

# ASCEND: A Spontaneous Chinese-English Dataset for Code-switching in Multi-turn Conversation

Holy Lovenia<sup>\*</sup>, Samuel Cahyawijaya<sup>\*</sup>, Genta Indra Winata<sup>\*†</sup>, Peng Xu, Xu Yan, Zihan Liu, Rita Frieske, Tiezheng Yu, Wenliang Dai, Elham J. Barezi, Qifeng Chen, Xiaojuan Ma, Bertram E. Shi, Pascale Fung

The Hong Kong University of Science and Technology

{hlovenia, scahyawijaya, giwinata}@connect.ust.hk

## Abstract

Code-switching is a speech phenomenon occurring when a speaker switches language during a conversation. Despite the spontaneous nature of code-switching in conversational spoken language, most existing works collect code-switching data from read speech instead of spontaneous speech. ASCEND (A Spontaneous Chinese-English Dataset) is a high-quality Mandarin Chinese-English code-switching corpus built on spontaneous multi-turn conversational dialogue sources collected in Hong Kong. We report ASCEND’s design and procedure for collecting the speech data, including annotations. ASCEND consists of 10.62 hours of clean speech, collected from 23 bilingual speakers of Chinese and English. Furthermore, we conduct baseline experiments using pre-trained wav2vec 2.0 models, achieving a best performance of 22.69% character error rate and 27.05% mixed error rate.

**Keywords:** code-switching, corpus, bilingual, speech, dialogue, Mandarin Chinese, English, low-resource

## 1. Introduction

Most of our knowledge about speech recognition and speech generation technologies comes from monolingual read speech data collected in a controlled setting (Panayotov et al., 2015). Monolingual read speech data allows researchers to exercise tight control over the linguistic backgrounds of speakers and the linguistic material (e.g. reading or repeating sounds, words, or sentences). While being highly informative, these monolingual read speech samples do not capture particular actualities of spoken speech (Howell and Kadi-Hanifi, 1991; Blaauw, 1994; Batliner et al., 1995; Li, 2002; Yang and Esposito, 2013; Haynes et al., 2015).

Code-switching is a phenomenon typical of spoken speech, characterized by alternating use of more than one language. It may occur within a single utterance, which is known as intra-sentential code-switching, or between utterance boundaries, which is commonly referred to as inter-sentential code-switching. To address this phenomenon, code-switching corpora in many different languages pairs have been introduced, including but not limited to Indonesian-English (Rizal and Stymne, 2020), Filipino-Spanish (Bautista, 2004), Latin-Irish (Horst, 2017), Spanish-English (García et al., 2018), Hindi-English (Si, 2011; Dey and Fung, 2014), and Chinese-English (Lyu et al., 2010).

Since code-switching occurs mostly during spontaneous conversational speech, building models using spontaneous speech should be more beneficial than using read speech. While the frequency of code-switching itself in read speech could be manually adjusted by modifying the transcription, spontaneous and read speech still have many other differences characterized by certain factors. For instance, reduced spectral space and increased spectral variance are observed in Japanese spontaneous speech (Nakamura et al.,

2008). Increased spectral variance has also been observed in French spontaneous speech (Rouas et al., 2010). Other studies observe a reduction in phoneme duration (Liu et al., 2010) and word duration (Spilkov et al., 2010) in spontaneous speech. Different patterns from the variance of GMM supervector are also shown to be able to discriminate spontaneous and read speech data (Asami et al., 2014). Read speech possesses different acoustic properties, and reliance on them in code-switching task might lead to a distributional shift, which consequently compromises the overall performance of the acoustic model in a real-setting. For this reason, building code-switching ASR using spontaneous speech is preferable to read speech.

In this work, we introduce ASCEND<sup>1</sup>, a spontaneous multi-turn conversational dialogue Mandarin Chinese-English code-switching corpus, to bridge the gap between the real-setting of code-switching speech utterances and the existing code-switching speech corpora. ASCEND comprises 10.62 hours of clean spontaneous Chinese-English code-switching data collected from dialogues between two people. To allow more variety in the utterances, speakers are diversified based on their English proficiency level and their Chinese dialects, covering Hong Kong, Taiwan, and various regions in Mainland China. In order to build a rich and diverse vocabulary, dialogues on various topics are incorporated into the corpus. These include education, persona, philosophy, sports, and technology. Overall, we collect 26 dialogue sessions with a total of 23 speakers. Our corpus is equally split between the genders.

## 2. Related Work

Code-switching has been widely studied in both text and speech modalities for multiple language pairs: 1) code-switching in Hindi-English, Bengali-English, Gujarati-English, and Tamil-English (Banerjee et al., 2018); 2) code-switching in Spanish-English and Modern Standard

<sup>\*</sup> These authors contributed equally.

<sup>†</sup> The work was done when the author was studying in The Hong Kong University of Science and Technology.

<sup>1</sup>We release ASCEND at <https://huggingface.co/datasets/CAiRE/ASCEND>.

Topic	Sample question
Technology	你使用任何社交媒体吗? (Do you use any social media?)
Sports	是谁鼓励你参加这项运动? (Who inspired you to play this sport?)
Education	你们的course project是什么? (What is your course project?)
Philosophy	你听说过火车电车问题吗? (Have you heard of the train trolley problem?)

Table 1: Examples of topic ideas and questions for the conversation in Session 2–4.

