



Improving the Gap in Visual Speech Recognition Between Normal and Silent Speech Based on Metric Learning

Sara Kashiwagi^{1*}, Keitaro Tanaka^{1*}, Qi Feng¹, Shigeo Morishima²

¹Waseda University, Japan

²Waseda Research Institute for Science and Engineering, Japan

sara.kashiwagi@moegi.waseda.jp, phys.keitaro1227@ruri.waseda.jp, fengqi@ruri.waseda.jp, shigeo@waseda.jp

Abstract

This paper presents a novel metric learning approach to address the performance gap between normal and silent speech in visual speech recognition (VSR). The difference in lip movements between the two poses a challenge for existing VSR models, which exhibit degraded accuracy when applied to silent speech. To solve this issue and tackle the scarcity of training data for silent speech, we propose to leverage the shared literal content between normal and silent speech and present a metric learning approach based on visemes. Specifically, we aim to map the input of two speech types close to each other in a latent space if they have similar viseme representations. By minimizing the Kullback-Leibler divergence of the predicted viseme probability distributions between and within the two speech types, our model effectively learns and predicts viseme identities. Our evaluation demonstrates that our method improves the accuracy of silent VSR, even when limited training data is available.

Index Terms: visual speech recognition, silent speech, metric learning

1. Introduction

Visual speech recognition (VSR), also known as lipreading, aims to recognize spoken words based solely on visual interpretation of lip movements [1, 2]. Since VSR is independent of background noise or sound quality, it has attracted significant attention as an alternative to traditional speech recognition, which relies on audio input. VSR also serves as a valuable communication tool for individuals with hearing or speech impairments. Consequently, VSR has already been implemented in practical systems such as the silent speech interface (SSI) [3, 4], which enables users to input text or operate smartphones through non-vocalized (i.e., silent) speech.

One of the major challenges in VSR is the difference in speech format, especially between vocalized (i.e., normal) and silent speech. Studies indicate that in silent speech, individuals tend to move their mouths more extensively than they do in normal speech [5, 6]. Currently, most VSR datasets contain normal speech data, as they are readily obtainable from TV programs [7] or the internet [8]. Consequently, VSR models intended for silent speech are typically trained on normal speech data, resulting in sub-optimal performance [9, 10].

A common approach to address the issue of data variance in VSR is to augment the training data to cover all possible variants by collecting or creating them automatically. This method has proven successful in speaker-independent [11] and pose-invariant lipreading [12, 13]. However, it is challenging to apply this approach to silent speech recognition due to the scarcity

*Equal contribution.

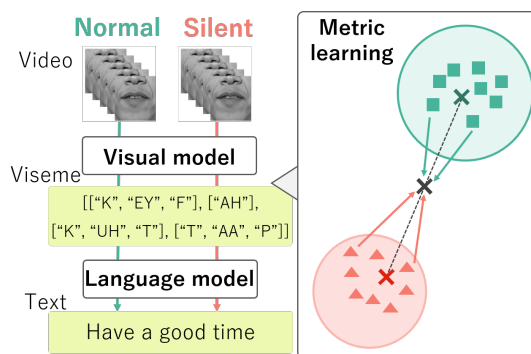


Figure 1: Our method employs metric learning in a latent space to establish relationships between normal and silent speech instances through shared visemes.

of public datasets and the need for manual data collection. In practice, the largest available dataset for normal speech contains over 150,000 samples [8], while the largest dataset for silent speech has only up to 4,600 samples [10].

In this paper, we present a novel metric learning method based on visemes. We observe that normal and silent speech may visually differ, but share the same literal contents (i.e., underlying texts). In the context of VSR for silent speech, it is desirable for a model to recognize input videos from normal and silent speech as similar when they have the same literal representation (or expected outputs). Visemes, the smallest units of visually indistinguishable lip movements [14] that are uniquely derived from phonemes and words, can capture the shared information between normal and silent speech[†]. Thus, we propose to focus on visemes to develop our novel metric learning method.

