



# MuAViC: A Multilingual Audio-Visual Corpus for Robust Speech Recognition and Robust Speech-to-Text Translation

Mohamed Anwar, Bowen Shi, Vedanuj Goswami, Wei-Ning Hsu, Juan Pino, Changhan Wang

Meta AI

{anwarvic, bshi, vedanuj, wnhsu, juancarabina, changhan}@meta.com

## Abstract

We introduce MuAViC, a multilingual audio-visual corpus for robust speech recognition and robust speech-to-text translation providing 1200 hours of audio-visual speech in 9 languages. It is fully transcribed and covers 6 English-to-X translation as well as 6 X-to-English translation directions. To the best of our knowledge, this is the first open benchmark for audio-visual speech-to-text translation and the largest open benchmark for multilingual audio-visual speech recognition. Our baseline results show that MuAViC is effective for building noise-robust speech recognition and translation models. We make the corpus available at <https://github.com/facebookresearch/muavic>.

**Index Terms:** corpus, audio-visual, multimodal, multilingual, speech recognition, speech-to-text translation

## 1. Introduction

Visual perception of lip movement plays a crucial role in supplementing audio signals for a better understanding of spoken languages [1]. Its immunity to acoustic noise makes it a valuable asset in building noise-robust speech processing systems. Audio-visual speech recognition (AVSR), which transcribes spoken utterances using both audio and visual inputs, has been shown to be effective in improving the robustness of speech recognition [2, 3]. The application of deep learning has greatly improved the performance of AVSR systems, as evidenced by the state-of-the-art (SOTA) models in the widely used LRS3-TED public benchmark. Typically, SOTA systems [3] are able to lower word error rate (WER) by 3 times compared to its audio-only counterpart under challenging conditions with 0 dB babble noise.

In spite of the encouraging progress mentioned earlier, modern AVSR models are heavily benchmarked under a monolingual setup, with the majority being based on English. Non-English audio-visual datasets are orders of magnitude smaller, with a common dataset, CMU-MOSEAS [4] providing less than 20 hours of audio-visual speech for each of the 4 non-English languages it covers, which is at least 20 times smaller than the English-based LRS3-TED corpus. The lack of large-scale audio-visual datasets also impedes the development of audio-visual speech-to-text translation (AVST), which is only simulated using synthetic data in prior works, but is expected to bring a noise-robust effect similar to AVSR, as will be demonstrated in this paper. With the increasing popularity of self-supervised audio-visual pre-training approaches [5], proper benchmarking of multilingual AVSR and AVST models will enable multilingual pre-training, which allows more effective use of large-scale multilingual audio-visual data such as AV-Speech [6].

In this paper, we present MuAViC, a multilingual audio-

Table 1: Duration and speaker statistics for MuAViC. <sup>†</sup> Train set example is either human-labeled (“H”) or pseudo-labeled (“P”) by a machine translation model, M2M-100 418M [7].

	Hours			Speakers		
	Train (H+P) <sup>†</sup>	Dev	Test	Train	Dev	Test
<i>Audio-Visual Speech Recognition</i>						
En [8, 3]	436 + 0	1.1	0.8	4.7K	0.9K	0.4K
Ar	16 + 0	1.5	1.2	95	7	7
De	10 + 0	1.6	1.5	53	9	9
El	25 + 0	2.3	2.0	113	10	8
Es	178 + 0	1.6	1.7	987	16	12
Fr	176 + 0	1.8	1.5	948	12	10
It	101 + 0	1.9	1.9	487	8	8
Pt	153 + 0	1.5	1.8	810	9	13
Ru	49 + 0	1.7	1.8	238	7	9
<i>En-X Audio-Visual Speech-to-Text Translation</i>						
En-El	17 + 420	0.2	0.5			
En-Es	21 + 416	0.2	0.6			
En-Fr	21 + 416	0.2	0.6	4.7K	0.9K	0.4K
En-It	20 + 417	0.2	0.6			
En-Pt	18 + 419	0.2	0.5			
En-Ru	20 + 417	0.2	0.6			
<i>X-En Audio-Visual Speech-to-Text Translation</i>						
El-En	8 + 17	2.3	2.0	113	10	8
Es-En	64 + 114	1.6	1.7	987	16	12
Fr-En	45 + 131	1.8	1.6	948	12	10
It-En	48 + 53	1.9	1.9	487	8	8
Pt-En	53 + 100	1.5	1.8	810	9	13
Ru-En	8 + 41	1.7	1.8	238	7	9

