

Prabhupadavani: A Code-mixed Speech Translation Data for 25 Languages

Jivnesh Sandhan¹, Ayush Daksh², Om Adideva Paranjay³,
Laxmidhar Behera^{1,4} and Pawan Goyal²

¹IIT Kanpur, ²IIT Kharagpur, ³University of Pennsylvania, ⁴IIT Mandi
jivnesh@iitk.ac.in, pawang@cse.iitkgp.ac.in

Abstract

Nowadays, the interest in code-mixing has become ubiquitous in Natural Language Processing (NLP); however, not much attention has been given to address this phenomenon for Speech Translation (ST) task. This can be solely attributed to the lack of code-mixed ST task labelled data. Thus, we introduce Prabhupadavani, which is a multilingual code-mixed ST dataset for 25 languages. It is multi-domain, covers ten language families, containing 94 hours of speech by 130+ speakers, manually aligned with corresponding text in the target language. The Prabhupadavani is about Vedic culture and heritage from Indic literature, where code-switching in the case of quotation from literature is important in the context of humanities teaching. To the best of our knowledge, Prabhupadavani is the first multi-lingual code-mixed ST dataset available in the ST literature. This data also can be used for a code-mixed machine translation task. All the dataset can be accessed at: <https://github.com/frozentoad9/CMST>.

1 Introduction

Speech Translation (ST) is a task in which speech is simultaneously translated from source language to a different target language.¹ It aids to overcome the language barriers across different communities for various applications such as social media, education, tourism, medical etc. Earlier attempts to build a robust speech translation system mainly focused on a cascaded approach where two separate architectures for Automatic Speech Recognition (ASR) and Machine Translation (MT) are used in pipeline mode (Cho et al., 2013; Post et al., 2013; Tsvetkov et al., 2014; Ruiz et al., 2015; Sperber et al., 2017). However, these approaches mainly suffer from cascading effect of error propagation. Thus, attention shifted to end-to-end approaches (Bérard et al., 2018; Duong et al., 2016; Weiss et al.,

2017; Bansal et al., 2017) due to their ability to obviate error propagation and ease of maintaining a single architecture. However, these end-to-end approaches could not match the performance of cascaded systems due to the lack of sufficiently large data (Niehues et al., 2021). Notably, with the recent upsurge in ST datasets (Di Gangi et al., 2019; Zanon Boito* et al., 2020; Iranzo-Sánchez et al., 2020; Wang et al., 2020), this gap has been closed (Niehues et al., 2021; Ansari et al., 2020).

Nowadays, most users prefer to communicate using a mixture of two or many languages on platforms such as social media, online blogs, chatbots, etc. Thus, code-mixing has become ubiquitous in all kinds of Natural Language Processing (NLP) resources/tasks (Khanuja et al., 2020; Chakravarthi et al., 2020; Singh et al., 2018a,b; Dhar et al., 2018). However, the existing NLP tools may not be robust enough to address this phenomenon of code-mixing for various downstream NLP applications (Srivastava and Singh, 2021). Therefore, there has been a surge in creating code-mixed datasets: (1) to understand reasons for the failure of existing models, and (2) to empower existing models for overcoming this phenomenon. Nevertheless, it is challenging to find natural resources that essentially capture different aspects of code-mixing for creating datasets for a wide range of NLP tasks. Although there has been considerable research in generating copious data and novel architectures for the ST task, we find that not much attention has been given to address the code-mixing phenomenon on the ST task. Possibly, this can be justified due to the lack of a code-mixed ST dataset. To the best of our knowledge, no such sufficiently large, multi-lingual, naturally occurring, code-mixed dataset is available for the ST.

