



A Novel Interpretable and Generalizable Re-synchronization Model for Cued Speech based on a Multi-Cuer Corpus

Lufei Gao^{1†}, Shan Huang^{1†}, Li Liu^{2(✉)}

¹Shenzhen Research Institute of Big Data, China

²The Hong Kong University of Science and Technology (Guangzhou), China

avrillliu@hkust-gz.edu.cn

Abstract

Cued Speech (CS) is a multi-modal visual coding system combining lip reading with several hand cues at the phonetic level to make the spoken language visible to the hearing impaired. Previous studies solved asynchronous problems between lip and hand movements by a cuer¹-dependent piecewise linear model for English and French CS. In this work, we innovatively propose three statistical measures on the lip stream to build an interpretable and generalizable model for predicting hand preceding time (HPT), which achieves cuer-independent by a proper normalization. Particularly, we build the first Mandarin CS corpus comprising annotated videos from five speakers including three normal and two hearing impaired individuals. Consequently, we show that the hand preceding phenomenon exists in Mandarin CS production with significant differences between normal and hearing impaired people. Extensive experiments demonstrate that our model outperforms the baseline and the previous state-of-the-art methods.

Index Terms: Mandarin Cued Speech, Asynchronous multimodality, Hand preceding time, Hearing-impaired

1. Introduction

Hearing-impairment is one of the social-concerned issues worldwide. According to the World Health Organization (WHO), more than 1.5 billion people worldwide are affected by some form of hearing impairment [1], presenting significant challenges for both society as a whole and the communities affected by this issue. **Cued Speech (CS)** is a visual coding system for the spoken language invented by Cornett in 1967 [2]. CS differs from sign language (SL) [3, 4, 5] in that it follows the rules of a normal language system, using hand gestures to code phonemes and supplementing the limitations of lip reading.

Currently, CS has been adapted to about 65 languages and dialects all over the world [6]. Liu et al. proposed a Chinese cued speech system, specifically for Mandarin cuers [7, 8]. In practice, we have found that for adults either normal hearing or hearing-impaired, who master Chinese *Pinyin*, it only takes 24 hours for them to master the rules of Mandarin CS and use it for simple expressions. Similar to other languages, hand cues in Mandarin CS also tend to reach the target state before the lip cues during expression. In the literature, it is called hand preceding phenomenon or lip-hand asynchronous problem [9, 10, 11]. Comprehending and dissecting this phenomenon holds great significance in understanding the cognitive prin-

ciples of CS expression, improving automatic CS recognition [12, 13, 14], and developing realistic CS synthesis systems.

There are three main contributions in this work:

- A new multi-cuer mandarin CS corpus is developed with accurate manual annotation for the lip and hand movements. It contains CS videos produced by three normal cuers and two hearing loss cuers. As far as we know, this is the first Mandarin CS corpus that contains data from hearing-loss people which is incredibly valuable for the academic research of CS.
- A novel method is proposed to address the lip-hand multi-modal asynchronous problem in the CS system and to deal with cuer adaptation. The method utilizes three lip stream measures to achieve lip-hand alignment for temporal segmentation and selects proper strategies for cuer normalization.
- Extensive experimental results show that our proposed re-synchronization method archives better performance compared with the previous methods concerning the re-synchronization effect, interpretability, and generalization ability for multiple cuers. Besides, we discuss the hearing-impaired performance of their mandarin CS expression with some noteworthy observations.

2. Related Work

2.1. Hand preceding phenomenon

The temporal organization of hand and speech coordination during French CS production was studied in [9, 10]. It was found that hand movement reaches the target position 239 ms before the onset of the acoustic vowel and hand shape formation accounts for a large fraction of hand transitions [9]. This relationship is not occasional but was also found on other proficient CS cuers [10]. The study [15] measured a 144.5 ms long HPT for syllables extracted from French continuous sentences and affirmed the importance of the instant at which the hand reaches the target position.

2.2. Lip-hand re-synchronizations in CS recognition

To resolve the issue of lip-hand asynchrony, Liu et al. [11, 13, 16] analyzed the correlation between HPTs and their lip target temporal positions, where a piece-wise linear function is used for HPM. It is presumed that the HPT remains constant until the turning point is achieved, after which it quickly declines. As the model is based on limited English/French CS videos, it is difficult to generalize to new cuers producing Mandarin CS, thus less generalizable.

