



MAVD: The First Open Large-Scale Mandarin Audio-Visual Dataset with Depth Information

Jianrong Wang¹ Yuchen Huo² Li Liu³(✉) Tianyi Xu¹ Qi Li⁴ Sen Li¹

¹College of Intelligence and Computing, Tianjin University, Tianjin, China

²Tianjin International Engineering Institute, Tianjin University, Tianjin, China

³The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

⁴School of Electrical and Information Engineering, Tianjin University, Tianjin, China

avrillliu@hkust-gz.edu.cn

Abstract

Audio-visual speech recognition (AVSR) gains increasing attention from researchers as an important part of human-computer interaction. However, the existing available Mandarin audio-visual datasets are limited and lack the depth information. To address this issue, this work establishes the MAVD, a new large-scale Mandarin multimodal corpus comprising 12,484 utterances spoken by 64 native Chinese speakers. To ensure the dataset covers diverse real-world scenarios, a pipeline for cleaning and filtering the raw text material has been developed to create a well-balanced reading material. In particular, the latest data acquisition device of Microsoft, Azure Kinect is used to capture depth information in addition to the traditional audio signals and RGB images during data acquisition. We also provide a baseline experiment, which could be used to evaluate the effectiveness of the dataset. The dataset and code will be released at <https://github.com/SpringHuo/MAVD>.

Index Terms: Audio-Visual Speech Recognition, Mandarin Audio-Visual Corpus, Azure Kinect, Depth Information

1. Introduction

Speech recognition is a key component of the human-computer interaction (HCI), serving as an indispensable bridge that connects human-machine communications. Traditionally, early speech recognition systems are capable of recognizing input audio signals and converting them into corresponding text units [1], which can assist in HCI [2], language learning [3], and aiding people with physical disabilities [4, 5]. Previous research on speech recognition has obtained significant achievements. Currently, it was reported that automatic speech recognition (ASR) systems performed almost as accurately as humans in relatively quiet laboratory environments [6].

However, audio-based speech recognition systems still cannot achieve satisfactory performance in noisy environments. This is because human speech recognition involves the acquisition of information from multiple modalities and an extensive analysis for a relatively accurate result [7, 8]. To address this issue, visual signals that remain unaffected by noise were presented to provide supplementary information to the speech signal [9]. The research results indicated that visual information is advantageous to ASR, particularly in the presence of noisy or unavailable audio [10].

Moreover, the quality and quantity of data used can greatly affect the performance and accuracy of deep neural networks (DNNs). In addition to traditional datasets containing audio

and color images, researchers wanted additional information to complement the visual feature. Proposed solutions included using multi-view capture devices and stereo cameras, which have been shown to improve the audio-visual speech recognition (AVSR) model based on experimental evidence [11, 12, 13, 14]. However, the availability of Mandarin open-source audio-visual datasets appears to be quite limited as it is time and labor-consuming. In conclusion, existing datasets still suffer from the following challenges:

- There is a shortage of Mandarin reading materials that are relevant to daily life and gathered from a wide range of domains.
- The annotations of the previous datasets are unitary, and thus cannot be directly adapted to various scenarios (e.g., sentence-level or phoneme-level).
- Due to the lack of the Mandarin audio-visual dataset with depth information, the speaker's head or body movements, and the noise caused by the recording environment may have a negative impact on feature extraction.

To solve the above-mentioned challenges, we establish a large-scale Mandarin Audio-Visual dataset with Depth information, named MAVD. The corpus contains 12,484 utterances spoken by 64 speakers. The distribution of speakers covers 24 different provinces of China, and all speakers are asked to record in Mandarin. The corpus of raw text comprises content from a variety of popular open-source media in China (e.g. Weibo, People's Daily, etc.) to better fit the words used in daily life. All raw text is obtained through an automated pipeline and then these texts are filtered manually to avoid some immoral or violent sentences. Azure Kinect is utilized to collect data, which includes depth images, color images, and audio. Information about the speaker is given along with multimodal data, including gender, age, hometown, etc. We annotate the audio at the phoneme-level and the *pinyin* sequences, tone sequences, part-of-speech combinations, di-phoneme, and tri-phoneme sequences of each sentence are obtained by using the automatic language dictionary-based rule to enlarge the annotations.