Arabic-Egyptian (Aguilar et al., 2018), Irish-Latin code-switching (Lynn and Scannell, 2019); 3) code-switching in Arabic-English and Arabic-French (Chowdhury et al., 2021); and 4) code-switching in Chinese-English (Lin et al., 2021; Lyu et al., 2010). Furthermore, many solutions specific to code-switching have been established, such as multitask and meta learning for code-switching (Yu and Chen, 2020; Song et al., 2017; Winata et al., 2018), code-switched data augmentation method (Qin et al., 2020; Winata et al., 2019), and adaptation method from large multilingual models for code-switching setting (Winata et al., 2021; Winata, 2021).

Despite the gradual progression of code-switching research, existing code-switching solutions merely reach a decent level of performance, which is several times inferior to that of their monolingual counterparts, especially in the automatic speech recognition (ASR) task. For instance, for the traditional ASR task, word error rate (WER) of  $\sim 2\%$  (Gulati et al., 2020; Baevski et al., 2020) and character error rate (CER) of  $\sim 5\%$  (Zhang et al., 2020) have been reported for LibriSpeech (English) (Panayotov et al., 2015) and AiShell-1 (Chinese) (Bu et al., 2017) corpora, respectively. Code-switching ASR, on the other hand, has much poorer state-of-the-art performance, with 24.2% mixed error rate (MER) (Li and Vu, 2019) and 29.30% CER (Winata et al., 2020) for Chinese-English, 26.4% WER for Arabic-English (Chowdhury et al., 2021), and 37.70% WER for Arabic-French (Chowdhury et al., 2021). We argue that these performance gaps occur due to the limitation of existing code-switching corpora in comparison with monolingual corpora, notably for high-resource languages, e.g., English and Mandarin Chinese.

In recent years, many speech corpora for Chinese-English code-switching have been introduced. CECOS corpus (Shen et al., 2011) is a collection of 12.1 hours of read Chinese-English code-switching corpus by Taiwanese speakers. SEAME corpus (Lyu et al., 2010) consists of 30 hours of spontaneous intra-sentential code-switching speech utterances collected from 92 speakers, covering Chinese-English code-switching within Singaporean and Malaysian populations. OC16-CE80 (Wang et al., 2016) is a Chinese-English code-switching corpus that consists of 80 hours of read speech collected from more than

Session	Average duration (minutes)	Done by
1	11.07	13 pairs
2	13.78	13 pairs
3	14.45	13 pairs
4	13.85	10 pairs

Table 2: Statistics of each recording session of our ASCEND corpus.

1,400 speakers from the Mainland Chinese population. ASRU 2019 (Shi et al., 2020) is a large-scale Chinese-English code-switching corpus with 740 hours of utterances, 240 hours of Chinese-English code-switching read utterances and 500 hours of monolingual Chinese utterances.<sup>2</sup> The dataset is collected from multiple speakers from 30 provinces in Mainland China. Finally, (Li et al., 2012) introduce 36 hours of spontaneous Chinese-English code-switching speech recordings, mainly in Chinese, English, and Cantonese with a small proportion of German and French. They report that only a fraction of this corpus is transcribed.

Despite the abundance of Chinese-English code-switching resources, many are no longer publicly available. OC16-CE80 and ASRU 2019 were only available for past competition purposes and are no longer publicly available.<sup>3</sup> Moreover, CECOS and (Li et al., 2012) are also no longer publicly available.<sup>4</sup> Hence, there is no publicly available Chinese-English code-switching corpus as of now, except for SEAME.

### 3. Corpus Collection

In this section, we describe the setup and procedure for the audio recording used for collecting ASCEND’s multi-turn conversational code-switched speech dialogues.

#### 3.1. Recording setup

ASCEND is collected through recording an informal conversation between two speakers. The recordings are made in a quiet classroom. Both speakers are seated across one another at a distance of  $\sim 1$  meter. Each speaker is equipped with a RODE SmartLav+ clip microphone as the recording device. The microphone is mounted on the speaker’s shirt collar. The audio recording is set to a mono channel with a sample rate of 16 kHz, and the audio signal is encoded as 16-bit pulse-code modulation (PCM), producing a total bit rate of 256 kbps. The resulting audio file is stored in an uncompressed WAVE (.wav) file format.

<sup>2</sup>The paper makes no explicit mention whether the code-switching corpus uses read speech. We gathered this information from the competition website <https://www.datatang.com/competition>. There is also no indication of whether the Chinese corpus is read or spontaneous. (Access date: 12 November 2021)

<sup>3</sup>Some steps of the procedures required to obtain the dataset given by the affiliated institution are missing.

<sup>4</sup>Dataset status was confirmed by contacting the authors and/or the affiliated institution.

<b># Speakers</b>	23 speakers
<b># Sessions</b>	49 sessions
<b># Raw recordings</b>	98 recordings
<b>Avg. utterances</b>	128.27 per speaker per session
<b>English speaking rate</b>	152.31 words/minute
<b>Chinese speaking rate</b>	262.33 characters/minute

Table 3: The overview of the collected raw audio data of our ASCEND corpus.

### 3.2. Recording procedure

We collect the conversational audio recording data in the form of a casual one-on-one conversation. Both speakers take turns to ask a question, answer, or talk about a certain topic however they would like to, maintaining the natural course of the conversation. Short pauses, coughs, laughter, incomplete sentences, and other spontaneous responses that usually do not come up in a formal or organized setting are allowed in the conversation. Both speakers are encouraged to use code-switching at all times during the recording, as long as the utterance feels natural to the speaker. The task description, along with written consent and information that the conversation will be recorded and the resulting audio data published, is provided to all speakers before the recording begins.

