

Zero-resource Speech Translation and Recognition with LLMs

Karel Mundnich, Xing Niu, Prashant Mathur, Srikanth Ronanki, Brady Houston, Veera Raghavendra Elluru, Nilaksh Das, Zejiang Hou, Goeric Huybrechts, Anshu Bhatia, Daniel Garcia-Romero, Kyu J. Han, Katrin Kirchhoff

AWS AI Labs

{kmundnic, xingniu, ronanks}@amazon.com

Abstract—Despite recent advancements in speech processing, zero-resource speech translation (ST) and automatic speech recognition (ASR) remain challenging problems. In this work, we propose to leverage a multilingual Large Language Model (LLM) to perform ST and ASR in languages for which the model has never seen paired audio-text data. We achieve this by using a pre-trained multilingual speech encoder, a multilingual LLM, and a lightweight adaptation module that maps the audio representations to the token embedding space of the LLM. We perform several experiments both in ST and ASR to understand how to best train the model and what data has the most impact on performance in previously unseen languages. In ST, our best model is capable to achieve BLEU scores over 23 in CoVoST2 for two previously unseen languages, while in ASR, we achieve WERs of up to 28.2%. We finally show that the performance of our system is bounded by the ability of the LLM to output text in the desired language.

Index Terms—Zero-resource, ASR, speech translation, LLMs.

I. INTRODUCTION

The goal of zero-resource speech processing is to learn from unlabeled audio (i.e., audio that is not paired to text), similar to what humans do in their infancy before acquiring the ability to read or write [1]. Even with a good deal of efforts over the past decades, zero-resource ASR or ST are still challenging tasks in the more general case (learning from unpaired text and audio), leading to work that splits the general zero-resource problem into simpler tasks [1]. In this work, we focus on the problem of cross-lingual transfer between *seen* and *unseen* languages in the audio data during training, to perform speech processing tasks in *unseen* languages at test time by leveraging cross-lingual transfer and multilingual LLMs.

Recently, speech LMs (SLMs) have shown remarkable performance in speech processing tasks [2]–[6] by leveraging both pre-trained audio encoders [7] as well as the tremendous progress and generalization capabilities of LLMs [8]. The success of these new models can be partly explained by the fact that pre-training techniques are enabling bridging the gap between the vast amounts of unlabeled audio and text. This has led to state-of-the-art performance in several tasks including multilingual ASR [4], [6] and speech translation [9].

Despite the successes of SLMs, most of the work has focused on understanding the generalization capabilities of instruction-tuned LLMs in supervised speech processing tasks, where the spoken languages used at test time are seen during the training of the SLM. Therefore, little to no effort has been put into understanding zero-resource cross-lingual transfer to *unseen* audio languages when the SLM is an end-to-end model composed of a multilingual pre-trained speech encoder and a fine-tuned LLM. For example, in [10] the authors study the zero-shot capabilities of the model in ST for unseen *pairs* of languages at training time, but all the input languages were used at training time for at least one of the tasks (so the model has seen labeled data for all the input languages). Two other works focus on zero-shot translation from English to other languages, so that the *pairs* of languages are previously unseen (but the input is fixed to English). In [9], the authors propose a technique to fine-tune a pre-trained speech encoder using CTC and an optimal transport loss to

map the speech representations directly into the text token embedding space, achieving state-of-the-art results in ST using a multilingual LLM, while in [11], the authors perform zero-shot translation from English to several other languages by proposing a model architecture to find a shared semantic space for audio and text representations. Finally, in [12] the authors propose a method for zero-resource ASR using multilingual transfer, by pre-training an audio encoder and using a universal phone recognizer trained on a different set of languages than the test languages and a deciphering algorithm.

In this paper, we study the capacity of a multilingual LLM to perform zero-resource ST and ASR on languages for which we use no labeled audio data. To achieve this, we pre-train a speech encoder with unlabeled data from over 130 languages, and train an SLM with a lightweight adapter that maps the audio representations to the token embedding space of the LLM. We show that with this approach, it is possible to perform zero-resource ST from previously-unseen languages to English, and to perform ASR in these same unseen languages.