By making our corpus open-source and completely free for research use, we intend to contribute to the Mandarin audio-visual speech recognition community. As far as we know, this is the first open large-scale Mandarin audio-visual dataset with depth information.

2. Related work

In the literature, there were some audio-visual datasets in multiple languages. Table 1 outlines some of the most prominent and widely utilized datasets.

This work was supported by the National Natural Science Foundation of China (No. 61977049) and National Natural Science Foundation of China (No. 62101351).

Table 1: Overview of recent audio-visual datasets which are related with ours. For the different modalities, A for audio, V for Visual (RGB) and D for Depth information.

DataSet	Language	Modalities	Speakers	Samples	Recoding Angle	Scenario
AV-Letters [15]	English	A, V	10	780	Front	Indoor
GRID [18]	English	A, V	34	1,000	Front	Indoor
CUAVE [27]	English	A, V	36	7,000	0°, 90°, -90°	Indoor
GAMVA [28]	Japanese	A, V	20	500	Eleven different angles	Indoor
BAVCD [30]	English	A, V, D	21	6,700	Front	Indoor
RGB-D [31]	English	A, V, D	53	1,060	Front	Indoor
CAVSR [19]	Mandarin	A, V	20	3,120	Front	Indoor
LRW-1000 [23]	Mandarin	A, V	2,000	718,018	Front	Wild
MAVD (ours)	Mandarin	A, V, D	64	12,484	Front	Indoor

2.1. Audio-Visual Datasets

For traditional audio-visual datasets, early AVSR corpora usually only asked speakers to read some simple and short phrases, such as numbers and letters. The AV-Letters [15] was one of the most used datasets for training the alphabetic AVSR system. It consisted of 780 utterances spoken by 10 subjects. Every English letter had been repeated 3 times by each speaker. In addition, their team released AV-Letters-2 [16], which was smaller in size but has been upgraded in terms of recording conditions (primarily improved video quality). Like DAVID [17], GRID [18] collected 1,000 sentences spoken by 34 speakers. Its sentences were simple, syntactically identical phrases. And those phrases consisted of some English words and numbers that are commonly used in daily life. In particular, GRID was recorded in a quiet and noise-free professional studio. These studies demonstrated that improving the recording quality may be beneficial for multimodal speech recognition systems.

CAVSR was the earliest Chinese audio-visual corpus in Mandarin [19]. This corpus was based on isolated syllables and did not consider continuous pronunciation cases. HIT Bi-CAVDDatabase was a dataset mainly for lipreading tasks [20, 21], and its subsequent release of the utterance-level continuous lipreading dataset HIT-AVDB II [22]. In addition, there was a large-scale lip reading dataset called LRW-1000 [23], which contained 718,018 samples from more than 2,000 individual speakers. This was a “speak in the wild” (SITW) dataset [24] with most of the data collected from dialogue-based television broadcasts. However, the differences in data sources, including shooting environments and equipment, may cause databases to become more challenging to work with compared to data collected under controlled laboratory conditions [25, 26].

In realistic scenarios, the speaker does not always face the camera but may speak from any angle with complicated phrases. Therefore, it was a common way to set up multi-view devices to capture data from different angles. CUAVE [27], containing 7,000 utterances spoken by 36 speakers, with the material of connected or isolated digits. In particular, this dataset had 3 perspectives from the front (0°) and two sides (90° and -90°). GAMVA [28] contained 20 Japanese male subjects, each reading 25 sentences of common Japanese phrases. Specifically, GAMVA used 12 cameras to capture information from all possible angles (including pitch and elevation angles).

2.2. Audio-Visual Datasets with Depth Information

Using a stereo camera to record depth information is a more common way to eliminate errors caused by different lighting conditions, as well as unconscious head movements of the




Figure 1: Pipeline of reading material acquisition

speaker. WAPUSK-20 [29] consisted of 20 subjects reading sentences with the same format as GRID using a Bumblebee stereo-camera capture. Designing and building a 3D-AVSR dataset is always a difficult task due to the complex data synchronization and the expensive cost. This has led researchers to focus on Kinect, which is a more efficient and cheaper device. BAVCD [30] used a first-generation Kinect to capture depth images. This dataset included 6700 utterances read by 15 English-speaking subjects and 6 Greek-speaking subjects. And then in its subsequent work, a new database with multiple views captured the conversation between two subjects. RGB-D [31] used a second-generation Kinect as a 3D information acquisition device to collect data from 20 speakers each reading 20 common English phrases. And through the Kinect, rich 3D data were released, including facial contours, depth images, head orientation, etc. Compared with a large number of multimodal datasets in English, however, available open-source mandarin AVSR databases with depth information are very scarce.