The recording is split into several sessions. During the first session, both speakers get to know each other by exchanging information on personal topics, such as nicknames, family, favorite pastimes, and recent activities. This session is intended to gradually make them feel at ease around one another to spark a more interactive and dynamic conversation in upcoming sessions. In the next two or three (depending on the remaining time) sessions, the conversation takes off on a broader subject to encourage a larger variety of vocabulary. To facilitate this, we provide a list of topic ideas and questions for the speakers to gather inspiration from. A few examples from this list can be seen in Table 1. For each session, speakers can choose one topic they are comfortable with and begin the conversation based on it. To ensure a natural conversation flow, no restriction is enforced to keep the conversation in-topic; speakers are free to deviate from the determined topic at any point of the conversation.

The recording takes approximately one hour. It includes 5 minutes of instructions, 40–55 minutes of mixed-language conversation and breaks in between each session. We



Figure 1: Speaker split in ASCEND corpus.

record 13 casual one-on-one conversations with 13 speaker pairs, collecting a total of 49 sessions (Table 2). Three speakers participate twice, with a different conversation partner in each round. On average, the first session goes on for 11 minutes, while the later sessions for around 14–15 minutes. For each session, we obtain two recordings from each speaker, which sum to 98 raw audio files. Table 3 presents the overall statistics of the raw data collected from our ASCEND corpus.

## 4. Annotation

The raw audio recordings of the sessions are split into utterances by a professional annotation company based on a natural semantic boundary or a long pause between utterances. Utterances corresponding to a speaker are obtained from the audio file recorded by the respective microphone. The utterances are manually transcribed in Chinese characters, English letters, or a mix of both, depending on the language in use. For consistency and accuracy of the annotation results, we formulate guidelines for the transcription annotation, as follows:

1. Numbers are written as words instead of numerals. For example, "24 hours" is transcribed as "twenty four hours" in the corpus.
2. Abbreviations are transcribed as capital letters or separated by a space.
3. Contractions and shortened versions of words (e.g., "can't", "won't", and "it's") are not expanded. We keep contractions as-is because of the possible difference in phoneme.
4. Fillers or discourse particles are annotated as either: *ah*, *oh*, or *um*.
5. Punctuation symbols, such as period (.), comma (,), question mark (?), and exclamation mark (!), are not used to mark the utterances.
6. Unintelligible speech is marked with an [UNK] placeholder token.
7. Repetitions are preserved. Annotators write the words down as what they hear from the speech data. For example, "I don't (I don't) think they should be in the Olympic Games" is transcribed as "I don't I don't think they should be in the Olympic Games".

Gender	# Utterance				Duration (hr)			
	Train	Val	Test	Total	Train	Val	Test	Total
<b>Female</b>	4,591	484	861	5,936	4.04	0.46	0.48	4.98
<b>Male</b>	5,278	646	454	6,378	4.74	0.46	0.44	5.63
<b>Total</b>	9,869	1,129	1,315	12,314	8.78	0.92	0.92	10.62

Table 4: Train, validation, and test split in ASCEND corpus.

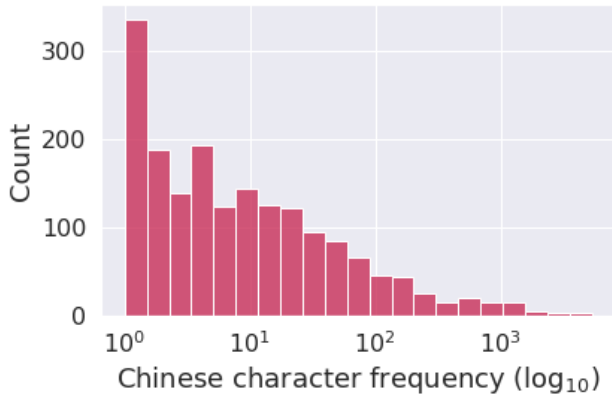


Figure 2: The distribution of the log character frequency in Chinese speech in ASCEND.

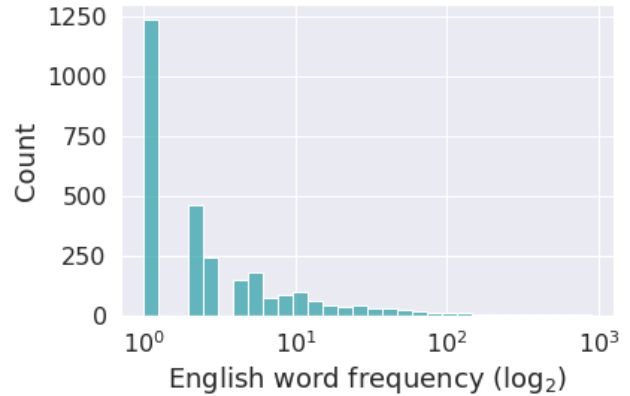


Figure 3: The distribution of the log word frequency in English speech in ASCEND.

**Post-annotation processing** To ensure the quality of our speech data, we do a second round of processing with a mix of manual and automatic checking. We inspect the transcriptions and remove the unnecessary symbols, whitespace, and annotation inconsistency. We exclude utterances that only contain [UNK] from the corpus. We re-format utterance audio files that do not follow the recording standards mentioned in Section 3.1.

**Corpus splitting** We divide the utterances into train, validation, and test sets. The sets have disjoint combinations of speakers (as presented by Figure 1) to enable this corpus to be used for the speaker-independent speech recognition task. Within each split, we balance the total duration of the audio data from each gender. At the end of this process, ASCEND is formed with the approximate ratio of 8:1:1 for its train, validation, and test sets respectively. This ratio is derived from both the audio duration and the number of utterances in each split. Table 4 describes the statistics of ASCEND’s train, validation, and test sets.

