



Comparison of Multilingual Self-Supervised and Weakly-Supervised Speech Pre-Training for Adaptation to Unseen Languages

Andrew Rouditchenko¹, Sameer Khurana¹, Samuel Thomas^{2,3}, Rogerio Feris^{2,3}, Leonid Karlinsky^{2,3}, Hilde Kuehne^{3,4}, David Harwath⁵, Brian Kingsbury^{2,3}, James Glass¹

¹MIT, USA ²IBM Research AI, USA ³MIT-IBM Watson AI Lab, USA ⁴Goethe University Frankfurt, Germany ⁵UT Austin, USA

roudi@mit.edu

Abstract

Recent models such as XLS-R and Whisper have made multilingual speech technologies more accessible by pre-training on audio from around 100 spoken languages each. However, there are thousands of spoken languages worldwide, and adapting to new languages is an important problem. In this work, we aim to understand which model adapts better to languages unseen during pre-training. We fine-tune both models on 13 unseen languages and 18 seen languages. Our results show that the number of hours seen per language and language family during pre-training is predictive of how the models compare, despite the significant differences in the pre-training methods.

Index Terms: speech recognition, multilingual, self-supervised, weakly-supervised

1. Introduction

Multilingual speech processing has seen rapid progress thanks to the availability of more data and larger models. Models like XLS-R [1] and Whisper [2] are pre-trained on hundreds of thousands of hours of audio from around 100 spoken languages. Downstream users can adapt these publicly available models for Automatic Speech Recognition (ASR) on different languages and domains [3–5]. XLS-R follows Wav2Vec2.0’s [6] self-supervised pre-training process, which only requires unlabeled speech, and it can be adapted for ASR with a small amount of labeled data. In contrast, Whisper was pre-trained with labeled speech-text pairs acquired by scraping the web and can be directly applied to ASR in many languages without requiring fine-tuning.

Although these models are pre-trained on many languages, they still omit several thousand spoken languages during pre-training, which limits their applicability to these unseen languages. Many of these unseen languages have over a million speakers, making them sufficiently widespread to be of importance. We believe more audio in these languages will become available through projects like CommonVoice [7]. However, it is infeasible to repeat pre-training each time more data is released since the compute required is only accessible to the largest corporations. For example, XLS-R [1] was pre-trained on 128 GPUs for smaller models and 200 GPUs for larger models. Therefore, we find it essential to study the adaptation of these pre-trained models to unseen languages.

In this work, we compare XLS-R and Whisper for adaptation to languages unseen during pre-training. We fine-tuned the models on 13 unseen languages in the FLEURS dataset [8], on which neither model was pre-trained. As shown in Figure 1, the models have seen different amounts of languages during pre-training, and we discuss how this impacts the performance on the unseen languages. We also fine-tuned on 18 seen lan-




Figure 1: Hours of audio used to pre-train Whisper [2] and XLS-R [1] per language family.

guages in FLEURS and found that the number of hours seen per language during pre-training are predictive of how the models compare.

In summary, our contributions are: 1) We compare two state-of-the-art multilingual speech models with self-supervised and weakly-supervised pre-training strategies for adaptation to 13 unseen and 18 seen languages through fine-tuning; 2) We analyze the impact of the pre-training data and give suggestions for selecting either model for adaptation. Please see our ArXiv paper version for additional results in the Appendix.

2. Methods

Multilingual speech models have recently been trained on around 100 languages [9–15]. However, few models are publicly available. We use XLS-R [1] and Whisper [2] since they are publicly available, were trained on the most languages and hours of audio, and have different pre-training strategies.

2.1. Self-Supervised Learning: XLS-R

Self-Supervised Learning has proved to be effective for training ASR models. Wav2Vec2.0 [6] was trained with large amounts of unlabeled audio and fine-tuned on smaller, labeled datasets to achieve good performance. XLSR53, a multilingual version of Wav2Vec2.0, was trained on 50k hours of audio in 53 languages and fine-tuned for ASR in different languages. Multilingual self-supervised training usually outperformed monolingual training, especially for lower-resource languages. XLS-R [1] scaled the pre-training to 436k hours of data in 128 languages, using publicly available datasets such as VoxPopuli [16], CommonVoice [7], Babel [17], and Multilingual LibriSpeech [18].