Our contributions are the following: (1) we propose a new approach to perform zero-resource ST and ASR through cross-lingual transfer by leveraging multilingual LLMs; (2) we study the impact of multi-task training (using ST and ASR as tasks) and sequential training in the zero-resource performance; (3) we show that the ability of the LLMs to generate text in the target language is crucial for the ASR performance.

II. METHODOLOGY

Figure 1 shows the model architecture, which consists of a pre-trained speech encoder, an LLM for text generation, and a convolutional network that adapts the audio representations to the text embedding space of the LLM.

A. Speech encoder

We pre-train a large multilingual speech encoder based on the Conformer [13] architecture using BEST-RQ [14]. We use 16 individual codebooks, with a vocabulary size of 8192 and a dimension of 16. The encoder has 24 layers, each with 8 attention heads, and a feature dimension d_{AE} equal to 1024, for a total of 630M parameters. We use a pre-training mask span of 10 with total effective masking ratio of about 40%. We use the transformer learning rate scheduler with a peak value of 0.0005 and 50k steps for warm-up, together with AdamW [15] with weight decay of 0.01. The model is trained for 500k steps in total. The downsampling factor of the model is 4, and the frame length of the output representations is 40ms. The pre-training dataset contains 180K hours of audio from 133 languages, including all the languages that we use for testing in this work.

B. Multilingual LLM

We use the mT0 family of models [16], an encoder-decoder architecture based on mT5 [17], [18]. The mT0 models have been pre-trained and fine-tuned on over 100 languages [16]. We use

the checkpoints available in Huggingface [19] (mT0-XL¹: 3.7B parameters, hidden dimension $d_{LLM} = 2048$ and mT0-XXL²: 12.9B parameters, hidden dimension $d_{LLM} = 4096$, respectively).

C. Adapter module

We use a 2-layer 1-dimensional CNN between the speech encoder and the LLM. We use a kernel size $k = 3$ on both layers. The first layer performs the upsampling in the feature dimension to d_{LLM} and downsampling in the time dimension (40ms to 80ms in frame lengths). The second layer adds depth to the adapter. The module has 18.9M and 62.9M parameters (initialized randomly) when used with mT0-XL and mT0-XXL, respectively.

D. Full model

To connect the speech encoder representations to the LLM, we use a weighted average of the output \mathbf{h}_{AE}^l of all the layers in the speech encoder, such that:

$$AE(\mathbf{x}) = \frac{1}{L} \sum_{l=1}^L w_l \mathbf{h}_{AE}^l(\mathbf{x}_{l-1}) \in \mathbb{R}^{d_{AE} \times t}, \quad (1)$$

where \mathbf{x} is the audio, $\mathbf{x}_l = \mathbf{h}_{AE}^l(\mathbf{x}_{l-1})$ is the output of layer l , $\{w_l, l \in 1, \dots, L\}$ is a set of learnable parameters, and t is the length of the sequence. $AE(\mathbf{x})$ is then passed to the CNN module, such that $A(\mathbf{x}) = CNN(AE(\mathbf{x})) \in \mathbb{R}^{d_{LLM} \times t/2}$. We concatenate these audio sequences with the tokenized text that has been passed through the trained token embedding layer of the LLM, therefore soft-prompting the LLM with audio representations.

E. Multi-task training

Tasks To evaluate the cross-lingual performance of our models for zero-resource ST and ASR, we train our models using either ST and ASR data. When training the model with both ST and ASR data, we are implicitly requesting the model to learn how to perform both tasks at the same time. To achieve this, we use next-token prediction with cross-entropy loss as implemented by Huggingface for both tasks and change the prompt that we use to instruct the model depending on the task.

Train prompts The tokenized text represents an instruction for the model that is task-dependent. In preliminary experiments, we observed better robustness and improved performance when using several prompts at training time. Therefore, for ST we use 25 paraphrased prompts similar to the following instruction: “Perform speech translation into English using the preceding audio:”. For ASR, we use 25 paraphrased prompts similar to the following instruction: “Perform speech recognition in {language} using the preceding audio:”, where {language} is the language of the input speech. These prompts are sampled uniformly at random conditional on the task.

Test prompts We use the following prompts at test time: for ST, we use “Transcribe the content of this audio into English in textual form:”, and for ASR, we use “The preceding audio is in {language}. Perform speech recognition (in {language}):”.