Therefore, by taking the advantage of each of the above datasets, we use the latest generation of Microsoft Kinect, Azure Kinect, to capture this dataset containing audio, color images, and depth images in a professional soundproof studio.

3. Our Mandarin Audio-Visual Dataset with Depth Information

3.1. Reading Material

3.1.1. Text Acquisition Pipeline

The text acquisition pipeline consists of three parts, raw text collection, data cleaning, and data filtering. The entire acquisition process is shown in Figure 1. For raw text collection, to ensure the performance of the dataset, it is essential to select raw text from authentic and widely circulated sources with high impact, relevance to daily life, and coverage of a wide range of domains. Based on these principles, we have made a selection of material from various topics such as social media, news, novels, and poetry.

After this, we acquire a large number of raw texts which need to be cleaned first. The data cleaning process includes several steps such as sentence segmentation based on punctuation marks like question marks, semicolons, periods, and exclama-

Table 2: Number of sentences from different data sources, sentence types, and the percentage of the total number of sentences.

Material Categorization		#Sentences	Percentage (%)
Data Source	Social Media	6,229	49.90
	News	3,135	25.11
	Book/Novel	2,524	20.22
	Poetry	596	4.77
Sentence Types	Statement sentences	9,573	76.68
	Question sentences	1,714	13.73
	Exclamation sentences	1,197	9.59



Figure 2: The main acquisition interface of the C#-based data acquisition system, which simultaneously acquires RGB color images, depth images and audio signals.

tion points. Special symbols such as colons, square brackets, English letters, and sensitive words are filtered out. Roman and Arabic numerals are converted to Mandarin expressions, and sentences that are too long or too short are filtered. This process ensures that the raw texts are clean and ready for use in the after accurate analysis.

After data cleaning, we utilize an algorithm to calculate the score of each sentence and select data by sorting the scores. The algorithm combines some metrics commonly found in Mandarin corpora, such as di-phones and tri-phones coverage, tone combination, etc. To be specific, our algorithm weighs the frequency of each metric, attributing a lower score to those with more frequent occurrences. The final score of the sentence is derived from the sum of the factor scores of each sentence. In particular, the types of sentence is taken into account, as the speaker usually shows some natural emotion when reading exclamatory or interrogative sentences.

3.1.2. Manual Selection

After this, all the remaining texts are manually filtered to get the final required reading material. In this process, sentences that cannot be recognized by the machine for sensitive words are removed by manual screening, including sentences with uncommon names, immoral or containing extremely negative messages, etc. The distribution of material categories and the quantities used and the percentage of the total number of sentences are described in Table 2.

3.2. Speakers

This dataset consists of 64 speakers, comprising 35 males and 29 females, whose ages range from 19 to 27, with an average




Figure 3: Recording environment

age of 23 years. Our criteria for speaker selection are as follows. Firstly, the speaker must not have any constant speech or hearing impairment. Secondly, the gender ratio should be close to 1:1. Finally, the speaker should cover most of the Mandarin-speaking areas as much as possible. The division of speakers according to gender and birthplace is presented in Table 3.

For research purposes, speaker information including age, gender, hometown, and the presence or absence of an accent was included in our dataset. All speakers signed consent and agreed with the public use of recorded data for speech recognition related research.

3.3. Recording Setup

3.3.1. Recording Preparation

As the primary data collection device, Azure Kinect is specially designed for researchers or developers. In the data acquisition, the captured video is consisting of a 30 fps RGBA image stream of 1920×1080 pixels and a 30 fps depth image stream of 640×576 pixels of the narrow field of view (NFoV) unboxed modal. Also, the microphone array of Kinect was used as the audio device.