2.2. Weakly-Supervised Learning: Whisper

Although Wav2Vec2.0 was pre-trained with unlabeled data which is easier to acquire than labeled data, it needs to be fine-tuned on each dataset with potentially different hyperparameters. Radford et al. [2] addressed this issue by developing Whisper, a supervised model trained on multilingual speech that can perform “zero-shot” ASR on different datasets (without fine-tuning). The authors trained Whisper on 680k hours of speech from 99 languages. In their experiments on English datasets, Whisper outperformed Wav2Vec2.0 without any fine-tuning. Recent work also finds that Whisper can outperform other state-of-the-art ASR models when fine-tuned [19].

When tested on multilingual ASR, Whisper performed well on high-resource languages but performed poorly on languages it saw few hours of during pre-training. We extend this analysis by comparing Whisper to XLS-R on multilingual ASR. We also fine-tune Whisper, improving performance on seen languages and enabling adaptation to unseen languages.

2.3. Language Adaptation

Language adaptation aims to use representations learned from high-resource languages to improve performance on low-resource languages. No audio in the target language is available in the zero acoustic resource setting [17]. However, a more realistic setting is to assume a limited amount of transcribed audio. Transferring representations from fully-supervised English models to other languages is a common setup [7, 20], while others transfer from multilingual models. Early work transferred from models trained on three languages [21] while others scaled training to 100 languages [9]. As self-supervised models became more common, cross-lingual transfer from methods like Contrastive Predictive Coding [22] were also explored [23]. More recently, XLS-R has been shown to outperform Wav2Vec2.0 for multilingual cross-lingual transfer [3–5].

Several prior studies also fine-tune XLS-R and other self-supervised models on low-resource languages [5, 24]. Our work differs since we focus on unseen languages and distinguish them from seen languages, we analyze how the pre-training data impacts the results, and we compare with Whisper.

Finally, we focus on adaptation to new languages and not on the *language forgetting* phenomenon in the models. Language forgetting could be studied as future work, with approaches including adaptors [20, 25] and weight factorization [26].

3. Experimental Setup

Dataset. We use the FLEURS [8] dataset, which covers 102 languages. We selected this dataset since neither XLS-R nor Whisper was pre-trained on it, there is a sufficient number of unseen languages, the data is approximately equal per language, and the data was verified by paid annotators. The dataset contains around 2k sentences per language, and each sentence is spoken by up to 3 different speakers for a total of around 12 hours of audio per language. The text was originally sourced from English Wikipedia and translated into other languages as the FLoRes-101 benchmark for machine translation [27]. For each language, we fine-tune the models on the training set, select the best checkpoint on the validation set, and report the performance on the test set. We evaluate models with the Character Error Rate (CER) since languages such as Chinese, Japanese, Thai, and Burmese do not use regular spacing, which is consistent with prior work [7, 8]. We show WER results and results on the CommonVoice [7] dataset in our ArXiv paper version.

Languages. We used the 13 unseen languages in FLEURS which are common to XLS-R and Whisper. We also used 18 seen languages. We picked a roughly equal number of seen languages for which XLS-R saw more hours of audio during pre-training than Whisper, and vice-versa. The languages use different scripts; the Latin script is used by 12 of the unseen languages and 8 of the seen languages.

Model Sizes. XLS-R ranges in size from 24 layers / 317M parameters to 48 layers / 2.16B parameters, while Whisper ranges in size from 4 layers / 39M parameters to 32 layers / 1.55B parameters. We compare XLS-R 317M with Whisper models with the closest number of parameters: 12 layers / 244M parameters (Whisper-Small) and 24 layers / 769M parameters (Whisper-Medium). We show the results for XLS-R 965M and Whisper-Large-V2 in the Appendix of our ArXiv paper version.

Tokenization. For XLS-R, we used the language-specific characters from the FLEURS text as the outputs for ASR training. For Whisper, we used the token IDs from the byte-level BPE tokenizer for ASR training, which can encode any UTF-8 text by converting characters to bytes. This results in no out-of-vocabulary characters, although the encoding might not be efficient since one character of multiple bytes could be encoded as multiple tokens. Whisper also expects a language token in the decoder input that is either predicted by its audio-based language ID system or manually specified. Since no token corresponds to the unseen languages, we tried setting the token equal to English or a language similar to each unseen language. Both strategies resulted in similar results, so we used English as the language token for all unseen languages.

