

# LEARNING ASR PATHWAYS: A SPARSE MULTILINGUAL ASR MODEL

Mu Yang<sup>1,†</sup>, Andros Tjandra<sup>2</sup>, Chunxi Liu<sup>2</sup>, David Zhang<sup>2</sup>, Duc Le<sup>2</sup>, Ozlem Kalinli<sup>2</sup>

<sup>1</sup>Center for Robust Speech Systems (CRSS), University of Texas at Dallas, USA  
<sup>2</sup>Meta AI, USA

mu.yang@utdallas.edu, androstj@fb.com

## ABSTRACT


Neural network pruning compresses automatic speech recognition (ASR) models effectively. However, in multilingual ASR, language-agnostic pruning may lead to severe performance drops on some languages because language-agnostic pruning masks may not fit all languages and discard important language-specific parameters. In this work, we present *ASR pathways*, a sparse multilingual ASR model that activates language-specific sub-networks (“pathways”), such that the parameters for each language are learned explicitly. With the overlapping sub-networks, the shared parameters can also enable knowledge transfer for lower-resource languages via joint multilingual training. We propose a novel algorithm to learn ASR pathways, and evaluate the proposed method on 4 languages with a streaming RNN-T model. Our proposed ASR pathways outperform both dense models and a language-agnostically pruned model, and provide better performance on low-resource languages compared to the monolingual sparse models.

**Index Terms**— Multilingual, speech recognition, sparse, pruning.

## 1. INTRODUCTION

Automatic speech recognition (ASR) technologies play an important role in customer smart devices such as smartphones, smart speakers, smart watches, virtual reality glasses, and more. Due to the limited computation resources and storage, hosting a high-performance yet dense ASR model on on-device hardware is challenging. In order to reduce on-device model size and boost run-time efficiency, recent works have attempted to compress the ASR models through knowledge distillation [1, 2], model sparsification [3, 4, 5, 6], quantization [7, 8, 9], etc. In this study, we compress ASR models via neural network pruning (i.e. zeroing out a certain amount of dense model weights by a learned pruning mask). We specifically target at multilingual ASR scenario [10, 11, 12, 13, 14], with the aim that sparsifying a multilingual ASR model does not significantly degrade recognition performance on certain languages.

Pruning a multilingual ASR model poses unique challenges. Given a dense multilingual model, simply doing **language-agnostic** pruning leads to large performance drops on certain languages. This is because language-agnostic pruning only learns a single language-agnostic mask and shares parameters for all languages, which may not fit for certain languages [15, 16]. As a result, some languages will see significant word error rate (WER) degradation. Alternatively, one can do **language-specific** pruning on monolingual data [3, 4], which learns different model parameters and different pruning masks across languages. However, the resulting models lose their



**Fig. 1.** Demonstration of ASR pathways: each language activates its own sparse sub-network.

multilingual ability: we end up with one sparse monolingual model for each language, which complicates on-device deployment.

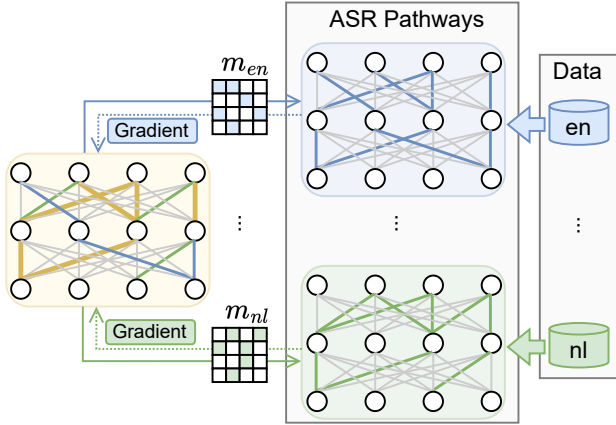
To tackle the challenges above, we propose that the sparse multilingual model should have language-specific masks while sharing the same set of model parameters. In other words, as shown in Figure 1, we aim to obtain a multilingual model that is sparsely activated, forming sparse “pathways” [17] (sub-networks) which can be dedicated to different languages. For this purpose, we propose a novel 2-step training scheme, where the first step identifies the language-specific masks via Iterative Magnitude Pruning (IMP) [18] or Lottery Ticket Hypothesis (LTH) [19], and in the second step, we fix the masks and update the model parameters via multilingual training, where each language governs the updates of the parameters under the mask that corresponds to the input language. Through the learning process, the model can learn the optimal pathway for each language and also learn the shared parameters across languages - such that when the language-specific masks applied, the model can recognize each language independently.