To capture video and audio effects, we have developed a related capture system based on C#. The system consists of two parts, the information acquisition page, and the data recording page as shown in Figure 2. On the information acquisition page, the age, gender, hometown, and whether the speaker has an accent are collected. In the formal recording, the speaker reads 200 sentences of text randomly selected from the reading material.

3.3.2. Environment Setting

The data recording process takes place in a professional soundproof studio that effectively eliminates sound echoes. The ac-

Table 3: Gender and regional distribution of speakers.

Gender	Hometown Area	#People	Total
Male	North China	22	35
	South China	13	
Female	North China	23	29
	South China	6	

quisition equipment used consisted of a monitor, the Kinect placed on top of the monitor via a stand, as well as a main-frame running the acquisition system. A white curtain is set up as a background and a fixed speaker position (0.6m-0.7m) to ensure that the depth images are captured correctly. The specific environmental setup is shown in Figure 3.

3.4. Data Annotation

First, we use the open-source dictionary generation model grapheme to phoneme (G2P)¹ to generate a dictionary of text-to-phoneme sequences. After this, a companion TextGrid file with recorded phoneme durations are generated for all captured audio data using the speech-forced alignment tool Montreal Forced Aligner (MFA)².

Since the mkv file with RGB image stream, depth image stream, and IR image stream is captured during the data recording process, the image frames need to be extracted and the images need to be named using their timestamps to align with the audio. In particular, in Figure 4(b), a Look Up Table (LUT) is used to transform the depth image for a more intuitive presentation.

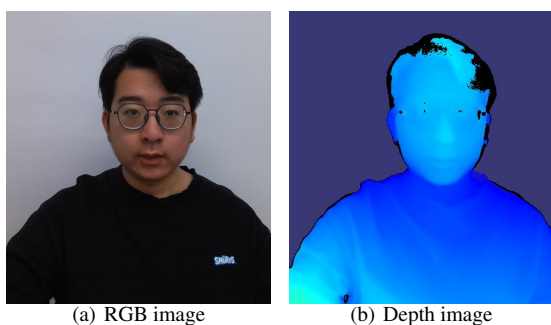


Figure 4: Visual information captured by the Kinect.

4. Baseline Experiment

In this section, we conducted several experiments to assess the efficacy of various modalities in our dataset for speech recognition tasks, by employing a model based on the open-source Wav2Vec 2.0 [32] and ResNet-18 [33].

4.1. Data Pre-processing

For audio signals, we resample the captured 48kHz audio data to 16kHz, 16-bit, mono-channel. For visual signals, the video is fixed at 25 frames per second to synchronize with the audio and also to avoid the risk of possible frame loss during recording. In

¹<https://github.com/cmuspinx/g2p-seq2seq>

²<https://mfa-models.readthedocs.io/en/latest/>

Table 4: The CER for different input modals on test set.

Features	CER (%)
Audio-Only	11.53
Audio-Visual (RGB only)	8.96
Audio-Visual (RGB + Depth)	8.78

addition, the lips are selected as the ROI to extract visual features. By using the pre-trained face landmark detection model in Dlib to detect face landmarks [34], the lip part of each image is cropped. Since the speaker is required not to move significantly during the recording, the lips are positioned every half second (15 frames) to speed up the interception. Finally, both the depth images and the color images, are resized to 64×64 pixels.

4.2. Baseline Model

We employ the Mandarin-Wav2Vec2³ open-source model, which was pre-trained on AISHELL-2 [35] and evaluated on AISHELL-1 [36], as the acoustic model to extract audio features. For visual signals, The ResNet-18, which was pre-trained on ImageNet, is used as an image recognition model to extract features of both depth and color images.

The audio features and visual features are then input into a Bi-directional Long Short-Term Memory (BiLSTM) network [37] and collocate the output features. And then as an input to Connectionist Temporal Classification (CTC) decoder to transcribe the text [38].

4.3. Result and Analysis

We divided the training, validation, and test set by random selection in the ratio of 8:1:1, under the condition of ensuring gender-balanced. We trained the Wav2Vec-based multimodal on the train set of our data with three different settings by controlling the combination of image features and audio features: (i) audio-only, (ii) audio-visual (RGB), and (iii) audio-visual (RGB + Depth). The character error rate (CER) is used as an evaluation metric, which is calculated by adding the number of substituted (S), inserted (I), and deleted (D) characters together and then dividing them by the total number of characters (N): $CER = \frac{S+I+D}{N} \times 100$. As shown in Table 4, the performance of the audio-only speech recognition is not good mainly due to challenging reading material, and the error rate decreases when combining color and depth images features, which also represents that our proposed corpus is effective in the AVSR task. Importantly, the use of depth information can also help the effectiveness of the Mandarin AVSR task.

