

MAESTRO: Matched Speech Text Representations through Modality Matching

Zhehuai Chen, Yu Zhang, Andrew Rosenberg, Bhuvana Ramabhadran,
Pedro Moreno, Ankur Bapna, Heiga Zen

Google, Inc.

{zhehuai,ngyuzh,rosenberg,bhuv,pedro,ankurbpn,heigazen}@google.com

Abstract

We present Maestro, a self-supervised training method to unify representations learnt from speech and text modalities. Self-supervised learning from speech signals aims to learn the latent structure inherent in the signal, while self-supervised learning from text attempts to capture lexical information. Learning aligned representations from unpaired speech and text sequences is a challenging task. Previous work either implicitly enforced the representations learnt from these two modalities to be aligned in the latent space through multi-tasking and parameter sharing or explicitly through conversion of modalities via speech synthesis. While the former suffers from interference between the two modalities, the latter introduces additional complexity. In this paper, we propose Maestro, a novel algorithm to learn unified representations from both these modalities simultaneously that can transfer to diverse downstream tasks such as Automated Speech Recognition (ASR) and Speech Translation (ST). Maestro learns unified representations through sequence alignment, duration prediction and matching embeddings in the learned space through an aligned masked-language model loss. We establish a new state-of-the-art (SOTA) on VoxPopuli multilingual ASR with a **8%** relative reduction in Word Error Rate (WER), multi-domain SpeechStew ASR (**3.7%** relative) and 21 languages to English multilingual ST on CoVoST 2 with an improvement of **2.8** BLEU averaged over 21 languages.

Index Terms: Speech Recognition, Speech Translation, Speech-Text, Self-supervised, Representation learning

1. Introduction

In recent years, self-supervised learning (aka pretraining) has become an effective and scalable paradigm to learn representations from speech and text. These representations of language learnt without labeled data have shown impressive performance in multiple speech and natural language processing tasks [1–4]. Auto-regressive, contrastive and multi-task learning objectives have been proposed in the literature to pre-train such models [3, 5–9], followed by fine-tuning with labeled data to the task of interest, namely, speech recognition, translation, voice conversion and spoken language understanding. These diverse tasks require that the learned representations capture acoustic, prosodic, speaker and linguistic characteristics as well as the semantics of the spoken language. Thus, joint pre-training of representations from both speech and text modalities is a natural extension for improved generalization.

Central to the problem of jointly learning shared speech and text representation is the challenging task of combining two very different modalities. Recent methods proposed in literature have addressed this challenge through two very

different approaches. One approach uses multi-task training which trains a single model with different objectives for each modality [5, 10, 11], while another uses modality conversion where text is directly converted to speech [12] via text-to-speech (TTS) approaches or viceversa [13] prior to self-supervised training.

Maestro (Section 3) lies at the intersection of these two lines of work. It retains the advantage of a common latent space and alignment between the speech and text representations without explicitly converting text to speech or viceversa, thereby learning a computationally efficient and aligned representation of the two modalities. The main contributions of this paper are:

- A novel modality matching algorithm, Maestro, that can effectively use small amounts of transcribed speech data to unify representations learnt from unlabeled speech and text.
- A new algorithm to use additional sources of supervision, namely, Machine Translation (MT) and Speech Translation (ST), to improve multilingual joint representations learnt during pre-training.
- Establishing a new state-of-the-art (SOTA) on the VoxPopuli multilingual ASR task with a **8%** relative reduction in WER and on CoVoST-2 21 languages-to-English speech translation (ST) with an absolute improvement of **2.8** BLEU.

2. Related Work

With the convergence of architectures [14–16] and objectives [9, 17] used for learning representations in speech and natural language processing, there has been growing interest in learning joint speech and text representations. One set of approaches pursue multi-task training [5, 10]. However, given the very different natures of the two modalities and their training objectives, these approaches can suffer from interference and capacity limitations. In contrast, alternative approaches utilize Text-To-Speech (TTS) synthesis to augment natural speech data, avoiding inter-modal interference and improving downstream performance, at the cost of requiring a complete TTS model [12, 18] despite utilizing only the intermediate feature representations between text and speech for augmentation [12, 19–21]. Maestro bridges these approaches by implicitly aligning text representations with speech in the latent representation space. Utilizing recent advances in end-to-end speech synthesis that allow explicit duration modeling and control [22, 23], we show that we can use the duration predictions for alignment while enforcing similarity to natural speech representations.

