



# AS-70: A Mandarin stuttered speech dataset for automatic speech recognition and stuttering event detection

Rong Gong<sup>1\*</sup>, Hongfei Xue<sup>2\*</sup>, Lezhi Wang<sup>1</sup>, Xin Xu<sup>3</sup>, Qisheng Li<sup>4</sup>, Lei Xie<sup>2</sup>, Hui Bu<sup>3</sup>, Shaomei Wu<sup>4</sup>, Jiaming Zhou<sup>5</sup>, Yong Qin<sup>5</sup>, Binbin Zhang<sup>6</sup>, Jun Du<sup>7</sup>, Jia Bin<sup>1</sup>, Ming Li<sup>8</sup>

<sup>1</sup>StammerTalk. <sup>2</sup>ASLP@NPU, Northwestern Polytechnical University. <sup>3</sup>AIShell Inc  
<sup>4</sup>AImpower. <sup>5</sup>Nankai University. <sup>6</sup>WeNet Open Source Community  
<sup>7</sup>University of Science and Technology of China. <sup>8</sup>Duke Kunshan University

rong.gong@stammertalk.net, hfxue@mail.nwpu.edu.cn

## Abstract

The rapid advancements in speech technologies over the past two decades have led to human-level performance in tasks like automatic speech recognition (ASR) for fluent speech. However, the efficacy of these models diminishes when applied to atypical speech, such as stuttering. This paper introduces AS-70, the first publicly available Mandarin stuttered speech dataset, which stands out as the largest dataset in its category. Encompassing conversational and voice command reading speech, AS-70 includes verbatim manual transcription, rendering it suitable for various speech-related tasks. Furthermore, baseline systems are established, and experimental results are presented for ASR and stuttering event detection (SED) tasks. By incorporating this dataset into the model fine-tuning, significant improvements in the state-of-the-art ASR models, e.g., Whisper and Hubert, are observed, enhancing their inclusivity in addressing stuttered speech.

**Index Terms:** mandarin stuttered speech dataset, speech recognition, stuttering event detection

## 1. Introduction

Stuttering is a speech impediment that affects around 1% of the world's population [1]. The impact of the stuttering on people who have such a problem could be on both social functioning and mental aspects. Stuttering hinders daily oral communication with speech repetition, prolongation, blocks, and secondary behaviors, such as body movements and facial grimaces [2].

Early children who stutter (CWS) have around 80% of a chance to recover within four years of the onset [3]. Early intervention by speech-language therapists (SLP) is crucial for increasing the recovery rate [4]. Compared with North American or Western European countries where the practice of SLP has been evolving for quite some time, the related practices in Mainland China are still in their infancy. Due to the lack of professionals, many families of children who stutter cannot receive timely diagnosis. While adding more professionals, automatic stuttering diagnosis can also meet the needs of some families.

When a person's stuttering continues into adulthood, the chance of full recovery becomes very low, and there is a high probability that stuttering will accompany them throughout their lives [5]. To alleviate their communication barriers or, in a broad sense, eliminate social discrimination, relevant products need to be designed to be inclusive to meet their needs.

With the rise of smart home devices and chatbot technology, e.g., Alexa and ChatGPT, voice-user interfaces have become indispensable tools in many of our lives. Although current ASR

systems can handle fluent speech well [6], they encounter difficulties in recognizing stuttered speech [7, 8]. The underlying reasons for this challenge may involve various factors, including insufficient data or a lack of awareness regarding the necessity to develop systems tailored for PWS.

### 1.1. Related works

Due to the highly sensitive and personal nature of the data, large-scale stuttering speech datasets collected by companies are typically not made openly accessible to the research community [7, 9, 10, 8, 11]. On the other hand, openly accessible datasets curated by academic researchers, such as FluencyBank [12] and UCLASS [13], are often small and primarily intended for use in SLP resources. Larger open datasets, while available, may require additional completeness in terms of annotation and authenticity. For instance, the Sep-28k dataset [14], comprising 28k 3-second podcast audio clips, only provides stuttering labels for the SED task but lacks accompanying text transcription. Meanwhile, the LibriStutter [15] dataset includes text transcriptions, but it is artificially generated from recordings of fluent speech. Furthermore, it is noteworthy that all these open stuttering speech datasets are limited to two Western languages [16].