visual corpus for robust speech recognition and robust speech-to-text translation. This corpus is sourced from TED and TEDx talks including total 1200 hours of transcribed audio-visual speech from over 8000 speakers in 9 languages: English (En), Arabic (Ar), German (De), Greek (El), Spanish (Es), French (Fr), Italian (It), Portuguese (Pt) and Russian (Ru). This makes MuAViC the largest open benchmark so far for multilingual audio-visual speech recognition (AVSR) and lipreading. In addition, we provide text translations and establish baselines for 6 English-to-X translation as well as 6 X-to-English translation directions. To the best of our knowledge, MuAViC is the first publicly available corpus for audio-visual speech-to-text translation.

## 2. Related Works

**Audio-visual speech recognition** Audio-visual speech recognition (AVSR) has been a popular research topic for the speech

Table 2: Results for English speech recognition and En-X speech translation in a clean environment (A: audio, AV: audio+video).

ID	Model	Mode	Target							Avg
			En	El	Es	Fr	It	Pt	Ru	
<i>Speech Recognition, Test WER ↓</i>										
A1	Monolingual AV-HuBERT [3]	A	2.5	-	-	-	-	-	-	-
		AV	2.3	-	-	-	-	-	-	-
<i>Text-To-Text Translation, Test BLEU ↑</i>										
A2	Bilingual Transformer base	-	-	25.8	29.5	27.0	22.6	23.9	17.2	24.3
A3	M2M-100 418M, multilingual, unconstrained data [7]	-	-	24.5	28.7	25.6	21.8	22.2	15.8	23.1
<i>Speech-To-Text Translation, Test BLEU ↑</i>										
A1+A2		A	-	25.1	28.9	26.7	21.9	23.8	16.9	23.9
		AV	-	25.5	29.1	26.5	22.3	23.6	16.8	24.0
A1+A3		A	-	23.8	27.8	25.4	21.3	22.0	15.2	22.6
		AV	-	24.1	28.2	25.2	21.4	22.0	15.3	22.7
Bilingual AV-HuBERT		A	-	23.0	27.5	25.1	20.7	20.1	14.7	21.9
		AV	-	23.4	26.6	25.3	20.7	20.5	14.6	21.9

community for a long time. Most of the early AVSR datasets, such as GRID [9], were collected in a controlled environment where subjects were required to recite pre-defined phrases in a laboratory-like setting (for example, with uniform illumination conditions). These datasets also have limitations in terms of visual perspectives and often consist of small-scale data with limited vocabulary size and number of speakers. The first large-scale “in-the-wild” AVSR dataset was LRW [10], which contains 173 hours of speaking video clips of 500 isolated words from BBC broadcasts. In recent years, there has been a growing trend towards collecting sentence-level audio-visual speech data, such as LRS3-TED [8, 2], which is more natural and widely accessible. LRS3-TED is the largest publicly available AVSR dataset and contains approximately 437 hours of English speaking videos from over 5000 speakers.

The majority of the widely-used AVSR datasets are in English, while non-English AVSR datasets [4, 11, 12] are significantly smaller in size. For example, the largest AVSR dataset in Mandarin Chinese, CMLR [11], comprises only 86 hours of utterances from 11 speakers. The current largest multilingual AVSR dataset, CMU-MOSEAS [4], contains total 68 hours of speech in Spanish, Portuguese, German, and French. Furthermore, it lacks cross-lingual translations, hindering its use for audio-visual speech-to-text translation (AVST).