Representation learning approaches can be extended to multilingual settings for text [2, 4], speech [24] and joint models [11]. Learnt representations can also be improved

by utilizing additional supervised data, joint unsupervised and supervised training on transcribed speech [25] or paired Masked Language Modeling (MLM) objectives on Machine Translation (MT) [26] or Speech Translation (ST) data [27]. We apply Maestro to train both, monolingual and massively multilingual models of speech and text. We also develop approaches to utilize cross-lingual supervision from MT and ST data during pre-training. Our work improves upon existing work on utilizing TTS models for speech representation learning by: (i) Developing an iterative and self-supervised approach for duration modeling, (ii) Circumventing the need for explicit speaker and prosody modeling by aligning resampled text representations with speech, (iii) Extending joint speech and text representation learning to massively multilingual datasets and tasks.

3. Proposed method

3.1. Architecture

The proposed framework to pretrain one model from untranscribed speech, unspoken text and any available labeled data (paired speech and text) is presented in Figure 1. It comprises of separate speech and text encoders to encode speech and text input signals, and a shared, multi-modal latent space encoder to learn the joint representation from these two modalities. The modality matching objective, \mathcal{L}_{MM} is used to explicitly enforce speech-text modality matching while the \mathcal{L}_{A-MLM} objective in the joint embedding space is used to learn representations from unspoken text. The framework allows for task-dependent additional RNN-T (or any other) decoders that predict grapheme, phoneme, word-piece or word targets from the shared representations.

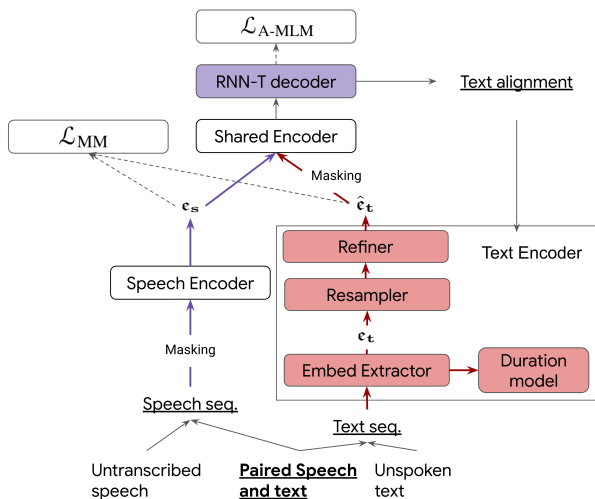


Figure 1: Proposed architecture of Maestro to learn unified representations from speech and text. The purple and red boxes denote differences from mSLAM [11]. The Text Encoder block utilizes alignments to explicitly resample text representations \hat{e}_t to match speech encoder output e_s .

3.2. Modality Matched Training

In this Section, we describe the overall training process using available paired speech and text data, unspoken text and modality matching. We propose a *modality matching mechanism* (MM) to explicitly unify the representation space

learned from text and speech. Specifically, we try to link the two modalities in the phonetic level and learn a mapping between text units and speech representations independent of speaker, prosody and other acoustic characteristics.

3.2.1. Learning Initial Embeddings

The speech encoder θ_s is trained on speech input using a within-utterance contrastive loss focusing on local acoustic and phonetic information [3, 9]. In contrast, the text embedding extractor θ_t captures the embedded lexical information from the input text signal.

3.2.2. Learning Aligned Embeddings

The initial speech e_s and text e_t embeddings learned are not expected to be aligned. Merging the two independently learned representations without attention to alignment can result in poor knowledge sharing and require extra model capacity. We hypothesize that this also limits the ability to jointly learn from the two tightly-coupled modalities representing any language. The self-alignment process described below learns alignments from the model itself in an iterative fashion.

