervised learning paradigm (Petridis et al. 2018; Zhou et al. 2019). Additionally, self-supervised representation learning has been explored to capture the correlations

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Research problem. “SNR” denotes the signal-to-noise ratio, and α is the threshold that modality-invariant representations lose the effectiveness compared with visual modality-specific representations.

between audio and visual lip movements for AVSR task, which has brought remarkable performance gain in terms of word error rate (WER) metric (Shi et al. 2022).

Mainstream AVSR methods focus on learning modality-invariant representations by fusing audio and visual modalities into a common subspace (Song, Sun, and Li 2022). However, such a fusion manner is prone to over-reliance on the audio modality, as it is much easier to recognize than video stream in clean conditions (Mittal et al. 2020). With the rise of noise levels, the importance of the video stream is increasingly underestimated in AVSR systems, and leads to sub-optimal performance since audio modality has already been corrupted by noise. We further illustrate this research problem in Fig 1. Though modality-invariant representations (green line) outperform audio modality-specific representations (yellow line) by a large margin, it is still vulnerable to low signal-to-noise ratio (SNR) conditions. It is worth noting that when SNR is lower than α , the modality-invariant representations even perform worse than visual modality-specific representations (blue line) which are completely unaffected by noise. We argue that this problematic situation can be avoided by reasonable coordination of modality-invariant and modality-specific representations, which is shown as the oracle system (red dashed line).

Although the significance of modality-specific representations has been emphasized in other multi-modal tasks, such as emotion recognition (Hazarika, Zimmermann, and Poria 2020), it still remains challenging to integrate them into AVSR system for several reasons. Firstly, the real-world noises have dynamic and non-stationary temporal distributions, which confuse the recognizer to estimate the importance of visual modality-specific representations during auto-regressive decoding. Secondly, due to the natural distinction of input modalities, a uniform training schedule probably results in the vanishing gradient when we add a further sub-network to extract visual modality-specific representations (Yao and Mihalcea 2022). Finally, with parameter growth of neural networks, the existing integration strategies for new representations are prone to over-fit to specific types of noise distribution (Fu, Wu, and Boulet 2022), thereby failing to adapt to unseen noises in the wild.

In this work, we aim to improve the AVSR system by leveraging visual modality-specific representations that carry noise-invariant information from the visual stream. To this end, we employ a pre-trained vision model that takes lip movement information as input and generates independent probability distribution for sequence generation. This idea is inspired by the language model rescoring that has been widely applied in popular ASR methods (Song et al. 2022; Xu et al. 2022; Chen et al. 2022a). However, instead of applying a typical integration approach, *e.g.*, late fusion (Inaguma et al. 2019), we propose a reinforcement learning (RL) based method to dynamically harmonize the integration process in terms of the task-specific metric. RL is appropriate to play this integration role for: 1) The auto-regressive decoding of AVSR can be modeled as an RL formulation (Bahdanau et al. 2016), where the agent can consider multiple information for reasonable token prediction. 2) The customized reward function of RL bridges the training criterion and WER, thus encouraging it to effectively improve the model performance. 3) The beam search of inference step can provide a set of hypotheses for RL sampling, which allows the agent to explore the optimal policy on token level (Chen et al. 2022c).

The main contributions of MSRL are:

- We propose MSRL—a novel AVSR system that utilizes visual modality-specific representations to dynamically remedy the noise-corrupted audio modality.
- MSRL adopts an RL-based integration method, where a new reward function is designed to encourage the agent to efficiently explore the optimal strategy in terms of the task-specific metric.
- Experimental results on the largest public LRS3 dataset show that MSRL is effective and achieves state-of-the-art performance in both clean and various noisy conditions. Furthermore, the comparative experiments on unseen noises demonstrate that MSRL has better generalization and adaptability than a strong baseline.

Related Work

Audio-visual speech recognition. Recently, AVSR has been attracting increasing research interest as it provides a po-

tential solution for noise-robust speech recognition. To process and fuse audio-video modalities, TM-seq2seq (Afouras, Chung, and Zisserman 2018a) applies a separated Transformer encoder for two modalities and fuses them before decoding. (Ma, Petridis, and Pantic 2021) presents a hybrid CTC/Attention model based on Resnet-18 and Conformer (Gulati et al. 2020), which can be trained in an end-to-end manner. (Tao and Busso 2018) demonstrate the importance of aligning two modalities before fusing them. Moreover, the AV-HuBERT (Shi et al. 2022) learn the correspondence of audio and video modalities in a self-supervised manner, which is further augmented in (Shi, Hsu, and Mohamed 2022) to improve noise-robustness.

Modality-invariant and modality-specific representations. Despite the advanced fusion techniques in multi-modal tasks (Chen et al. 2022b), prior works suggest that the model can benefit from modality-specific representations which capture some desirable properties (Xiong et al. 2020). Nevertheless, how to effectively utilize it is still an open question to be explored. MISA (Hazarika, Zimmermann, and Poria 2020) maps the multi-modal inputs into two subspaces for modality-invariant and modality-specific representations and then fuses them for final classification. MuSE (Yao and Mihalcea 2022) employs separated encoders for multiple modalities, then harmonize them using different learning rates and late-fusion. Similarly, (Feng, Lai, and Xie 2019) constructs an individual network for each modality, as well as designing a modality-shared identity loss to facilitate the extraction of modality-invariant representation. The integration of modality-specific representations is particularly difficult in an AVSR system because the single decoder is hard to dynamically weight the representations in a sequential decision process.