The primary focus of most ASR research on stuttered speech is directed towards predicting semantic rather than verbatim transcription. Lea et al. [7] observed enhanced performance by modifying the ASR decoder, specifically by increasing the language model weight and imposing a higher penalty for word insertions. Shonibare et al. [8] introduced a frame-level stuttering classifier, termed Detect-and-Pass, designed for the RNNT ASR model. Zhang et al. [11] demonstrated performance improvement through fine-tuning the ASR model using synthesized stuttering speech. Alharbi et al. [17] utilized ASR to generate orthographic transcriptions that include stuttering words.

SED involves identifying instances of stuttering in speech audio clips. Research in this domain predominantly focused on four openly available datasets: Sep-28k [14], FluencyBank [12], LibriStutter [15], and KSoF [16]. Different neural network architectures, including ConvLSTM [14], StutterNet [18], and SE-Resnet [15], were developed specifically for this task. Additionally, machine learning techniques like multi-task learning [14, 19] were employed to enhance the accuracy of stuttering detection.

The subsequent sections of this paper showcase our primary contribution - the AS-70 dataset, which can be accessed from

\*Equal contribution.

this download location<sup>1</sup>. In Section 2, we elaborate on the data collection, annotation process, and conduct a descriptive analysis that compares AS-70 with other openly available stuttered speech datasets. In Section 3, we establish baseline systems for the ASR and SED tasks using the AS-70 dataset.

## 2. Dataset

### 2.1. Participants and recording sessions

The AS-70 dataset was recorded between January 2023 and October 2023 by two of the authors, both adults who stutter (AWS) and native Mandarin speakers. A total of 70 native Mandarin AWS took part in the recording sessions, with 24 of them being female, establishing a male-to-female ratio of 1.9:1. There was no gender selection process as all willing candidates were accepted to participate in the recording sessions.

Each participant engaged in a recording session lasting up to one hour, comprising two parts: conversation and voice command reading. Conversations were conducted through online interviews using platforms like Zoom or Tencent Meet, aiming to capture spontaneous speech on diverse topics. The interviewer, one of the two authors, posed questions based on a prepared list, with the flexibility to introduce impromptu questions as needed.

In the voice command reading part, participants were tasked with reading a set of 200 commands, categorized into car navigation and smart home device interaction. To ensure variety, a new set of 200 commands was introduced for every 25 participants, resulting in a dataset featuring a total of 600 unique commands. Participants were encouraged to employ the Voluntary Stuttering technique, deliberately introducing stuttering.

### 2.2. Annotation

The dataset annotation process involved 15 non-PWS annotators and 5 quality controllers (QCs). Among the 5 QCs, 4 are non-PWS, and 1 is the PWS author (QC0). All annotators and QCs are experienced Mandarin speech annotators, with some having prior experience annotating atypical speech. To ensure annotation quality, QC0 conducted a training session prior to the commencement of the annotation process. In this session, the 19 non-PWS annotators and QCs were initially tasked with annotating a few examples of stuttered speech. Following this, QC0 provided feedback, and the 19 annotators were then asked to annotate new examples, incorporating the lessons learned from the feedback. This iterative process was repeated multiple times until no further feedback was deemed necessary.

Five types of stuttering were specified by the annotation guidelines, including:

- **[]: Word/phrase repetition.** Designated for marking entire repeated character or phrase.
- **/b: block.** Gasps for air or stuttered pauses.
- **/p: prolongation.** Elongated phoneme.
- **/r: sound repetition.** Repeated phoneme that do not constitute an entire character.
- **/i: interjections.** Filler characters due to stuttering e.g., ‘嗯’, ‘啊’, or ‘呃’. Notably, naturally occurring interjections that don’t disrupt the speech flow are excluded.