F. Adapter Training

To align speech and text representations, we employ a CNN adapter for speech and incorporate LoRA [20] for text LLMs, as outlined in [21]. These components are jointly trained on speech translation and ASR data. To further improve the alignment, prior to the joint training of adapters, we pre-train the CNN adapter on ASR data. This pre-training in turn facilitates a better initialization for joint training. In

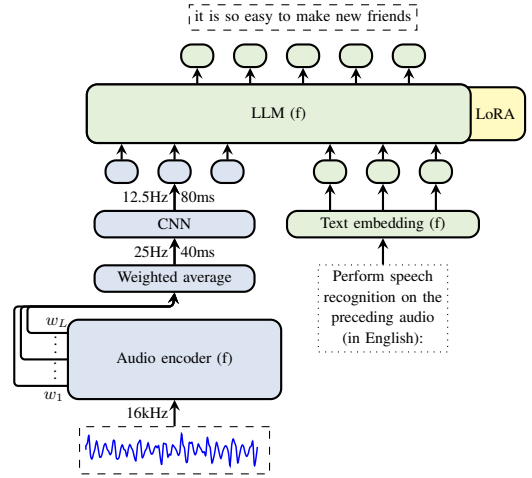


Fig. 1. Model architecture. The audio components are in blue, while the text components are in green. We feed a weighted average of the output of all the layers of the pre-trained audio encoder (Eq. 1) to a 1-dimensional CNN to adapt them to the token embedding space. In terms of trainable parameters, we use (f) to represent frozen weights. We do not freeze the (optional) LoRA weights in the LLM if we use them.

the following section, we assess the impact of training these adapters in a multi-stage manner. In LoRA, we use $\alpha = 10$ and rank $r = 16$, resulting in 9.4M and 18.8M additional parameters for mT0-XL and mT0-XXL models.

III. ZERO-RESOURCE SPEECH TRANSLATION

A. Datasets

For Speech Translation (ST), we used both ASR and ST data from CoVoST2 [22] and Europarl-ST [23]. The training data was limited to the following languages: Portuguese (pt), French (fr), Italian (it), German (de), Spanish (es), and English (en). The X \rightarrow en ($X \in \text{PFIGS}$) ST training data comprised approximately 800 hours of data. Additionally, the ASR partition of the same data added another 800 hours, totaling around 1600 hours of training data. The evaluations were conducted solely on the CoVoST2 dataset, focusing on the Catalan-English (ca \rightarrow en) and Dutch-English (nl \rightarrow en) language pairs. These language pairs are strategically selected for zero-resource testing due to their linguistic characteristics. Catalan, being a Romance language similar to French, Italian, and Spanish, and Dutch, sharing similarities with German, allowed us to effectively assess cross-lingual transfer capabilities of the models. We report BLEU scores [24] to evaluate speech translation quality.

B. Experiments

Prior to assessing the performance of SLMs for speech translation on zero-resource language pairs, we establish an upper bound (i.e., text-only translation) for this work by utilizing the same LLMs employed for training the SLMs. Two different models are evaluated: mT0-XL and mT0-XXL. Furthermore, we include a fine-tuned version of mT0-XL, denoted as mT0-XL-FT, which is trained (with LoRA) on the same text translation data as used for training the SLMs.

We trained several variants of SLMs to evaluate the effectiveness of the parameterization techniques, and pre-training strategies for zero-resource cross-lingual transfer learning in ST.

¹<https://hf.co/bigscience/mt0-xl>

²<https://hf.co/bigscience/mt0-xxl>

TABLE I

SPEECH TRANSLATION RESULTS. THE NUMBERS SHOW BLEU SCORES IN THE CoVoST2 DATASET. THE MODEL HAS NOT SEEN PAIRED AUDIO/TEXT DATA FOR DUTCH (NL) NOR CATALAN (CA). **BOLD FACE** SHOWS THE BEST PERFORMANCE IN ZERO-RESOURCE FOR THE TWO DIFFERENT MODEL SIZES.