## 5. ASCEND: A Spontaneous Chinese-English Dataset

In this section, we report statistical findings regarding the Chinese-English code-switching of ASCEND. We also provide the statistics of the speakers who have participated in the corpus collection.

### 5.1. Corpus profile

ASCEND comprises 10.62 hours and  $\sim 12.3$ K utterances of spontaneous speech, with an average duration of 3.10 seconds per utterance. ASCEND includes a total of 145,146 tokens (i.e., words in English and characters in Chinese) with 1,795 types of Chinese characters and 2,860 types of

English words. An utterance is approximately 11.78 tokens long. In both languages, we find that a small portion of the vocabulary (e.g., particles, pronouns, affirmations, etc.) appears much more frequently than the rest. The distribution of the token frequency in ASCEND is depicted in Figure 2 and Figure 3.

ASCEND is collected from multiple speakers from different locations, including Taiwan, Hong Kong, and various provinces in Mainland China. Section 5.2. will discuss more details about our speakers. In terms of code-switching characteristics, the dialogues in ASCEND encompass both inter-sentential code-switching (from monolingual Chinese to monolingual English utterance or vice versa) and intra-sentential code-switching (mixed Chinese-English). Table 5 describes the proportion of the languages used in the speech data.

### 5.2. Speaker distribution

We hire 23 university students as our speakers, all of whom are native Chinese speakers who converse using English on a daily basis. Their personal information is obtained from an online form we provide during the speaker registration. Of the speakers, 13 identify as female and the other 10 identify as male. The speakers’ ages range from 19 to 30 years old, with a mean of 24 and a standard deviation of 2.24.

In addition to gender and age demographics, we also collect information that is indicative of their English proficiency to ensure the quality of the acquired code-switching utterances (Table 6). Most speakers have been studying English for 10 years or more, except for two whose experience has just passed the five year count. We also collect their speaking scores (according to IELTS or TOEFL iBT) to measure their fluency in English as a second language. The speaking scores among the speakers are then standardized to IELTS

Language	# Utterance	Duration (hr)
<b>Chinese</b> (49.85%)	6,139	5.32
<b>English</b> (23.14%)	2,850	2.42
<b>Mixed</b> (27.01%)	3,325	2.88

Table 5: Utterance distribution per language.

English study	Chinese	English	Mixed
< 10 years	53.38%	23.78%	22.84%
10-15 years	50.57%	21.07%	28.37%
> 15 years	47.87%	27.77%	24.37%

7262 Table 6: Language usage by years of English study.



Figure 4: Example of inter-sentential code-switching in ASCEND.

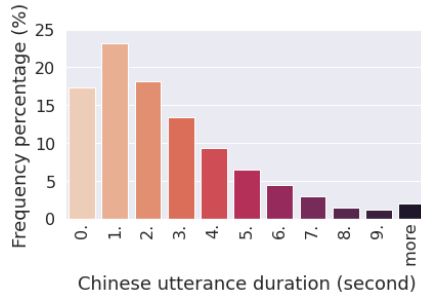


Figure 5: Distribution of Chinese utterance duration in ASCEND.




Figure 6: Distribution of English utterance duration in ASCEND.

band criteria. We find that all speakers reach or surpass the 5.5 mark, with an average score of 6.5.

### 5.3. Topic and code-switching

As mentioned in Section 3.2., each session uses one topic as a conversation starter. The topic in the first session is always persona, which covers both speakers' backgrounds such as name, hobbies, and age. The topics for the later sessions adhere to the speakers' choice, which is either education, philosophy, sports, or technology. From the total of 49 sessions, 12 correspond to education, 13 correspond to persona, 4 correspond to philosophy, 7 correspond to sports, and 13 correspond to technology.

In general, around half of the utterances (44.78%–55.09%) spoken for all of the topics consists of code-switching. Although the proportions of utterances with inter-sentential and the ones with intra-sentential code-switching are quite balanced, as shown in Table 7, the usage of intra-sentential code-switches increases for topics involving many widely-known English terms. One of these topics is technology, where intra-sentential code-switching makes up 31.93% of the utterances. The other is philosophy, which is composed of the highest overall percentage of code-switching. We also find that, despite code-switches using monolingual English utterances tending to be more occasional, their occurrence frequency increases along with the speakers' familiarity and knowledge about the conversation subject. For example, talking about communication devices during the technology topic or themselves during the persona topic triggers inter-sentential code-switches slightly more often among the speakers.

### 5.4. Common English phrases used in ASCEND

While the lexical resources used during code-switching from Chinese to English vary, some come up more frequently than others in the conversations. According to Ta-

Topic	Chinese	English	Mixed
Education	51.57%	23.16%	25.27%
Persona	48.76%	25.85%	25.40%
Philosophy	44.91%	26.54%	28.55%
Sports	55.22%	21.85%	22.94%
Technology	48.06%	20.01%	31.93%

Table 7: Language usage by conversation topic.

ble 8, the types of phrases that often occur in our corpus are related to asking a question (e.g., "do you think" and "what do you") and giving or thinking of a response (e.g., "how to say", "want to do", and "you know"). Aside from these, speakers quite frequently exchange phrases that are used to describe an idea (e.g., "like", "in the", "you can", and "this kind of"). A few topic-related phrases, such as "smart phone" for technology and "meaning of life" for philosophy, are also mentioned a lot during the discussions.

### 5.5. Inter-sentential code-switching in ASCEND

Our ASCEND corpus contains several inter-sentential code-switching instances. Inter-sentential code-switching differs from intra-sentential in a way that its language switch occurs between utterances. For example, in Figure 4, the second speaker completes the first utterance in Chinese then switches to English for the entire second utterance. As a result, all the involved utterances are still monolingual despite a language switch occurring. As shown by Figure 5 and Figure 6, we find that the monolingual utterances in ASCEND have a similar duration distribution for both Chinese and English utterances.