Thus, in this work, we introduce **Prabhupadavani**, a multi-lingual, multi-domain, speech translation dataset for 25 languages containing 94 hours of speech by 130 speakers. The Prabhupadavani is about Vedic culture and heritage from Indic litera-

¹We refer to speech translation as a speech-to-text task.

ture, where code-switching in the case of quotation from literature is important in the context of humanities teaching. The multiple domains cover utterances from public lectures, conversations, debates, and interviews on various social issues. This is the first code-mixed data for speech translation to the best of our knowledge. It is code-mixed with English, Bengali and Sanskrit. From the typological point of view, the languages covered vary over ten language families. All the audios files have been manually aligned and translated. We believe that our work will ignite research in this direction to understand- (1) How to make existing systems robust for handling this phenomenon effectively? (2) Can multi-lingual training help to improve performance on code-mixed speech translation? (3) Will the gap between the cascade and end-to-end systems be closed? (4) Can we train a single model for all languages using parallel nature of the dataset?²

2 Related Work

Speech Translation: Recently, there have been increased efforts for creating large speech translation data sets for many languages. Nevertheless, the available datasets are limited to certain languages or only underpaid licenses for non-English languages. Most of the relatively larger datasets are English-centric (Di Gangi et al., 2019; Bérard et al., 2018), domain-specific (Zanon Boito* et al., 2020; Iranzo-Sánchez et al., 2020) or with limited speech hours (Zanon Boito* et al., 2020; Zhang et al., 2021). Recently, there has been upsurge in ST datasets in the literature (Di Gangi et al., 2019; Iranzo-Sánchez et al., 2020; Wang et al., 2020; Salesky et al., 2021). To the best of our knowledge, there is no such naturally occurring, sufficiently large and multi-lingual ST dataset that contains a code-switching phenomenon. We fill this gap by contributing a code-mixed speech translation dataset for 25 languages.

Code-mixing: Code-mixing is ubiquitous and well addressed on variety of downstream NLP tasks. However, majority of code-mixed datasets are synthetically generated (Gonen and Goldberg, 2018; Khanuja et al., 2020); therefore, they may not be able to capture the different aspects of code-mixing. The code-mixed datasets for an ASR task are either limited to only one/two languages or contain

²Prabhupadavani has parallel translations available in all the 25 languages for all the utterances.

only a few hours of speech data (Nakayama et al., 2019; Lyu et al., 2015). Synthetic code-mixed datasets may not capture the different aspects of code-mixing. To the best of our knowledge, Prabhupadavani is the first code-mixed dataset available for 25 languages on speech translation task.

3 Data Description

Resource: Vanimedia’s Multi-language Subtitle Project³ has created 1,080 audio mini-clips of Śrīla Prabhupāda’s lectures, conversations, debates, interviews and is now transcribing them in multiple languages.⁴ 700+ translators participate in creating subtitles for all 1,080 mini-clips for 108+ languages. Currently, this work has been completed for 25 languages. To procure subtitles for each clip in multiple languages, translators are provided with mini-clips and English subtitles that are manually aligned with each utterance. They use a third-party software named Dotsub.com⁵ and the task of translators is to provide translation for the corresponding utterance with the help of given transcription. On average, there are 3-4 translators for each language, and each clip takes more or less one hour for translation. Each translator has invested an average of 6 hours every day in translating these clips. Collectively, the time taken to translate 1,080 clips into 25 languages is over 45 weeks. Current release of the dataset contains 25 languages (including transcription) for which transcription/translations are available. For these languages, we have over 53K utterances of transcription and translations.

Preprocessing: In this section, we describe preprocessing steps followed to arrive at the final version of the data. First, we scrape transcription and their translations in 25 languages from Dotsub and extracted the corresponding Youtube links of all audio clips. We use Selenium⁶ web crawler to automatize the process of downloading. We convert those videos into MP3 audio clips using a third-party application⁷. We chop the converted audio files based on the timestamps provided in the subtitle (.srt) files.⁸ This process boils down to 53,000