Paired Speech and Text input: We utilize an RNN-T decoder (See Figure 1) and the available labeled data to first learn an RNN-T model. When learning from paired speech and text, the *Text Encoder* block in the figure uses this RNN-T model to generate alignments between the predicted text targets \mathbf{t} and the speech encoder output e_s . The *Resampler* and *Refiner* layers replicate the initially learned text embeddings to match the duration of the speech embedding using this alignment information and a Mean-Squared Error (MSE) training objective given in Equation 1 below.

$$e_s = \theta_s(\mathbf{s}), e_t = \theta_t(\mathbf{t}), (\mathbf{t}, \mathbf{s}) \in \mathcal{X}_{\text{paired}}$$

$$\hat{e}_t = \theta_{\text{Refiner}}(\text{Resample}(e_t, \text{Align}_{\text{RnnT}}(e_s, \mathbf{t}))) \quad (1)$$

$$\mathcal{L}_{MM} = \text{MSE}(e_s, \hat{e}_t) + \mathcal{L}_{\text{RnnT}}(\mathbf{t} | e_s)$$

Unpaired text input: When learning from unspoken text, the speech-text alignment information is unavailable. We substitute it with durations predicted from a duration prediction model in a fashion similar to speech synthesis [22]. This model is trained on any available paired data to predict the duration of each token. The predicted duration on unspoken text is subsequently used to resample the initially learned text embeddings.

3.2.3. Aligned Masked Language Model training objective

With the introduction of speech-aligned and resampled text embeddings (\hat{e}_t), we can use the same RNN-T loss for both, unpaired text (aligned embedding \hat{e}_t , text) and paired (speech embedding e_s , text) data. We replace the original MLM/BERT loss used in prior work [1, 11] with the *aligned masked language model* training objective (\mathcal{L}_{A-MLM}). This is the RNN-T loss applied over the masked, resampled text embeddings with masking in frequency and time domain similar to SpecAugment [28]. This new objective allows for the use of the same RNN-T objective on speech embedding or unspoken text with no associated speech embedding. Motivated by [18] to enforce consistency between text and speech modalities, we also include the A-MLM loss when training with paired speech and text input. We summarize this training in Equation 2.

$$\begin{aligned} \mathbf{e}_t &= \theta_t(\mathbf{t}), \hat{\mathbf{e}}_t = \theta_{\text{Refiner}}\left(\text{Resample}(\mathbf{e}_t, \theta_{\text{duration}}(\mathbf{e}_t))\right) \\ \mathcal{L}_{\text{A-MLM}} &= \mathcal{L}_{\text{Rnnt}}\left(\mathbf{t} \mid \text{Mask}(\hat{\mathbf{e}}_t)\right), \quad \mathbf{t} \in \mathcal{X}_{\text{text}} \end{aligned} \quad (2)$$

3.3. Untranscribed Speech

Similar to W2v-BERT [9], we use contrastive loss on the speech encoder outputs and a masked language model (MLM) loss on the shared encoder output to pretrain our model on untranscribed speech data.

4. Data and Experimental Setup

4.1. Pre-training Data

We explore pre-training Maestro on both monolingual (English) and multilingual scenarios as well as ASR and translation tasks.

Monolingual ASR: We use unlabeled English-only speech from the LibriLight corpus [29] to pre-train monolingual models. Our unspoken text corpus consists of the text corpora from TEDLIUM [30] and Librispeech [31]. We use SpeechStew [32] as the labeled, paired speech and text corpus.

Multilingual ASR: Following [11], we use 429k hours of public unlabeled speech corpora: VoxPopuli [33], Common-Voice [34], MLS [35] and BABEL [36] to pre-train multilingual ASR models. Our transcribed data draws from VoxPopuli transcribed dataset (VP-P), MLS and Babel, following [11]. We use three different setups to study the effectiveness of unlabeled text corpora: (1) VoxPopuli text dataset (VP-T, 3GB) (2) mC4 [2] spanning 101 languages (15TB) and (3) VP-T and mC4.

Speech Translation: We explore the use of MT and ST supervised data in pretraining for downstream tasks such as ST (Section 4.4). We use the source speech and text data from CoVoST 21 languages to English ST corpus as transcribed speech. We concatenate paired MT text sequences from CoVoST, WMT and TED datasets [11] and use them as unspoken text for pretraining.

4.2. Architecture Details

Speech encoder and shared encoder : The encoders in Figure 1 are a stack of ‘‘Conformer blocks’’. We use the Conformer XL architecture described in [12] with 24 layers of full-context Conformer blocks (600M parameters) where 6 of the lower layers are assigned as speech encoder and the upper 18 layers are the speech-text shared encoder.