Asynchronous multi-modal fusion is a fast-growing research area [17, 18]. Prior researches have explored the use of neural networks to tackle this issue in CS. In [19, 20], RNNs trained with a CTC loss were adopted to address the

[†] Equal contribution. This work was supported by the National Natural Science Foundation of China (No. 62101351)

¹The people who perform Cued Speech are called the cuer.

asynchronous problem. In [21], a transformer based on cross-modal mutual learning was proposed by achieving feature fusion based on the aligned modalities with the long-time dependencies. However, the black-box methods lack interpretability, which limits our ability to identify the specific factors related to the issue.

In the following sections, we will describe the corpus construction and define three measurements on the acoustic stream and employ statistical-based display methods to address the asynchronous issue, thereby enhancing both the interpretability and generalization ability of the multi-modal fusion process.

3. Methodology

3.1. Mandarin CS Corpus

3.1.1. Data collection

To collect the Mandarin CS videos, five volunteers, including three normal hearing (2 female and 1 male, named NF1, NF2, and NM1) and two hearing-impaired (1 female and 1 male, named DF1 and DM1) learned to use Mandarin CS in advance within 24 hours of training. A 1000-sentence Chinese corpus consisting of 5 themes was constructed for normal-hearing cuers, while a smaller corpus consisting of simpler sentences was provided for the hearing impaired cuers. Each sentence contains 4 to 27 words, 10.6 on average. The volunteers recorded 1000 videos at home, using the front-facing cameras on their personal mobile phones with the following standards: 1) white background with no decoration; 2) 720 HD at 30 fps; 3) turning off the camera mirror effect. Finally, we collected 3000 CS videos by the normal and 979 by the hearing-impaired².

3.1.2. Data annotation

Each video contains two annotation documents of temporal segments corresponding to phonemes, one based on hand gestures and the other based on acoustics. Hand annotations were completed by ELAN [22], while acoustic annotations by Praat [23]. The annotations were all manually done and double-checked.

3.2. Hand preceding analysis

For an input video, the audio signal and visual signal are regarded as the lip stream and hand stream, respectively. Due to the fact that one vowel corresponds to one syllable in Mandarin, only vowel segments are taken into our consideration. The required measurements related to the analysis are drawn in Fig. 1. Specifically, HPT is defined as the time difference between two instants, i.e., the hand target instant T_i^{mid} and lip target instant t_i^{mid} of the i^{th} vowel in a sentence. The target instants are the mid-time points between the start and the end time of the vowel segments extracted from temporal annotations. Therefore, we denote the i^{th} HPT by $\Delta_i = t_i^{mid} - T_i^{mid}$, $0 \leq i \leq N - 1$, where N is the number of syllables in a sentence.

3.2.1. Lip vowel to end

Lip vowel to end (LVE) refers to the time difference between the lip target instant of a vowel and the sentence end. It is observed that HPT decreases sharply as the cuer is about to finish a sentence. This phenomenon can be explained from the perspective of cognitive function in the brain. The articulatory process of pronunciation lags behind the hand gestures during




Figure 1: *Hand preceding analysis illustration. Green:* the box refers to the hand temporal segment corresponding to the vowel at t_i^{mid} ; the arrow refers to HPT. *Black:* the vertical line is the sentence end; the arrow refers to LVE. *Red:* two vertical lines are the lip target time instants of two adjacent vowels; the arrow refers to LVI. *Orange:* the arrow refers to LVD.

the expression, but as the expression is nearing completion, the two streams will gradually approach synchronization. Therefore, we define the variable that measures the time difference between the sentence end and the lip target instant of a vowel segment as lip vowel to end (LVE), denoted by δ_i . The mapping relationship from δ_i to Δ_i is written as $\hat{\Delta}_i = f_0(\delta_i)$.

3.2.2. Lip vowel interval

Pauses within an utterance may play a similar role for HPT at the end of the sentence. In other words, HPT decreases sharply to a small value when there is a large time gap between two syllables. Besides, a Mandarin syllable may have zero, one or two consonants preceding its vowel. When the hand moves from a vowel position to the next, a hand-shape transition occurs during Mandarin CS production if it is decorated with two consonants. The transition can increase the time duration between two syllables, thus lengthening the HPT. Therefore, we define the variable that measures the time difference between two adjacent lip target instants as lip vowel interval (LVI), denoted by α_i , i.e. $\alpha_i = t_i^{mid} - t_{i-1}^{mid}$. The last LVI of a sentence is set to the maximum value among previous LVI, i.e. $\max_{0 < i < N-1} \alpha_i$. The mapping relationship from α_i to Δ_i is written as $\hat{\Delta}_i = f_1(\alpha_i)$.