Reinforcement learning in sequence generation. Extensive existing works have demonstrated that RL is suitable to play an optimizing role in sequence generation tasks. In captioning tasks like image captioning (Rennie et al. 2017) and audio captioning (Mei et al. 2021), a self-critical training approach based on RL can optimize the trained model in terms of non-differentiable metrics. Such an idea is also expanded to ASR tasks with a customized reward function (Tjandra, Sakti, and Nakamura 2019; Chen et al. 2022c). Additionally, actor-critic based RL optimization algorithms are designed to improve task-specific score (*e.g.* BLEU) in machine translation task (Williams 1992; Bahdanau et al. 2016). Compared with previous work, the proposed MSRL commits to the stability of RL training, where 1) we utilize pre-trained models to provide learned representations as state space, and 2) we design a reward function to encourage the policy network to explore in trust region (Schulman et al. 2015).

Methodology

We first introduce the main structure of the proposed MSRL system. Then we model the auto-regressive decoding of AVSR as an RL formulation, and also give mathematical derivation for the optimization. Finally, the training schedule of MSRL is illustrated by algorithm pseudo.

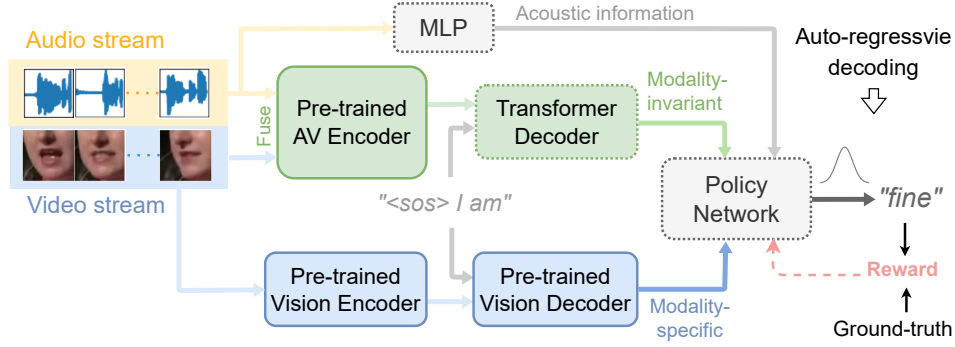


Figure 2: The block diagram of the proposed MSRL system. The solid box denotes such module is fixed during training, while the dashed box denotes it is trainable. The red dashed arrow denotes the process of back-propagation. Policy network considers multiple information in auto-regressive decoding and predicts the current token “fine”.

Main Structure

Given the acoustic utterance A and its paired l -frame video stream $V = (v_1, v_2, \dots, v_l)$, the neural network of AVSR intends to predict a hypothesis sequence $Y = (y_1, y_2, \dots, y_T)$. As shown in Fig. 2, the audio and video streams are fed into a pre-trained AV encoder to extract hidden representation, where a ResNet block (He et al. 2016) and a linear layer are respectively served as front-end to obtain the audio and visual features. After concatenating them, a Transformer encoder (Vaswani et al. 2017) with self-attention mechanism is employed to extract hidden representations. Subsequently, we utilize a learnable Transformer decoder including cross-attention mechanism to produce modality-invariant representations F_i corresponding to the probability distribution of predicted tokens. Meanwhile, we further utilize a pre-trained vision model which similarly consists of ResNet front-end, Transformer encoder, and Transformer decoder. It consumes the video stream as input and independently produces visual modality-specific representations F_v with the same shape as F_i . To harmonize the modality-specific and modality-invariant representations, a linear layer-based policy network are designed in the auto-regressive decoding process. Besides the F_i and F_v themselves, we argue that policy networks should also be aware of audio quality which is useful to estimate the importance of representations. To this end, an MLP block with 2 linear layers is used for downsampling and provides acoustic information I_a for policy network. Finally, a combined distribution is generated to predict the current token (“fine” in Fig. 2).

It is noted that all pre-trained models are fixed in the whole training process. The pre-trained AV encoder is initialized from AV-HuBERT (Shi, Hsu, and Mohamed 2022), which captures cross-modal correlations between audio and video modalities by a self-supervised approach. The vision model (Shi et al. 2022) including encoder and decoder is pre-trained via lip-reading task on LRS3 dataset, where only the video stream is required to generate target sequence.

RL Policy

Basic Reinforcement Learning is typically formulated as a Markov Decision Process (MDP) that includes a tuple of

trajectories $\langle S, \mathcal{A}, \mathcal{R}, \mathcal{T} \rangle$. For each time step t , the agent consider state $s_t \in S$ to generate an action $a_t \in \mathcal{A}$ which interacts with environment. The transition dynamics $\mathcal{T}(s_{t+1}|s_t, a_t)$ is defined as transition probability from current state s_t to next state s_{t+1} , and gain an instant reward $r_t(s_t, a_t)$. The objective of RL is to learn optimal policy to maximize the cumulative reward \mathcal{R} :

$$\mathcal{R} = \max_{a_t \in \mathcal{A}} \sum_{t=1}^T r_t \quad (1)$$

In AVSR task, we summarize the MDP tuple as:

- State \mathcal{S} should contain the comprehensive information or learned patterns for decision-making. Therefore, we denoted \mathcal{S} as a combination of F_i , F_v , and I_a defined in Section , as they are related to predict current token.
- Action \mathcal{A} aims to interact with the environment and update the \mathcal{S} . In this work, \mathcal{A} is defined as a probability distribution P_a for the current predicted token.
- Reward \mathcal{R} is an instant feedback signal to evaluate the performance of \mathcal{A} . we define a token-level reward function for each hypothesis Y as follows:

$$\mathcal{R}(Y, Y^*) = -D_{ED}(Y||Y^*) - \lambda_1 \sum_{t=0}^T D_{KL}(P_a^t||F_i^t) - \lambda_2 \sum_{t=0}^T D_{KL}(P_a^t||F_v^t) \quad (2)$$

where $D_{ED}(\cdot||\cdot)$ denotes the edit distance between two sequence, and Y^* denotes the ground-truth sequence. It is noted that such distance is directly related to WER. The $D_{KL}(\cdot||\cdot)$ denotes the KL-divergence between two distributions, which are used to constrain the policy network to explore in trust region (Schulman et al. 2015). λ_1 and λ_2 are the weights to balance them.

- Transition dynamics $\mathcal{T}(s_{t+1}|s_t, a_t)$ denotes that the predicted token a_t will influence next state s_{t+1} , since the decoding of AVSR is auto-regressive generation process.

In order to maximize the cumulative reward \mathcal{R} , the training objective of policy network is defined as:

$$\begin{aligned} \mathcal{L}_\theta(\langle A, V \rangle, Y^*) &= -\mathbb{E}[\mathcal{R}(Y, Y^*)] \\ &= \sum_Y \mathcal{P}(Y | \langle A, V \rangle, \theta) \mathcal{R}(Y, Y^*) \end{aligned} \quad (3)$$

where θ denotes the neural network, $\mathcal{P}(Y | \langle A, V \rangle, \theta)$ is the probability of hypothesis Y determined by input $\langle A, V \rangle$ and θ . The reward function $\mathcal{R}(Y, Y^*)$ is defined in E.q (2).

Since $\sum_Y \mathcal{P}(Y | \langle A, V \rangle, \theta)$ involves a summation over all possible sequences, we employ the REINFORCEMENT algorithm (Williams 1992) to approximate the expected \mathcal{R} and calculate the gradient $\nabla_\theta \mathcal{L}_\theta$:

$$\nabla_\theta \mathcal{L}_\theta = -\mathbb{E}_{Y^n \sim \mathcal{P}(Y^n | \langle A, V \rangle, \theta)} [\mathcal{R}(Y^n, Y^*) \nabla_\theta \log \mathcal{P}(Y^n | \langle A, V \rangle, \theta)] \quad (4)$$

Where Y^n is the sampling hypothesis drawn from the current model distribution. Different from other sampling methods, we directly utilize the beam search algorithm during decoding to select the N -best hypothesis, which indicates the number of sampling hypotheses is equal to the beam size N . Furthermore, we introduce the baseline to normalize the reward as follows:

$$\nabla_\theta \mathcal{L}_\theta = -\frac{1}{N} \sum_{Y^n \in \text{Beam}} \nabla_\theta \log \mathcal{P}(Y^n | \langle A, V \rangle, \theta) [\mathcal{R}(Y^n, Y^*) - \bar{\mathcal{R}}] \quad (5)$$

where $\bar{\mathcal{R}}$ is the baseline defined as the average of reward of all hypotheses in a beam set. Subtracting $\bar{\mathcal{R}}$ does not influence the gradient, but importantly, it can reduce the variance of the gradient estimation, thus stabilizing the training process. To simplify the calculation, we assume that the probability mass is concentrated on the N -best list only. Consequently, the loss function can be approximated as:

$$\mathcal{L} \approx - \sum_{Y^n \in \text{Beam}} \log \hat{\mathcal{P}}(Y^n | \langle A, V \rangle, \theta) [\mathcal{R}(Y^n, Y^*) - \bar{\mathcal{R}}] \quad (6)$$

where $\hat{\mathcal{P}}(Y^n | \langle A, V \rangle, \theta) = \frac{\mathcal{P}(Y^n | \langle A, V \rangle, \theta)}{\sum_{Y^n \in \text{Beam}} \mathcal{P}(Y^n | \langle A, V \rangle, \theta)}$ represents the re-normalized distribution over the N -best hypotheses. Accordingly, in one Beam set, those hypotheses with a higher reward than average are encouraged to be selected by increasing their possibilities. Conversely, the hypothesis that obtains a lower reward will be suppressed. By minimizing the criterion of E.q (6), the current model intends to pursue higher reward by effective exploration in a beam set.

Training Schedule of MSRL

The training process contains two stages (shown in Algorithm 1). We first use typical cross-entropy criterion to train the randomly initialized decoder (shown from step 1 - 4). The best model is selected by a valid set for subsequent sampling. Then the RL optimization is applied to integrate the visual modality-specific representations based on the reward function in step 4. Considering the continuity of utterance, we adopt an online training manner and the gradient is calculated after the completion of beam search. Consequently, to achieve higher reward, the policy network will be updated to the direction which optimizes the posterior metric.