Text encoder : The text embedding extractor θ_t (Section 3.2) includes 3 convolutional layers of 512 filters with kernel size (5, 1), followed by a 6-layer Transformer with positional embedding. The upsampling is done by copying the original text embedding to the target length of specified duration with positional embeddings to capture frame positions within text units as described in [22]. The Refiner θ_{Refiner} includes 2 layers of 8-headed self-attention blocks with 17×1 lightweight convolutions [37]. The duration model includes four blocks of 3×1 lightweight convolutions taking the original text embedding to predict the duration.

RNN-T decoder : We use a 2-layer, 1280-dim LSTM with a joint network of 640 dims as the RNN-T decoder. While we use an RNN-T decoder in this paper, the proposed framework

Table 1: *Multi-domain ASR: SpeechStew results on 5 domains. All the models use 0.6B parameters.*

| Method | LS-test | | AMI | | TED | SWB | | CV |
|----------------|--------------|--------------|------------|-------------|------------|------------|------------|------------|
| | <i>clean</i> | <i>other</i> | <i>ihm</i> | <i>sdm</i> | | <i>swb</i> | <i>chm</i> | |
| Wav2vec2 | 1.7 | 3.3 | 9.6 | 23.8 | 5.7 | 4.9 | 10.8 | 8.5 |
| W2v-BERT | 1.6 | 3 | 9.1 | 23.1 | 5.4 | 4.5 | 9 | 8.6 |
| + LM | 1.5 | 2.8 | | | | | | |
| SLAM | 1.6 | 3.1 | 9.3 | 23.5 | 5.6 | 4.6 | 9.1 | 8.6 |
| TTS4Pretrain2 | 1.6 | 2.8 | 8.7 | 21.9 | 5 | 4.5 | 8.5 | 8.4 |
| Maestro | 1.5 | 2.8 | 8.5 | 21.9 | 4.9 | 4.3 | 8 | 8.1 |
| + LM | 1.5 | 2.7 | | | | | | |

allows for the use of any decoder (CTC, LAS, etc.). We use both phoneme and grapheme decoders for monolingual pre-training. When switching to multilingual pre-training, we use a single decoder with 4k sentence-pieces as targets to cover 101 languages in the mC4 dataset. We also consider variants using grapheme and phoneme targets with vocabulary sizes of 100 and 256 respectively on the VoxPopuli corpora for comparison.

4.3. Pretraining hyper-parameters

We include untranscribed speech, unspoken text, transcribed speech in each batch with a fixed ratio. To stabilize training, we use, (1) exponential-moving-averaged (EMA) with decay rate 0.9999 to all prediction steps during alignment of transcribed speech, duration prediction and resampling of unspoken text; (2) a curriculum learning schedule to start from untranscribed speech-only training, include transcribed speech after 500k steps and unspoken text after another 15k steps. The joint training of three types of data lasts for another 300K steps with a learning rate schedule and optimizer given in [7]. We use an effective batch size of (256, 512, 256) for (untranscribed speech, unspoken text and paired data) in SpeechStew. It increases to (1024, 1024, 256) in VoxPopuli and (1024, 8192, 512) when including mC4 to increase text throughput.

4.4. Downstream Tasks

Speech Recognition We evaluate our monolingual Maestro on multi-domain ASR using the SpeechStew [32] task. All fine-tuning data and parameters follow [7]. We also initialize the RNN-T decoder from Maestro, which means both encoder and decoder are pretrained. Multilingual versions of Maestro are fine-tuned and evaluated on the 14-language VoxPopuli [33] multilingual ASR task using grapheme targets [11].

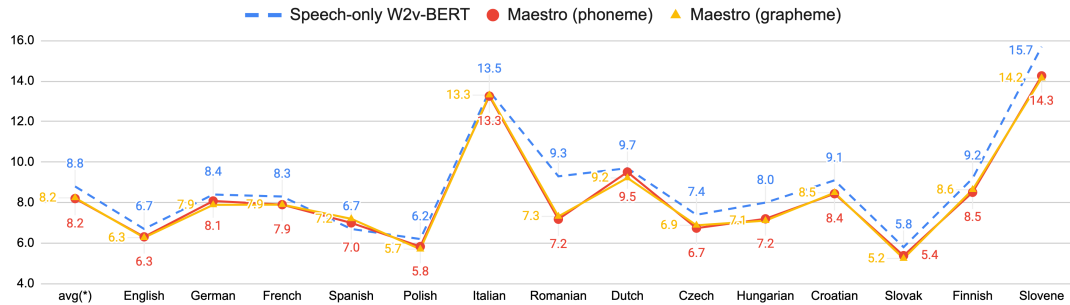
Speech Translation Our multilingual speech translation evaluation on the CoVoST-2 multilingual XX-to-English task that covers translation from 21 source languages into English [38] closely follows [11]. For joint fine-tuning we utilize the text translation datasets composed of CoVoST 2, WMT and TED.

