



XLSR-TRANSDUCER: STREAMING ASR FOR SELF-SUPERVISED PRETRAINED MODELS

Shashi Kumar Srikanth Madikeri
Zuluaga-Gomez. Juan Esaú VILLATORO-TELLO^a
Nigmatulina Iuliia Petr Motlicek Manjunath K E
Aravind Ganapathiraju

Idiap-RR-08-2024

AUGUST 2024

^aIdiap

XLSR-Transducer: Streaming ASR for Self-Supervised Pretrained Models

Shashi Kumar^{1,4}, Srikanth Madikeri¹, Juan Zuluaga-Gomez^{1,4}, Esaú Villatoro-Tello¹, Iuliia Nigmatulina^{1,2}, Petr Motlicek^{1,3}, Manjunath K E⁵, Aravind Ganapathiraju⁵

¹Idiap Research Institute, Switzerland ²University of Zurich, Switzerland

³ Faculty of Information Technology, Brno University of Technology, Czech Republic

⁴LIDIAP, Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland

⁵Uniphore, India

Abstract

Self-supervised pretrained models exhibit competitive performance in automatic speech recognition on finetuning, even with limited in-domain supervised data for training. However, popular pretrained models are not suitable for streaming ASR because they are trained with full attention context. In this paper, we introduce XLSR-Transducer, where the XLSR-53 model is used as encoder in transducer setup. Our experiments on the AMI dataset reveal that the XLSR-Transducer achieves 4% absolute WER improvement over Whisper large-v2 and 8% over a Zipformer transducer model trained from scratch. To enable streaming capabilities, we investigate different attention masking patterns in the self-attention computation of transformer layers within the XLSR-53 model. We validate XLSR-Transducer on AMI and 5 languages from CommonVoice under low-resource scenarios. Finally, with the introduction of attention sinks, we reduce the left context by half while achieving a relative 12% improvement in WER.

Index Terms: streaming ASR, self-supervised learning, XLSR, transformer transducer

1. Introduction

End-to-End (E2E) modeling approaches for automatic speech recognition (ASR) have become ubiquitous in literature. There are three popular architecture for achieving E2E models: encoder-decoder (AED) based models [1, 2, 3, 4], using Connectionist Temporal Classification (CTC) loss [5], and neural Transducer [6, 7] based models. The transducer models, shown in Figure 1a, consist of an encoder, predictor and joint networks. Using Transformers [8] encoder, termed as Transformer Transducer (TT) [9, 10], is a popular choice for streaming ASR [11, 12, 13] because of its inherently streaming nature. It is common to train TT models from scratch which requires large amount of supervised data [12, 11]. Self-supervised (SSL) models have shown strong ASR performance when trained with small amount of in-domain supervised data [14, 15]. Most of the recent works using SSL models for ASR use encoder-decoder setup [16, 17, 18] or CTC loss [14, 15, 19, 20]. In this paper, we present XLSR-Transducer where we use a pretrained XLSR-53 model as encoder in the TT architecture. This opens the door to streaming TT systems for low-resource applications.

In streaming ASR, partial hypotheses are generated for each audio chunk sequentially [12, 21] to produce the transcript for the full audio, whereas the entire audio segment is available for non-streaming decoding. Depending on the latency requirements, the chunk size may vary from few hundred milliseconds to few seconds [21]. Typically, a drastic degradation in word error rate (WER) is observed when non-streaming models are decoded in streaming fashion [13], because only a limited con-