5. Conclusions

In the work, we present a new multi-modal audio-visual corpus for Mandarin named MAVD, including about 12k sentences read by 64 speakers. Our corpus focuses on contributing to the Mandarin audio-visual corpus community. As far as we know, this is the first open large-scale Mandarin audio-visual dataset with depth information. In the future, we will continue to expand this work such as using several devices to form a multi-views corpus, working with more recording environments, etc.

³<https://github.com/kehanlu/Mandarin-Wav2Vec2>

6. References

- [1] T. Sakai and S. Doshita, "The Automatic Speech Recognition System for Conversational Sound," *IEEE Transactions on Electronic Computers*, vol. EC-12, no. 6, pp. 835–846, 1963.
- [2] P. Sanderson, "Cognitive work analysis and the analysis, design, and evaluation of human-computer interactive systems," in *Proceedings 1998 Australasian Computer Human Interaction Conference. OzCHI'98 (Cat. No.98EX234)*, 1998, pp. 220–227.
- [3] E. M. Golonka, A. R. Bowles, V. M. Frank, D. L. Richardson, and S. Freynik, "Technologies for foreign language learning: A review of technology types and their effectiveness," *Computer assisted language learning*, vol. 27, no. 1, pp. 70–105, 2014.
- [4] M. H. Raskind and E. L. Higgins, "Speaking to read: The effects of speech recognition technology on the reading and spelling performance of children with learning disabilities," *Annals of Dyslexia*, vol. 49, pp. 251–281, 1999.
- [5] L. Liu and G. Feng, "A pilot study on mandarin chinese cued speech," *American Annals of the Deaf*, vol. 164, no. 4, p. 496–518, 2019.
- [6] S. P. Panda, "Automated speech recognition system in advancement of human-computer interaction," in *2017 ICCMC*, 2017, pp. 302–306.
- [7] H. MCGURK and J. MACDONALD, "Hearing lips and seeing voices," *Nature*, vol. 264, no. 5588, pp. 746–748, Dec. 1976. [Online]. Available: <https://doi.org/10.1038/264746a0>
- [8] L. Liu, G. Feng, D. Beutemps, 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.
- [9] C. Ding and D. Tao, "Robust Face Recognition via Multimodal Deep Face Representation," *IEEE Transactions on Multimedia*, vol. 17, no. 11, pp. 2049–2058, 2015.
- [10] S. Dupont and J. Luetin, "Audio-visual speech modeling for continuous speech recognition," *IEEE Transactions on Multimedia*, vol. 2, no. 3, pp. 141–151, 2000.
- [11] S. Themos and G. Potamianos, "Audio-visual speech activity detection in a two-speaker scenario incorporating depth information from a profile or frontal view," in *2016 IEEE Spoken Language Technology Workshop (SLT)*, 2016, pp. 579–584.
- [12] C. Xu, Y. Wang, T. Tan, and L. Quan, "Depth vs. intensity: which is more important for face recognition?" in *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, vol. 1, 2004, pp. 342–345 Vol.1.
- [13] L. Liu, G. Feng, and D. Beutemps, "Automatic temporal segmentation of hand movements for hand positions recognition in french cued speech," in *2018 IEEE ICASSP*. IEEE, 2018, p. 3061–3065.
- [14] L. Liu, T. Hueber, G. Feng, and D. Beutemps, "Visual recognition of continuous cued speech using a tandem cnn-hmm approach," in *Interspeech*, 2018, p. 2643–2647.
- [15] I. Matthews, T. Cootes, S. Cox, R. Harvey, and J. A. Bangham, "Lipreading using shape, shading and scale," in *AVSP'98 International Conference on Auditory-Visual Speech Processing*, 1998.
- [16] S. Cox, R. Harvey, and Y. Lan, "The Challenge of Multispeaker Lip-Reading," in *Proc of International Conference on Auditory-visual Speech Processing*, 2008.