The annotation was performed in a verbatim manner, with stuttering labels embedded as markups. An example of the

annotated transcription can be “嗯/i/p, 我[我我]的名/b字是 小/r明。”. The character ‘嗯’ is an interjection and also is prolonged. The character ‘我’ is repeated twice additionally. A block annotation is applied to the character ‘名’. Furthermore, sound repetition is identified in one of the phonemes of ‘小’. It is noteworthy that a single character may carry multiple labels. Additional examples of audio and annotations can be accessed through this link<sup>2</sup>.

The conversation and command reading parts of each session were annotated separately by different annotators. Subsequently, one of the 4 non-PWS QCs conducted a cross-check of the annotations. To ensure consistent annotation, QC0 then performed a final cross-check of all session annotations. The annotators noted that this annotation process required approximately three times more time compared to annotating fluent speech.

Personal Identifiable Information (PII) such as name, address, birth date, and workplace has been redacted from the transcription. The audio segments containing PII have been muted. A data collection agreement, adhering to Guidance for personal information security impact assessment [20], was signed between the data collectors and the participant.

### 2.3. Descriptive analysis

Table 1 presents a comparison of dataset scales between AS-70 and 4 other open datasets. A total duration of 48.8-hours speech data were included in AS-70 dataset from 70 recording sessions. Excluding two interviewers, each participant contributed on average 457 utterances and 33.0 minutes of total speech, with an average of 141 utterances and 17.8 minutes for conversations and an average of 316 utterances and 15.23 minutes for voice command reading.

Table 1 indicates that the AS-70 dataset is the largest among open datasets in terms of duration and participants. When compared to the previously largest open dataset, Sep-28k [14], which is annotated at a 3-second clip level without transcription, the AS-70 dataset is twice as large, annotated at the character level, and includes verbatim transcription. Moreover, the AS-70 dataset encompasses both spontaneous conversational and voice command reading speech tasks, and stands out as the only non-Western language open dataset.

Table 2 depicts the distribution of the stuttering events across four datasets. The *Average Stuttering Rate* is calculated by dividing the total count of stuttering events by the duration of the speech. The *Event Type Distribution* is calculated by the formula “number of labels of a certain stuttering event type / overall number of all stuttering event labels”. It’s important to note that a direct comparison between Sep-28k and AS-70 may not be entirely equitable. This is because AS-70 dataset is annotated on the character level, offering finer granularity compared to the clip level annotation of Sep-28k. This difference may partially account for the higher average stuttering rate observed in the AS-70 conversation part compared to Sep-28k.

Additionally, there is a significant shift towards more word and phrases repetitions and fewer sound repetitions in the AS-70 dataset, signaling potential phonological differences between stuttering in Chinese and in English. Meanwhile we notice that SEP-28k dataset contains almost twice more interjections than in AS-70 conversations, which could be attributed to different definitions of interjections in these two datasets. While SEP-28k includes any filler words, such as “um,” “uh,” and

<sup>1</sup>[https://www.aishelltech.com/AISHELL\\_6A](https://www.aishelltech.com/AISHELL_6A)

<sup>2</sup><https://stammertalk.github.io/interspeech2024-page>

Table 1: *Dataset scale and scope as characterized by speech duration (Duration), the number and types of speakers (Speakers), whether it provides speech transcription (Transcription), types of speaking tasks (Tasks), and Language. \* Limited to the transcribed portion of the dataset. \*\* Including two interviewers. Abbreviations: AWS - adults who stutter; CWS - children who stutter; PWS - people who stutter.*

Dataset	Duration	Speakers	Transcription	Tasks	Language
LibriStutter [15]	20 hrs	50 non-PWS	Yes	audiobook	English
UCLASS* [13]	53 mins	25 CWS	Yes	conversation	English
SEP-28k [14]	23 hrs	not reported	No	podcast	English
FluencyBank* [12]	3.5 hrs	32 AWS	Yes	conversation and reading	English
KSoF [16]	4.6 hrs	37 PWS	No	spontaneous and reading	German
<b>AS-70</b>	<b>48.8 hrs</b>	<b>72 AWS**</b>	<b>Yes (verbatim)</b>	<b>conversation, voice commands</b>	<b>Mandarin</b>

Table 2: *Distribution of annotated stuttering events. FluencyBank’s annotation was done in another work [14]. AS-70 Conversation includes two interviewers.*