Row	Experiment	Model	Pretraining	Trainable Params	Tasks	Trained language pairs					Zero-resource pairs			
						fr→en (14,760)	de→en (13,511)	es→en (13,221)	it→en (8,951)	pt→en (4,023)	Mean (PFIGS)	nl→en (1,699)	ca→en (12,730)	Mean
1	Text-only LLMs	mT0-XL	None	-	-	31.70	29.53	32.76	29.89	40.55	32.89	35.61	24.81	30.21
2		mT0-XXL		-	-	35.70	33.80	36.73	34.10	46.50	37.37	39.44	29.09	34.27
3		mT0-XL-FT		LoRA	MT	46.90	40.48	46.62	43.15	54.00	46.23	43.79	36.78	40.28
4	SLMs (Ours)	mT0-XL	✗	CNN only	ST & ASR	29.46	23.07	29.40	25.49	30.67	27.62	18.01	20.42	19.22
5		mT0-XL	✗	CNN + LoRA	ST & ASR	31.81	25.91	32.05	27.82	30.88	29.69	16.02	21.75	18.88
6		mT0-XL	✓	CNN + LoRA	ST & ASR	33.05	27.41	33.27	28.62	33.60	31.19	18.23	22.89	20.56
7		mT0-XL	✓	CNN + LoRA	ST	33.78	28.10	34.10	29.60	34.66	32.05	17.86	23.50	20.68
8		mT0-XL-FT	✓	CNN + LoRA	ST	34.42	29.14	34.61	30.58	35.75	32.90	17.49	23.81	20.65
9	mT0-XXL	✓	CNN + LoRA	ST & ASR	35.57	32.41	37.08	34.23	45.17	36.89	23.26	25.78	24.52	



Fig. 2. Language confusions for the ASR test languages in Common Voice.

The work conducted by authors in [4] demonstrated that a 2-layer transformer adapter with random subsampling achieves comparable BLEU scores to previous state-of-the-art work in ST. Inspired by this approach, we use light-weight adapters for training SLMs but disentangle the adaptation from representation learning. We utilize CNN parameters for both subsampling and encoding speech representations into textual space, while LoRA parameters are employed for joint learning and decoding. Additionally, we perform pre-training of the CNN layers before fine-tuning both CNN and LoRA parameters together. This pre-training stage aims to enhance the encoding of speech features into textual space and stabilize the training process before introducing LoRA parameters.

C. Results

Table I presents the speech translation results from both text-only LLMs and our spoken language models (SLMs). For the text-only LLMs, we utilized the original ground-truth transcriptions as input. Firstly, comparing rows 1 and 2, the performance of mT0-XXL significantly outperforms mT0-XL across all language pairs, likely due to its increased model size and capacity. Additionally, we observe that mT0-XL-FT, which underwent additional fine-tuning, achieved a 14% absolute improvement across trained language pairs and a 10% improvement across zero-resource pairs. The performance of these models is presented solely to serve as a topline reference that can be achieved with these LLMs. However, it is important to note that these public text-only models may have already encountered the data for zero-resource language pairs, and their naming conventions do not carry significant implications.

Impact of CNN subsampling with LoRA adaptation We compare the performance of various spoken language models trained following the methodology we presented in Section II. Rows 4 and 5 of Table I show the results of trained language pairs with and without LoRA adaptation. We observe a notable improvement with the introduction of LoRA in addition to CNN subsampling. However, the performance

of zero-source pairs degraded slightly, likely due to LoRA parameters being fine-tuned before the CNN features could properly encode speech features, resulting in suboptimal performance for unseen languages. Comparing rows 5 and 6, where we performed additional pre-training of CNN layers before fine-tuning CNN and LoRA, the performance of both trained language pairs and zero-resource language pairs improved significantly. This clearly demonstrates that while LoRA adaptation benefits spoken language models, sequential training with pre-training of CNN layers is crucial for optimal performance.

Multi-task training By comparing the results from rows 6 and 7, we evaluate the performance of speech translation as a standalone task by omitting the ASR task. In contrast to the findings reported in [10], our observations indicate that incorporating ASR data to train the model in a multi-task fashion does not improve performance across all language pairs. Instead, we notice a slightly lower average performance for both seen and unseen languages during training when ASR data is included. This result suggests that if the speech translation (ST) data is sufficient to train the adapter to map the audio representations to the token embedding space, additional ASR data may not be necessary. The results in row 8 indicate that further fine-tuning of the text-only LLM did not lead to significant performance improvements in the spoken language model (SLM) task.