Hyperparameters. XLS-R is fine-tuned with the CTC objective [28], and Whisper is fine-tuned with the Cross-Entropy objective. We used Fairseq [29] and HuggingFace [30] for fine-tuning XLS-R and Whisper, respectively, and adopted most of the default hyperparameters, with the following changes. We fine-tuned XLS-R for 20k steps with a learning rate of $5e-5$ and a gradient accumulation factor of 5. We used a batch of around 1 minute of transcribed speech. Unlike [1], we did not find it necessary to freeze the Transformer [31] encoder’s parameters at the beginning of training. For Whisper, we used a batch of 48 recordings for the small model and 24 recordings for the medium model, and fine-tuned for 2k steps with a learning rate of $1e-5$. We used a single V100 32GB GPU, and each experiment took around 12 hours.

Decoding. We used greedy search for ASR inference with XLS-R. We tried beam search without a language model (LM) but found no improvement. We do not train LMs so that our results are independent of external modules used in the decoding process, and also since FLEURS contains little text to train them. For inference with Whisper, we used greedy search for a fair comparison. However, as discussed in Section 4.2, beam search decoding gives a slight improvement.

Text Normalization. Since Whisper was trained with unnormalized transcripts, it generates transcripts with punctuation. Hence, Radford et al. [2] introduced a multilingual text normalizer to standardize both the Whisper transcripts and evaluation transcripts by removing punctuation. The text in FLEURS was already normalized; however, we still found some punctuation and irrelevant characters. Therefore, we applied the Whisper normalizer to all model-generated and ground-truth transcripts during evaluation to make a fair comparison between models with and without fine-tuning. The Whisper normalizer erroneously introduces spaces to transcripts in a few languages; we discuss the impact of this in the Appendix of our ArXiv paper version.

Table 1: Results on the FLEURS dataset. XLS-R 317M, Whisper-Small 244M, and Whisper-Medium 769M are fine-tuned on each language individually. CER=Character Error Rate (\downarrow is better); PT=Pre-Train; FT=Fine-Tune; ZT=Zero-Shot (no fine-tuning).

Language	Language Family	XLS-R PT Hours	Whisper PT Hours	XLS-R FT CER	Whisper-S FT CER	Whisper-M FT CER	Whisper-S ZT CER	Whisper-M ZT CER
Seen Languages								
<i>XLS-R PT Hours < Whisper PT Hours</i>								
English	Indo-European	69,493	438,218	6.45	3.63	2.72	3.18	2.81
Mandarin Chinese	Sino-Tibetan	90	23446	20.25	<u>9.55</u>	7.77	25.84	16.06
Korean	Koreanic	61	7993	15.92	9.31	5.27	7.56	<u>5.78</u>
Japanese	Japonic	49	7054	24.84	12.10	7.22	14.47	<u>9.18</u>
Indonesian	Austronesian	41	1014	3.57	4.19	2.85	6.83	<u>3.33</u>
Arabic	Afro-Asiatic	95	739	<u>5.99</u>	6.86	4.80	9.64	6.08
Ukrainian	Indo-European	72	697	4.16	4.81	3.20	6.39	<u>3.89</u>
Vietnamese	Austro-Asiatic	96	691	10.54	9.42	6.42	11.05	<u>7.01</u>
Thai	Kra-Dai	20	226	<u>10.29</u>	12.65	8.82	24.05	17.75
<i>XLS-R PT Hours \approx Whisper PT Hours</i>								
Azerbaijani	Turkic	47	47	<u>5.84</u>	8.05	5.25	19.62	12.54
<i>XLS-R PT Hours > Whisper PT Hours</i>								
Czech	Indo-European	18514	192	<u>5.00</u>	8.51	4.69	11.30	6.43
Maltese	Afro-Asiatic	9120	1.1	4.08	8.80	<u>5.78</u>	100.26	88.24
Bengali	Indo-European	100	1.3	6.04	11.08	<u>7.67</u>	114.56	111.30
Swahili	Atlantic-Congo	91	5.4	<u>4.88</u>	6.91	4.71	44.40	37.67
Afrikaans	Indo-European	87	4.1	<u>9.31</u>	13.83	9.14	29.80	21.24
Hindi	Indo-European	65	12	5.94	8.21	<u>6.13</u>	45.46	24.16
Khmer	Austro-Asiatic	33	1.3	12.49	27.18	<u>18.63</u>	133.97	112.02
Burmese	Sino-Tibetan	33	0.1	11.58	30.46	<u>27.58</u>	152.30	178.60
<i>Average CER for Seen Languages</i>				<u>9.29</u>	10.86	7.70		
Unseen Languages								
Asturian	Indo-European	0	0	<u>5.29</u>	7.09	5.02	N/A	N/A
Kabuverdianu	Indo-European	0	0	<u>4.65</u>	5.64	4.00	N/A	N/A
Sorani Kurdish	Indo-European	0	0	7.75	12.57	<u>9.89</u>	N/A	N/A
Oromo	Afro-Asiatic	0	0	16.23	18.95	<u>17.57</u>	N/A	N/A
Fula	Atlantic-Congo	0	0	16.07	20.87	<u>16.26</u>	N/A	N/A
Kamba	Atlantic-Congo	0	0	12.87	19.20	<u>16.50</u>	N/A	N/A
Sotho	Atlantic-Congo	0	0	7.35	13.90	<u>11.16</u>	N/A	N/A
Nyanja	Atlantic-Congo	0	0	8.64	<u>13.37</u>	14.69	N/A	N/A
Wolof	Atlantic-Congo	0	0	14.59	17.60	<u>14.71</u>	N/A	N/A
Xhosa	Atlantic-Congo	0	0	6.34	12.48	<u>8.52</u>	N/A	N/A
Igbo	Atlantic-Congo	0	0	12.56	19.34	<u>19.64</u>	N/A	N/A
Umbundu	Atlantic-Congo	0	0	18.13	<u>24.69</u>	<u>19.58</u>	N/A	N/A
Luo	Nilo-Saharan	0	0	5.96	8.55	<u>6.54</u>	N/A	N/A
<i>Average CER for Unseen Languages</i>				10.49	14.94	<u>12.62</u>		