	Avg. Stuttering Rate (per minute)	Total Stuttering Events	Event Type Distribution %				
			[]	/b	/p	/r	/i
LibriStutter [15]	12.5	15,000	20.0	20.0	20.0	20.0	20.0
SEP-28k [14]	12.26	17,267	16.0	19.5	16.3	13.6	34.6
FluencyBank[12]	13.88	2,875	15.0	14.9	11.8	19.1	39.3
KSoF [16]	13.05	3,602	6.0	32.2	18.7	22.9	20.2
<b>AS-70 Conversation</b>	15.58	29,017	42.7	6.9	19.8	8.7	21.9
<b>AS-70 Command</b>	8.11	8,636	53.0	8.4	16.9	15.8	5.9

Table 3: *Stuttering event distribution of SED task datasets. no-dis: no-disfluency. AS-70 Conversation includes two interviewers.*

	Avg. length (s)	Number of clips	Event Type Distribution %					no-dis
			[]	/b	/p	/r	/i	
SEP-28k [14]	3.00	28,177	9.8	12.0	10.0	8.3	21.2	56.9
FluencyBank [12]	3.00	4,144	10.38	10.33	8.16	13.25	27.27	54.10
KSoF [16]	3.00	5,597	3.88	20.74	12.02	14.76	12.97	24.75
<b>AS-70 Conversation</b>	5.69	19,654	38.67	8.45	22.12	10.77	24.80	36.31
<b>AS-70 Command</b>	2.87	22,299	14.09	2.87	5.31	5.08	1.76	79.61

“you know” as stuttering interjections, our annotation excludes natural interjections that blend into the speech flow.

In line with previous studies [14, 16], we split the utterances of the AS-70 dataset into short clips for the SED task. Table 3 provides a comparison between the AS-70 dataset and other three datasets. Notably, AS-70 dataset contains 49% more clips than the Sep-28k dataset.

To compute the *Event Type Distribution*, we employ the consistent formula as in the work [14]: “number of clips containing a certain stuttering event type / total number of clips”. An interesting observation is that the conversation part of the AS-70 exhibits almost 4 times word/phrase repetitions and 2.2 times prolongation compared to Sep-28k. As a relatively easy task for our participants, the command reading task is reflected by a high no-disfluency rate, low interjection, and block rate.

### 3. Experiments

The objective of the work in this section is not to attain state-of-the-art results but to highlight the limitations of specific pre-trained models through an evaluation of the AS-70 dataset. Additionally, we aim to demonstrate the enhancements achievable by integrating our dataset into the model training process.

#### 3.1. Data partition

As highlighted in the study of Bayerl et al. [21], the data partition method can heavily impact the reliability of experimental results. A data partition that neglects speaker exclusivity can cause a model to learn speaker-specific traits, resulting in overly

optimistic results. We ensure that speakers with varying degrees of stuttering severity are distributed across each partition.

The stuttering severity of a speaker is measured using the stuttering rate (*SR*), calculated by dividing the number of stuttering events by the number of non-stuttering characters in the transcription. E.g., the annotation “嗯/i/p, 我[我我]的名/b字是小/r明。” would lead to a stuttering rate of 71.42% as there are 5 stuttering events and 7 non-stuttering characters. The stuttering severity is classified as mild:  $SR \leq 7\%$  for 45 participants, moderate:  $7\% < SR \leq 12\%$  for 16 participants, and severe:  $SR > 12\%$  for 9 participants. The speaker numbers of train/development/test partitions can be found in the link<sup>2</sup>. The data of two PWS interviewers is included in the train partition.

#### 3.2. ASR

We conduct an in-depth evaluation of ASR performance using the AS-70 dataset, focusing on three distinct models: Conformer [22], HuBERT [23], and Whisper [6]. These models are selected to represent a comprehensive range of approaches to ASR, including supervised end-to-end ASR, self-supervised pretraining, and large-scale semi-supervised methods. To emphasize semantic transcription, the annotations undergo preprocessing to exclude stuttering event labels, stuttering characters, and punctuation.