To facilitate the proposed learning method and address the accuracy imbalance between normal and silent speech, we propose to utilize a two-stage VSR model [16, 17], comprising a video-to-viseme (i.e., visual) model and a viseme-to-text (i.e., language) model. Our method enforces the visual model to map input videos of normal and silent speech close to each other in a latent space, provided they share the same viseme (see Fig. 1). Specifically, each point in the latent space represents a probability distribution as estimated visemes, and we minimize the Kullback-Leibler (KL) divergence between the two probability distributions of the two types of speech. Additionally, our method extends to visemes within the same speech type, but presenting in different positions within the same text. This allows the model to learn viseme identities, as well as to estimate visemes independently.

[†]This is because visemes are uniquely decided based on a group of phonemes, the smallest units of speech sounds [15], which are also uniquely derived from words or sentences.

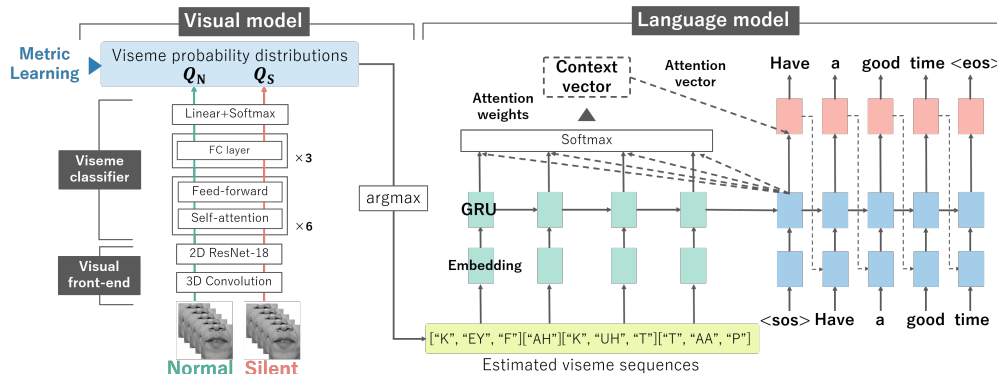


Figure 2: Two-stage VSR model incorporating a visual model and a language model. Our approach incorporates metric learning into the visual model outputs, representing viseme probability distributions of normal and silent speech data inputs.

The main contribution of this study is to propose a novel learning method to leverage already existing yet limited amounts of silent speech data more efficiently. Our method achieves better performance without collecting single additional speech data, even that of normal speech. In addition, we demonstrate that our method’s effectiveness is further enhanced by accounting for data imbalance between the two speech types. Experimental results show that our method improves the performance of the model by nearly 3% from a baseline. Furthermore, the model trained with 400 silent speech data using our method performs comparably to the baseline model trained with twice the amount of data.

2. Related work

This section provides a summary of the existing literature regarding the differences in lip movements between normal and silent speech. Additionally, we explore the advantages of incorporating visemes as an intermediate feature in lipreading models as compared to alternative approaches.

2.1. Normal speech and silent speech

Research examining lip movements during normal and silent speech has utilized facial electromyography (EMG) signals [5] and ultrasound images [6] to establish significant differences between the two types of speech. Specifically, studies have shown that the absence of auditory feedback in silent speech leads speakers to rely more on somatosensory feedback to ensure proper pronunciation, resulting in larger lip movements, particularly when articulating bilabial consonants such as “p”, “b”, and “m”. Notably, two additional studies [9, 10] have reported a decline in performance during silent speech recognition, a challenge that remains unresolved. Petridis et al. [10] observed that a model trained on normal speech achieved 8.5% lower accuracy when tested on silent speech compared to normal speech using their original dataset.