We applied the proposed method to an efficient memory transformer (Emformer) based [20, 21] low-latency streaming RNN-T architecture [22] with a 70.6% block-wise structured sparsity [23]. Our empirical results on a 4-language dataset show that the proposed multilingual ASR pathways model outperforms both dense model (-5.0% average WER) and language-agnostic pruning (-21.4% average WER), and has better performance on low-resource languages compared to the monolingual sparse models. By analyzing the language-specific masks, we reveal that with the help of a regularization that is tailored to structured pruning, LTH masks are superior to IMP masks, even with less total effective parameters.

## 2. RELATED WORKS

**Multilingual ASR.** Among recent works on multilingual ASR, language-specific modeling is usually adopted to achieve universally decent performance on all languages. This includes language-aware encoding [24], parameterizing language-specific attention heads [25], adapter modules [13], decoders [11] and Mixture-of-experts [16]. In [15], LTH was used to identify language-specific sub-networks inside the pre-trained multilingual XLSR model [26].

† Work was done when Mu Yang was an intern at Meta.



**Fig. 2.** Multilingual training with language-specific masks.  $m_{en}$  and  $m_{nl}$  denotes English (*en*) and Dutch (*nl*) mask, respectively. Blue and green segments represent trainable weights for *en* and *nl*. Bold yellow weights are shared between *en* and *nl*. Grey weights are set to zero during training and inference.

These sub-networks were shown to be able to pre-train an enhanced XLSR model that improves downstream multilingual ASR fine-tuning. In contrast to [15], we directly identify the language-specific sparse structures in a multilingual ASR model, which can potentially learn masks that are more tailored to ASR tasks.

**Model compression.** Prior works on ASR model compression can be grouped into 3 categories: (1) knowledge distillation [1, 2] which transfers knowledge from a larger model to a smaller model; (2) model sparsification [3, 4, 5, 6, 27] which prunes (zeros out) a subset of model parameters to reduce the representational complexity and improve run-time efficiency; (3) parameter quantization [7, 8, 9] which represents the trained model parameters with fewer bits for model size reduction. In this study, we adopt model pruning approaches (IMP and LTH) to find language-specific sub-networks. In [5], LTH was shown to learn a better sparse monolingual ASR than IMP. We study how the monolingual sub-networks could be used for multilingual scenario. A similar scheme was also investigated in [28] for multitask sequence-labeling.

### 3. PROPOSED METHOD

Firstly, we introduce 2 pruning methods: IMP and LTH, which we use to learn language-specific masks (Section 3.1). We then show how these identified masks are utilized to train our proposed ASR pathways (Section 3.2). Finally, we describe an essential regularization technique to improve the pruning performance (Section 3.3).

#### 3.1. Learning language-specific pruning masks

##### 3.1.1. Iterative Magnitude Pruning (IMP)

For a network  $f(x; \theta)$  with input samples  $x$  and model weights  $\theta$ , we denote a sub-network as  $f(x; m \odot \theta)$ , where  $m \in \{0, 1\}^{|\theta|}$  is a binary pruning mask, and  $\odot$  represents element-wise product. In this work, we consider a structured  $8 \times 1$  block-wise pruning strategy [4] which means that weight matrices will be zeroed out in  $8 \times 1$  groups. We prune the encoder Emformer layers and the predictor LSTM layer, with a uniform sparsity across all prunable

layers [4, 6]. Given the pre-trained dense weights  $\theta_0$ , we initialize  $\theta = \theta_0$  and  $m = 1^{|\theta_0|}$ , and proceed as follows:

##### Repeat

1. Train  $f(x; m \odot \theta)$  for  $T$  steps, leading to  $f(x; m \odot \theta_T)$ .
2. Prune  $p\%$  of the remaining weights that have the smallest magnitudes. Update  $m$  by setting the pruned positions to 0.
3. Set  $\theta = \theta_T$  (start from the updated weights in next iteration).