It is noteworthy that there exist large-scale, unlabeled multilingual audio-visual speech datasets, such as VoxCeleb2 [13] (2442 hours) and AV-Speech [6] (around 4700 hours). They have been widely utilized for various non-transcription tasks, such as speaker verification [13, 14] and source separation [6], or as a source of pre-training data for self-supervised audio-visual representation learning [5, 15, 16]. However, they are not directly applicable to AVSR and AVST due to the lack of transcriptions.

**Speech-to-text translation** Speech-to-text translation (ST) or spoken language translation (SLT) has become an increasingly popular topic of speech research with the recent emergence of open ST benchmarks. Almost all the ST corpora were created by adding translations of transcriptions to pre-existing speech recognition corpora, where speech data was collected from various sources, such as phone call recordings [17], read speech recordings [18, 19, 20] and public speech recordings [21, 22, 23, 24]. Despite the occasional availability of video tracks in the original data sources of these corpora, their

creators have focused on only the audio modality and included only audio data in the corpora. Due to the lack of labeled audio-visual corpora, audio-visual speech-to-text translation remains a largely uncharted research area.

### 3. Corpus Creation

MuAViC sources data from TED and TEDx talk recordings, where native or non-native speakers (only one speaker most of the time) deliver public speech on stage and cameras capture stage scenes switching among different viewpoints. We collect both audio and video tracks from the recordings, and align them with human transcriptions as well as text translations.

For English talks, we reuse the audio-visual data from LRS3-TED [8] and follow the original data split. We find human translations for these talks from a machine translation corpus, TED2020 [26] by matching transcriptions in LRS3-TED and source sentences in TED2020. Matched LRS3-TED examples are then paired with the corresponding target sentences in TED2020 for translation labels. We apply exact text matching for development set and test set examples to ensure the best accuracy. To improve matching recall on the train set, we develop a fuzzy text matching strategy: we first segment TED2020 source and target sentences by punctuation if both sides of the sentence pair contain the same amount of segments. Then we normalize TED2020 and LRS3-TED texts by punctuation removal and lowercasing. Finally, we conduct exact text matching between the two corpora. For LRS3-TED train set examples without a match from TED2020, we acquire pseudo-translation labels from a machine translation model, M2M-100 418M [7] with default decoding hyper-parameters.

For non-English talks, we reuse the audio-only data, transcriptions and text translations collected by mTEDx [24]. Our data split also follows mTEDx. We acquire video tracks of the original recordings, and align processed video data with the audio one to form audio-visual data, similar to LRS3-TED [8]. Although all the audio data in mTEDx is transcribed, only a subset of it is translated. We acquire pseudo-translation labels from M2M-100 418M [7] for the untranslated train set examples with default decoding hyper-parameters.

In Table 1, we provide duration and speaker statistics for each data split in MuAViC. MuAViC covers 9 languages providing total 1200 hours of transcribed audio-visual data from

Table 3: Results for non-English speech recognition and X-En speech-to-text translation in a clean environment (A: audio, AV: audio+video)