4. Results

4.1. Seen Languages

In Table 1, we show the results of fine-tuning XLS-R and Whisper on the seen and unseen languages in FLEURS. We also show the Whisper results on seen languages without fine-tuning. Both Whisper-Small and Whisper-Medium usually improved after fine-tuning. The improvements are significant for languages with only a few hours of audio seen during pre-training; for example, CER improved from 88.24 to 5.78 for Whisper-Medium on Maltese (1.1 hours seen during pre-training). However, for high-resource languages like English and Korean, the improvement was small, and the performance slightly decreased in some cases. Whisper-Medium outperformed Whisper-Small on all seen languages both with and without fine-tuning.

Comparing XLS-R to Whisper, the results tend to depend on the number of hours seen per language during pre-training. Whisper-Medium outperformed XLS-R both with and without fine-tuning on languages for which it saw more hours of audio during pre-training. Whisper-Medium fine-tuned also outperformed XLS-R on some languages for which it saw fewer

hours of audio during pre-training, such as Czech (4.69 vs 5.00) and Swahili (4.71 vs 4.88). However, XLS-R outperformed Whisper-Medium on languages that it had seen far more of during pre-training, such as Hindi (5.94 vs 6.13) and Khmer (12.49 vs 18.63). XLS-R outperformed Whisper-Small on most languages including those that Whisper had seen more hours of during pre-training, such as Indonesian (3.57 vs 4.19) and Arabic (5.99 vs 6.86). However, Whisper-Small outperformed XLS-R on languages that it had seen far more hours of during pre-training, such as Mandarin (9.55 vs 20.25) and Japanese (12.10 vs 24.84).

In Figure 2, we compare the performance of XLS-R to Whisper-Small and XLS-R to Whisper-Medium based on the hours of pre-training data for each language. The y-axis measures the relative CER and the x-axis measures the relative number of pre-training hours per language on a log-log scale. Based on both plots, we observe a medium-to-strong correlation between the log of the relative CER and the log of the relative number of pre-training hours, which could be helpful for predicting whether XLS-R or Whisper will perform better on other seen languages. The first line of best fit predicts that Whisper-Small will outperform XLS-R on languages for which it has




Figure 2: Comparing the performance of XLS-R 317M to Whisper Small 244M (left) and XLS-R 317M to Whisper Medium 769M (right) on seen languages in FLEURS. X-axis: $\log_{10}\left(\frac{\text{XLS-R \#PT Hours}}{\text{Whisper \#PT Hours}}\right)$. Y-axis: $\log_{10}\left(\frac{\text{XLS-R CER}}{\text{Whisper CER}}\right)$.

seen at least 3.47 times more hours of audio than XLS-R during pre-training. The second line of best fit predicts that XLS-R will outperform Whisper-Medium on languages for which it has seen at least 9.23 times more hours of audio during pre-training. This result is intuitive since it predicts that a smaller model would outperform a larger model on languages for which it has been trained with far more hours of audio. It is also interesting that the raw number of hours seen per language is predictive of the models’ relative performance, despite the significant differences between XLS-R’s self-supervised pre-training and Whisper’s weakly supervised pre-training.