Zero-resource performance As expected, we observe lower zero-resource performance for the models compared to seen languages. However, the larger-scale experiment using mT0-XXL demonstrates that speech translation quality can be comparable to text-based translation in certain scenarios: for Catalan, the ST BLEU score with mT0-XXL (row 9) is remarkably only 4% absolute lower than the text-only translation BLEU score (row 2). This result, however, does not translate to Dutch, where we observe a greater degradation of 16%. We hypothesize that two factors contribute to the worse results for Dutch: (1) Dutch has less phonetic overlap with the languages in the training set (max Jaccard similarity for phonemes in PFIGS+en

languages and Dutch is 0.48 for German, while Catalan’s is 0.79 with Spanish), making generalization more challenging, and (2) as observed in multilingual models, different languages require different training times to be fitted adequately (e.g., in [25], the authors discuss how Dutch is under-fitted while Portuguese, a Romance language, is already fitted).

IV. ZERO-RESOURCE ASR

A. Datasets

For the ASR task, we utilized data from various sources, including Common Voice v14.0 [26], VoxPopuli [27] (labeled portion only), FLEURS [28], LibriVox-MLS [29], and the ASR data from Europarl-ST (since Common Voice and CoVoST2 have significant overlap). For the experiments described in Section IV-B, we created several variants of the ASR datasets, each containing different sets of languages seen during training. However, we use Catalan (ca), Dutch (nl), and Romanian (ro) as the test languages for evaluation. For metrics, we use word error rate (WER). We perform basic normalization of the references (removing punctuation, lowercase, etc.), and directly compare to the output of the LLM.

B. Experiments

We begin by presenting a topline system as a reference, which includes all zero-resource languages during training. Subsequently, we describe a set of experiments to investigate the model’s behavior in the zero-resource scenario under the following conditions: (1) using sequential training, (2) increasing the amount of training data, and (3) varying the number of languages present in the training sets. For all the ASR experiments, we utilize the mT0-XXL model.

ASR topline Similar to the setup in ST, we use PFIGS+en languages for training, sampling up to 500hrs per language from Common Voice, VoxPopuli, FLEURS, and including CoVoST2, Europarl-ST, and Europarl-ST’s ASR data. For the topline, we add up to 350hrs of Catalan, Dutch, and Romanian data, totaling 4,400hrs of speech data.

ASR experiment 1 (ASR-E1) To study how the model performs in the zero-resource setting, we remove Catalan, Dutch, and Romanian data from the topline dataset, totaling 3,500hrs of audio. We train models with and without LoRA adaptation.

ASR experiment 2 (ASR-E2) In preliminary results, we observe that the model does not always output the correct language in the ASR task. Therefore, to understand the role of the dataset size and number of training languages, we create a dataset with 36 languages from Common Voice, Voxpopuli, FLEURS, and LibriVox-MLS. We cap the number of hours for each language between 350hrs and 20hrs, for a total of 50 languages. Out of these 50 languages, we keep 36 that were used in the training of mT0 (not including Catalan, Dutch, and Romanian). We also add the ST data from CoVoST2 and ST+ASR data from Europarl-ST, for a total of about 9,000hrs of data.

ASR experiment 3 (ASR-E3) To further study the impact of the number of languages in performance and language confusion, we follow the same steps as in ASR-E2, but create a dataset with 84 languages by capping the maximum and minimum number of hours to 350hrs and 10hrs, respectively. We obtain 67 languages that overlap with mT0 training. After adding ST data, we obtain about 9,000hrs of training data.

C. Results

Table II presents the results on zero-resource speech recognition across different experiments.

TABLE II
ZERO-RESOURCE ASR RESULTS IN COMMON VOICE. THE LANGUAGE ACCURACY SHOWS THE PERCENTAGE OF HYPOTHESES GENERATED IN THE PROMPTED LANGUAGE. [†] REPRESENTS THE WER CALCULATED BY FIRST FILTERING OUT THE HYPOTHESES THAT ARE IN THE WRONG LANGUAGE.