### 5.6. Intra-sentential code-switching in ASCEND

Aside from inter-sentential code-switching, ASCEND also consists of numerous intra-sentential code-switching utterances. An utterance is considered to have intra-sentential code-switching when a switch from one language to another happens within the utterance at least once. We refer

Top	English phrases		
	1-gram	2-gram	3-gram
1	the	do you	do you think
2	you	in the	what do you
3	to	you can	how to say
4	like	kind of	in hong kong
5	and	smart phone	this kind of
6	is	to do	you want to
7	in	hong kong have	so do you
8	so	you have	want to do
9	of	want to	you are not
10	for	you know	meaning of life

Table 8: Top 10 English 1-gram, 2-gram, and 3-gram phrases.

Top	Language turn	
	zh → en	en → zh
1	个 project	school 的
2	读 phd	phd 的
3	个 topic	ok 的
4	做 research	smartphone 的
5	的 major	phone 的

Table 9: Top 5 code-switches in language turns between Chinese and English.

to this language switching phenomenon as language turn. In the example in Figure 7, the utterance begins in Chinese, switches to English, goes back to Chinese, and so on until the language turns sum to six. In practice, most utterances tend to have a lower number of language turns. In the intra-sentential code-switching utterances in our ASCEND corpus, language turns appear 2.18 times per utterance on average, with a maximum of 14 times in a single utterance.



Figure 7: Intra-sentential code-switching utterance with six language turns.

**Language turn within utterances** As the speech data in our ASCEND corpus is spontaneous, all code-switches, including those used in language turns, occur on the speaker’s own accord without any fixed or predefined rules. Nevertheless, people tend to follow certain lexical patterns during code-switching, so a few mixes of Chinese and English phrases get used in language turns more frequently than others. We select one Chinese character and one English word from every language turn and sort them based on their occurrence frequency. Table 9 reports the five most common language turns for code-switching from Chinese to English and vice versa in ASCEND.

**Utterance as multiple monolingual segments** As shown in Figure 7, the presence of language turns causes the corresponding intra-sentential code-switching utterance to be

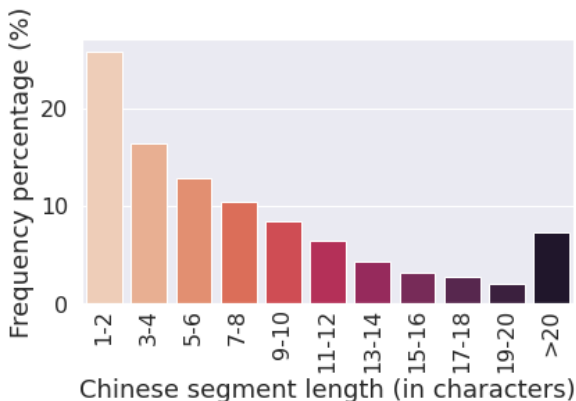


Figure 8: The distribution of Chinese segment length.

Top	Chinese segments		English segments	
	1-char	2-char	1-word	2-word
1	的	就是	ai	smart phone
2	啊	然后	phd	social media
3	是	所以	ok	hong kong
4	对	这个	so	i think
5	吗	那个	and	it’s like

Table 10: Top 5 short monolingual segments in intra-sentential code-switching.

composed of multiple monolingual segments. Depending on the language usage and the speaker, these segments vary in length. To observe the style of intra-sentential code-switching in spontaneous conversations, we separate the Chinese segments from the English ones, then calculate the number of segments found in each utterance. We find that an intra-sentential code-switching utterance typically comprises 1.75 Chinese segments and 1.38 English segments. In addition to the number of segments, we also calculate the occurrence frequency for each segment length. We report the overall distribution of the number of Chinese characters per segment in Figure 8 and the number of English words per segment in Figure 9.

Despite having a similar number of segments in an utterance, the characteristics of the Chinese segment length distribution differs from the English, with the former having a more even length distribution than the latter. Short Chinese segments (1–4 characters long) make up approximately 35% of the population, while the percentage doubles for English. Around 70% of English segments found in intra-sentential code-switching utterances consist of one or two words. Although language turns can occur in both languages, we can see that people tend to speak in longer Chinese segments (7.96 characters per segment on average) then switch to a shorter English segment (2.96 words per segment on average) in between. This phenomenon is expected, considering that all the speakers have Chinese as their first language. This speaking pattern aligns with the characteristics of code-switching in Hong Kong, Taiwan, Singapore, and Malaysia reported by (Chan et al., 2005), (Lyu et al., 2006), and (Lyu et al., 2010). English-

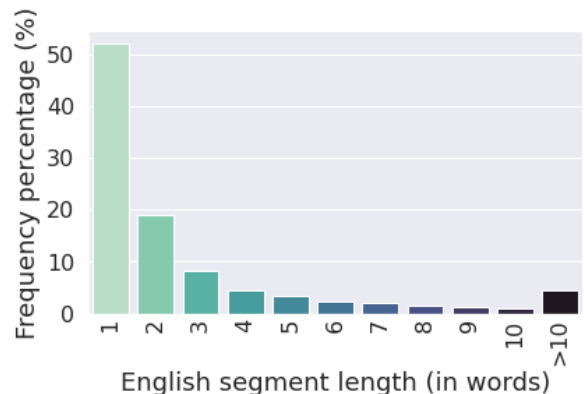


Figure 9: The distribution of English segment length.