Algorithm 1: Pseudocode for MSRL Training

Input: The paired audio A , video V , and corresponding ground-truth sequence $Y^* = (y_1^*, y_1^*, \dots, y_T^*)$.

1: Initialize the pre-trained parameters for AV encoder θ_{av} and vision model θ_v .

2: Initialize the random parameters for Transformer decoder θ_d , MLP block θ_m , and RL policy network θ_p .

3: **while** TRUE **do**

Freeze the parameters of encoder θ_{av}

Obtain the hidden feature $h_{av} = \theta_{av}(A, V)$

Train the decoder using cross-entropy loss \mathcal{L}_{ce} :

$$\mathcal{L}_{ce} = \sum_{t=1}^T -\log \mathcal{P}_{\theta_d}(y_t^* | y_{t-1}^*, \dots, y_1^*, h_{av}) \quad (7)$$

end while

4: **while** TRUE **do**

for hypothesis Y^n in N -best list **do**

Freeze the encoder θ_{av} and vision model θ_v

for t in $1, 2, \dots, T$ **do**

Obtain representations F_i and F_v :

$F_i = \theta_d(h_{av})$

$F_v = \theta_v(V)$

Compute current action probability:

$P_a^t = \theta_p(F_i^t, F_v^t, \theta_m(A))$

end for

Compute probability $\mathcal{P}(Y^n) = \prod_{t=1}^T P_a^t$

Determine accumulative reward $\mathcal{R}(Y^n, Y^*)$

end for

Train the policy network using E.q (6)

end while

Experiment Setting

Database

We conduct the experiments on LRS3 (Afouras, Chung, and Zisserman 2018b), which is the largest publicly available dataset for audio-visual speech recognition task. It includes face tracks from over 400 hours of TED and TEDx videos from more than 5,000 speakers, along with the corresponding subtitles and word alignment boundaries. The original training set is divided into 2 partitions: pretrain (403 hours) and trainval (30 hours), which are both from the same sources with test set (1452 utterances, 1 hour). In this paper, we randomly choose 1,200 utterances (1 hour) as a valid set for hyper-parameter tuning and best model selection.

For the noisy test set, we follow the categories and mixing strategy from prior work (Shi, Hsu, and Mohamed 2022). The seen noises contains categories of “*babble*”, “*music*” and “*natural*” sampled from MUSAN dataset (Snyder, Chen, and Povey 2015), and “*speech*” noise is sampled from utterances in LRS3. These four categories of noises are seen by both pre-trained models and the training process. For the unseen noises, we select 4 categories of “*Cafe*”, “*Meeting*”, “*River*”, and “*Resto*” from DEMAND noise set (Thiemann, Ito, and Vincent 2013) and mix them with test set. The detailed data pre-processing strategy is illustrated in appendix.

AV Pre-trained Enc. (# Enc.)	Decoder (# Dec.)	V Pre-trained Model (# Enc./ # Dec.)	Data (hour)
Small (12)	Small (6)	Large (24/16)	30
Small (12)	Small (6)	Large (24/16)	433
Large (24)	Large (9)	Large (24/16)	30
Large (24)	Large (9)	Large (24/16)	433

Table 1: Different settings of MSRL. “# Enc.” and “# Dec.” denotes the numbers of encoder and decoder blocks.

MSRL Setup

We develop several MSRL frameworks with different settings, as shown in Table 1. The small transformer block has 768/3072/12 of embedding dimension/feed-forward dimension/attention heads, and the large transformer block increases to 2034/4096/16 respectively. Considering the task difficulty of lip reading, the encoder and decoder of vision model adopt large blocks.

The labeled data is first used for the pre-trained models (Shi et al. 2022), then it is reused for the training of the decoder and RL module. According to labeled data amount, we define it as two modes. The normal-resource contains 433 hours of full training data (pretrain subset and trainval subset), and the low-resource only contains 30 hours of training data (trainval subset).

Result and Analysis

In this section, we conduct extensive experiments and answer the following questions:

- What is the effect of modality-specific representations in AVSR task? We display the experimental results to prove that MSRL addresses the research problem in Fig. 1, and the problematic situation will not happen at low SNR conditions.
- What is the effect of the RL integration? We conduct comparative experiments including other integration strategies to show the superiority of RL method.
- How does the MSRL performance against other competitive methods? We carry out a series of experiments in various conditions to compare our method with previously published works.
- How is the generalization of MSRL to unseen noises? We directly test the MSRL on various conditions with unseen noise to demonstrate its ability of generalization.

Effect of Modality-specific Representations

In this part, we first quantitatively analyze the effect that MSRL utilizes the visual modality representations. To this end, we construct three baseline systems that leverage different representations. **Audio-only (A-only)** baseline only consumes audio modality as input to generate target sequence. **Visual-only (V-only)** baseline is trained as a lip-reading task that consumes visual modality as input. **Modality-invariant (Inv.)** baseline is trained as a vanilla AVSR task that consumes both audio and visual modality as input. Considering the intensity, the *babble* noise is selected to simulate the noisy condition with different SNR levels. The WER