Experiment	LoRA	Eval.	Language Accuracy				WER			
			ca	nl	ro	Mean	ca	nl	ro	Mean
Topline	✗	All	0.92	0.89	0.98	0.93	25.6	13.9	26.9	22.1
ASR-E1	✗	All	0.91	0.80	0.89	0.87	44.7	37.8	64.2	48.9
	✗	Subset [†]	–	–	–	–	41.3	31.6	60.4	44.4
	✓	All	0.57	0.28	0.73	0.53	51.8	75.5	57.9	61.8
	✓	Subset [†]	–	–	–	–	31.6	39.0	45.8	38.8
ASR-E2	✗	All	0.91	0.83	0.97	0.90	48.7	32.3	52.5	42.4
	✗	Subset [†]	–	–	–	–	41.4	29.6	37.1	36.0
ASR-E3	✗	All	0.92	0.86	0.98	0.92	53.6	30.5	50.1	42.6
	✗	Subset [†]	–	–	–	–	38.1	28.2	43.6	36.6

Language confusion We initially observed that the model does not consistently output text in the desired language. Figure 2 presents a confusion matrix for all three zero-resource languages against all other languages used in training. While Catalan is confused with Spanish, Italian, and Portuguese, Dutch is highly confused with Afrikaans. To quantify this behavior, we employ `langdetect` [30] to detect the output language and report the language accuracy for each test language and experiment in Table II. Even in the topline model, where we included the test languages during training, the language accuracy ranges from 89% (Dutch) to 98% (Romanian). In the zero-resource scenarios, the language accuracy is further reduced. To decouple the effect of language confusions from the ASR capabilities, we filter out the hypotheses in the wrong language and re-score, presenting both the unfiltered and filtered results.

Sequential training In ST task, we observed that sequential training with pre-training of CNN layers is crucial for optimal performance. However, for the experimental results presented in ASR-E1, we observe that using LoRA adaptation in addition to CNN subsampling degrades the performance of the model. This can be explained by the model losing generalization capabilities to unseen languages, giving the worst results across our experiments.

Larger scale We finally show the results for the experiments where we use more hours of data and more languages in the training set (36 languages in ASR-E2 and 67 languages in ASR-E3). In both of these experiments we observe WER improvements over the previous experiments. In ASR-E3, we observe that the language confusion problem is reduced, reaching the same language accuracy for both Catalan and Romanian than the baseline. However, this is not reflected in the WER, where the model trained in ASR-E2 shows the best WER performance.

V. CONCLUSIONS

In this paper, we propose a method to perform zero-resource speech translation and automatic speech recognition. We use a pre-trained speech encoder and a multilingual LLM, and show that training a lightweight adapter between the output of the speech encoder and the token embedding space of the LLM allows to perform speech tasks in languages for which the model has never seen paired audio-text data.

REFERENCES

- [1] E. Dunbar, N. Hamilakis, and E. Dupoux, “Self-supervised language learning from raw audio: Lessons from the zero resource speech chal-