3.2.3. Lip vowel duration

Another factor that may influence Δ_i is the amount of time that is spent pronouncing the vowel. Since half of the vowel segment is included in the HPT measurement, if the vowel is pronounced longer, HPT will be prolonged accordingly. Thereby, we refer to this measurement as lip vowel duration (LVD) and denote it by β_i . The mapping from β_i to Δ_i is denoted by $\hat{\Delta}_i = f_2(\beta_i)$.

3.3. Cuer normalization

The corpus adopted in this work consists of five cuers with different ages, genders, hearing conditions, and proficiencies of CS production. Besides, the number of words and the expression speed vary substantially across different videos. Hence, selecting appropriate normalization is necessary in order to address subject differences, i.e. the issue of cuer adaptation.

²All cuers have consented to the public release of the data. The corpus can be accessed from: <https://github.com/lufei321/ReSync-CS>.




Figure 2: Linear regressions of HPT with respect to (a) LVE, (b) LVI and (c) LVD in the normalized scales.

3.3.1. Z-score normalization

Since the statistical distributions of HPT and LVD are approximately normal distributions, Z-score normalization is applied. The normalized variables are denoted by Δ'_i and β'_i , respectively. The descriptive statistics in the unit of milliseconds for each cuer, average on all cuers, average on normal cuers, and average on deaf cuers are shown in Table 1. The total number of vowel segments is listed in the last column of the table.

Table 1: Descriptive statistics of the data set

Cuers	HPT		LVD		Syllable number
	μ_Δ	σ_Δ	μ_β	σ_β	
NF1	242	177	338	79	10556
NF2	246	150	389	94	10556
NM1	352	137	464	144	10556
DF1	154	110	365	97	4584
DM1	164	105	397	101	4440
NORMAL	280	164	397	121	31668
DEAF	159	108	369	99	9024
ALL	253	161	390	117	40692

*N/D stands for normal/deaf; F/M stands for female/male. NF1 represents “Normal Female 1”, NF2 represents “Normal Female 2”, and so forth. NORMAL refers to the corpus subset containing NF1, NF2 and NM1. DEAF refers to the corpus subset containing DF1 and DM1. ALL refers to the corpus containing five cuers.

3.3.2. Log-scale normalization

Log-scale normalization is used for LVI as it shows a left-skewed distribution, denoted by $\alpha'_i = \lg \alpha_i$. Due to the different lengths and number of characters among videos, the statistical distribution of LVE is also skewed to left. So log-scale normalization is suggested for LVE as well, denoted by $\delta'_i = \lg \delta_i$.

3.4. Temporal organization

Fig.2 shows the linear regressions (LRs) [24] $f_0(\delta')$, $f_1(\alpha')$ and $f_2(\beta')$ based on ALL. Due to the presence of a distinct inflection point towards the end of the sentence, a piece-wise LR is employed for f_0 , where the segment point is denoted by γ . It is observed that HPT is dominated by LVE when δ' is smaller than γ . Otherwise, when the lip target instant is away from the sentence end, we propose a linear combination of $f_1(\alpha')$ and $f_2(\beta')$ to predict HPT. As in Eq.(1), $f(\cdot)$ is a segmentation function taking in the normalized LVE, LVI, and LVD as input variables to calculate HPT, where $\lambda_1 + \lambda_2 = 1$ and $\lambda_1 > 0, \lambda_2 > 0$. Intuitively, If LVI is larger for a syllable in a sentence, it is logical to assume that $f_1(\alpha')$ carries a greater weight and the same holds true when LVD is larger. Therefore we define λ_1 and λ_2 as in Eq. (2), where $\bar{\alpha}$ and $\bar{\beta}$ are the mean

values of α and β in a sentence, respectively.

$$f(\delta', \alpha', \beta') = \begin{cases} f_0(\delta'), & \delta' \leq \gamma \\ \lambda_1 f_1(\alpha') + \lambda_2 f_2(\beta'), & \delta' > \gamma \end{cases} \quad (1)$$

$$\lambda_1 = \frac{\alpha}{\alpha + C\beta}, \lambda_2 = \frac{C\beta}{\alpha + C\beta}, C = \frac{\bar{\alpha}}{\bar{\beta}}. \quad (2)$$

In the following experiments, γ is empirically set to -0.34 and LR analyses are conducted for ALL, NORMAL and DEAF separately, named A-LR, N-LR and D-LR for short, respectively.