4.2. Unseen Languages

The bottom half of Table 1 shows the results of fine-tuning XLS-R and Whisper on the unseen languages. We also attempted to test Whisper without fine-tuning but found that the transcripts were incomprehensible. XLS-R outperformed Whisper-Small on all languages and outperformed Whisper-Medium on all languages except for Asturian and Kabuverdianu. The models tend to perform worse on the unseen languages compared to the seen languages based on the average CER, but XLS-R’s average CER only decreased by 12.9% relative, while Whisper-Medium’s average CER decreased by 63.9% relative.

To understand the impact of the models’ pre-training data on performance, we plotted the hours of audio seen during pre-training per language family in Figure 1. Language families consist of languages that have a common ancestral language. Whisper saw more audio in high-resource language families such as Indo-European (Whisper: 506k hours vs XLS-R: 379k hours), while XLS-R saw more audio in lower-resource language families such as Afro-Asiatic (Whisper: 1.4k hours versus XLS-R: 10k hours) and Atlantic Congo (Whisper: 5.4 hours versus XLS-R: 1.5k hours). In our experiments, the model which saw more data in the corresponding language family usually performed better on the unseen language. XLS-R did better on the Atlantic-Congo, Afro-Asiatic, and Nilo-Saharan languages, consistently outperforming the larger Whisper-Medium model. Although neither model was pre-trained on Nilo-Saharan languages, the family is similar to Atlantic-Congo which explains XLS-R’s better performance on Luo. Whisper-Medium outperformed XLS-R on Indo-European languages, a

family that Whisper had seen more data from, except for Sorani Kurdish. Overall, these results show that the distribution of hours seen per language family can be useful to make predictions about how XLS-R and Whisper compare on unseen languages, despite the different pre-training methods.

For seen languages, the hours seen per language during pre-training appears to be a more important factor than the hours seen per family. For example, XLS-R outperformed Whisper-Medium on languages that XLS-R had seen more hours of during pre-training, such as Hindi (Indo-European), Khmer (Austro-Asiatic), and Burmese (Sino-Tibetan), even though Whisper saw more hours in those families. Similarly, XLS-R outperformed Whisper-Small on Indo-European languages that XLS-R had seen more hours of during pre-training, such as Czech, Bengali, and Afrikaans, even though Whisper saw more hours from the family in total.

Finally, we attempted to improve Whisper’s performance on the unseen languages by using beam search instead of greedy search. With a beam size of 5, Whisper’s performance improved by around 3.34% relative or 0.42 absolute on average, but the improvement was never large enough to outperform XLS-R.

5. Conclusion

We compared XLS-R and Whisper for adaptation to unseen languages not included during pre-training. XLS-R largely outperformed Whisper on unseen languages, which we attribute to more data seen in the related language families during pre-training. The hours per language seen during pre-training was predictive of how the models performed on seen languages. Our results confirm that diverse pre-training data is essential for generalization and should ideally include more languages. Since unlabeled audio is easier to collect, future self-supervised models like XLS-R may have greater potential to outperform supervised models like Whisper on other unseen languages, with the limitation that self-supervised models can’t perform ASR without fine-tuning. Future work could try ensembles of the two models or different tasks such as speech translation.

6. Acknowledgements

This research was supported by the MIT-IBM Watson AI Lab, the MIT SuperCloud, and an NDSEG Fellowship to A.R.