5. Experiments and Results

5.1. Monolingual Multi-domain Speech Recognition

Table 1 presents the results from fine-tuning monolingual Maestro on multi-domain ASR. The first two blocks present baselines from speech-only pretraining methods, Wav2vec2 and W2v-BERT and previous speech-text pretraining methods, TTS4Pretrain 2.0 and SLAM [10] respectively. Maestro clearly outperforms speech-only pretraining and further improves over previous speech-text pretraining, TTS4Pretrain 2.0 (upto 6% relative) and SLAM (upto 12% relative) reduction in WER.

Figure 2: Breakdown of the improvement from Maestro on VoxPopuli and the comparison between phoneme and grapheme pretrain. The languages are sorted by the amount of data from high to low.



Comparing Maestro with TTS4Pretrain 2.0, the text encoder here is learnt from all the SpeechStew data while TTS model used in TTS4Pretrain 2.0 is only trained on Librispeech. We believe this is the reason we see better performance. On Librispeech, Maestro closes the gap with LM shallow fusion (Row 3) and yields additional wins with LM fusion (last row). Both use a neural Conformer language model trained from Librispeech text corpora.

5.2. Multilingual Speech Recognition

Table 2: Multilingual ASR: VoxPopuli results on 14 languages. All systems are finetuned with the same 1.3k hours of multilingual VoxPopuli data (VP-P). In pretrain, 2.4k hours refer to VP-P+MLS+BABEL. Maestro uses 4k SPM during pretrain. (XLS-R uses 8k extra hours of untranscribed speech from VoxLingua [39].)

| Method | Model size | Paired (hours) | Speech (hours) | Text | Avg WER |
|----------------|------------|----------------|----------------|-------|------------|
| XLS-R [24] | 1B | - | 437k | - | 10.6 |
| W2v-BERT [9] | 0.6B | - | 429k | - | 8.8 |
| mSLAM [11] | 0.6B | 2.4k | 429k | mC4 | 9.2 |
| mSLAM [11] | 2B | 2.4k | 429k | mC4 | 9.1 |
| Maestro | 0.6B | 1.3k | 429k | VP-T | 8.2 |
| Maestro | 0.6B | 2.4k | 429k | mC4 | 8.3 |
| Maestro | 0.6B | 2.4k | 429k | +VP-T | 8.1 |

We evaluate Maestro on multilingual speech recognition on the VoxPopuli benchmark in Table 2 and include several best-performing systems as comparison. Using in-domain unspoken text (VP-T) and paired data (VP-P) provided by VoxPopuli release, the proposed method Maestro in the first row of the third block outperforms all the previous systems using either in-domain or out-domain unspoken data. A slight degradation (Row 6) is seen with the inclusion of massive amounts of out-of-domain unspoken text such as mc4 [2]. Combining both mC4 and VP-T corpora (last row) yields the best performance, improving over speech-only W2v-BERT by 8% relative and create a new state of the art.

To understand the improvement w.r.t. data sizes, Figure 2 illustrates the language-wise breakdown corresponding to Row 5 in Table 2. Most of languages show improvements including *Slovenian* with only 10 hours of supervised speech. This shows that the proposed method can learn meaningful speech-text representations even from small amounts of supervised data. Figure 2 also compares the use of phoneme and grapheme

Table 3: Speech Translation (ST): CoVoST 2 X→En results on 21 language pairs.