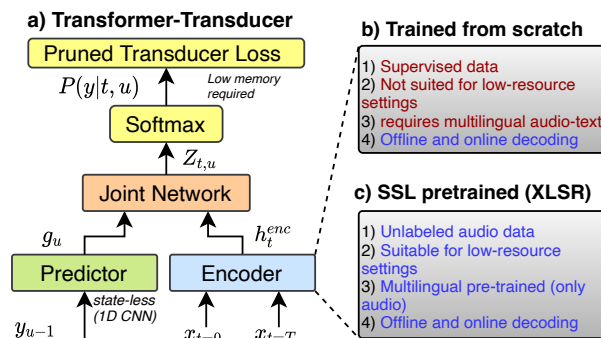


Figure 1: Current state-of-the-art a) Transducer ASR includes state-less predictor, pruned transducer loss and b) Transformer-based encoder, trained from scratch. We replace the encoder by XLSR-53, an SSL model suitable for low-resource applications. Our contributions lead to the c) XLSR-Transducer.

text is available. In this work, we propose a variety of attention masking patterns that enables streaming training and decoding of our XLSR-Transducer model. We also study the importance of chunk sizes and past context by varying them during inference. Increasing past context typically enhances ASR performance [21], at the expense of increased latency. Recently, it was shown that the transformer layers learn to assign disproportionate attention scores to few initial tokens for streaming language models [22], termed as attention sinks. We study the effects of attention sinks for the first time in streaming ASR. Formally, at decoding time, we allow the transformer layers in XLSR to attend to few initial frames in addition to designated frames in chunk and past context. In theory, this reduces total computation required for processing an audio chunk during streaming decode.

Our contributions are covered below:

- Introduction of the XLSR-Transducer, a multilingual SSL encoder based transducer model, demonstrating significant WER improvement on the AMI dataset compared to large speech foundational models and other open-source ASR models;
- extension to streaming XLSR-Transducer and a systematic study of chunk size and past context on training and inference;
- to author's knowledge, this is the first work that explores the attention sink [22] phenomenon for streaming ASR which leads to improved WER; and
- Evaluation of the XLSR-Transducer on AMI and five languages of CommonVoice dataset in low resource settings.




Figure 2: *Masking strategies for streaming XLSR-Transducer. Multi-chunk training allows decoding with variable chunk size (blue) and left context (orange). Each square denotes “n” frames. Attention sink (yellow) allows context from the first n frames. Our results show that attention sink frames offer a better trade-off w.r.t increasing left context alone, leading to lower WERs.*

2. XLSR-Transducer

In a typical Transformer-Transducer (TT) ASR model (Figure 1a), there are three networks: the encoder, predictor, and joiner. The encoder processes audio frames to produce acoustic embeddings. The predictor generates token embeddings in an autoregressive manner, taking previous non-blank tokens as input. Lastly, the joiner combines the outputs from the encoder and predictor to predict a probability distribution over the tokens in the vocabulary. In this work, we utilize a stateless predictor [23] composed of an embedding layer and one 1-D CNN layer, and the joiner network consists of one linear layer. Typically, the encoders [24, 25, 26, 27] are trained from scratch and require a large amount of in-domain supervised data to achieve decent WER, which may not always be feasible. We train the TT model using the pruned-transducer loss [7] from k2¹ toolkit.

2.1. Non-Streaming XLSR-Transducer

In contrast to encoders trained from scratch, recent advancements in SSL pretrained models demonstrate competitive performance [15, 14] when fine-tuned with a limited amount of labeled data for ASR. Previously, the ASR models employing SSL pretrained models have utilized CTC loss [5], encoder-decoder based architecture [2, 3], and Lattice-Free MMI loss [28] (hybrid approach) for training. In this paper, we integrate pretrained models as encoders in the TT setup, as illustrated in Figure 1c. One notable advantage is the ability to achieve strong ASR performance with relatively low amounts of training data. We select XLSR-53 [14] as our encoder model, which takes raw audio as input and outputs audio frames with a frame duration of 25ms and a stride of 20ms. The selection of XLSR-53 is driven by its large-scale pre-training on multilingual audio, which has demonstrated competitive ASR performance [14] in the low-resource across multiple languages.