2.2. Viseme-level lipreading

Current research on VSR has primarily focused on developing end-to-end models [2, 7], which necessitate significant amounts of paired video and text data to capture the context of the inputted utterances. Alternatively, lipreading models that generate other tokens, such as sub-words [18], phonemes [19, 20], or visemes [16, 17], have also been investigated. These models employ an external language model that predicts spoken texts based on the output tokens, facilitating the prediction of unseen

words during training. Moreover, several studies [16, 17] have identified a common challenge in both end-to-end and viseme-level two-stage models, namely, differentiating homophones: words that possess unique meanings but similar lip movements. This issue arises due to the fact that the number of visemes is significantly smaller than the number of phonemes. In the latter models, language models can have decoding systems that correspond to homophones, leading to more efficient utilization of visual information from small datasets.

3. Proposed method

We first explain the model architecture and then introduce two kinds of metric learning, inter-speech and within-speech types, applied to the outputs of the visual model (Fig. 2).

3.1. Two-stage VSR model

As a pioneering work for mitigating the imbalance between normal and silent speech, we aim to explore the potential of viseme information. In our work, we develop a two-stage model for VSR that explicitly incorporates this information following Fenghour et al. [16,21,22]. It consists of a visual model that predicts from input videos viseme sequences and a language model that generates from the predicted viseme sequences the output texts. Note that the input video is a sequence of colored images. We convert them to grayscale and crop around their lips using Dlib [23]. It turns into a tensor with the shape of $T \times 112 \times 112$, where T is the number of frames.

The visual model consists of a visual front-end that extracts spatial and temporal features from the input tensor and a viseme classifier that outputs the predicted viseme sequences. In the visual front-end, a 3D convolutional layer captures the changes in lip shapes, and then 2D ResNet-18 follows to output a 512-dimensional feature vector for each video frame. For the viseme classifier, we employ an encoder-decoder architecture. Specifically, we utilize a transformer’s encoder with repeated self-attention and feed-forward networks and three fully connected layer blocks with 1,024, 2,048, and 1,024 nodes, respectively. The last layer of the classifier consists of a linear layer and a softmax function, which yields probability distributions for visemes. Each frame is classified into 17 classes, including 13 viseme classes and four special characters for space, start of sentence (SoS), end of sentence (EoS), and padding.

The language model is based on an attention architecture and two layers of the gated recurrent unit (GRU) with 256 nodes. Specifically, the first layer of the GRUs serves as an encoder, where attention weights are calculated against its out-

put, and the second one does as a decoder. The language model takes as inputs the viseme sequences predicted by the visual model, where each combination of visemes is regarded as a distinct class, and every possible word is considered a target text label. During the training of the language model, the decoder uses the ground truth as input. In contrast, during inference, the decoder uses a context vector calculated from attention weights that access the encoder’s last hidden state to predict the next word based on all previously predicted words. The visual model and the language models are trained so that cross-entropies between their output and the ground truth are minimized. We denote the cross-entropy loss function for normal speech data as \mathcal{L}_{NCE} and that for silent speech data as \mathcal{L}_{SCE} .

3.2. Metric learning for inter-speech type

Although normal and silent speech exhibits different lip movements, the predicted viseme probability distributions should be similar regardless of speech type, as long as they have the same viseme. To fill this gap, we propose two loss functions based on KL divergence, derived from different handling manners of the two types of speech. Let P_N (P_S) be the target distribution of normal (silent) speech and Q_N (Q_S) be the predicted distribution of normal (silent) speech. For both speech types, the predicted distribution $Q = \{q_1, \dots, q_L\} \in \mathbb{R}^{C \times L}$ is generated by the visual model, where L is the length of the viseme sequence and C is the number of viseme classes. The ground truth of the viseme sequence corresponding to Q is given as $Y = \{y_1, \dots, y_L\} \in \mathbb{N}^L$, where $(y_l)_{l=1}^L \in \{1, \dots, C\}$. We use the average of all distributions classified into the same viseme within each speech type as the representative distribution for the viseme label. The alignment of distributions is represented as $S = \{s_1, \dots, s_C\} \in \mathbb{R}^{C \times C}$, where $(s_c)_{c=1}^C \in \mathbb{R}^C$ is the representative distribution for viseme class c defined as:

$$s_c = \frac{\sum_{y_l=c} q_l}{m}, \quad (1)$$

where m represents the number of times the viseme label c appears in the true viseme sequence Y , i.e., $m = \#\{y_l = c\}$. By aligning each distribution included in S with the viseme label of Y , we create the target distribution $P = \{p_1, \dots, p_L\} \in \mathbb{R}^{C \times L}$ (for example, if $y_2 = 8$, s_8 is allocated to p_2).