Method	Size	<i>Babble</i> Noise, SNR=						Clean ∞
		-15	-10	-5	0	5	avg	
Normal-resource (433 hrs of labeled data)								
<i>A-only</i>	S	99.1	98.1	82.7	32.6	11.9	64.9	2.53
<i>A-only</i>	L	98.6	97.4	75.8	24.6	9.01	61.1	1.95
<i>V-only</i>	L	26.9						26.9
<i>Inv.</i>	S	55.6	38.0	19.1	7.24	4.02	24.8	1.84
<i>Inv.</i>	L	43.4	30.3	13.5	4.90	2.50	18.9	1.45
MSRL	S	26.1	24.7	14.8	5.92	3.19	14.9	1.44
MSRL	L	25.5	22.3	11.3	4.51	2.31	13.2	1.33
Low-resource (30 hrs of labeled data)								
<i>A-only</i>	S	*	*	84.2	36.1	13.9	66.8	4.69
<i>A-only</i>	L	*	98.0	77.0	25.9	14.2	63.0	3.51
<i>V-only</i>	L	27.8						27.8
<i>Inv.</i>	S	53.0	39.5	21.4	10.2	5.92	26.0	4.10
<i>Inv.</i>	L	44.8	32.3	16.4	7.37	4.87	21.1	3.27
MSRL	S	27.4	25.8	16.7	7.24	5.20	16.5	3.38
MSRL	L	26.5	24.9	13.0	6.36	3.97	14.9	2.82

Table 2: The WER (%) results in *babble* noise and clean conditions. “avg” denotes the average performance across all SNR. “*” denotes the input modality can not be recognized. “S” denotes small size of block, “L” denotes large size.

results of MSRL and three baselines with different Transformer blocks and resource modes are shown in Table 2.

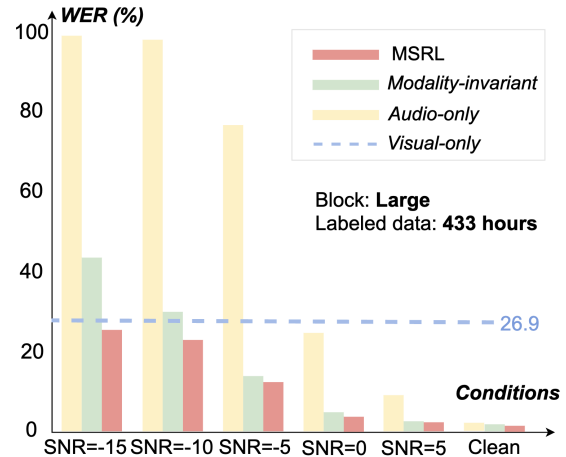


Figure 3: The visualization WER(%) results of *Audio-only* baseline, *Modality-invariant* baseline, and proposed MSRL with large block and 433 hours of labeled data.

We observe that except *Visual-only* baseline, the performance of other methods degrades obviously with the decrease of SNR. When SNR is lower than -5, the performance of *Inv.* baseline is even worse than *V-only* baseline. However, such a problematic situation does not happen in MSRL, since the visual modality-specific representations have been increasingly effective if audio quality becomes hard to recognize. Furthermore, we observe that MSRL system achieves up to 17.6% relative WER reduction than *Inv.* baseline in clean conditions. It is out of intuition because the visual modality-specific representations are usually con-

Method	Babble Noise, SNR=						Clean ∞
	-15	-10	-5	0	5	avg	
Normal-resource & Large Transformer block							
<i>Inv.</i>	43.4	30.3	13.5	4.90	2.50	18.9	1.45
<i>Ef</i>	38.2	25.8	12.6	5.07	2.96	16.9	1.58
<i>Lf</i>	36.7	26.2	12.2	4.78	2.70	16.5	1.68
<i>Me</i>	31.6	23.4	11.8	5.36	3.15	15.1	2.26
MSRL	25.5	22.3	11.3	4.51	2.31	13.2	1.33
Low-resource & Large Transformer block							
<i>Inv.</i>	44.8	32.3	16.4	7.37	4.87	21.1	3.27
<i>Ef</i>	40.1	25.6	15.5	7.40	5.01	18.7	3.26
<i>Lf</i>	38.4	25.9	13.3	6.40	4.19	17.6	3.31
<i>Me</i>	33.7	25.4	13.6	6.77	4.18	16.7	3.35
MSRL	26.5	24.9	13.0	6.36	3.97	14.9	2.82

Table 3: The WER (%) results of MSRL and other integration methods in *babble* noise and clean conditions. “*Inv.*” serves as the baseline without any fusion manner. Best results are in bold.

sidered trivial when audio quality is high. We reason that 1) visual modality-specific representations add the diversity of information, and it might be helpful when some ambiguous acoustic pronunciations have similar probabilities. 2) The training objective in E.q.(6) is related to WER, thus ameliorating the mismatch problem between training and testing modes. In general, the proposed MSRL system not only guarantees the lower-bound performance in noisy conditions but also improves the upper-bound performance in clean conditions.

In order to visualize the effect of modality-specific representations, we draw the histogram of *A-only* baseline, *Modality-invariant* baseline, and MSRL system in Fig. 3. The *V-only* baseline is shown as the blue dashed line as the WER keeps invariant (26.9%) in all conditions. It is noted that Fig. 3 roughly reproduces the situation in Fig. 1. The *Inv.* baseline loses its effectiveness in low SNR setting, while MSRL system performs similarly to the oracle line in all conditions.