- lenge,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1211–1226, 2022.
- [2] H. Gao, J. Ni, K. Qian, Y. Zhang, S. Chang, and M. Hasegawa-Johnson, “WavPrompt: Towards Few-Shot Spoken Language Understanding with Frozen Language Models,” in *Proc. Interspeech 2022*, 2022, pp. 2738–2742.
 - [3] Y. Fathullah, C. Wu, E. Lakomkin, J. Jia, Y. Shangguan, K. Li, J. Guo, W. Xiong, J. Mahadeokar, O. Kalinli *et al.*, “Prompting large language models with speech recognition abilities,” *arXiv preprint arXiv:2307.11795*, 2023.
 - [4] M. Wang, W. Han, I. Shafran, Z. Wu, C.-C. Chiu, Y. Cao, N. Chen, Y. Zhang, H. Soltau, P. K. Rubenstein *et al.*, “SLM: Bridge the thin gap between speech and text foundation models,” in *ASRU 2023 Workshop*. IEEE, 2023, pp. 1–8.
 - [5] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, “Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models,” *arXiv preprint arXiv:2311.07919*, 2023.
 - [6] G. Team, R. Anil, S. Borgeaud, Y. Wu, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
 - [7] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
 - [8] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa, “Large language models are zero-shot reasoners,” *Advances in neural information processing systems*, vol. 35, pp. 22 199–22 213, 2022.
 - [9] I. Tsiamas, G. I. Gállego, J. A. Fonollosa, and M. R. Costa-jussà, “Pushing the limits of zero-shot end-to-end speech translation,” *arXiv preprint arXiv:2402.10422*, 2024.
 - [10] P. K. Rubenstein, C. Asawaroengchai, D. D. Nguyen, A. Bapna, Z. Borsos, F. d. C. Quiry, P. Chen, D. E. Badawy, W. Han, E. Kharitonov *et al.*, “AudioPaLM: A large language model that can speak and listen,” *arXiv preprint arXiv:2306.12925*, 2023.
 - [11] C. Wang, Y. Liu, B. Chen, J. Zhang, W. Luo, Z. Huang, and C. Zong, “Discrete cross-modal alignment enables zero-shot speech translation,” *arXiv preprint arXiv:2210.09556*, 2022.
 - [12] O. Klejch, E. Wallington, and P. Bell, “Deciphering Speech: a Zero-Resource Approach to Cross-Lingual Transfer in ASR,” in *Proc. Interspeech 2022*, pp. 2288–2292.
 - [13] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Proc. Interspeech 2020*, 2020, pp. 5036–5040.
 - [14] C.-C. Chiu, J. Qin, Y. Zhang, J. Yu, and Y. Wu, “Self-supervised learning with random-projection quantizer for speech recognition,” in *ICML*, 2022, pp. 3915–3924.
 - [15] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*, 2019.
 - [16] N. Muennighoff, T. Wang, L. Sutawika, A. Roberts, S. Biderman, T. L. Scao, M. S. Bari, S. Shen, Z.-X. Yong, H. Schoelkopf *et al.*, “Crosslingual generalization through multitask finetuning,” *arXiv preprint arXiv:2211.01786*, 2022.
 - [17] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
 - [18] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Siddhant, A. Barua, and C. Raffel, “mT5: A massively multilingual pre-trained text-to-text transformer,” in *Proceedings of the 2021 Conference of NAACL-HLT*, 2021, pp. 483–498.
 - [19] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, “Huggingface’s transformers: State-of-the-art natural language processing,” *arXiv preprint arXiv:1910.03771*, 2019.
 - [20] E. J. Hu, yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” in *International Conference on Learning Representations*, 2022.
 - [21] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. Glass, “Listen, think, and understand,” *International Conference on Learning Representations*, 2024.
 - [22] C. Wang, A. Wu, and J. Pino, “CoVoST 2 and Massively Multilingual Speech-to-text Translation,” *arXiv preprint arXiv:2007.10310*, 2020.
 - [23] 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*. IEEE, 2020, pp. 8229–8233.
 - [24] M. Post, “A call for clarity in reporting BLEU scores,” in *Proceedings of the Third Conference on Machine Translation: Research Papers*, 2018, pp. 186–191.
 - [25] J. Bai, B. Li, Q. Li, T. N. Sainath, and T. Strohman, “Efficient adapter finetuning for tail languages in streaming multilingual asr,” *arXiv preprint arXiv:2401.08992*, 2024.
 - [26] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” *arXiv preprint arXiv:1912.06670*, 2019.
 - [27] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talmnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, “VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” in *Proc. of ACL and the 11th IJCNLP (Volume 1: Long Papers)*, 2021, pp. 993–1003.
 - [28] A. Conneau, M. Ma, S. Khanuja, Y. Zhang, V. Axelrod, S. Dalmia, J. Riesa, C. Rivera, and A. Bapna, “Fleurs: Few-shot learning evaluation of universal representations of speech,” in *SLOT 2022 Workshop*. IEEE, 2023, pp. 798–805.
 - [29] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, “MLS: A large-scale multilingual dataset for speech research,” *arXiv preprint arXiv:2012.03411*, 2020.
 - [30] M. Danilk, “langdetect,” 2024, <https://pypi.org/project/langdetect/>. Accessed: March 1, 2024.