3.2.1. Metric learning with single KL divergence

In this method, we regard normal speech as the primary speech type and silent speech as its variant. Specifically, we leverage the KL divergence to bring the predicted distribution of silent speech close to the target distribution of normal speech by minimizing the following \mathcal{L}_{KL} :

$$\mathcal{L}_{\text{KL}} = \mathcal{D}_{\text{KL}}(P_N || Q_S) = P_N \log \frac{P_N}{Q_S}. \quad (2)$$

3.2.2. Metric learning with weighted KL divergence

Next, we regard both types of speech as variants of speech in general (i.e., we regard them as of equal importance). Such a point of view lets us consider a weighted distribution that depends on the number of training data used for each speech type to the target distribution of the KL divergence. The weighted distribution M can be defined as:

$$M = \frac{N_N P_N + N_S P_S}{N_N + N_S}, \quad (3)$$

where N_N and N_S represent the number of normal and silent speech data, respectively. The predicted distributions of normal

and silent speech approach to the distribution M by minimizing the following sum of KL divergences:

$$\mathcal{L}_{\text{WKL}} = \mathcal{D}_{\text{KL}}(M || Q_N) + \mathcal{D}_{\text{KL}}(M || Q_S). \quad (4)$$

Note that M becomes equal to P_N if $N_N/N_S \rightarrow \infty$, which represents the situation of the current VSR for silent speech.

3.3. Metric learning for within-speech type

Even when people pronounce the same word in the same speech type, differences in speaker [24] and location [25] can cause variations in their lip movements. Inspired by this, we consider within-class variances in viseme probability distributions. By reducing such differences through metric learning within each speech type, the model can classify the distributions into their true visemes more accurately. To achieve this, we bring all distributions close to a representative distribution if they show the same viseme through minimizing the KL divergence \mathcal{L}_{NKL} (\mathcal{L}_{SKL}) for normal (silent) speech data as follows:

$$\mathcal{L}_{\text{NKL}} = \mathcal{D}_{\text{KL}}(P_N || Q_N) = P_N \log \frac{P_N}{Q_N}, \quad (5)$$

$$\mathcal{L}_{\text{SKL}} = \mathcal{D}_{\text{KL}}(P_S || Q_S) = P_S \log \frac{P_S}{Q_S}. \quad (6)$$

4. Evaluation

This section describes experiments conducted to evaluate the performance of the proposed method for silent speech VSR.

4.1. Database

To conduct our experiments, we used the AV Digits database [10], the only publicly available dataset that contains both normal and silent speech. This limited dataset comprises 39 speakers who repeated 10 brief phrases (including “Excuse me”, “Goodbye”, “Hello”, “How are you”, “Nice to meet you”, “See you”, “I am sorry”, “Thank you”, “Have a good time”, “You are welcome”) five times each in both speaking modes. Following the methodology of [10], we split the dataset into training, validation, and test sets of 1,000 utterances (from 20 speakers), 400 utterances (from 8 speakers), and 550 utterances (from 11 speakers), respectively. Additionally, we augmented the amount of training data for normal speech by incorporating the OuluVS2 database [12], which includes 1560 utterances (from 52 speakers) with the same vocabulary as the AV Digits database. In total, we utilized 2560 utterances (from 72 speakers) for normal speech and 1,000 utterances (from 20 speakers) for silent speech during the training phase by combining these two datasets. Since the datasets only include text-based labels, we transformed them into viseme sequences by first translating texts to phoneme sequences using the Carnegie Mellon Pronouncing (CMU) Dictionary, and then mapping the phonemes to visemes according to Lee and Yook’s approach [14].