We also conduct a case study to observe how the visual modality-specific representations help the MSRL system. To this end, we sample a divergent step in decoding, where the *Inv.* baseline predicts a wrong token but MSRL predict the correct one. As shown in Fig 4, three probability distributions are drawn from *Inv.* baseline, *V-only* baseline, and MSRL system. The x-axis is the vocabulary size and each value denotes a BPE (Sennrich, Haddow, and Birch 2015) token. The y-axis denotes the probability of the corresponding token. For better visualization, the improbable tokens (probability < 0.05) are not included in the figure. It is observed that *Inv.* baseline predict a wrong token (ID=48) in this decoding step, but with help of visual modality-specific representations (*i.e.*, *V-only* baseline), the MSRL predicts the correct token ‘*que*’ (ID=582).




Figure 4: Case study of a divergent decoding step, where the ground-truth token is ‘*que*’ (ID=582). The probabilities higher than 0.05 are displayed.

Effect of RL Module

In this part, we examine the effect of RL module by replacing it using other integration methods, *i.e.*, *Early fusion (Ef)*, *Late fusion (Lf)*, and *Model ensemble (Me)*. Since the pre-trained vision model also has the Transformer-based encoder-decoder pipeline, the *Early fusion* adds the hidden features from the final encoder layer of pre-trained vision model to the corresponding layer of pre-trained AV encoder. The *Late fusion* (Inaguma et al. 2019) executes a similar operation but add features at the final layer of decoder. Both early and late fusion strategies are applied in the cross-entropy training (step 3 in Algorithms 1), where the decoder is trainable to fit the new features. The *Model ensemble* method directly computes the average of probabilities from *Modality-invariant* baseline and *V-only* baseline for token prediction in auto-regressive decoding, without any tuning operation.

From the WER results of Table 3, in noisy conditions, all integration methods can benefit from visual modality-specific representations compared with *Inv.* baseline. MSRL achieves best performance in all SNR levels. Surprisingly, the untrained *Model ensemble* baseline beats the *Early fusion* and *Late fusion* baselines on average in both normal-resource and low-resource modes. When SNR is -15, except MSRL, three other baselines are not able to avoid the problematic situation that perform worse than *V-only* baseline. In clean conditions, however, the visual modality-specific representations might be redundant. We observe that *Model ensemble* baseline overestimates the importance of visual modality-specific representations, thus suffering 55.9% of performance deterioration (1.45% \rightarrow 2.26%) in normal-resource mode. *Early fusion* and *Late fusion* can dilute it by tunable parameters, thereby obtaining comparable WER results with *Inv.* baseline. In general, MSRL can reasonably balance the importance of modality-specific and modality-invariant representations, as the policy network always considers acoustic information in auto-regressive decoding.

Benchmark against Other Methods

We then report the WER performance of MSRL in various conditions, as well as comparing it with other compet-

Method	Hr	<i>Babble</i> , SNR=					<i>Natural</i> , SNR=					<i>Music</i> , SNR=					<i>Speech</i> , SNR=					Clean ∞
		-10	-5	0	5	avg	-10	-5	0	5	avg	-10	-5	0	5	avg	-10	-5	0	5	avg	
RNN-T	34K	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	4.5
TM-s2s	595	-	-	42.5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	7.2
AE-MSR	1.4K	38.6	31.1	25.5	24.3	29.9	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	6.8
AV-HBT	30	35.1	18.4	8.3	4.9	16.7	11.6	6.5	4.6	4.0	6.7	12.4	7.4	4.7	4.1	7.2	11.5	6.8	5.0	4.2	6.9	3.3
AV-HBT	433	34.9	16.6	5.8	2.6	15.0	9.4	4.3	2.4	1.9	4.5	10.9	4.6	2.6	1.8	5.0	11.4	4.6	2.9	2.2	5.3	1.4
MSRL	30	24.9	13.0	6.4	4.1	12.1	9.8	5.6	3.7	3.4	5.6	10.8	6.5	4.0	3.3	6.2	8.6	5.5	4.0	3.5	5.4	2.8
MSRL	433	22.4	11.3	4.5	2.3	10.1	8.0	4.1	2.3	1.6	4.0	8.9	4.4	2.4	1.7	4.4	7.2	3.4	2.3	1.8	3.7	1.3

Table 4: The WER (%) results of MSRL and prior works on LRS3 dataset. “Hr” denotes the the amount of labeled audio-visual speech data used in each system.

Method	<i>Cafe</i> , SNR=					<i>Meeting</i> , SNR=					<i>River</i> , SNR=					<i>Resto</i> , SNR=					Clean ∞
	-10	-5	0	5	avg	-10	-5	0	5	avg	-10	-5	0	5	avg	-10	-5	0	5	avg	
Low-resource (30 hrs of labeled data) & Large Transformer block																					
AV-HBT	16.4	7.5	4.7	4.0	8.2	13.6	7.3	4.9	4.1	7.5	23.6	11.0	5.9	4.4	11.2	36.8	19.9	8.3	5.1	17.5	3.3
MSRL	13.0	6.1	3.9	3.1	6.5	11.1	6.4	4.4	3.4	6.3	18.5	9.5	5.0	3.7	9.2	24.5	16.3	7.0	4.3	13.0	2.8
Normal-resource (433 hrs of labeled data) & Large Transformer block																					
AV-HBT	13.1	4.8	2.6	1.9	5.6	12.4	5.4	3.0	2.2	5.8	21.0	8.3	3.6	2.4	8.8	35.9	17.4	5.9	2.8	15.5	1.4
MSRL	11.2	4.2	2.3	1.7	4.9	10.4	4.5	2.6	1.8	4.8	17.8	7.8	3.2	1.9	7.7	23.9	13.9	5.1	2.4	11.3	1.3