**Until**  $m$  reaches the target sparsity.

After  $n$  pruning iterations,  $(1 - p\%)^n$  weights survive. When the target sparsity is reached,  $m$  will be fixed afterward and we continue training until convergence. Note that depending on the training data, IMP can be language-agnostic (multilingual data) or language-specific (monolingual data). If trained on monolingual data, we consider the final mask  $m$  as the language-specific mask. We set  $p = 20$ , i.e. 20% of the remaining weights will be pruned in each iteration. Pruning interval  $T$  depends on training data size.

##### 3.1.2. Lottery Ticket Hypothesis (LTH)

LTH aims to identify a sub-network from a random or pre-trained initialization  $\theta_0$ , which can be trained or transferred in isolation and obtain matching performance as the original dense network [19, 29]. The algorithm for identifying such a sub-network is identical to IMP, except that there is a “rewinding” step replacing step 3 of the pruning loop in Section 3.1.1. In LTH, the rewinding step always sets  $\theta = \theta_0$ , as opposed to  $\theta = \theta_T$  in IMP, in each iteration. After rewinding, LTH may learn a set of model parameters that are different from IMP after training on the same data for certain steps, which decide different weights to be pruned. Hence, the final mask  $m$  identified by LTH may also be different from that identified by IMP. In [5], the authors show that the sparse network identified by LTH can outperform IMP on a monolingual ASR task. In this work, we also compare the performance of the multilingual ASR pathways trained using masks identified by IMP and LTH, respectively.

#### 3.2. Multilingual ASR pathways training

With the language-specific masks separately learned on each language, the ASR pathways are trained in a multilingual manner. Figure 2 demonstrates a 2-language case. Given the pre-trained dense multilingual ASR model weights  $\theta_0$ , and the language-specific masks  $m_{en}$  and  $m_{nl}$  learned via IMP or LTH, we identify the language-specific sub-networks  $f(x; m_{en} \odot \theta_0)$  and  $f(x; m_{nl} \odot \theta_0)$ . During training, for each step, we sample a monolingual batch from the multilingual dataset, identify a language-specific sub-network using the mask corresponding to the input language, do a forward and backward pass in the sub-network, and only update the parameters inside the sub-network. After training, we end up with a set of universal weights  $\theta^*$ , and  $f(x; m_{en} \odot \theta^*)$  and  $f(x; m_{nl} \odot \theta^*)$  form the ASR pathways for English and Dutch, respectively.

Compared to language-agnostic pruning which degrades performance on certain languages due to a single sub-network, we expect the sparse sub-networks that are tailored to specific languages can mitigate the performance loss. In addition, since the masks can overlap with each other, the ASR pathways share some common parameters between languages, which potentially benefit similar languages (e.g. with shared vocabularies), especially for the low-resource ones.

#### 3.3. Regularization: group lasso weight decay

We adopt structured  $8 \times 1$  block-wise pruning. However, the trained dense weights may not follow this pattern. As a result, block-wise

Model	Monolingual or Multilingual?	Sparsity (%)	WER (%)				
			en	fr	it	nl	Avg.
Dense (100M)	Multilingual	0	14.03	11.74	18.90	17.69	15.59
Dense (30M)	Monolingual	0	16.13	14.64	24.66	21.64	19.27
LAP	Multilingual	70.6	14.73	13.30	25.66	21.67	18.84
LSP - IMP	Monolingual	70.6	<b>11.34</b>	11.88	18.85	18.34	15.10
LSP - LTH	Monolingual	70.6	11.80	<b>11.85</b>	<b>18.38</b>	<b>18.16</b>	<b>15.05</b>
ASR Pathways - Random mask	Multilingual	70.6	13.76	13.07	20.60	19.78	16.80
ASR Pathways - IMP mask	Multilingual	70.6	12.94	11.98	19.91	17.65	15.62
ASR Pathways - LTH mask	Multilingual	70.6	<b>12.74</b>	<b>11.59</b>	<b>17.79</b>	<b>17.12</b>	<b>14.81</b>

**Table 1.** Test set WERs on MLS. All sparse models have 70.6% sparsity. The last column shows the average WER on the 4 languages.