4.2. Training settings

Following previous studies [16, 21, 22], the visual and language models were trained separately while fixing the visual front-end with weights pre-trained by [7]. Note that we conducted a thorough exploration of various parameters, including the weights for loss terms, and adopted the parameters that achieved the best performance. To train the visual model, we used the Adam optimizer [26] with an initial learning rate of 10^{-5} and a mini-batch size of 64 for 50 epochs, and decreased the learning rate

Table 1: Performance comparison of different loss functions during training. A lower score indicates higher accuracy.

Loss function	VER (%)		WER (%)	
	Normal	Silent	Normal	Silent
\mathcal{L}_{NCE}	5.87	12.49	11.58	22.55
$\mathcal{L}_{\text{NCE}} + \mathcal{L}_{\text{SCE}}$ (Baseline)	5.30	7.48	10.21	12.91
+ \mathcal{L}_{NKL}	5.24	7.77	7.36	13.45
+ \mathcal{L}_{SKL}	5.10	7.20	8.70	12.42
+ $\mathcal{L}_{\text{NKL}} + \mathcal{L}_{\text{SKL}}$	5.22	7.07	8.48	12.24
$\mathcal{L}_{\text{NCE}} + \mathcal{L}_{\text{SCE}} + \mathcal{L}_{\text{KL}}$	5.21	7.17	8.12	11.06
+ \mathcal{L}_{NKL}	5.15	6.92	7.85	10.85
+ \mathcal{L}_{SKL}	5.13	7.35	9.36	10.71
+ $\mathcal{L}_{\text{NKL}} + \mathcal{L}_{\text{SKL}}$	5.11	6.96	7.85	10.73
$\mathcal{L}_{\text{NCE}} + \mathcal{L}_{\text{SCE}} + \mathcal{L}_{\text{WKL}}$	5.16	6.89	8.70	10.73
+ \mathcal{L}_{NKL}	5.09	6.71	8.00	10.88
+ \mathcal{L}_{SKL}	5.10	6.89	8.76	11.00
+ $\mathcal{L}_{\text{NKL}} + \mathcal{L}_{\text{SKL}}$	5.03	6.66	8.20	9.97

by 50% if the validation loss did not improve over five epochs. For the language model, we employed the same settings, except for the initial learning rate and mini-batch size, which were set to $5 \cdot 10^{-4}$ and 10, respectively. During all experiments, the language model was trained on the viseme sequences predicted by the visual model using normal speech data.

4.3. Evaluation metrics

For evaluation, we used two metrics: viseme error rate (VER) for the visual model outputs and word error rate (WER) for the language model outputs. We calculated VER as follows:

$$\text{VER} = \frac{V_S + V_D + V_I}{V_N}, \quad (7)$$

where V_N is the total number of visemes in the ground truth. V_S , V_I , and V_D represent the number of visemes substituted for wrong classifications, inserted for those not to be present, and deleted for those to be present, respectively. WER refers to the same concept as VER yet is against words instead of visemes.

4.4. Experimental results

We conducted 13 experiments with different combinations of loss functions to evaluate the accuracy of our proposed method, and the results are summarized in Table 1. Initially, to confirm the issue with silent speech, we trained the model solely on normal speech data using the cross-entropy \mathcal{L}_{NCE} . This resulted in the largest performance drop for silent speech in both VER and WER, with values of 6.62% and 10.97%, respectively, compared to normal speech. To serve as a baseline for comparison, we used the model trained only with the classification loss function $\mathcal{L}_{\text{NCE}} + \mathcal{L}_{\text{SCE}}$, which achieved higher accuracy than the adaptation approach pre-trained with normal speech and then fine-tuned with silent speech. Additionally, we incorporated two kinds of metric learning for inter-speech type, namely \mathcal{L}_{KL} and \mathcal{L}_{WKL} . Both approaches improved accuracy for silent speech, leading to a 1.85% and 2.18% decrease in WER, respectively, compared to the baseline. These results demonstrate the effectiveness of using metric learning between normal and silent speech, and treating them with equal consideration. For further improvement, we added metric learning for within-speech type, i.e., using only \mathcal{L}_{NKL} , only \mathcal{L}_{SKL} ,