Pre-training language	Vocabulary size	
	Pre-trained only	With ASCEND
Chinese	3503	3593 (+90)
English	33	1833 (+1800)
Multilingual	9913	9920 (+7)

Table 11: Vocabulary size of the models before and after the additions from ASCEND

dominated utterances with Chinese code-switches also appear in ASCEND, albeit more occasionally. Table 10 presents one-token and two-token segments that are commonly utilized as code-switches in ASCEND.

## 6. Baseline Experiment

In this section, we conduct an experiment on ASCEND to show its reliability and validity as a code-switching speech corpus.<sup>5</sup> For the experiment, a state-of-the-art speech recognition model architecture, namely wav2vec 2.0 (Baeovski et al., 2020), is employed. As no code-switching ASR model is available, we utilize wav2vec 2.0 models with different initializations as the baselines. The first two models are pre-trained on the English corpus and the Chinese corpus of Common Voice, respectively. Motivated by (Winata et al., 2021) who use multilingual models to approach the code-switching task, the third model is initialized with a multilingual wav2vec 2.0 pre-trained on 53 languages of the Common Voice corpus.

**Preprocessing.** Before we fine-tune both models on ASCEND, we omit unnecessary characters and symbols from the transcription data. The resulting texts are used to build ASCEND-specific vocabulary, which we leverage to extend the pre-trained tokenizer that comes with the model. Table 11 shows the vocabulary size of each model with and without ASCEND-specific vocabulary.

For the audio data, we normalize the audio data and apply SpecAugment (Park et al., 2019) to increase the robustness of the model. Specifically, we apply time masking and frequency masking, with a time masking probability of 0.065, a time masking length of 2, a frequency masking probability of 0.004, and a frequency masking length of 2. No time warping is applied to the audio data.

<sup>5</sup>The baseline experiment code can be found at <https://github.com/HLTCHKUST/ASCEND>.




Figure 10: Loss on ASCEND train set in the baseline experiments.




Figure 11: MER on ASCEND validation set in the baseline experiments.

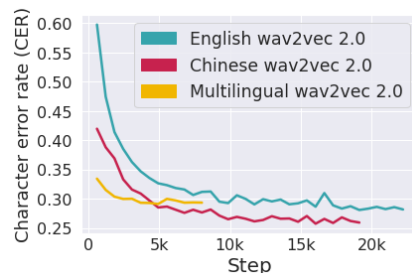


Figure 12: CER on ASCEND validation set in the baseline experiments.

Pre-training language	Validation		Test	
	MER (%)	CER (%)	MER (%)	CER (%)
Chinese	<b>30.37</b>	<b>25.72</b>	<b>27.05</b>	<b>22.69</b>
English	35.77	28.07	28.72	22.78
Multilingual	35.30	28.68	29.35	24.31

Table 12: Baseline experiment results on ASCEND validation and test set. **Bold** denotes the best performance over different models.

**Training details.** During the training, we employ Adam (Kingma and Ba, 2015) to optimize the wav2vec 2.0 model. For the objective function, we use the Connectionist Temporal Classification (CTC) loss. The model is fine-tuned on a single GeForce GTX 3090 GPU with a learning rate of  $5e-5$  and a batch size of 16. We train the model up to 100 epochs with early stopping of 5 epochs.

**Evaluation.** During the evaluation, we apply CTC decoding for generating the transcription. As for evaluation metrics, considering the character-based nature of the Chinese transcriptions and the word-based nature of the English transcriptions in ASCEND, we measure the models' performance using character error rate (CER) and mixed error rate (MER) (Schultz et al., 2011; Hu et al., 2020; Qiu et al., 2020). The CER is computed as the total of substitutions, deletions, and insertions divided by the number of characters in the reference, while the MER is calculated by measuring the CER for Chinese characters and word error rate (WER) for other characters.

### 6.1. Result and analysis

The evaluation results of the English, Chinese, and multilingual pre-trained models are shown in Table 12. The experiment results suggest that the Chinese pre-trained model outperforms both the English and the multilingual pre-trained models. While the Chinese pre-trained model is slower at converging compared to the multilingual model, it converges much faster than the English one, as shown by the training loss curve in Figure 10. Furthermore, Figure 11 and Figure 12 denote that the multilingual pre-trained model reaches a plateau earlier on both CER and MER in the ASCEND validation set. However, the Chinese pre-trained model ultimately yields better performance (30.37% MER and 25.72% CER) than the multilingual pre-trained model (35.30% MER and 28.68% CER) and the English pre-trained model (35.77% MER and

28.07% CER). This result is expected for three reasons: 1) almost 50% of the language distribution in ASCEND is Chinese, 2) as presented by Table 11, there is a huge vocabulary overlap between the Chinese pre-trained model and ASCEND-specific vocabulary, and 3) its pre-training solely focuses on learning Chinese instead of multiple languages at once like the multilingual model, in which the Chinese and English language only make up 0.95% and 28.48% of all the pre-training audio data.

Compared to other works on code-switching datasets (Banerjee et al., 2018; Chowdhury et al., 2021; Lynn and Scannell, 2019; Lyu et al., 2010; Winata et al., 2020), the baseline experiment on ASCEND yields a comparable performance with  $\sim 28\%$  MER and  $\sim 23\%$  CER on the test set. Additionally, in terms of dataset size, the number of tokens, and word distribution, ASCEND is on par with other existing Chinese-English spontaneous code-switching datasets, such as CECOS (Shen et al., 2011) and SEAME (Lyu et al., 2010). These results indicate that ASCEND is reliable for training and evaluating Chinese-English code-switching ASR.