Table 5: The WER (%) results of MSRL on unseen noises. “Cafe”, “Meeting”, “River”, and “Resto” are drawn from DEMAND.

itive methods. Four recent published methods are selected as strong baselines, which are RNN-T (Makino et al. 2019), TM-seq2seq (abbreviated as TM-s2s) (Afouras, Chung, and Zisserman 2018a), AE-MSR (Xu et al. 2020), and AV-HuBERT (abbreviated as AV-HBT) (Shi et al. 2022). Since RNN-T and TM-seq2seq methods focus on clean conditions, and the AE-MSR is only evaluated on *babble* noise, we only report the available results from their respective papers. For AV-HuBERT, the “*babble*”, “*speech*”, and “clean” columns present the WER results from original paper. The “*natural*” and “*music*” columns were reproduced using the official code as they are not available in original paper. The comparison of WER results is shown in Table 4.

In clean conditions, we observe that MSRL achieves 5% (1.4% \rightarrow 1.33%) relative WER reduction over the best baseline of AV-HuBERT in normal-resource mode. In low-resource mode, such superiority increases to 14.5% (3.3% \rightarrow 2.82%). It indicates that visual information is particularly important when training data is limited. Furthermore, the MSRL using 30 hours of labeled data even performs better than RNN-T employs 34k hours of labeled data, which shows better data efficiency.

In noisy conditions, MSRL achieves the best performances in all kinds of noises and SNR levels. Specifically, for the “*babble*”, “*natural*”, “*music*” and “*speech*” noises, MSRL respectively surpasses AV-HuBERT baseline by 32.5%/27.5%, 11.1%/15%, 12.0%/19.5% and 30.2%/21.7% relatively in normal-resource/low-resource mode. It is worthy noted that the *speech* noise is the utterance drawn from the same source of LRS3 which might confuse the recognizer, while the MSRL can address it well without any separation module.

Generalization on Unseen Noise

Finally, we evaluate the generalization of MSRL method, as the AVSR model usually encounters unseen noises in practical applications. We test the AV-HuBERT and MSRL models on a customized test set which contains 4 types of unseen noises, and the WER results are shown in Table 5.

We observe MSRL system has better generalization on all 4 kinds of noises. The visual modality-specific representations are still effective as they are unaffected by the domain shift of audio modality. Consequently, MSRL respectively surpasses the AV-HuBERT baseline by 20.7%/12.5%, 16.4%/17.2%, 17.9%/12.5% and 25.7%/27.1% relatively in low-resource/normal-resource mode. Furthermore, we notice that models show distinct adaptability to different unseen noises. Since the “*cafe*” and “*meeting*” noises mainly consist of human voice, both AV-HuBERT and MSRL adapt them well and achieve comparable WER results with seen “*speech*” noise. However, the WER performance degrades obviously on strange “*river*” and “*resto*” noises.

Conclusion

In this paper, we propose a reinforcement learning-based method MSRL to leverage the modality-specific representations into AVSR task. MSRL employs a pre-trained vision model to provide the visual modality-specific and a policy network to explore the optimal integrated strategy in autoregressive decoding process. We design the experiments to examine the effects of both visual modality-specific representations and RL integration module. WER results demonstrate that MSRL achieves state-of-the-art performance on LRS3 dataset in clean and noisy conditions, as well as showing better generalization on unseen noises.

Acknowledgements

This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-PhD-2021-01-002). The computational work for this article was (fully/partially) performed on resources of the National Supercomputing Centre, Singapore (<https://www.nsc.sg>)