³https://vanimedia.org/wiki/Multi-language_Subtitle_Project

⁴https://vanimedia.org/wiki/Table:_Clips_to_subtitle

⁵<https://dotsub.com>

⁶<https://www.selenium.dev/>

⁷<https://ytmp3.cc/uul00cc/>

⁸<https://pypi.org/project/audioclipextractor/>

Languages	# Types	# Tokens	Types per line	Tokens per line	Avg. token length
English	40,324	601,889	10.58	11.27	4.92
French (France)	50,510	645,651	11.38	12.09	5.08
German	50,748	584,575	10.44	10.95	5.57
Gujarati	41,959	584,989	10.37	10.95	4.46
Hindi	29,744	716,800	12.36	13.42	3.74
Hungarian	84,872	506,608	9.13	9.49	5.89
Indonesian	39,365	653,374	11.54	12.23	6.14
Italian	52,372	512,061	9.23	9.59	5.37
Latvian	70,040	477,106	8.69	8.93	5.72
Lithuanian	75,222	491,558	8.92	9.20	6.04
Nepali	52,630	570,268	10.03	10.68	4.88
Persian (Farsi)	51,722	598,096	10.61	11.20	4.10
Polish	71,662	494,263	8.99	9.25	5.86
Portuguese(Brazil)	50,087	608,432	10.80	11.39	5.12
Russian	72,162	490,908	8.96	9.19	5.79
Slovak	73,789	520,465	9.39	9.75	5.37
Slovenian	68,619	516,649	9.35	9.67	5.30
Spanish	49,806	608,868	10.75	11.40	5.07
Swedish	48,233	581,751	10.31	10.89	5.00
Tamil	84,183	460,678	8.37	8.63	7.65
Telugu	72,006	464,665	8.34	8.70	6.56
Turkish	78,957	453,521	8.27	8.49	6.35
Bulgarian	60,712	564,150	10.10	10.56	5.24
Croatian	73,075	531,326	9.58	9.95	5.28
Danish	50,170	587,253	10.40	11.00	4.98
Dutch	42,716	595,464	10.52	11.15	5.05

Table 1: Statistics of the **Prabhupadavani** dataset

Language	Tokens	Types	Percentage
English	500,136	6,312	83.6
Bengali	46,933	3,907	7.84
Sanskrit	51,246	7,202	8.56
Total	598,315	17,421	100

Table 2: Statistics of code-mixing in **Prabhupadavani**

utterances with their transcription and translation in 25 languages. Table 4 illustrates the example from Prabhupadavani. In order to oblivate time and efforts needed for data pre-processing, we provide train, dev and test set splits using stratified sampling. There are 51,000, 1,000 and 1,000 utterances in train, dev and test set, respectively. We consider the following dimensions for stratified sampling: (1) different speakers (2) proportion of intra-sentential and inter-sentential code-mixing.

Code-mixing: Prabhupadavani is code-mixed across three languages: English, Sanskrit, and Bengali. Table 2 reports the overall statistics of the code-mixing present in Prabhupadavani. If we consider the number of tokens in utterances, then it is mainly dominated by English (83.0%) tokens; however, it is not the case in terms of a number of types. This attributes to the contrasting nature of morphology (English vs Sanskrit/Bengali). Code-mixing is categorized into two classes- (1) inter-

sentential: speaker chooses to switch the language after completion of utterance (2) intra-sentential: speaker switches the language within an utterance. Table 3 illustrates examples of code-mixing from Prabhupadavani. Table 5 shows the statistics of both these types of code-mixing. Mainly speaker explains Sanskrit verses to the English audience; therefore, transitions between the Sanskrit-English pair is more. However, sometimes speaker also use Bengali literature to illustrate the points. Thus, we observe Bengali-English code-switching. Notably, there is no code-switching between Bengali-Sanskrit because the audience is English speaking.