**Conformer.** This system is based on Wenet [24]. U2++ is a unified two-pass framework with bidirectional attention decoders, which includes the future contextual information by a right-to-left attention decoder to improve the representative

Table 4: The CER(%) results of different ASR model.

Model	Pre-trained Dataset	Fine-tuned Dataset	mild	moderate	severe	conversation	command	all
Conformer	WenetSpeech	-	11.41	17.81	33.21	11.82	20.40	15.00
	WenetSpeech	AS-70	5.20	<b>8.00</b>	9.25	<b>7.96</b>	3.32	<b>6.24</b>
Hubert-large	WenetSpeech (unlabel)	AISHELL-1	16.28	22.86	26.76	16.80	23.79	19.06
	WenetSpeech (unlabel)	AS-70	6.25	9.67	<b>7.85</b>	9.75	<b>2.03</b>	7.25
Whisper-large-v2	Whisper	-	14.50	28.50	95.33	17.83	46.85	27.20
	Whisper	AS-70	<b>5.18</b>	13.46	18.92	10.19	5.74	8.75

ability of the shared encoder and the performance during the rescoring stage. We show the results of the u2++ conformer model pre-trained using WenetSpeech dataset [25] and the results of the fine-tuned model with AS-70 train partition.

**Hubert.** HuBERT has demonstrated remarkable efficacy in ASR through self-supervised learning. We employ the Chinese Hubert Large model as our pre-trained experiment system [26]. We use the results of AISHELL-1 [27] fine-tuning as a baseline to compare with the model fine-tuned by the AS-70 dataset.

**Whisper.** Whisper exhibits exceptional proficiency across multiple languages, developed through semi-supervised training on a large-scale dataset. Our baseline employs the Whisper large-v2 model for direct inference. Subsequent efforts are channeled towards fine-tuning the Whisper model, specifically with the AS-70 dataset.

**Results and discussion.** Table 4 indicates the above models’ Character Error Rate (CER) results against the AS-70 test set. These results communicate a clear message: even though the above models have performed impressively with generic data, their proficiency in dealing with stuttering data seemed to falter. However, after fine-tuning using AS-70 data, all models exhibit commendable improvement. The conformer with AS-70 fine-tuning takes the lead, possibly owing to extensive labeled data in the conformer pretrain stage. On the other hand, Whisper seems to falter on severe data. Upon examination, we notice that Whisper tends to generate an overabundance of repeat labels when decoding utterances from severe-level PWS.

An interesting finding from these experiments is the relatively low CER registered on command data processed by the AS-70 fine-tuned models, possibly due to the training set containing identical command text as that in the test set, with only variations being the speakers. We plan to address this limitation in the future work by considering the factor of command text in the data partition process. Examples of ASR results can be found in the link<sup>2</sup>.

### 3.3. SED

The evaluation in SED is undertaken by establishing a fundamental benchmark predicated on random guessing, thus providing a comparative baseline for subsequent analyses. We employ some reputed SED methods, such as StutterNet [18], ConvLSTM [14], Conformer and wav2vec2.0 [28], to assess the effectiveness of these methods.

**StutterNet.** StutterNet [18] uses a time-delay neural network suitable for capturing contextual aspects of the disfluent utterances, which is trained on the MFCC input features. We reproduce the structure in StutterNet [18] with 12.2M parameters. We use multi-task learning with two output branches: a fluent/dysfluent prediction and a soft prediction for each of the five event types.

**ConvLSTM.** ConvLSTM [14] input is a set of 40-dimensional mel-filterbank energy features. Feature maps from

the convolution layer are combined after batch normalization and fed to three LSTM layers, which result in a 1.6M parameter size. We use the same multi-task learning as in StutterNet.

**Conformer.** The same Wenet conformer [22] encoder architecture mentioned in the section 3.2 is applied. However, due to the small amount of training data, we use 3 conformer blocks, which result in 9.7M parameter size. We use single-task learning, which is to predict the five event types. The multi-label soft margin loss is used for the model training.

**Wav2Vec2.0.** Referring to Bayerl et al.’s approach [28], we fine-tune the wav2vec2.0 base model. The model we used for our initial experiments (Chinese-wav2vec2-base) [29] is pre-trained in an unsupervised manner on the WenetSpeech corpus. We use the same single-task learning as in Conformer.