- [17] C. C. Chibelushi, F. Deravi, and J. Mason, "A review of speech-based bimodal recognition," *IEEE Transactions on Multimedia*, vol. 4, no. 1, pp. 23–37, 2002.
- [18] M. Cooke, J. Barker, S. 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 Pt 1, p. 2421–2424, November 2006.
- [19] Y. Xu, L. Du, G. Li, P. Wu, and X. Zhang, "Chinese audiovisual bimodal speech database CAVSR1.0," in *Proc. Int. Symp. Chin. Spoken Lang. Process.*, 2000, pp. 98–101.
- [20] X. Lin, H. Yao, and Hong, "A sentence-level lip-reading-based corpus and its slice and dice algorithm," *Computer Engineering and Applications*, vol. 41, no. 3, pp. 174–177, 2005.
- [21] L. Liu, G. Feng, and D. Beutemps, "Automatic dynamic template tracking of inner lips based on cnf," in *2017 ICASSP*. IEEE, 2017, p. 5130–5134.
- [22] X. Lin, H. Yao, X. Hong, and Q. Wang, "HIT-AVDB-II: A new multi-view and extreme feature cases contained audio-visual database for biometrics," in *11th Joint International Conference on Information Sciences*. Atlantis Press, 2008, pp. 357–363.
- [23] S. Yang, Y. Zhang, D. Feng, M. Yang, C. Wang, J. Xiao, K. Long, S. Shan, and X. Chen, "LRW-1000: A Naturally-Distributed Large-Scale Benchmark for Lip Reading in the Wild," in *2019 14th IEEE International Conference on Automatic Face Gesture Recognition (FG 2019)*, 2019, pp. 1–8.
- [24] M. McLaren, L. Ferrer, D. Castan, and A. Lawson, "The Speakers in the Wild (SITW) Speaker Recognition Database," in *Interspeech 2016*, 2016.
- [25] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A Large-Scale Speaker Identification Dataset," in *Interspeech 2017*. ISCA, aug 2017.
- [26] C. Zhang, G. Huang, L. Liu, S. Huang, Y. Yang, X. Wan, S. Ge, and D. Tao, "Webuav-3 m: A benchmark for unveiling the power of million-scale deep uav tracking," *IEEE TPAMI*, 2022.
- [27] E. Patterson, S. Gurbuz, Z. Tufekci, and J. Gowdy, "CUAVE: A new audio-visual database for multimodal human-computer interface research," in *2002 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 2, 2002, pp. II–2017–II–2020.
- [28] S. Isobe, R. Hirose, T. Nishiwaki, T. Hattori, S. Tamura, Y. Gotoh, and M. Nose, "GAMVA: A Japanese Audio-Visual Multi-Angle Speech Corpus," in *2021 24th Conference of the Oriental CO-COSDA International Committee for the Co-ordination and Standardisation of Speech Databases and Assessment Techniques (O-COCOSDA)*, 2021, pp. 134–139.
- [29] A. Vorwerk, X. Wang, D. Kolossa, S. Zeiler, and R. Orglmeister, "WAPUSK20-A Database for Robust Audiovisual Speech Recognition," in *LREC*. Citeseer, 2010.
- [30] G. Galatas, G. Potamianos, D. Kosmopoulos, C. McMurrough, and F. Makedon, "Bilingual corpus for AVASR using multiple sensors and depth information," in *Auditory-Visual Speech Processing 2011*, 2011.
- [31] N. Ahmed, "RGB-D dynamic facial dataset capture for visual speech recognition," in *2019 International Conference on Image and Video Processing, and Artificial Intelligence*, vol. 11321. SPIE, 2019, pp. 42–46.
- [32] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [34] D. E. King, "Dlib-ml: A machine learning toolkit," *The Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.
- [35] J. Du, X. Na, X. Liu, and H. Bu, "Aishell-2: Transforming mandarin asr research into industrial scale," *arXiv preprint arXiv:1808.10583*, 2018.
- [36] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHELL-1: An open-source Mandarin speech corpus and a speech recognition baseline," in *2017 20th O-COCOSDA*, 2017, pp. 1–5.
- [37] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *2013 IEEE ICASSP*. Ieee, 2013, pp. 6645–6649.
- [38] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.