References

- Afouras, T.; Chung, J. S.; and Zisserman, A. 2018a. The conversation: Deep audio-visual speech enhancement. *arXiv preprint arXiv:1804.04121*.
- Afouras, T.; Chung, J. S.; and Zisserman, A. 2018b. LRS3-TED: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*.
- Bahdanau, D.; Brakel, P.; Xu, K.; Goyal, A.; Lowe, R.; Pineau, J.; Courville, A.; and Bengio, Y. 2016. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086*.
- Chen, C.; Hou, N.; Hu, Y.; Shirol, S.; and Chng, E. S. 2022a. Noise-robust speech recognition with 10 minutes unparallelized in-domain data. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4298–4302. IEEE.
- Chen, C.; Hou, N.; Hu, Y.; Zou, H.; Qi, X.; and Chng, E. S. 2022b. Interactive Audio-text Representation for Automated Audio Captioning with Contrastive Learning. *arXiv preprint arXiv:2203.15526*.
- Chen, C.; Hu, Y.; Hou, N.; Qi, X.; Zou, H.; and Chng, E. S. 2022c. Self-Critical Sequence Training for Automatic Speech Recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3688–3692. IEEE.
- Feng, Z.; Lai, J.; and Xie, X. 2019. Learning modality-specific representations for visible-infrared person re-identification. *IEEE Transactions on Image Processing*, 29: 579–590.
- Fu, Y.; Wu, D.; and Boulet, B. 2022. Reinforcement learning based dynamic model combination for time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6639–6647.
- Gulati, A.; Qin, J.; Chiu, C.-C.; Parmar, N.; Zhang, Y.; Yu, J.; Han, W.; Wang, S.; Zhang, Z.; Wu, Y.; et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Hazarika, D.; Zimmermann, R.; and Poria, S. 2020. Misa: Modality-invariant and-specific representations for multimodal sentiment analysis. In *Proceedings of the 28th ACM international conference on multimedia*, 1122–1131.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, Y.; Hou, N.; Chen, C.; and Chng, E. S. 2022a. Dual-Path Style Learning for End-to-End Noise-Robust Speech Recognition. *arXiv preprint arXiv:2203.14838*.
- Hu, Y.; Hou, N.; Chen, C.; and Chng, E. S. 2022b. Interactive feature fusion for end-to-end noise-robust speech recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6292–6296. IEEE.
- Inaguma, H.; Cho, J.; Baskar, M. K.; Kawahara, T.; and Watanabe, S. 2019. Transfer learning of language-independent end-to-end ASR with language model fusion. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6096–6100. IEEE.
- Ma, P.; Petridis, S.; and Pantic, M. 2021. End-to-end audio-visual speech recognition with conformers. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7613–7617. IEEE.
- Makino, T.; Liao, H.; Assael, Y.; Shillingford, B.; Garcia, B.; Braga, O.; and Siohan, O. 2019. Recurrent neural network transducer for audio-visual speech recognition. In *2019 IEEE automatic speech recognition and understanding workshop (ASRU)*, 905–912. IEEE.
- Mei, X.; Huang, Q.; Liu, X.; Chen, G.; Wu, J.; Wu, Y.; Zhao, J.; Li, S.; Ko, T.; Tang, H. L.; et al. 2021. An encoder-decoder based audio captioning system with transfer and reinforcement learning. *arXiv preprint arXiv:2108.02752*.
- Mittal, T.; Bhattacharya, U.; Chandra, R.; Bera, A.; and Manocha, D. 2020. M3er: Multiplicative multimodal emotion recognition using facial, textual, and speech cues. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 1359–1367.
- Noda, K.; Yamaguchi, Y.; Nakadai, K.; Okuno, H. G.; and Ogata, T. 2015. Audio-visual speech recognition using deep learning. *Applied Intelligence*, 42(4): 722–737.
- Petridis, S.; Stafylakis, T.; Ma, P.; Tzimiropoulos, G.; and Pantic, M. 2018. Audio-visual speech recognition with a hybrid ctc/attention architecture. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, 513–520. IEEE.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7008–7024.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Sennrich, R.; Haddow, B.; and Birch, A. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Shi, B.; Hsu, W.-N.; Lakhotia, K.; and Mohamed, A. 2022. Learning audio-visual speech representation by masked multimodal cluster prediction. *arXiv preprint arXiv:2201.02184*.
- Shi, B.; Hsu, W.-N.; and Mohamed, A. 2022. Robust Self-Supervised Audio-Visual Speech Recognition. *arXiv preprint arXiv:2201.01763*.
- Snyder, D.; Chen, G.; and Povey, D. 2015. Musan: A music, speech, and noise corpus. *arXiv preprint arXiv:1510.08484*.

- Song, Q.; Sun, B.; and Li, S. 2022. Multimodal Sparse Transformer Network for Audio-Visual Speech Recognition. *IEEE Transactions on Neural Networks and Learning Systems*.
- Song, T.; Xu, Q.; Ge, M.; Wang, L.; Shi, H.; Lv, Y.; Lin, Y.; and Dang, J. 2022. Language-specific Characteristic Assistance for Code-switching Speech Recognition. *arXiv preprint arXiv:2206.14580*.
- Tao, F.; and Busso, C. 2018. Aligning audiovisual features for audiovisual speech recognition. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Thiemann, J.; Ito, N.; and Vincent, E. 2013. The diverse environments multi-channel acoustic noise database (demand): A database of multichannel environmental noise recordings. In *Proceedings of Meetings on Acoustics ICA2013*, volume 19, 035081. Acoustical Society of America.
- Tjandra, A.; Sakti, S.; and Nakamura, S. 2019. End-to-end speech recognition sequence training with reinforcement learning. *IEEE Access*, 7: 79758–79769.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3): 229–256.
- Xiong, H.; Ou, W.; Yan, Z.; Gou, J.; Zhou, Q.; and Wang, A. 2020. Modality-specific matrix factorization hashing for cross-modal retrieval. *Journal of Ambient Intelligence and Humanized Computing*, 1–15.
- Xu, B.; Lu, C.; Guo, Y.; and Wang, J. 2020. Discriminative multi-modality speech recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14433–14442.
- Xu, Q.; Song, T.; Wang, L.; Shi, H.; Lin, Y.; Lv, Y.; Ge, M.; Yu, Q.; and Dang, J. 2022. Self-Distillation Based on High-level Information Supervision for Compressing End-to-End ASR Model. In *Proc. Interspeech 2022*, 1716–1720.
- Yao, Y.; and Mihalcea, R. 2022. Modality-specific Learning Rates for Effective Multimodal Additive Late-fusion. In *Findings of the Association for Computational Linguistics: ACL 2022*, 1824–1834.
- Zhou, P.; Yang, W.; Chen, W.; Wang, Y.; and Jia, J. 2019. Modality attention for end-to-end audio-visual speech recognition. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6565–6569. IEEE.