## 7. Conclusion

In this paper, we introduce ASCEND, a spontaneous multi-turn conversational dialogue Chinese-English code-switching corpus. ASCEND consists of 10.62 hours of spontaneous speech with a total of  $\sim 12.3\text{K}$  utterances. The corpus is split into three sets: training, validation, and test with a ratio of 8:1:1 and a balanced gender proportion on each set. We further conduct a deeper analysis of the speech data to show the statistical distribution of both inter-sentential and intra-sentential code-switching utterances in ASCEND. Lastly, we experiment with the Chinese pre-trained wav2vec 2.0 model, English pre-trained wav2vec 2.0 model, and the multilingual pre-trained wav2vec 2.0 model to establish some baselines on ASCEND. Based on our experiment, the Chinese pre-trained model achieves the best code-switching performance (22.69% CER and 27.05% MER) on ASCEND’s test set.

## 8. Acknowledgements

This work is funded by ITS/353/19FP of the Innovation Technology Commission, The Hong Kong SAR Government, School of Engineering Ph.D. Fellowship Award, The Hong Kong University of Science and Technology, and the Hong Kong Fellowship Scheme from the Hong Kong Research Grants Council (RGC).

## 9. References

Aguilar, G., AlGhamdi, F., Soto, V., Diab, M., Hirschberg, J., and Solorio, T. (2018). Overview of the CALCS 2018 Shared Task: named entity recognition on code-switched data. In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, Melbourne, Australia, July. Association for Computational Linguistics.

Asami, T., Masumura, R., Masataki, H., and Sakauchi, S. (2014). Read and spontaneous speech classification based on variance of gmm supervectors. In *INTER-SPEECH*.

Baevski, A., Zhou, Y., Mohamed, A., and Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. In H. Larochelle, et al., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12449–12460. Curran Associates, Inc.

Banerjee, S., Moghe, N., Arora, S., and Khapra, M. M. (2018). A dataset for building code-mixed goal oriented conversation systems. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3766–3780, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.

Batliner, A., Kompe, R., Kießling, A., Nöth, E., and Niemann, H. (1995). Can you tell apart spontaneous and read speech if you just look at prosody? In *Speech Recognition and Coding*, pages 321–324. Springer Berlin Heidelberg.

Bautista, M. L. S. (2004). Tagalog-english code switching as a mode of discourse. *Asia Pacific Education Review*, 5(2):226–233, June.

Blaauw, E. (1994). The contribution of prosodic boundary markers to the perceptual difference between read and spontaneous speech. *Speech Communication*, 14(4):359–375, September.

Bu, H., Du, J., Na, X., Wu, B., and Zheng, H. (2017). Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline. In *Oriental COCODA 2017*, page Submitted.

Chan, J. Y., Ching, P., and Lee, T. (2005). Development of a cantonese-english code-mixing speech corpus. In *Ninth European Conference on Speech Communication and Technology*.

Chowdhury, S. A., Hussein, A., Abdelali, A., and Ali, A. (2021). Towards one model to rule all: Multilingual strategy for dialectal code-switching arabic asr. *ArXiv*, abs/2105.14779.

Dey, A. and Fung, P. (2014). A Hindi-English code-switching corpus. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, Reykjavik, Iceland, May. European Language Resources Association (ELRA).

García, P. B., Leibold, L., Buss, E., Calandruccio, L., and Rodriguez, B. (2018). Code-switching in highly proficient spanish/english bilingual adults: Impact on masked word recognition. *Journal of Speech, Language, and Hearing Research*, 61(9):2353–2363, September.

Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., and Pang, R. (2020). Conformer: Convolution-augmented Transformer for Speech Recognition. In *Proc. Interspeech 2020*, pages 5036–5040.

Haynes, R. M., White, L., and Mattys, S. L. (2015). What do we expect spontaneous speech to sound like? In Maria Wolters, et al., editors, *18th International Congress of Phonetic Sciences, ICPhS 2015, Glasgow, UK, August 10-14, 2015*. University of Glasgow.

Horst, T. t. (2017). Codeswitching in the irish-latin leabhar breac: mediæval homiletic culture.