2.2. Streaming XLSR-Transducer

XLSR is typically trained and decoded using entire audio sample. This makes the proposed XLSR-Transducer non-streaming despite the use of stateless predictor and linear joiner that are inherently streaming. The main challenge to port SSL models to streaming is the use of self-attention in the transformer layers [8], i.e., computed over entire acoustic frames of an utterance. In this paper, we present multiple masking patterns [21] to limit the frame context over which self-attention is computed, simulating streaming ASR within the XLSR-Transducer setup.

Chunked masking For a typical streaming ASR, decoding partial hypotheses should occur after receiving a few audio frames, known as the chunk size. As depicted in Figure 2a,

we implement chunk-wise decoding by masking frames outside a specific chunk during the forward pass from the XLSR model. The mask is applied after dot-product computation during self-attention, ensuring that each frame inside a chunk has access to all the frames within that chunk. Note that the XLSR model also includes a CNN front-end, which takes raw audio as input. Thus, we feed chunk-size equivalent raw waves to the CNN front-end in a sequential manner and concatenate them across the time dimension to obtain all the frames for an utterance. In this paper, we explore chunk sizes of 16, 32, 64, and 128, translating to approximately 320ms, 640ms, 1280ms, and 2560ms, respectively, for XLSR.

Chunked masking with variable left context chunks In practice, when decoding chunk “n”, we have access to all the previous chunks, which can be utilized as left context. As illustrated in Figures 2b and 2c, a variable number of left context chunks can be utilized during the self-attention computation of a chunk, with the possibility of using the full left context. The number of frames in the left context is a multiple of the chunk size, as this can be efficiently implemented to store past chunks in the cache.

Streaming training and decoding The use of non-streaming XLSR-Transducer for streaming decoding with the described masking patterns presents a challenge. The model has been trained on full context, creating a train-test mismatch. To address this challenge, we train the model in a streaming fashion using a fixed chunk size and left context. Flexibility in our chunked mask implementation allow us to perform both streaming and non-streaming decoding using a single model. The advantage of our method is that it only affects the fine-tuning stage and we can avoid the computationally prohibitive pre-training.

Multi-chunk Training In many practical use-cases of streaming ASR, varying the chunk size at decoding time is often desirable depending upon latency requirements. However, a streaming XLSR-Transducer model trained with a fixed chunk size may not yield optimal WERs when decoded with different chunk sizes. Also, training multiple models for varying chunk sizes may be infeasible. To address this limitation, we propose randomly selecting the chunk size from the predefined list mentioned above for each batch during training.

2.3. Attention sinks for streaming ASR

In a recent work on streaming language models (LM) [22, 29], it was shown that a surprisingly large amount of attention scores during self-attention computation inside transformer layers is directed towards the initial tokens, termed as *attention sinks*. This was attributed to the Softmax operation, which mandates attention scores to sum up to one and in autoregressive LMs, all subsequent tokens have access to the initial tokens. Consequently, the model may find it easier to learn to assign large

¹<https://github.com/k2-fsa/k2>

scores to these initial tokens. In our streaming model training, where we utilize full left context, we employ a similar setup. This leads us to introduce the first utilization of the attention sinks in the context of streaming ASR during inference. Specifically, as depicted in Figure 2d, we enable self-attention to focus on not only frames within a chunk and left context chunks, but also on the initial few frames.

3. Experimental setup

3.1. Datasets

AMI We train and evaluate the XLSR-Transducer model on the individual head microphone (IHM) split from the AMI dataset [30] containing audios with a sampling rate of 16kHz. We use the default recipe for AMI from lhotse² toolkit to prepare the train, dev and eval sets containing 80hr, 8.8hr, 8.5hr of audios respectively. In all our experiments on the AMI dataset, we use WER on the dev set to select the best epoch and report the results on the eval set in this paper.