Language diversity: From typological point of view, Prabhupadavani covers 25 languages (including ASR) from 10 language families. They are listed as follows: (1) *Indo-European*:- (a) *Romance*: Italian, Portuguese, Spanish, French (b) *Germanic*: German, Swedish, Danish, Dutch, English (c) *Baltic*: Lithuanian, Latvian (d) *Slavic*: Croatian, Polish, Slovak, Russian, Slovenian, Bulgarian, (e) *Indo-Aryan*: Bengali, Hindi, Sanskrit, Nepali (f) *Indo-Iranian*: Persian (Farsi) (2) *Uralic*:- (a) *Finno-Ugric*: Hungarian (3) *Austronesian*:- (a) *Malayo-Polynesian*: Indonesian (4) *Dravidian*:- Tamil, Telugu. Prabhupadavani contains fusional languages: Indo-European, Uralic and Agglutinative languages: Austronesian (Indonesian), Dravid-

Language	Type	Examples
English-Sanskrit	Inter Intra	Kṛṣṇa is assuring. ahaṁ tvāṁ sarva-pāpebhyo mokṣayiṣyāmi Sense gratification means udara-upastha-jihvā
Sanskrit-English	Inter Intra	Īśāvāsyam idaṁ sarvam . Everything belongs to God andhā yathāndhair upaniyamānāḥ and people, leaders.
English-Bengali	Inter Intra	Give up their. Asat-saṅga-tyāga ei vaiṣṇava ācāra Therefore Caitanya Mahāprabhu said guru-kṛṣṇa-kṛpāya
Bengali-English	Inter Intra	Guru-kṛṣṇa-kṛpāya pāya bhakti-latā-bīja . Then our devotional service is perfect tānhāra nāhika doṣa means he is not faulty.

Table 3: Examples of code-mixing in **Prabhupadavani**. English, Sanskrit and Bengali are indicated by black, red and violet color, respectively.

	Language	Translations
Source	English	You can become Brahman. Brahma-bhūyāya kalpate
Target	Bulgarian	Можете да станете Брахман. Брахма бхуяя калпате.
	Hindi	तुम ब्रह्म बन सकते हो । ब्रह्मभूयाय कल्पते ।
	Russian	Вы можете стать "Брахманом". "Брахма-бхуйайа калпате".
	Tamil	நீ ப்ரம்மன் ஆகலாம். ப்ரம்ம-பூயாய-கல்பதே.
	Gujarati	તમે બ્રહ્મ બની શકો છો. બ્રહ્મ-ભૂયાય-કલ્પતે.

Table 4: Sample data point from **Prabhupadavani**. For a code-mixed utterance, we show its English transcription (source) and the corresponding translations for 5 languages.

	Inter-Sentential	Intra-Sentential
English-Sanskrit	2,356	2,338
Sanskrit-English	2366	851
English-Bengali	339	124
Bengali-English	339	0

Table 5: Code-Mixing type for our dataset

ian and Turkic (Turkish). In the former languages, grammatical markers bear several meanings and for the latter ones, they exhibit only one meaning at the same time. Except Hindi (Indo-Aryan), all the languages in our dataset use nominative-accusative marking. Hindi uses ergative-absolutive marking upto a limited extent. We can also categorize languages based on the number of grammatical genders. Some languages pose (1) three genders: German, Russian, Swedish, etc. (2) two genders: French, Spain, Hindi, etc. (3) no genders: English, Nepali, Persian, Turkish, etc. Based on a syntactic construct, we can categorize the languages present in Prabhupadavani- (1) SVO word order: English, Italian, French, Indonesian, etc. (2) SOV word order: Indo-Aryan, Dravidian and Turkic (3) flexible word order: Russian, Hungarian.

Thus, the diversified language coverage of **Prabhupadavani** along with its code-mixed nature makes it a suitable dataset to investigate various linguistic phenomena for the speech translation task, ASR and machine translation.

Applications of dataset: In this section, we throw some light on possible applications of Prabhupadavani: (1) The full dataset of Prabhupadavani contains 2,400 hours of speech. The current release of Prabhupadavani dataset can be utilized to facilitate automatic subtitling of the remaining part of data in the different languages. In this way, it will also be helpful to generate relatively larger dataset. (2) This dataset can provide a fertile soil to investigate on- How to make existing systems robust to code-mixing phenomenon? Will the gap between cascade and end-to-end to approaches be closed? Can multi-lingual training help to address code-mixing phenomenon? Can we train single ST system for all languages? How robust will be these models on another domain?