4. Experimental Setup

4.1. Data pre-processing

To evaluate the effectiveness of the proposed HPT model, we conduct the vowel recognition experiment and calculate the average distance between the hand coordinates at the target time point and the predicted time point. Therefore, video clips containing vowels are segmented based on the lip annotations, hand annotations and the predicted hand target instants. The regions of interest (ROIs) of lip and hand are then extracted from each clip using the facial landmark detector included in the dlib library [25, 26] and hands solution in mediapipe [27], respectively. Once the key points of the lips and hands have been extracted, a 140×140 square is obtained around the lips, while a 320×320 square is obtained around the hands, with the exception of DM1, whose hand square is set to 400×400 due to larger hand size. To do the vowel recognition, the training and testing sets are divided randomly in a 4:1 ratio based on sentences.

4.2. Vowel recognition

We utilize the Slowfast video classification network [28] as the backbone for feature extraction and concatenate the lip and hand features for multi-modal feature fusion. The fused feature is then classified using a two-layered multilayer perceptron into 16 vowel categories.

The network is initialized randomly and trained using the SGD optimizer with an initial learning rate of 0.01, which is adjusted using the CosineAnnealing policy [29]. The mini-batch size is set to 64, and the entire network is trained for 100 epochs.

4.3. Metrics

Three metrics are adopted to evaluate the performance of HPT predictions. The first is mean squared error (MSE) between the HPT predictions and their ground truth in the normalized scale, denoted by e_{HPT} . The second is mean hand coordinate distances (MHCD) between the hand coordinates at T_i^{mid} and $t_i^{\text{mid}} - \hat{\Delta}_i$, denoted by $d_{\text{HPT}} = \frac{1}{N} \sum_{i=0}^{N-1} \sqrt{(x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2}$, where (x_i, y_i) and (\hat{x}_i, \hat{y}_i) are the hand coordinates³ of two time instants. The third is the vowel recognition accuracy, defined as the ratio of correctly recognized vowels to the total number of vowels.

5. Results and Discussion

The hand temporal segmentation is done through five HPT calculations: ground truth by hand annotations, named GT; baseline by audio annotations, named audio-based; mean HPT value

³The hand coordinates are obtained either by taking the fingertip or by taking the midpoint between the index finger and the middle finger. When one of the two fingers is not extended, the only extended finger (either the middle finger for hand shape 3 or the index finger for hand shapes 1 and 6) is taken.

Table 2: The results of vowel recognition.

LR data	Cuers	Mono-Modal		Multi-Modal Fusion (Lip + Hand)				
		Lip	Hand	Ground truth	Audio-based	Mean-based [11]	LVE	LVE + LVI + LVD
ALL	NF1	86.52%	81.19%	97.02%	89.92%	95.87%	96.35%	96.50%
	NF2	93.91%	77.16%	98.27%	96.31%	98.27%	98.27%	98.39%
	NM1	81.57%	71.79%	96.40%	84.69%	95.11%	95.30%	95.30%
	DF1	86.20%	80.54%	96.84%	91.96%	96.09%	96.09%	96.41%
	DM1	76.59%	73.18%	91.29%	85.76%	90.59%	90.71%	90.82%
	Average	86.06%	76.78%	96.56%	90.01%	95.76%	95.95%	96.06%
NORMAL	Average	87.33%	76.71%	97.23%	90.31%	96.42%	96.51%	96.61%
DEAF	Average	81.58%	77.01%	94.18%	88.98%	93.45%	93.56%	93.90%

*N/D stands for normal/deaf; F/M stands for female/male. NF1 represents “Normal Female 1”, NF2 represents “Normal Female 2”, and so forth.




Figure 3: Comparison of MSE and MHCD: (a) MSE across different corpus subsets calculated on different LRs. (b) Boxplot of MHCD across different cuers calculated on different LRs.

of the dataset, named mean-based; HPT by $f_0(\delta')$, named LVE; HPT by Eq.(1), named LVE + LVI + LVD. Here, the other combinations such as LVE + LVI and LVE + LVD are omitted since they exhibited inferior results compared to the selected ones.

5.1. Evaluation

The vowel recognition results are shown in Table 2. LVE + LVI + LVD achieves the closest accuracy to GT. It even exceeds GT for NF2 by A-LR, reaching as much as 98.39% accuracy rate, which may be due to inaccurate manual annotations. The results of MSE and MHCD are shown in Fig. 3. Comparing with mean-based, LVE performs better on both metrics, while LVE + LVI + LVD leads to the minimum MSE and MHCD. Besides, Fig. 3(b) shows that LVE + LVI + LVD has better stability.