CommonVoice We validate XLSR-Transducer on five non-English languages from CommonVoice-v11 [31].³ This includes Catalan (CA), Belarusian (BE), Spanish (ES), French (FR) and Italian (IT). To keep experimentation under the low-resource domain, we extract randomly a 100-hr subset from the training data per language. Later, we train streaming and non-streaming models. We report WERs on the full official test sets.

3.2. Zipformer-Transducer Baseline

We establish strong baselines by training non-streaming and streaming Zipformer transducer models [26] from scratch, following the AMI recipe⁴ from Icefall toolkit (we only use the IHM set). The state-less Predictor [23] consists of an embedding layer followed by one 1-D CNN layer and joiner consists of 1 linear layer. We use the default hyper-parameters from the recipes and train for 30 epochs. See more training details in [26]. Beam search with width of 4 is used across all ASR decoding in this paper.

3.3. XLSR-Transducer Training

The XLSR-transducer model is constructed from the Icefall’s Transducer recipe for AMI dataset adapted with the XLSR model from fairseq [32]. The fine-tuning uses Scaled Adam [33] and a learning rate scheduler that consists of a 500-step warmup phase [8] followed by a decay phase directed by number of steps and epochs. Furthermore, the model is optimized with pruned rnn-t loss [7, 6]. The learning rate is set to $lr=1.25e^{-3}$ for AMI and $lr=5.0e^{-3}$ for CommonVoice. We train AMI and CommonVoice models for 10 and 20 epochs, respectively.

4. Results & Discussion

Non-streaming ASR We begin by benchmarking the XLSR-Transducer model for non-streaming ASR on the AMI dataset and the results are reported in the Table 1 (*full-attm*). We compare against large open source foundational speech models. It can be seen that the proposed XLSR-Transducer model achieves significant improvement in WERs. Specifically, it achieves a relative improvement of 19% in WER when compared to the

Table 1: WERs on the AMI eval set. On non-streaming decoding[†] XLSR-Transducer yields significant WER reduction. On streaming decoding, the multi-chunk training (multiple) provides significant gain w.r.t encoders trained from scratch with minimal degradation in non-streaming performance. [‡]encoder-decoder model. [¶]decoding chunk size 2000ms.

Encoder	Chunk Size training	Chunk Size decoding		
		320 ms	1280 ms	full-att [†]
decoding: non-streaming ASR				
Whisper large-v2 (1.6B) [‡] [1]	-	-	-	16.9
FastConformer (1.1B) [27]	-	-	-	15.6
Zipformer (70M)	-	-	-	21.0
decoding: streaming ASR				
FastConformer (114M) [¶] [12]	-	-	24.2	-
Zipformer (70M)	multiple	28.5	24.6	23.2
XLSR (300M)	full-att	35.3	17.8	12.7
XLSR (300M)	320 ms	17.1	15.0	14.2
XLSR (300M)	1280 ms	19.7	14.5	13.1
XLSR (300M)	multiple	17.7	14.2	12.9

best open source large foundational ASR model. Next, we train a Zipformer encoder based transducer model from scratch and observe that XLSR-Transducer yields 39% relative improvement in WER. It is clear that there are significant advantages of using pretrained encoders in TT setup for low resource ASR. **Streaming ASR** First, we decode the non-streaming trained XLSR-Transducer model in streaming fashion by applying different masks (see §2.2). Results are reported in the Table 1, where full left context is used during decoding. The XLSR-Transducer achieves significant improvement over Zipformer and FastConformer based transducer models. Despite the improvements, there is a significant degradation from non-streaming performance because the model was not trained for streaming. When model was trained for streaming with full left context and decoded using chunk size of 320ms, the performance improves (35.3% → 17.1% WER) because of the train-test matched chunk size setting. We also train models with larger chunk sizes, but it degrades the performance showcasing the importance of context during self attention computation inside a chunk which the model may have learned during training.