ID	Model	Mode	Source								Avg
			Ar	De	El	Es	Fr	It	Pt	Ru	
<i>Speech Recognition, Test WER ↓</i>											
	Transformer, monolingual [24]	A	104.4	111.1	109.5	46.4	45.6	48.0	54.8	74.7	74.3
	Hybrid LF-MMI, monolingual [24]	A	80.8	42.3	25.0	16.2	19.4	16.4	20.2	28.4	31.1
	Whisper, multilingual, unconstrained data [25]	A	91.5	24.8	25.4	12.0	12.7	13.0	15.5	31.1	28.2
B1	Monolingual AV-HuBERT	A	99.3	61.1	35.1	16.5	24.4	19.3	23.0	33.3	39.0
		AV	98.5	52.4	26.3	15.9	23.7	18.5	19.4	30.0	35.6
B2	Multilingual AV-HuBERT	A	67.7	46.5	40.4	30.6	27.0	19.3	19.8	37.6	36.1
		AV	69.3	47.2	41.2	16.2	19.0	19.8	19.9	38.0	33.8
<i>Text-To-Text Translation, Test BLEU ↑</i>											
B3	Bilingual Transformer <i>base</i>	-	-	-	13.9	28.2	33.8	27.2	31.9	14.2	27.0
B4	M2M-100 <i>418M</i> , multilingual, unconstrained data [7]	-	-	-	25.9	29.5	34.9	30.7	33.6	19.1	29.0
<i>Speech-To-Text Translation, Test BLEU ↑</i>											
	Whisper, multilingual, unconstrained data [25]	A	-	-	24.2	28.9	34.5	29.2	32.6	16.1	29.9
B1+B3		A	-	-	10.7	21.2	24.2	20.3	23.5	10.7	18.4
		AV	-	-	11.3	21.1	24.3	20.8	24.9	11.0	18.9
B2+B3		A	-	-	7.6	9.2	12.8	19.3	23.1	8.7	13.4
		AV	-	-	7.5	20.2	24.4	18.6	23.2	8.6	17.1
B1+B4		A	-	-	15.9	23.1	25.7	23.4	24.2	13.6	21.0
		AV	-	-	18.4	23.0	26.1	23.8	26.3	13.9	21.9
Bilingual AV-HuBERT		A	-	-	7.3	20.4	26.4	19.5	24.2	9.2	17.8
		AV	-	-	7.2	20.2	25.0	19.4	24.1	9.0	17.5
Multilingual AV-HuBERT		A	-	-	9.3	21.0	26.3	21.2	24.3	9.3	18.6
		AV	-	-	7.6	20.5	25.2	20.0	24.0	8.1	17.6

more than 8000 speakers (number of speakers estimated based on TED/TEDx talk IDs). It also provides text translations into or out of English for 6 of the 9 languages. All the AVSR and AVST test sets are human-labeled, with no less than 0.5 hours of audio-visual speech from at least 7 speakers.

## 4. Experiments

### 4.1. Experimental Setup

For both AVSR and AVST, we use an English AV-HuBERT *large* pre-trained model [3]<sup>1</sup>, which is trained on the combination of LRS3-TED [8] and the English portion of Vox-Celeb2 [27]. We follow [3] for fine-tuning hyper-parameters, except that we fine-tune our bilingual models to 30K updates and our multilingual AVSR model to 90K updates. We freeze the pre-trained encoders for the first 4K and 24K updates for X-En AVST and En-X AVST models, respectively. Also, we randomly augment 25% of the input samples with multiple types of additive noises with a SNR (signal-to-noise ratio) of 0. The noise audio clips in the categories of “natural”, “music” and “babble” are sampled from MUSAN dataset [28], while the overlapping “speech” noise samples are drawn from LRS3-TED. In creating “speech” and “babble” noise sets, we ensure there are no speaker overlap among different partitions. We remove extremely short utterances (less than 0.2 seconds) and long utterances (more than 20 seconds) for better training stability. Besides end-to-end AVST models, we also build cascaded systems composed of a AVSR model and a text-based machine

translation model (MT). For bilingual MT, we train Transformer *base* [29] models using all the transcription-translation pairs in MuAViC train set. For multilingual MT, we leverage M2M-100 *418M* [7], which is trained on large-scale open-domain mined bitext data. For speech recognition and X-En speech translation, we also evaluate Whisper *Large V2* [25] as a SOTA baseline.

For inference, we use the best checkpoint by validation accuracy for AVSR/AVST and the best checkpoint by validation BLEU for MT. We use a beam size of 5 and default values for the other beam search decoding hyper-parameters. For AVSR, we normalize texts by punctuation removal and lowercasing [30] before calculating WER (word error rate). For AVST and MT, we use SacreBLEU [31] with default options, where texts are processed by its built-in *l3a* tokenizer before BLEU [32] calculation.