4 Conclusion and Discussion

In this work, we focused on a code-mixed speech translation dataset. Although code-mixing is a spoken language phenomenon, not much attention has been given to address this phenomenon due to the unavailability of such a dataset. Thus, we released Prabhupadavani, a high-quality multilingual multi-domain code-mixed ST dataset containing 94 hours of speech data, 130+ speakers, for 25 languages covering 10 language families. The dataset is code-mixed with three languages: English, Bengali, and Sanskrit. In order to reduce the efforts needed for

pre-processing, we provide stratified data splits for the dataset. The same dataset can be utilized for the code-mixed machine translation task. We believe that these efforts will (1) set a fertile soil for investigating the applicability of existing solutions, (2) help to analyze the kind of errors existing systems are making, and (3) facilitate researchers to propose a novel solution to make existing systems robust for code-mixing. We plan to extend this dataset for 108+ languages.

In code-switching conversations, speakers prefer a specific communication in a specific language of choice. In that context, the interesting factor is often when the boundary between languages becomes fuzzy, as in cases where an English verb stem is used with Spanish morphology. In the case of Prabhupadavani, the audience does not necessarily speak Sanskrit or Bengali, and the code-switching is primarily quotation or explanation. That is interesting as a phenomenon, but it is not the same as dual bilingual code-switching. We believe future work may help to get deep insights into this phenomenon.

Ethics Statement: We do not foresee any ethical concerns with the work presented in this manuscript. We have taken the consent of the Vanipedia team and Bhaktivedanta Book Trust International to use translations and audio in our dataset.

Acknowledgements

We would like to thank the Vanipedia team (<https://vanipedia.org/>) of 700+ translators for establishing this multi-lingual database for us to develop. We thank the Bhaktivedanta Book Trust International for permitting us to use Prabhupadavani audio in our dataset. We are grateful to Manish Gupta, Microsoft, for helping us with insightful discussions. We would like to thank the anonymous reviewers for their constructive feedback towards improving this work. The TCS Fellowship supports the first author’s work under Project TCS/EE/2011191P.

References

Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, et al. 2020. Findings of the iwslt 2020 evaluation campaign. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34.

Sameer Bansal, Herman Kamper, Adam Lopez, and Sharon Goldwater. 2017. [Towards speech-to-text translation without speech recognition](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 474–479, Valencia, Spain. Association for Computational Linguistics.

Alexandre Bérard, Laurent Besacier, Ali Can Kocabiyikoglu, and Olivier Pietquin. 2018. End-to-end automatic speech translation of audiobooks. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6224–6228. IEEE.

Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, Ruba Priyadharshini, and John P McCrae. 2020. Corpus creation for sentiment analysis in code-mixed tamil-english text. *arXiv preprint arXiv:2006.00206*.

Eunah Cho, Christian Fügen, Teresa Herrmann, Kevin Kilgour, Mohammed Mediani, Christian Mohr, Jan Niehues, Kay Rottmann, Christian Saam, Sebastian Stüker, et al. 2013. A real-world system for simultaneous translation of german lectures. In *INTER-SPEECH*, pages 3473–3477.

Mrinal Dhar, Vaibhav Kumar, and Manish Shrivastava. 2018. Enabling code-mixed translation: Parallel corpus creation and mt augmentation approach. In *Proceedings of the First Workshop on Linguistic Resources for Natural Language Processing*, pages 131–140.

Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [MuST-C: a Multilingual Speech Translation Corpus](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2012–2017, Minneapolis, Minnesota. Association for Computational Linguistics.

Long Duong, Antonios Anastasopoulos, David Chiang, Steven Bird, and Trevor Cohn. 2016. [An attentional model for speech translation without transcription](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 949–959, San Diego, California. Association for Computational Linguistics.

Hila Gonen and Yoav Goldberg. 2018. Language modeling for code-switching: Evaluation, integration of monolingual data, and discriminative training. *arXiv preprint arXiv:1810.11895*.

J. Iranzo-Sánchez, J. A. Silvestre-Cerdà, J. Jorge, N. Roselló, A. Giménez, A. Sanchis, J. Civera, and A. Juan. 2020. Europarl-st: A multilingual corpus for speech translation of parliamentary debates. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233.