As we increase the chunk duration during decoding (320ms → 1280ms), the performance improves monotonically [13], which is expected due to larger context available for frame attention score computation. Now, the streaming trained models are decoded in non-streaming which can serve as performance upper bound. We observe that increasing chunk size during training improves the non-streaming performance and even when chunk size of 320ms is used during training, the results only degrade by 1.5% in absolute WER. When random chunk sizes are used during training, the gap is 0.2% when compared with the best non-streaming ASR performance. Overall, a streaming trained XLSR-Transducer model using random chunk sizes shows best WER when decoded in streaming fashion with 1280ms chunk duration and performance gap for non-streaming decoding is minimal. Thus, a single model can be used for both streaming and non-streaming ASR.

Streaming ASR with variable left context Using full left context during training and decoding will incur additional computation which may not always be desirable. Limiting left context during training of streaming models resulted in significant degradation of results; therefore, we use full left context for all streaming XLSR-Transducer training. Figure 3 list WERs

²<https://github.com/lhotse-speech/lhotse>

³CommonVoice-v11: cv-corpus-11.0-2022-09-21)

⁴github.com/k2-fsa/icefall/tree/master/egs/ami




Figure 3: Plots of WERs on AMI eval set for XLSR-Transducer trained on three configurations (a, b and c) and decoded on multiple streaming scenarios. Note that adding one or more left-context chunks at decoding time reduces WERs dramatically.

Table 2: WERs of streaming XLSR-Transducer on five Common-Voice languages. Models are fine-tuned on random 100hr train subset and with multi-chunk training. *full-att*: non-streaming decoding. [†]CA is 28h long, while the rest 26h. [‡]median (duration) in seconds. [¶]non-streaming training and decoding.

Lang [†]	Test Set		Streaming model - chunk size [ms]				full-att [¶]	
	#utt	[50%-dur] [‡]	320	640	1280	2560	full-att	full-att
CA	16k	6.1	17.5	15.2	13.9	12.9	12.0	10.7
BE	15.8k	5.7	20.0	17.5	15.9	14.8	13.8	13.7
ES	15.5k	6.1	17.7	15.0	13.5	12.2	11.3	10.8
FR	16k	5.7	24.3	21.6	20.0	18.7	17.6	17.1
IT	15k	6.3	18.5	15.9	14.3	13.1	12.1	11.5

when the number of left context chunks are varied during inference. Note that the left context duration is in multiple chunk size. Increasing the left context improves the performance for all training scenarios and chunk duration. At the same time, a significant improvement is observed using just one chunk of left context. Further increase in left context improves the performance but the reduction in WER per left context chunk is lower. When the XLSR-Transducer is trained with multi-chunk streaming strategy and decoded with 1280ms chunk size, a relative improvement of only 4% in WER is observed using full left context instead of one left context chunk. Thus, a limited number of left context chunks should be enough for most real-world streaming ASR with XLSR-Transducer model.

XLSR-Transducer on multiple languages We also train the proposed model on five non-English languages of Common-Voice [31]. WERs are listed in Table 2 for multi-chunk streaming and full-attention non-streaming models. We see competitive WERs for models evaluated under different streaming conditions, with constant WERs improvement as chunk size increases; similar behavior is reported in [13]. The upper-bound WERs are obtained with a model trained and evaluated in non-streaming fashion, i.e., last column of Table 2. We note negligible WER degradation (up to 1.5% absolute WER, worse CA; best BE) for full-attention decoding (*full-att*) on streaming models vs. their non-streaming counterparts. This confirms the robustness of XLSR-Transducer on multiple languages.

Improving streaming ASR with attention sinks In theory, restricting left context chunks should lead to overall latency improvements for streaming ASR. The recent observation of attention sinks phenomena [22, 34], where the transformer models learn to assign relatively higher attention scores to initial tokens, may help in reducing the overall computation required to decode one chunk of audio in streaming ASR. WERs on the

Table 3: WERs on AMI eval set for varied decoding settings, including attention sink. Adding attention sinks offer a better trade off than larger left context. (*blue*) denotes relative WER reduction w.r.t no attention sink within same chunk and left context. [†]nb. of chunks. [‡]nb. attention sink frames.