Howell, P. and Kadi-Hanifi, K. (1991). Comparison



- of prosodic properties between read and spontaneous speech material. *Speech Communication*, 10(2):163–169, June.
- Hu, X., Zhang, Q., Yang, L., Gu, B., and Xu, X. (2020). Data augmentation for code-switch language modeling by fusing multiple text generation methods. In *Interspeech 2020*. ISCA, October.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Li, C.-Y. and Vu, N. T. (2019). Integrating knowledge in end-to-end automatic speech recognition for mandarin-english code-switching. *2019 International Conference on Asian Language Processing (IALP)*, pages 160–165.
- Li, Y., Yu, Y., and Fung, P. (2012). A Mandarin-English code-switching corpus. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2515–2519, Istanbul, Turkey, May. European Language Resources Association (ELRA).
- Li, A. (2002). Chinese prosody and prosodic labeling of spontaneous speech. In *Proc. Speech Prosody 2002*, pages 39–46.
- Lin, Z., Madotto, A., Winata, G. I., Xu, P., Jiang, F., Hu, Y., Shi, C., and Fung, P. (2021). Bitod: A bilingual multi-domain dataset for task-oriented dialogue modeling. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Liu, G., Lei, Y., and Hansen, J. H. L. (2010). Dialect identification: Impact of differences between read versus spontaneous speech. *2010 18th European Signal Processing Conference*, pages 2003–2006.
- Lynn, T. and Scannell, K. (2019). Code-switching in irish tweets: A preliminary analysis. In *Proceedings of the Celtic Language Technology Workshop*, pages 32–40.
- Lyu, D.-C., Lyu, R.-Y., Chiang, Y.-c., and Hsu, C.-N. (2006). Speech recognition on code-switching among the chinese dialects. In *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, volume 1, pages I–I. IEEE.
- Lyu, D.-C., Tan, T. P., Siong, C. E., and Li, H. (2010). Seame: a mandarin-english code-switching speech corpus in south-east asia. In *INTERSPEECH*.
- Nakamura, M., Iwano, K., and Furui, S. (2008). Differences between acoustic characteristics of spontaneous and read speech and their effects on speech recognition performance. *Computer Speech & Language*, 22(2):171–184.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. (2015). Librispeech: An asr corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Park, D. S., Chan, W., Zhang, Y., Chiu, C.-C., Zoph, B., Cubuk, E. D., and Le, Q. V. (2019). SpecAugment: A simple data augmentation method for automatic speech recognition. In *Interspeech 2019*. ISCA, September.
- Qin, L., Ni, M., Zhang, Y., and Che, W. (2020). Cosdaml: Multi-lingual code-switching data augmentation for zero-shot cross-lingual nlp. *ArXiv*, abs/2006.06402.
- Qiu, Z., Li, Y., Li, X., Metze, F., and Campbell, W. M. (2020). Towards context-aware end-to-end code-switching speech recognition. In *Proc. Interspeech 2020*, pages 4776–4780.
- Rizal, A. N. and Stymne, S. (2020). Evaluating word embeddings for Indonesian-English code-mixed text based on synthetic data. In *Proceedings of the The 4th Workshop on Computational Approaches to Code Switching*, pages 26–35, Marseille, France, May. European Language Resources Association.
- Rouas, J.-L., Beppu, M., and Adda-Decker, M. (2010). Comparison of spectral properties of read, prepared and casual speech in french. In *LREC*.
- Schultz, I. T., Fung, P., Gebhardt, J., and Schlippe, D. I. T. (2011). Speech recognition on english-mandarin code-switching data using factored language models.
- Shen, H.-P., Wu, C.-H., Yang, Y.-T., and Hsu, C.-S. (2011). Cecos: A chinese-english code-switching speech database. In *2011 International Conference on Speech Database and Assessments (Oriental CO-COSDA)*, pages 120–123.
- Shi, X., Feng, Q., and Xie, L. (2020). The asru 2019 mandarin-english code-switching speech recognition challenge: Open datasets, tracks, methods and results.
- Si, A. (2011). A diachronic investigation of hindi-english code-switching, using bollywood film scripts. *International Journal of Bilingualism*, 15(4):388–407, January.
- Song, X., Zou, Y., Huang, S., Chen, S., and Liu, Y. (2017). Investigating multi-task learning for automatic speech recognition with code-switching between mandarin and english. In *2017 International Conference on Asian Language Processing (IALP)*, pages 27–30.
- Spilkov, H., van Dommelen, W. A., et al. (2010). English of in 11 and 12 speakers' read and spontaneous speech. *Working papers/Lund University, Department of Linguistics and Phonetics*, 54:91–96.
- Wang, D., Tang, Z., Tang, D., and Chen, Q. (2016). Oc16-ce80: A chinese-english mixlingual database and a speech recognition baseline. *2016 Conference of The Oriental Chapter of International Committee for Coordination and Standardization of Speech Databases and Assessment Techniques (O-COCOSDA)*, pages 84–88.
- Winata, G. I., Madotto, A., Wu, C.-S., and Fung, P. (2018). Code-switching language modeling using syntax-aware multi-task learning. In *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*, pages 62–67, Melbourne, Australia, July. Association for Computational Linguistics.
- Winata, G. I., Madotto, A., Wu, C.-S., and Fung, P. (2019). Code-switched language models using neural based synthetic data from parallel sentences. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 271–280, Hong Kong, China, November. Association for Computational Linguistics.
- Winata, G. I., Cahyawijaya, S., Lin, Z., Liu, Z., Xu, P., and Fung, P. (2020). Meta-transfer learning for code-switched speech recognition. In *Proceedings of the 58th Annual Meeting of the Association for Computational*

- Linguistics*, pages 3770–3776, Online, July. Association for Computational Linguistics.
- Winata, G. I., Cahyawijaya, S., Liu, Z., Lin, Z., Madotto, A., and Fung, P. (2021). Are multilingual models effective in code-switching? In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages 142–153, Online, June. Association for Computational Linguistics.
- Winata, G. I. (2021). Multilingual transfer learning for code-switched language and speech neural modeling. *arXiv preprint arXiv:2104.06268*.
- Yang, L.-c. and Esposito, R. (2013). Understanding Mandarin prosody: Tonal and contextual variations in spontaneous conversation. In *International Journal of Computational Linguistics & Chinese Language Processing, Volume 18, Number 3, September 2013-Special Issue on Processing Lexical Tones in Natural Speech*, September.
- Yu, F.-H. and Chen, K.-Y. (2020). A preliminary study on leveraging meta learning technique for code-switching speech recognition. In *Proceedings of the 32nd Conference on Computational Linguistics and Speech Processing (ROCLING 2020)*, pages 136–147, Taipei, Taiwan, September. The Association for Computational Linguistics and Chinese Language Processing (ACLCLP).
- Zhang, B., Wu, D., Yao, Z., Wang, X., Yu, F., Yang, C., Guo, L., Hu, Y., Xie, L., and Lei, X. (2020). Unified streaming and non-streaming two-pass end-to-end model for speech recognition. *ArXiv*, abs/2012.05481.