Table 5: The F1-score(%) of different SED model. We report results for each label individually.

Model	/p	/b	/r	[]	/i
Random guess	15.65	8.07	15.34	31.7	17.03
StutterNet [18]	61.07	33.33	47.81	50.12	58.82
ConvLSTM [14]	33.30	18.22	30.19	64.02	46.70
Conformer [22]	66.77	30.94	46.84	65.49	73.10
Wav2vec2.0 [28]	<b>70.48</b>	<b>42.51</b>	<b>65.76</b>	<b>78.48</b>	<b>83.80</b>

**Results.** Table 5 indicates the F1-score of the above models for the five stuttering event types. All models exceed the random guess baseline in each event type. ConvLSTM underperformed in comparison, which is attributed to its limited parameter count of 1.6m. StutterNet and Conformer display intermediate results, whereas wav2vec 2.0 emerged as the most proficient, benefiting significantly from its pre-training on a large dataset. It is worth noting that the block (/b) performance is substantially worse than the other stuttering types, probably attributed to the fact that blocks are often represented solely by silence in the speech. Similar findings are reported in Lea et al. [7] and Bayerl et al. [21]. For a detailed SED results on stuttering severities and speech tasks, please refer to the link<sup>2</sup>.

## 4. Conclusion

This paper introduces the AS-70 dataset, which is the first publicly available Mandarin stuttered speech dataset and is notable for being the largest in its category. AS-70 consists of conversational and voice command reading speech recordings with verbatim manual transcription, making it suitable for various speech-related tasks. In addition, baseline systems have been established, and experimental results are presented for ASR and SED tasks, demonstrating significant improvements in state-of-the-art ASR models by incorporating this dataset into model fine-tuning processes. We hope that AS-70 can assist in detecting stuttering and aiding people who stutter in developing speech interaction systems.