| Method | Size | Pretrain | | | | | Avg BLEU |
|---|------|----------|------|-----|----|----|-------------|
| | | Speech | Text | ASR | ST | MT | |
| <i>Finetune: ST-only; mBART decoder init [40]</i> | | | | | | | |
| XLS-R [24] | 1B | 437k | × | × | × | × | 19.3 |
| XLS-R [24] | 2B | 437k | × | × | × | × | 22.1 |
| <i>Finetune: ST and MT jointly</i> | | | | | | | |
| W2v-BERT [9] | 0.6B | 429k | × | × | × | × | 21.0 |
| mSLAM [11] | 0.6B | 429k | mC4 | ✓ | × | × | 22.4 |
| mSLAM [11] | 2B | 429k | mC4 | ✓ | × | × | 24.8 |
| Maestro | 0.6B | 429k | mC4 | ✓ | × | × | 24.3 |
| Maestro | 0.6B | 429k | mC4 | ✓ | ✓ | ✓ | 25.2 |

targets during pretraining. Notably, regardless of different units used in pretraining, graphemes are always used in the supervised finetune. The overall performances are similar.

5.3. Speech Translation

We evaluate the potential of Maestro to serve as a foundation model for several speech recognition and understanding tasks. We use the same pretrained ASR encoder from Section 5.2 to evaluate speech translation performance. Table 3 demonstrates that Maestro also advances SOTA results on the speech translation CoVoST 2 benchmark. With 0.6B parameters and the same pretrain and finetune data setups as mSLAM, Maestro (second last row) outperforms most of the previous systems except the larger, 2B parameter mSLAM model. To minimize the mismatch in the text encoder outputs between pretraining and finetuning, we include MT and ST data into pretraining with the procedure described in Section 4.1 yielding the best performance (last row). Overall, the resulting system outperforms mSLAM (0.6B) by 2.8 BLEU scores and creates a new SOTA on this benchmark with 30% of the parameters. We hypothesize that the extra efficiency stems from improved knowledge sharing between the two modalities that can be attributed to the proposed modality matching.

6. Conclusions

We have described Maestro, a new technique for joint speech and text representation learning that outperforms the state of the art on speech recognition and speech translation tasks. Maestro addresses the joint representation problem by first aligning text to speech and then training a text representation to match a W2v-BERT speech representation. This results in significant improvements by 8% on ASR and 2.8 BLEU on ST.