Decoding settings		Decoding chunk-size	
Left-context [†]	attn-sink [‡]	320ms	640ms
full	none	17.7	15.5
1	none	25.9	18.1
	1	22.9 (+11.6)	17.4 (+4.1)
	4	21.4 (+17.4)	16.8 (+7.1)
	16	19.7 (+23.7)	16.3 (+10.0)
2	none	22.5	16.7
	1	20.9 (+7.3)	16.4 (+1.7)
	4	20.0 (+11.3)	16.2 (+3.2)
	16	18.9 (+16.1)	15.9 (+4.8)
4	none	19.8	15.9

AMI dataset are reported in Table 3. For different chunk sizes and left context, we observe that increased frames for attention sinks improve the performance monotonically. Specifically, for a smaller chunk of 320ms, using 1 left context chunk and 16 frames (i.e., 320ms) for attention sinks performs better than using 2 left context chunks by 12% in relative terms despite attending over same number of frames. We do not observe a significant reduction in WERs beyond a chunk size of 640ms. Overall, our results show that it is better to use attention sinks than increasing left context chunks beyond 1 for improved performance. We also run decoding with attention sinks on the CommonVoice dataset and observe similar trends but do not include a results table for brevity.

5. Conclusions

In this paper, we demonstrate that using a self-supervised trained model as encoder in the transducer framework, termed as XLSR-Transducer, leads to significant improvement in WER on AMI dataset. We explore various chunked masks and left context configurations to enable streaming decoding in XLSR. Our findings across 2 datasets and 6 languages shows that the proposed model achieves streaming performance comparable to non-streaming ASR. With introduction of attention sinks, we reduce left context during inference with improvements in WER.