To visualize the correcting effect of the method, polar coordinates are used to represent hand position, where the center point of the lip is set to be the coordinates origin. The spatial distributions of hand positions at the predicted hand target instants are shown in Fig. 4. LVE + LVI + LVD has a corrective effect on the hand positions at target instants by revising the HPT values. Intuitively, a better hand temporal segment corresponds to a more separable hand position distribution. It is obvious that the separability across five positions becomes more and more remarkable from (b) to (d). In summary, the proposed method can improve the performance of hand temporal segmentation by comprehensive evaluation.

5.2. Comparison between normal and hearing-impaired

As shown in Table 1, the statistic characteristics are considerably different between NORMAL and DEAF. A noteworthy observation is that the mean HPT of DEAF is significantly lower than that of NORMAL. This may be because hearing-impaired subjects have less accurate speech perception, leading to their hand movements arriving closer in time to the lip target in-

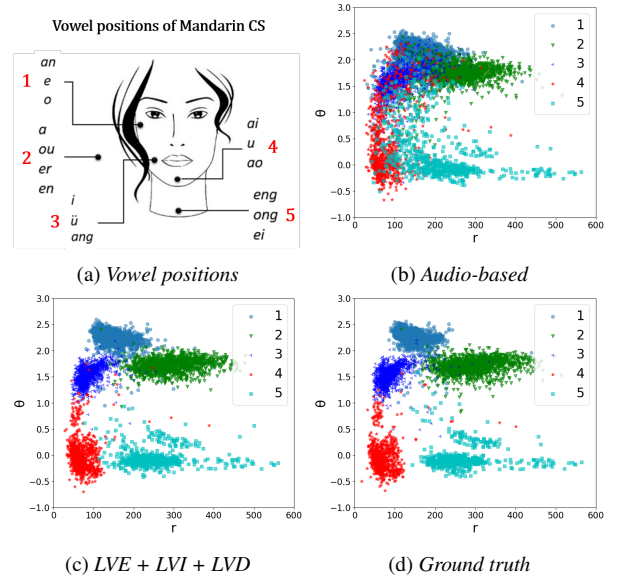


Figure 4: (a) Five hand positions and their corresponding vowels of the Mandarin CS system. (b)-(d) Spatial distributions of hand positions on the polar coordinates at $t_i - \hat{\Delta}_i$ of the vowels randomly selected from five cuers.

stant. Moreover, female cuers generally exhibit smaller mean HPT values than male cuers, indicating that women’s language proficiency extends to cued speech.

In Fig. 3(a), the MSE of NORMAL by A-LR are almost the same as NORMAL by N-LR. But the MSE reduction in DEAF by D-LR is significant. This is reasonable as the amount of DEAF data occupies less than 25% in the complete corpus, so that there is a larger bias compared with the NORMAL group. This finding also gives a hint that it may be better to distinguish between different hearing conditions when developing a CS recognition system.

6. Conclusions

We propose a novel re-synchronization method for the lip-hand asynchronous problem during CS production. A multi-cuer mandarin CS corpus is collected to build the statistical-based model. Extensive experiments illustrate the efficiency of our proposed method and the differences that exist between normal and hearing-impaired cuers are discussed. In the future, in addition to expanding the corpus, we will use the interpretable model to provide semi-supervised signals to train a pre-trained model for automatic CS segmentation and recognition.