- Simran Khanuja, Sandipan Dandapat, Sunayana Sitaram, and Monojit Choudhury. 2020. A new dataset for natural language inference from code-mixed conversations. *arXiv preprint arXiv:2004.05051*.
- Dau-Cheng Lyu, Tien-Ping Tan, Eng-Siong Chng, and Haizhou Li. 2015. Mandarin–english code-switching speech corpus in south-east asia: Seame. *Language Resources and Evaluation*, 49(3):581–600.
- Sahoko Nakayama, Takatomo Kano, Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. 2019. **Recognition and translation of code-switching speech utterances**. In *2019 22nd Conference of the Oriental COCOSDA International Committee for the Coordination and Standardisation of Speech Databases and Assessment Techniques (O-COCOSDA)*, pages 1–6.
- Jan Niehues, Elizabeth Salesky, Marco Turchi, and Matteo Negri. 2021. **Tutorial proposal: End-to-end speech translation**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Tutorial Abstracts*, pages 10–13, online. Association for Computational Linguistics.
- Matt Post, Gaurav Kumar, Adam Lopez, Damianos Karakos, Chris Callison-Burch, and Sanjeev Khudanpur. 2013. Improved speech-to-text translation with the fisher and callhome spanish–english speech translation corpus. In *Proc. IWSLT*.
- Nicholas Ruiz, Qin Gao, Will Lewis, and Marcello Federico. 2015. Adapting machine translation models toward misrecognized speech with text-to-speech pronunciation rules and acoustic confusability. In *Interspeech*. ISCA-International Speech Communication Association.
- Elizabeth Salesky, Matthew Wiesner, Jacob Bremerman, Roldano Cattoni, Matteo Negri, Marco Turchi, Douglas W. Oard, and Matt Post. 2021. Multilingual tedx corpus for speech recognition and translation. In *Proceedings of Interspeech*.
- Kushagra Singh, Indira Sen, and Ponnurangam Kumaraguru. 2018a. A twitter corpus for hindi-english code mixed pos tagging. In *Proceedings of the Sixth International Workshop on Natural Language Processing for Social Media*, pages 12–17.
- Vinay Singh, Deepanshu Vijay, Syed Sarfaraz Akhtar, and Manish Shrivastava. 2018b. Named entity recognition for hindi-english code-mixed social media text. In *Proceedings of the seventh named entities workshop*, pages 27–35.
- Matthias Sperber, Jan Niehues, and Alex Waibel. 2017. Toward robust neural machine translation for noisy input sequences. In *International Workshop on Spoken Language Translation (IWSLT)*, page 18.
- Vivek Srivastava and Mayank Singh. 2021. **Challenges and considerations with code-mixed nlp for multilingual societies**.
- Yulia Tsvetkov, Florian Metze, and Chris Dyer. 2014. **Augmenting translation models with simulated acoustic confusions for improved spoken language translation**. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 616–625, Gothenburg, Sweden. Association for Computational Linguistics.
- Changhan Wang, Juan Pino, Anne Wu, and Jiatao Gu. 2020. **CoVoST: A diverse multilingual speech-to-text translation corpus**. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4197–4203, Marseille, France. European Language Resources Association.
- Ron J Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. Sequence-to-sequence models can directly translate foreign speech. *arXiv preprint arXiv:1703.08581*.
- Marcely Zanon Boito*, William Havarad*, Mahault Garnerin, Eric Le Ferrand, and Laurent Besacier. 2020. **MaSS: A Large and Clean Multilingual Corpus of Sentence-aligned Spoken Utterances Extracted from the Bible**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6486–6493, Marseille, France. European Language Resources Association.
- Ruiqing Zhang, Xiyang Wang, Chuanqiang Zhang, Zhongjun He, Hua Wu, Zhi Li, Haifeng Wang, Ying Chen, and Qinfei Li. 2021. **Bstc: A large-scale chinese-english speech translation dataset**. *arXiv preprint arXiv:2104.03575*.