### 4.2. Audio-Visual Speech Recognition (AVSR)

#### 4.2.1. Clean Setup

We evaluate AVSR models in both audio-only (“A”) and audio-visual (“AV”) modes, where the former leverages only audio modality in fine-tuning and inference while the latter leverages both audio and visual modalities. As shown in Table 2, the English AVSR model [3] has a low test WER of 2.5 and 2.3 respectively for audio-only and audio-visual modes. For non-English AVSR, we fine-tune the pre-trained English AV-HuBERT model either separately on each language (8 monolingual models) or jointly on all the 8 non-English languages (a multilingual model), whose test WER can be found in Table 3. We observe that our monolingual AVSR models in the audio-

<sup>1</sup>Randomly initialized models perform poorly (see Appendix A.1).

Table 4: Results for speech recognition and speech translation in a noisy environment with multilingual babble noise (SNR=0. A: audio, AV: audio+video).

Model	Mode	Source/Target									Avg
		En	Ar	De	El	Es	Fr	It	Pt	Ru	
<i>Speech Recognition, Test WER ↓</i>											
Whisper, multilingual, unconstrained data [25]	A	202.4	197.9	244.4	113.3	116.3	172.3	172.4	223.6	126.2	174.3
Monolingual AV-HuBERT	A	25.9	102.3	83.5	74.6	65.1	63.6	71.0	77.3	68.1	70.2
	AV	5.7	101.0	74.6	53.1	45.1	48.1	51.6	47.2	51.9	53.2
Multilingual AV-HuBERT	A	-	86.1	76.0	73.7	59.7	54.7	61.0	61.9	61.9	66.9
	AV	-	82.2	66.9	62.2	40.7	39.0	44.3	43.1	43.1	52.7
<i>En-X Speech-to-Text Translation, Test BLEU ↑</i>											
Bilingual AV-HuBERT	A	-	-	-	15.9	19.2	17.1	12.9	14.4	10.3	15.0
	AV	-	-	-	22.7	24.8	23.8	20.0	20.0	13.7	20.8
<i>X-En Speech-to-Text Translation, Test BLEU ↑</i>											
Whisper, multilingual, unconstrained data [25]	A	-	-	-	0.1	0.4	0.7	0.1	0.1	0.2	0.3
Bilingual AV-HuBERT	A	-	-	-	3.9	8.6	13.1	7.7	9.5	5.7	8.1
	AV	-	-	-	4.2	12.4	16.5	12.4	15.4	6.1	11.2
Multilingual AV-HuBERT	A	-	-	-	2.9	8.4	12.4	8.1	8.6	0.9	6.9
	AV	-	-	-	4.2	12.8	15.0	12.5	14.8	4.6	10.7

visual mode outperform a comparable ASR baseline (Transformer, monolingual) by average 52% WER reduction. They underperform the hybrid LF-MMI monolingual ASR baselines by average 14% WER increase, which leverage 4-gram language models in decoding while our models do not. Our multilingual AVSR model on average slightly outperforms monolingual AVSR models, with gains observed on some low-resource languages (Ar and De) and degradation observed on the others (El and Ru). We also observe that our multilingual AVSR model falls behind Whisper by average 20% WER due to the large gap in the amount of training data (0.7K hours compared to 680K hours).

#### 4.2.2. Noisy Setup

The first section of Table 4 shows the test WER of our AVSR models in a noisy setup, where we simulate noisy environments by adding multilingual babble noises to clean speech inputs with a SNR (signal-to-noise ratio) of 0<sup>2</sup>. We observe that Whisper, a SOTA multilingual ASR model, performs catastrophically in this challenging setup, with a high average WER of 174.3 over the 9 languages. In contrast, our monolingual AVSR models in the audio-only mode have an average WER of 70.2 and 66.7 respectively. In the audio-visual mode, the average WER of our models drop significantly by 32%, suggesting their efficient use of visual information to alleviate the distraction of noisy environments. Our multilingual AVSR model outperforms the monolingual counterparts on every non-English language (except El) in both audio-only and audio-visual modes.