7. References

- [1] “Who: 1 in 4 people projected to have hearing problems by 2050,” who.int, March 2021. [Online]. Available: <https://www.who.int/news/item/02-03-2021-who-1-in-4-people-projected-to-have>
- [2] R. O. Cornett, “Cued speech,” *American annals of the deaf*, pp. 3–13, 1967.
- [3] S. K. Liddell and R. E. Johnson, “American sign language: The phonological base,” *Sign language studies*, vol. 64, no. 1, pp. 195–277, 1989.
- [4] C. Valli and C. Lucas, *Linguistics of American sign language: An introduction*. Gallaudet University Press, 2000.
- [5] W. C. Stokoe Jr, “Sign language structure: An outline of the visual communication systems of the american deaf,” *Journal of deaf studies and deaf education*, vol. 10, no. 1, pp. 3–37, 2005.
- [6] “Find your cued language,” cuedspeech.org. [Online]. Available: <https://cuedspeech.org/learn/find-your-cued-language/>
- [7] L. Liu and G. Feng, “A pilot study on mandarin chinese cued speech,” *American Annals of the Deaf*, vol. 164, no. 4, pp. 496–518, 2019.
- [8] L. Liu, G. Feng, X. Ren, and X. Ma, “Objective hand complexity comparison between two mandarin chinese cued speech systems,” in *2022 13th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2022, pp. 215–219.
- [9] V. Attina, D. Beaufemps, M.-A. Cathiard, and M. Odisio, “A pilot study of temporal organization in cued speech production of french syllables: rules for a cued speech synthesizer,” *Speech Communication*, vol. 44, no. 1-4, pp. 197–214, 2004.
- [10] V. Attina, M.-A. Cathiard, and D. Beaufemps, “Temporal measures of hand and speech coordination during french cued speech production,” in *Gesture in Human-Computer Interaction and Simulation: 6th International Gesture Workshop, GW 2005, Berder Island, France, May 18-20, 2005, Revised Selected Papers 6*. Springer, 2006, pp. 13–24.
- [11] L. Liu, G. Feng, D. Beaufemps, and X.-P. Zhang, “Resynchronization using the hand preceding model for multi-modal fusion in automatic continuous cued speech recognition,” *IEEE Transactions on Multimedia*, vol. 23, p. 292–305, 2020.
- [12] L. Liu, G. Feng, and D. Beaufemps, “Automatic dynamic template tracking of inner lips based on clnf,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, p. 5130–5134.
- [13] —, “Automatic temporal segmentation of hand movements for hand positions recognition in french cued speech,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, p. 3061–3065.
- [14] L. Liu, T. Hueber, G. Feng, and D. Beaufemps, “Visual recognition of continuous cued speech using a tandem cnn-hmm approach,” in *Interspeech*, 2018, p. 2643–2647.
- [15] N. Aboutabit, D. Beaufemps, and L. Besacier, “Hand and lip desynchronization analysis in french cued speech: Automatic temporal segmentation of hand flow,” in *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, vol. 1. IEEE, 2006, pp. I–I.
- [16] L. Liu, G. Feng, D. Beaufemps, and X.-P. Zhang, “A novel resynchronization procedure for hand-lips fusion applied to continuous french cued speech recognition,” in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.
- [17] M. Tang, X. Wu, P. Agrawal, S. Pongpaichet, and R. Jain, “Integration of diverse data sources for spatial pm2. 5 data interpolation,” *IEEE Transactions on Multimedia*, vol. 19, no. 2, pp. 408–417, 2016.
- [18] Y. Sun, M. Liu, and M. Q.-H. Meng, “Active perception for foreground segmentation: An rgb-d data-based background modeling method,” *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 4, pp. 1596–1609, 2019.
- [19] J. Wang, Z. Tang, X. Li, M. Yu, Q. Fang, and L. Liu, “Cross-modal knowledge distillation method for automatic cued speech recognition,” p. 2986–2990, 2021.
- [20] S. Sankar, D. Beaufemps, and T. Hueber, “Multistream neural architectures for cued speech recognition using a pre-trained visual feature extractor and constrained ctc decoding,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 8477–8481.
- [21] L. Liu and L. Liu, “Cross-modal mutual learning for cued speech recognition,” *arXiv preprint arXiv:2212.01083*, 2022.
- [22] H. Sloetjes, “305ELAN: Multimedia Annotation Application,” in *The Oxford Handbook of Corpus Phonology*. Oxford University Press, 05 2014. [Online]. Available: <https://doi.org/10.1093/oxfordhb/9780199571932.013.019>
- [23] P. Boersma, “Praat: doing phonetics by computer [computer program],” <http://www.praat.org/>, 2011.
- [24] X. Su, X. Yan, and C.-L. Tsai, “Linear regression,” *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 4, no. 3, pp. 275–294, 2012.
- [25] D. E. King, “Dlib-ml: A machine learning toolkit,” *The Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.
- [26] —, “Dlib-models,” 2020. [Online]. Available: <https://github.com/davisking/dlib-models>
- [27] Google, “Mediapipe,” 2020. [Online]. Available: <https://github.com/google/mediapipe>
- [28] C. Feichtenhofer, H. Fan, J. Malik, and K. He, “Slowfast networks for video recognition,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6202–6211.
- [29] I. Loshchilov and F. Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” *arXiv preprint arXiv:1608.03983*, 2016.