7. References

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [2] 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,” *arXiv preprint arXiv:2010.11934*, 2020.
- [3] A. Baevski *et al.*, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *arXiv preprint arXiv:2006.11477*, 2020.
- [4] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov, “Unsupervised cross-lingual representation learning at scale,” *arXiv preprint arXiv:1911.02116*, 2019.
- [5] J. Ao, R. Wang, L. Zhou, S. Liu, S. Ren, Y. Wu, T. Ko, Q. Li, Y. Zhang, Z. Wei *et al.*, “SpeechT5: Unified-modal encoder-decoder pre-training for spoken language processing,” *arXiv preprint arXiv:2110.07205*, 2021.
- [6] Y.-A. Chung, H. Tang, and J. Glass, “Vector-quantized autoregressive predictive coding,” *arXiv preprint arXiv:2005.08392*, 2020.
- [7] Y. Zhang *et al.*, “Pushing the limits of semi-supervised learning for automatic speech recognition,” *arXiv preprint arXiv:2010.10504*, 2020.
- [8] W.-N. Hsu, Y.-H. H. Tsai, B. Bolte, R. Salakhutdinov, and A. Mohamed, “HuBERT: How much can a bad teacher benefit ASR pre-training?” in *Proc. ICASSP*, 2021, pp. 6533–6537.
- [9] Y.-A. Chung *et al.*, “W2v-BERT: Combining contrastive learning and masked language modeling for self-supervised speech pre-training,” *arXiv preprint arXiv:2108.06209*, 2021.
- [10] A. Bapna, Y.-a. Chung, N. Wu, A. Gulati, Y. Jia, J. H. Clark, M. Johnson, J. Riesa, A. Conneau, and Y. Zhang, “SLAM: A unified encoder for speech and language modeling via speech-text joint pre-training,” *arXiv preprint arXiv:2110.10329*, 2021.
- [11] A. Bapna, C. Cherry, Y. Zhang, Y. Jia, M. Johnson, Y. Cheng, S. Khanuja, J. Riesa, and A. Conneau, “mSLAM: Massively multilingual joint pre-training for speech and text,” *arXiv preprint arXiv:2202.01374*, 2022.
- [12] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, G. Wang, and P. Moreno, “Injecting text in self-supervised speech pretraining,” *arXiv preprint arXiv:2108.12226*, 2021.
- [13] R. Kuhn and R. De Mori, “A cache-based natural language model for speech recognition,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 12, no. 6, pp. 570–583, 1990.
- [14] A. Vaswani *et al.*, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [15] A. Gulati, J. Qin, C.-C. Chiu *et al.*, “Conformer: Convolution-augmented Transformer for speech recognition,” *arXiv preprint arXiv:2005.08100*, 2020.
- [16] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, “Perceiver: General perception with iterative attention,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 4651–4664.
- [17] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” *arXiv preprint arXiv:2202.03555*, 2022.
- [18] Z. Chen, Y. Zhang, A. Rosenberg, B. Ramabhadran, P. Moreno, and G. Wang, “Tts4pretrain 2.0: Advancing the use of text and speech in ASR pretraining with consistency and contrastive losses,” in *Proc. ICASSP*, 2022.
- [19] A. Renduchintala, S. Ding, M. Wiesner, and S. Watanabe, “Multimodal data augmentation for end-to-end ASR,” *arXiv preprint arXiv:1803.10299*, 2018.
- [20] G. Wang, A. Rosenberg, Z. Chen, Y. Zhang, B. Ramabhadran, Y. Wu, and P. Moreno, “Improving speech recognition using consistent predictions on synthesized speech,” in *ICASSP*, 2020.
- [21] S. Ueno, M. Mimura, S. Sakai, and T. Kawahara, “Data augmentation for asr using TTS via a discrete representation,” *ASRU*, 2021.
- [22] I. Elias, H. Zen, J. Shen, Y. Zhang, Y. Jia, R. J. Weiss, and Y. Wu, “Parallel Tacotron: Non-autoregressive and controllable TTS,” in *Proc. ICASSP*, 2021, pp. 5709–5713.
- [23] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech: Fast, robust and controllable text to speech,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [24] 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.
- [25] J. Bai, B. Li, Y. Zhang, A. Bapna, and T. N. Sainath, “Joint unsupervised and supervised training for multilingual ASR,” *arXiv preprint arXiv:2111.08137*, 2021.
- [26] G. Lample and A. Conneau, “Cross-lingual language model pretraining,” *arXiv preprint arXiv:1901.07291*, 2019.
- [27] R. Zheng *et al.*, “Fused acoustic and text encoding for multimodal bilingual pretraining and speech translation,” *arXiv preprint arXiv:2102.05766*, 2021.
- [28] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Proc. Interspeech 2019*, pp. 2613–2617, 2019.
- [29] J. Kahn, M. Rivière, W. Zheng *et al.*, “Libri-Light: A benchmark for asr with limited or no supervision,” in *Proc. ICASSP*, 2020, pp. 7669–7673.
- [30] A. Rousseau, P. Deléglise, and Y. Esteve, “TED-LIUM: an automatic speech recognition dedicated corpus,” in *Proc. LREC*, 2012, pp. 125–129.
- [31] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proc. ICASSP*, 2015.
- [32] W. Chan *et al.*, “SpeechStew: Simply mix all available speech recognition data to train one large neural network,” *arXiv preprint arXiv:2104.02133*, 2021.
- [33] 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,” *arXiv preprint arXiv:2101.00390*, 2021.
- [34] R. Ardila *et al.*, “Common Voice: A massively-multilingual speech corpus,” *arXiv preprint arXiv:1912.06670*, 2019.
- [35] 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.
- [36] 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)*, 2014, pp. 16–23.
- [37] F. Wu, A. Fan, A. Baevski, Y. N. Dauphin, and M. Auli, “Pay less attention with lightweight and dynamic convolutions,” *arXiv preprint arXiv:1901.10430*, 2019.
- [38] C. Wang, A. Wu, and J. Pino, “CoVoST 2 and massively multilingual speech-to-text translation,” *arXiv preprint arXiv:2007.10310*, 2020.
- [39] J. Valk and T. Alumäe, “VoxLingua107: A dataset for spoken language recognition,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 652–658.
- [40] Y. Tang, C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, and A. Fan, “Multilingual translation from denoising pre-training,” in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021, pp. 3450–3466.