### 4.3. Audio-Visual Speech-to-Text Translation (AVST)

#### 4.3.1. Clean Setup

We report test BLEU for En-X AVST and X-En AVST models in Table 2 and Table 3, respectively. Besides end-to-end AV-HuBERT AVST models, we also set up bilingual and multilingual MT baselines, and build cascaded AVST baselines by

<sup>2</sup>See Appendix A.2 for model performance under different noise types and different SNRs

pipelining AV-HuBERT AVSR models with MT models. We see that our end-to-end AVST models are on par with the cascaded counterparts in the constrained data setup (“A1+A3” and “B1+B3”) for both En-X and X-En directions. Similar to the case in non-English AVSR, our 6-language multilingual X-En AVST model on average perform slightly better than the corresponding bilingual AVST models.

#### 4.3.2. Noisy Setup

We evaluate our En-X AVST and X-En AVST models in a noisy setup, whose test BLEU are shown in the second and third section of Table 4, respectively. We simulate noisy environments in the same approach as that for AVSR models, where multilingual babble noises are added to clean speech inputs with a SNR of 0. We observe that Whisper, a SOTA multilingual X-En speech-to-text translation model, has a catastrophic performance under this setup, with only 0.3 average BLEU over the 6 directions. Our bilingual and multilingual AVST models in the audio-only mode outperform it largely with 7.8 and 6.6 average BLEU improvement, respectively. AVST models in the audio-visual mode consistently outperform those in the audio-only mode, with an improvement of 3.2 and 3.3 on average BLEU for the bilingual and multilingual settings, respectively.

## 5. Conclusion

We introduce a multilingual audio-visual corpus, MuAViC, for 9 languages totaling 1200 hours of speech. It is the first open benchmark for audio-visual speech-to-text translation and the largest open benchmark for multilingual audio-visual speech recognition. The corpus is available at <https://github.com/facebookresearch/muavic>.

## 6. References

- [1] W. H. Sumby and I. Pollack, “Visual contribution to speech intelligibility in noise,” *Journal of the Acoustical Society of America*, vol. 26, pp. 212–215, 1954.
- [2] T. Afouras, J. S. Chung, A. W. Senior, O. Vinyals, and A. Zisser-