Figure 3: Impact of decreasing training data for silent speech on VER, from 1,000 to 200 by 200. The number of normal speech data remained constant at 1,000. The shaded area represents the range $m \pm \sigma$, where m is the mean and σ is the standard deviation of VER over five experiments.

or both. The best accuracy achieved for silent speech was 6.66% in VER and 9.97% in WER, using the loss function $\mathcal{L}_{\text{NCE}} + \mathcal{L}_{\text{SCE}} + \mathcal{L}_{\text{WKL}} + \mathcal{L}_{\text{NKL}} + \mathcal{L}_{\text{SKL}}$, which showed that metric learning within each speech type could reduce the variances in the output distributions.

In these experiments, we utilized 2,560 normal speech data and 1,000 silent speech data in the training process. However, recent major studies [7,8,27] have employed VSR datasets consisting of tens of thousands of normal speech data, emphasizing the need to evaluate the effectiveness of our proposed method in more realistic scenarios. To assess its robustness to an imbalanced amount of data used for each speech type, we reduced the number of utterances per speaker from 50 to 40, 30, 20, and 10 in silent speech data, by modifying N_S in Equation 3. Each experiment using the same number of data was conducted five times, with different training data excluded in each run. Figure 3 shows that the proposed method consistently achieved a lower VER for silent speech than the baseline method across all training data settings. Notably, the proposed model trained on 400 silent speech data exhibited comparable accuracy to the baseline model trained on twice the amount of data.

5. Conclusion

We proposed a novel approach to improve the accuracy of silent speech prediction using a VSR model trained on normal speech. Specifically, we presented a viseme-based metric learning method that addressed the limited availability of silent speech datasets by focusing on the common viseme representations between normal and silent speech. Our approach employed metric learning techniques to learn visemes inter- and within-speech types, allowing us to efficiently leverage silent speech data while accounting for variations within the same speech type. Through extensive experimentation with various loss functions, we demonstrated that our proposed method consistently outperformed the baseline approach, even under challenging conditions. In future work, we plan to extend our approach to larger normal speech datasets to facilitate VSR on silent speech data with unknown vocabulary, which was not included in the 10 phrases used in our experiments.

6. Acknowledgements

This work is supported in part by JSPS KAKENHI Nos. 21H05054, 22J22424, 22KJ2959.