7. References

- [1] A. Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino *et al.*, “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [2] 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.
- [3] S. Khurana, A. Laurent, and J. Glass, “Magic dust for cross-lingual adaptation of monolingual wav2vec-2.0,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6647–6651.
- [4] D. Krishna, P. Wang, and B. Bozza, “Using large self-supervised models for low-resource speech recognition,” in *Interspeech*, 2021, pp. 2436–2440.
- [5] J. Zhao and W.-Q. Zhang, “Improving automatic speech recognition performance for low-resource languages with self-supervised models,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1227–1241, 2022.
- [6] 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.
- [7] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” in *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 4218–4222.
- [8] 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 *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 798–805.
- [9] O. Adams, M. Wiesner, S. Watanabe, and D. Yarowsky, “Massively multilingual adversarial speech recognition,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 96–108.
- [10] V. Pratap, A. Sriram, P. Tomasello, A. Hannun, V. Liptchinsky, G. Synnaeve, and R. Collobert, “Massively Multilingual ASR: 50 Languages, 1 Model, 1 Billion Parameters,” in *Proc. Interspeech 2020*, 2020, pp. 4751–4755.
- [11] B. Li, R. Pang, Y. Zhang, T. N. Sainath, T. Strohmaier, P. Haghani, Y. Zhu, B. Farris, N. Gaur, and M. Prasad, “Massively multilingual asr: A lifelong learning solution,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6397–6401.
- [12] A. Tjandra, N. Singhal, D. Zhang, O. Kalinli, A. Mohamed, D. Le, and M. L. Seltzer, “Massively multilingual asr on 70 languages: Tokenization, architecture, and generalization capabilities,” *arXiv preprint arXiv:2211.05756*, 2022.
- [13] L. Lugosch, T. Likhomanenko, G. Synnaeve, and R. Collobert, “Pseudo-labeling for massively multilingual speech recognition,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7687–7691.
- [14] A. Bapna, C. Cherry, Y. Zhang, Y. Jia, M. Johnson, Y. Cheng, S. Khanuja, J. Riesa, and A. Conneau, “msslam: Massively multilingual joint pre-training for speech and text,” *arXiv preprint arXiv:2202.01374*, 2022.
- [15] Y. Zhang, W. Han, J. Qin, Y. Wang, A. Bapna, Z. Chen, N. Chen, B. Li, V. Axelrod, G. Wang *et al.*, “Google usm: Scaling automatic speech recognition beyond 100 languages,” *arXiv preprint arXiv:2303.01037*, 2023.
- [16] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, 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 *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*. Association for Computational Linguistics, Aug. 2021, pp. 993–1003.
- [17] M. J. Gales, K. M. Knill, A. Ragni, and S. P. Rath, “Speech recognition and keyword spotting for low-resource languages: Babel project research at cued,” in *Fourth International workshop on spoken language technologies for under-resourced languages (SLTU-2014)*. International Speech Communication Association (ISCA), 2014, pp. 16–23.
- [18] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, “MLS: A Large-Scale Multilingual Dataset for Speech Research,” in *Proc. Interspeech 2020*, 2020, pp. 2757–2761.
- [19] S. Gandhi, P. Von Platen, and A. M. Rush, “Esb: A benchmark for multi-domain end-to-end speech recognition,” *arXiv preprint arXiv:2210.13352*, 2022.
- [20] S. Kessler, B. Thomas, and S. Karout, “An adapter based pre-training for efficient and scalable self-supervised speech representation learning,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 3179–3183.
- [21] F. Grézl, M. Karafiát, and K. Veselý, “Adaptation of multilingual stacked bottle-neck neural network structure for new language,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 7654–7658.
- [22] A. v. d. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [23] M. Riviere, A. Joulin, P.-E. Mazaré, and E. Dupoux, “Unsupervised pretraining transfers well across languages,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7414–7418.
- [24] N.-Q. Pham, A. Waibel, and J. Niehues, “Adaptive multilingual speech recognition with pretrained models,” in *Proc. Interspeech 2022*, 2022, pp. 3879–3883.
- [25] W. Hou, H. Zhu, Y. Wang, J. Wang, T. Qin, R. Xu, and T. Shinzaki, “Exploiting adapters for cross-lingual low-resource speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 317–329, 2021.
- [26] N.-Q. Pham, J. Niehues, and A. Waibel, “Continually learning new languages,” *arXiv preprint arXiv:2211.11703*, 2022.
- [27] N. Goyal, C. Gao, V. Chaudhary, P.-J. Chen, G. Wenzek, D. Ju, S. Krishnan, M. Ranzato, F. Guzmán, and A. Fan, “The Flores-101 evaluation benchmark for low-resource and multilingual machine translation,” *Transactions of the Association for Computational Linguistics*, vol. 10, pp. 522–538, 2022.
- [28] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [29] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, “fairseq: A fast, extensible toolkit for sequence modeling,” in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [30] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. Rush, “Transformers: State-of-the-art natural language processing,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Oct. 2020, pp. 38–45.
- [31] 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.