- man, "Deep audio-visual speech recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, pp. 8717–8727, 2018.
- [3] B. Shi, W.-N. Hsu, and A. Mohamed, "Robust Self-Supervised Audio-Visual Speech Recognition," in *Proc. Interspeech 2022*, 2022, pp. 2118–2122.
- [4] A. Bagher Zadeh, Y. Cao, S. Hessner, P. P. Liang, S. Poria, and L.-P. Morency, "CMU-MOSEAS: A multimodal language dataset for Spanish, Portuguese, German and French," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Nov. 2020, pp. 1801–1812.
- [5] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," in *International Conference on Learning Representations*, 2021.
- [6] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. W. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party," *ACM Transactions on Graphics (TOG)*, vol. 37, pp. 1 – 11, 2018.
- [7] A. Fan, S. Bhosale, H. Schwenk, Z. Ma, A. El-Kishky, S. Goyal, M. Baines, O. Celebi, G. Wenzek, V. Chaudhary *et al.*, "Beyond english-centric multilingual machine translation," *Journal of Machine Learning Research*, vol. 22, pp. 1–48, 2021.
- [8] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," *arXiv preprint arXiv:1809.00496*, 2018.
- [9] M. Cooke, J. Barker, S. P. Cunningham, and X. Shao, "An audio-visual corpus for speech perception and automatic speech recognition." *The Journal of the Acoustical Society of America*, vol. 120 5 Pt 1, pp. 2421–4, 2006.
- [10] J. S. Chung and A. Zisserman, "Lip reading in the wild," in *Asian Conference on Computer Vision*, 2016.
- [11] Y. Zhao, R. Xu, and M. Song, "A cascade sequence-to-sequence model for chinese mandarin lip reading," *Proceedings of the ACM Multimedia Asia*, 2019.
- [12] D. Ivanko, A. Axyonov, D. Ryumin, A. Kashevnik, and A. Karpov, "RUSAVIC corpus: Russian audio-visual speech in cars," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, Jun. 2022, pp. 1555–1559.
- [13] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Interspeech*, 2018.
- [14] J. S. Chung, J. Huh, S. Mun, M. Lee, H.-S. Heo, S. Choe, C. Ham, S.-Y. Jung, B.-J. Lee, and I. Han, "In defence of metric learning for speaker recognition," *ArXiv*, vol. abs/2003.11982, 2020.
- [15] W.-N. Hsu and B. Shi, "u-hubert: Unified mixed-modal speech pretraining and zero-shot transfer to unlabeled modality," in *NeurIPS*, 2022.
- [16] T. Afouras, J. S. Chung, and A. Zisserman, "Asr is all you need: Cross-modal distillation for lip reading," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2143–2147, 2019.
- [17] M. Post, G. Kumar, A. Lopez, D. Karakos, C. Callison-Burch, and S. Khudanpur, "Improved speech-to-text translation with the fisher and callhome Spanish-English speech translation corpus," in *Proceedings of the 10th International Workshop on Spoken Language Translation: Papers*, Dec. 5-6 2013.
- [18] A. C. Kocabiyikoglu, L. Besacier, and O. Kraif, "Augmenting librispeech with French translations: A multimodal corpus for direct speech translation evaluation," in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, May 2018.
- [19] C. Wang, J. Pino, A. Wu, and J. Gu, "CoVoST: A diverse multilingual speech-to-text translation corpus," in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, May 2020, pp. 4197–4203.
- [20] C. Wang, A. Wu, J. Gu, and J. Pino, "CoVoST 2 and Massively Multilingual Speech Translation," in *Proc. Interspeech 2021*, 2021, pp. 2247–2251.
- [21] M. A. Di Gangi, R. Cattoni, L. Bentivogli, M. Negri, and M. Turchi, "Must-c: a multilingual speech translation corpus," in *2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, pp. 2012–2017.
- [22] J. Iranzo-Sánchez, J. A. Silvestre-Cerda, J. Jorge, N. Roselló, A. Giménez, A. Sanchis, J. Civera, and A. Juan, "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)*. IEEE, 2020, pp. 8229–8233.
- [23] R. Cattoni, M. A. Di Gangi, L. Bentivogli, M. Negri, and M. Turchi, "Must-c: A multilingual corpus for end-to-end speech translation," *Computer Speech & Language*, vol. 66, p. 101155, 2021.
- [24] E. Salesky, M. Wiesner, J. Bremerman, R. Cattoni, M. Negri, M. Turchi, D. W. Oard, and M. Post, "Multilingual tedx corpus for speech recognition and translation," in *Proceedings of Interspeech*, 2021.
- [25] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.
- [26] N. Reimers and I. Gurevych, "Making monolingual sentence embeddings multilingual using knowledge distillation," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 11 2020.
- [27] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *Proc. Interspeech 2018*, pp. 1086–1090, 2018.
- [28] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [30] C. Wang, Y. Tang, X. Ma, A. Wu, D. Okhonko, and J. Pino, "Fairseq S2T: Fast speech-to-text modeling with fairseq," in *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: System Demonstrations*, Dec. 2020, pp. 33–39.
- [31] M. Post, "A call for clarity in reporting BLEU scores," in *Proceedings of the Third Conference on Machine Translation: Research Papers*, Oct. 2018, pp. 186–191.
- [32] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.