## 5. References

- [1] E. Yairi, N. Ambrose, and N. Cox, "Genetics of stuttering," *Journal of Speech, Language, and Hearing Research*, vol. 39, pp. 771–784, 1996.
- [2] J. Prasse and G. Kikano, "Stuttering: An overview," *American family physician*, vol. 77, pp. 1271–6, 2008.
- [3] E. Yairi and N. G. Ambrose, "Early childhood stuttering i," *Journal of Speech, Language, and Hearing Research*, vol. 42, pp. 1097–1112, 1999.
- [4] M. Onslow and S. O'Brian, "Management of childhood stuttering," *Journal of Paediatrics and Child Health*, vol. 49, pp. E112–E115, 2013.
- [5] J. O. Boyce, V. E. Jackson, O. van Reyk, R. Parker, A. P. Vogel, E. Eising, S. E. Horton, N. A. Gillespie, I. E. Scheffer, D. J. Amor, M. S. Hildebrand, S. E. Fisher, N. G. Martin, S. Reilly, M. Bahlo, and A. T. Morgan, "Self-reported impact of developmental stuttering across the lifespan," *Developmental Medicine & Child Neurology*, vol. 64, pp. 1297–1306, 2022.
- [6] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*, vol. 202. PMLR, 2023, pp. 28 492–28 518.
- [7] C. Lea, Z. Huang, J. Narain, L. Tooley, D. Yee, D. T. Tran, P. Georgiou, J. P. Bigham, and L. Findlater, "From user perceptions to technical improvement: Enabling people who stutter to better use speech recognition," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2023.
- [8] O. Shonibare, X. Tong, and V. Ravichandran, "Enhancing asr for stuttered speech with limited data using detect and pass," *ArXiv*, vol. abs/2202.05396, 2022.
- [9] B. MacDonald, P.-P. Jiang, J. Cattiau, R. Heywood, R. Cave, K. Seaver, M. Ladewig, J. Tobin, M. Brenner, P. Q. Nelson, J. R. Green, and K. Tomanek, "Disordered speech data collection: Lessons learned at 1 million utterances from project euphonia," in *Interspeech*. ISCA, 2021, pp. 4833–4837.
- [10] V. Mitra, Z. Huang, C. Lea, L. Tooley, P. Georgiou, S. Kajarekar, and J. Bigham, "Analysis and tuning of a voice assistant system for dysfluent speech," in *Interspeech*. ISCA, 2021, pp. 4848–4852.
- [11] X. Zhang, I. Vallés-Pérez, A. Stolcke, C. Yu, J. Droppo, O. Shonibare, R. Barra-Chicote, and V. Ravichandran, "Stutter-tts: Controlled synthesis and improved recognition of stuttered speech," in *NeurIPS 2022 Workshop on SyntheticData4ML*, 2022.
- [12] N. B. Ratner and B. MacWhinney, "Fluency bank: a new resource for fluency research and practice," *Journal of Fluency Disorders*, vol. 56, pp. 69–80, 2018.
- [13] P. Howell, S. Davis, and J. Bartrip, "The university college london archive of stuttered speech (uclass)," *Journal of speech, language, and hearing research : JSLHR*, vol. 52 2, pp. 556–69, 2009.
- [14] C. Lea, V. Mitra, A. Joshi, S. Kajarekar, and J. Bigham, "Sep-28k: A dataset for stuttering event detection from podcasts with people who stutter," in *ICASSP*. IEEE, 2021, pp. 6798–6802.
- [15] T. Kourkounakis, A. Hajavi, and A. Etemad, "Fluentnet: End-to-end detection of stuttered speech disfluencies with deep learning," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2986–2999, 2021.
- [16] S. Bayerl, A. Wolff von Gudenberg, F. Hönig, E. Noeth, and K. Riedhammer, "Ksof: The kassel state of fluency dataset – a therapy centered dataset of stuttering," in *Proceedings of the Language Resources and Evaluation Conference*. European Language Resources Association, 2022, pp. 1780–1787.
- [17] S. Alharbi, A. J. H. Simons, S. Brumfitt, and P. D. Green, "Automatic recognition of children's read speech for stuttering application," in *Workshop on Child, Computer and Interaction*, 2017.
- [18] S. A. Sheikh, M. Sahidullah, F. Hirsch, and S. Ouni, "Advancing stuttering detection via data augmentation, class-balanced loss and multi-contextual deep learning," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, pp. 2553–2564, 2023.
- [19] S. P. Bayerl, D. Wagner, I. Baumann, F. Hönig, T. Bocklet, E. Nöth, and K. Riedhammer, "A stutter seldom comes alone - cross-corpus stuttering detection as a multi-label problem," in *INTERSPEECH*. ISCA, 2023, pp. 1538–1542.
- [20] "GB/T 39335-2020: Information security technology—Guidance for personal information security impact assessment," Standard, 2020.
- [21] S. P. Bayerl, D. Wagner, E. Nöth, T. Bocklet, and K. Riedhammer, "The influence of dataset partitioning on dysfluency detection systems," in *Text, Speech, and Dialogue*. Springer International Publishing, 2022.
- [22] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Interspeech*. ISCA, 2020, pp. 5036–5040.
- [23] W. Hsu, B. Bolte, Y. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [24] Z. Yao, D. Wu, X. Wang, B. Zhang, F. Yu, C. Yang, Z. Peng, X. Chen, L. Xie, and X. Lei, "Wenet: Production oriented streaming and non-streaming end-to-end speech recognition toolkit," in *Interspeech*. ISCA, 2021, pp. 4054–4058.
- [25] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. Zeng *et al.*, "Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition," in *ICASSP*. IEEE, 2022, pp. 6182–6186.
- [26] "chinese-hubert-large," <https://huggingface.co/TencentGameMate/chinese-hubert-large>, accessed: 2024-03-07.
- [27] 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 *Oriental COCODA*, 2017.
- [28] S. P. Bayerl, D. Wagner, E. Nöth, and K. Riedhammer, "Detecting dysfluencies in stuttering therapy using wav2vec 2.0," in *Interspeech*. ISCA, 2022, pp. 2868–2872.
- [29] "chinese-wav2vec2-large," <https://huggingface.co/TencentGameMate/chinese-wav2vec2-base>, accessed: 2024-03-07.