7. References

- [1] J. S. Chung, A. W. Senior, O. Vinyals, and A. Zisserman, "Lip reading sentences in the wild," *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3444–3453, 2016.
- [2] P. Ma, S. Petridis, and M. Pantic, "Visual speech recognition for multiple languages in the wild," *Nature Machine Intelligence*, vol. 4, pp. 930–939, 2022.
- [3] K. Sun, C. Yu, W. Shi, L. Liu, and Y. Shi, "Lip-interact: Improving mobile device interaction with silent speech commands," *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pp. 581–593, 2018.
- [4] L. Pandey and A. S. Arif, "Liptype: A silent speech recognizer augmented with an independent repair model," *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1–19, 2021.
- [5] M. Janke, M. Wand, and T. Schultz, "Impact of lack of acoustic feedback in emg-based silent speech recognition," in *Proceedings of the 11th Annual Conference of the International Speech Communication Association (Interspeech)*, 2010, pp. 2686–2689.
- [6] K. Bicevskis, J. de Vries, L. Green, J. M. Heim, J. Božič, J. D'Aquisto, M. Fry, E. Sadlier-Brown, O. Tkachman, N. Yamane, and B. Gick, "Effects of mouthing and interlocutor presence on movements of visible vs. non-visible articulators," *Canadian acoustics = Acoustique canadienne*, vol. 44, no. 1, pp. 17–24, 2016.
- [7] T. Afouras, J. S. Chung, A. W. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 8717–8727, 2018.
- [8] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," *arXiv preprint arXiv:1809.00496*, 2018.
- [9] V. M. Florescu, L. Crevier-Buchman, B. Denby, T. Hueber, A. Colazo-Simon, C. Pillot-Loiseau, P. Roussel-Ragot, C. Gendrot, and S. Quattrocchi, "Silent vs vocalized articulation for a portable ultrasound-based silent speech interface," in *Proceedings of the 11th Annual Conference of the International Speech Communication Association (Interspeech)*, 2010, pp. 450–453.
- [10] S. Petridis, J. Shen, D. Cetin, and M. Pantic, "Visual-only recognition of normal, whispered and silent speech," *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6219–6223, 2018.
- [11] J. S. Chung and A. Zisserman, "Lip reading in the wild," in *Asian Conference on Computer Vision*, 2016, pp. 87–103.
- [12] I. Anina, Z. Zhou, G. Zhao, and M. Pietikäinen, "Ouluvs2: A multi-view audiovisual database for non-rigid mouth motion analysis," *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, vol. 1, pp. 1–5, 2015.
- [13] J. S. Chung and A. Zisserman, "Lip reading in profile," in *British Machine Vision Conference*, 2017, pp. 1–11.
- [14] S. Lee and D. Yook, "Audio-to-visual conversion using hidden markov models," in *Proceedings of the 7th Pacific Rim International Conference on Artificial Intelligence: Trends in Artificial Intelligence*, 2002, pp. 563–570.
- [15] L. Cappelletta and N. Harte, "Viseme definitions comparison for visual-only speech recognition," *2011 19th European Signal Processing Conference*, pp. 2109–2113, 2011.
- [16] S. Fenghour, D. Chen, K. Guo, and P. Xiao, "Lip reading sentences using deep learning with only visual cues," *IEEE Access*, vol. 8, pp. 215 516–215 530, 2020.
- [17] J. Peymanfard, M. R. Mohammadi, H. Zeinali, and N. Mozayani, "Lip reading using external viseme decoding," *2022 International Conference on Machine Vision and Image Processing (MVIP)*, pp. 1–5, 2021.
- [18] K. Prajwal, T. Afouras, and A. Zisserman, "Sub-word level lip reading with visual attention," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5162–5172, 2021.
- [19] B. Shillingford, Y. Assael, M. W. Hoffman, T. L. Paine, C. Hughes, U. Prabhu, H. Liao, H. Sak, K. Rao, L. Bennett, M. Mulville, B. Coppin, B. Laurie, A. W. Senior, and N. de Freitas, "Large-scale visual speech recognition," *arXiv preprint arXiv:1807.05162*, 2018.
- [20] H. L. Bear and R. Harvey, "Alternative visual units for an optimized phoneme-based lipreading system," *Applied Sciences*, vol. 9, no. 18, pp. 3870–3881, 2019.
- [21] S. Fenghour, D. Chen, and P. Xiao, "Decoder-encoder lstm for lip reading," *Proceedings of the 8th International Conference on Software and Information Engineering*, pp. 162–166, 2019.
- [22] S. Fenghour, D. Chen, K. Guo, B. Li, and P. Xiao, "An effective conversion of visemes to words for high-performance automatic lipreading," *Sensors (Basel, Switzerland)*, vol. 21, no. 23, p. 7890, 2021.
- [23] V. Kazemi and J. Sullivan, "One millisecond face alignment with an ensemble of regression trees," *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1867–1874, 2014.
- [24] M. Wand and J. Schmidhuber, "Improving speaker-independent lipreading with domain-adversarial training," in *Proceedings of the 18th Annual Conference of the International Speech Communication Association (Interspeech)*, 2017, pp. 3662–3666.
- [25] J. Šimko, S. Benus, and M. Vainio, "Hyperarticulation in lombard speech: Global coordination of the jaw, lips and the tongue," *The Journal of the Acoustical Society of America*, vol. 139, no. 1, pp. 151–162, 2016.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [27] M. Cooke, J. Barker, S. P. Cunningham, and X. Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *The Journal of the Acoustical Society of America*, vol. 120, no. 5, pp. 2421–2424, 2006.