6. References

- [1] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International Conference on Machine Learning*. PMLR, 2023, pp. 28 492–28 518.
- [2] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [3] F. Mai, J. Zuluaga-Gomez, T. Parcollet, and P. Motlicek, "Hyper-Conformer: Multi-head HyperMixer for Efficient Speech Recognition," in *Proc. INTERSPEECH 2023*, 2023, pp. 2213–2217.
- [4] R. Prabhavalkar, T. Hori, T. N. Sainath, R. Schlüter, and S. Watanabe, "End-to-end speech recognition: A survey," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [5] A. Graves and A. Graves, "Connectionist temporal classification," *Supervised sequence labelling with recurrent neural networks*, pp. 61–93, 2012.
- [6] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [7] F. Kuang, L. Guo, W. Kang, L. Lin, M. Luo, Z. Yao, and D. Povey, "Pruned rnn-t for fast, memory-efficient asr training," *arXiv preprint arXiv:2206.13236*, 2022.
- [8] 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.
- [9] C.-F. Yeh, J. Mahadeokar, K. Kalgaonkar, Y. Wang, D. Le, M. Jain, K. Schubert, C. Fuegen, and M. L. Seltzer, "Transformer-transducer: End-to-end speech recognition with self-attention," *arXiv preprint arXiv:1910.12977*, 2019.
- [10] Q. Zhang, H. Lu, H. Sak, A. Tripathi, E. McDermott, S. Koo, and S. Kumar, "Transformer transducer: A streamable speech recognition model with transformer encoders and rnn-t loss," in *ICASSP*. IEEE, 2020, pp. 7829–7833.
- [11] B. Li, A. Gulati, J. Yu, T. N. Sainath, C.-C. Chiu, A. Narayanan, S.-Y. Chang, R. Pang, Y. He, J. Qin *et al.*, "A better and faster end-to-end model for streaming asr," in *ICASSP*. IEEE, 2021, pp. 5634–5638.
- [12] V. Noroozi, S. Majumdar, A. Kumar, J. Balam, and B. Ginsburg, "Stateful fastconformer with cache-based inference for streaming automatic speech recognition," *arXiv preprint arXiv:2312.17279*, 2023.
- [13] B. Li, S.-y. Chang, T. N. Sainath, R. Pang, Y. He, T. Strohm, and Y. Wu, "Towards fast and accurate streaming end-to-end asr," in *ICASSP*. IEEE, 2020, pp. 6069–6073.
- [14] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," *arXiv preprint arXiv:2006.13979*, 2020.
- [15] V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi *et al.*, "Scaling speech technology to 1,000+ languages," *arXiv preprint arXiv:2305.13516*, 2023.
- [16] J. Ao, Z. Zhang, L. Zhou, S. Liu, H. Li, T. Ko, L. Dai, J. Li, Y. Qian, and F. Wei, "Pre-training transformer decoder for end-to-end asr model with unpaired speech data," *arXiv preprint arXiv:2203.17113*, 2022.
- [17] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, "W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 244–250.
- [18] J. Ao, R. Wang, L. Zhou, C. Wang, S. Ren, Y. Wu, S. Liu, T. Ko, Q. Li, Y. Zhang *et al.*, "Speech5: Unified-modal encoder-decoder pre-training for spoken language processing," *arXiv preprint arXiv:2110.07205*, 2021.
- [19] T. Javed, S. Doddapaneni, A. Raman, K. S. Bhogale, G. Ramesh, A. Kunchukuttan, P. Kumar, and M. M. Khapra, "Towards building asr systems for the next billion users," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, 2022, pp. 10813–10821.
- [20] 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.
- [21] P. Swietojanski, S. Braun *et al.*, "Variable attention masking for configurable transformer transducer speech recognition," in *ICASSP*. IEEE, 2023, pp. 1–5.
- [22] G. Xiao, Y. Tian, B. Chen, S. Han, and M. Lewis, "Efficient streaming language models with attention sinks," *arXiv preprint arXiv:2309.17453*, 2023.
- [23] M. Ghodsi, X. Liu, J. Apfel, R. Cabrera, and E. Weinstein, "Rnn-transducer with stateless prediction network," in *ICASSP*. IEEE, 2020, pp. 7049–7053.
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *arXiv preprint arXiv:2005.08100*, 2020.
- [26] Z. Yao, L. Guo, X. Yang, W. Kang, F. Kuang, Y. Yang, Z. Jin, L. Lin, and D. Povey, "Zipformer: A faster and better encoder for automatic speech recognition," *arXiv preprint arXiv:2310.11230*, 2023.
- [27] D. Rekish, N. R. Koluguri, S. Krizan, S. Majumdar, V. Noroozi, H. Huang, O. Hrinchuk, K. Puvvada, A. Kumar, J. Balam *et al.*, "Fast conformer with linearly scalable attention for efficient speech recognition," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [28] A. Vyas, S. Madikeri, and H. Bourlard, "Comparing ctc and lfmmi for out-of-domain adaptation of wav2vec 2.0 acoustic model," in *Proceedings of Interspeech*, 2021. [Online]. Available: <https://arxiv.org/abs/2104.02558>
- [29] M. Zaheer, G. Guruganesh, K. A. Dubey, J. Ainslie, C. Alberti, S. Ontanon, P. Pham, A. Ravula, Q. Wang, L. Yang *et al.*, "Big bird: Transformers for longer sequences," *Advances in neural information processing systems*, vol. 33, pp. 17 283–17 297, 2020.
- [30] J. Carletta, S. Ashby *et al.*, "The ami meeting corpus: A pre-announcement," in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [31] 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*, 2020, pp. 4218–4222.
- [32] 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 the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, 2019, pp. 48–53.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [34] G. Attanasio, B. Savoldi, D. Fucci, and D. Hovy, "Multilingual speech models for automatic speech recognition exhibit gender performance gaps," *arXiv preprint arXiv:2402.17954*, 2024.