pruning may unexpectedly discard important weights, causing performance drop. Hence, it is necessary to have specific algorithmic designs for structured pruning [30]. Following [6], we use group lasso regularization to mitigate the undesired pruning. Specifically, we define a group as a  $8 \times 1$  block, and the regularization term is added to the loss function  $\mathcal{L}$ :

$$\min_W \mathcal{L} + \sum_{i=1}^l \lambda_i \sum_{g \in \mathcal{G}} \|W_g^{(i)}\|_2 \quad (1)$$

where  $l$  denotes the number of prunable layers,  $W_g^{(i)}$  denotes a weight group in the  $i$ -th layer,  $\lambda_i$  denotes regularization strength. With group lasso, L2-norms of the  $8 \times 1$  blocks will be suppressed, so the blocks become more pruning-friendly: pruning such blocks leads to minimal performance drop. Following [6], we dynamically set  $\lambda_i$  based on layer-wise averages of the group L2-norms. We add group lasso regularization in the dense model training stage, as well as in the pruning iterations. Once the model reaches the target sparsity, the regularization is disabled. Since no pruning is involved during or after the multilingual ASR pathways training (Section 3.2), we do not include the regularization term in this stage. We discuss the effect of group lasso in Section 5.3.

## 4. EXPERIMENTAL SETUP

### 4.1. Dataset

We use 4 languages, English (*en*), French (*fr*), Italian (*it*) and Dutch (*nl*) from MLS dataset [31] as our training data. We follow the train/validation/test splits in [31]. The sizes of the training audio for the 4 languages are 44.7k hrs, 1.1k hrs, 0.2k hrs, 1.6k hrs, respectively. We use 80-dims log-Mel with 25 ms window size and 10 ms step size as the input features. For all multilingual training, we use the sampling strategy in [32] with  $\alpha = 0.5$  to re-balance the training data.

### 4.2. Implementation details

For our dense multilingual model, we use a streaming RNN-T model with 30 Emformer layers with 512 input dims, 2048 feed-forward dims, GeLU non-linearity, and no memory banks [20], resulting in about 100M parameters. We use an input feature stride 6, Emformer center segment length=4 (240 ms) and right context length=1 (60 ms), which results in 300 ms streaming latency. For the target vocabulary, we generate 512 word pieces for each language, aggregate the word pieces across 4 languages, and remove the duplicated tokens, leading to 1548 items in total. For simplicity, we use the same

output layer size for all the multilingual and monolingual training. We use an Adam optimizer with a tri-stage learning rate schedule [33] and a peak learning rate 1e-3. For monolingual pruning, we set 350K maximum updates and separately tune the pruning interval  $T$  (Section 3.1.1) on the validation set of each language. We train the multilingual ASR pathways for 200K steps. 16 GPUs are used for monolingual training and 32 GPUs for multilingual training, with a per-GPU batch size of 64 samples for both.

## 5. RESULTS

We denote the multilingual language-agnostic pruning as *LAP*, and the monolingual language-specific pruning via IMP and LTH as *LSP - IMP* and *LSP - LTH*, respectively. The proposed ASR pathways trained with IMP masks or LTH masks are denoted as *ASR Pathways - IMP mask* and *ASR Pathways - LTH mask*, respectively. We also include an ASR pathways model trained with randomly initialized language-specific masks (denoted as *ASR Pathways - Random mask*), which serves as a baseline to compare different language-specific masks. All sparse models have 70.6% sparsity. We denote the pre-trained dense model as *Dense (100M)*. To understand the difference between a sparse model and a small dense model that has around the same number of parameters, we also include a monolingual small dense model baseline by reducing the Emformer layers, which is denoted as *Dense (30M)*. We show the test WERs on the 4 languages in Table 1.<sup>1</sup>

### 5.1. Language-agnostic vs. language-specific pruning

First, comparing *Dense (100M)* with *LAP*, we see more significant performance drops on *it* and *nl* than *en* and *fr*, indicating that a single language-agnostic mask does not fit all languages. Since *en* has the most training data, we hypothesize that the learned pruning mask may be biased to *en* and results in performance degradation on other languages.

Second, we find that both *LSP - IMP* and *LSP - LTH* can perform on par with or better than *Dense (100M)*, even at the high sparsity of over 70%, suggesting that finding language-specific structures and parameters from an over-parameterized multilingual model with minimal performance drop is possible. Besides, both language-specific pruning methods significantly outperform *Dense (30M)*, which indicates that it is better to prune a pre-trained over-parameterized model than to train a small dense model from scratch.

<sup>1</sup>Note that the results presented in this paper are all based on low-latency streaming ASR models, so the WERs fall behind [31], while our in-house non-streaming Transformer/Conformer models [34, 35] provide similar WERs compared to [31].

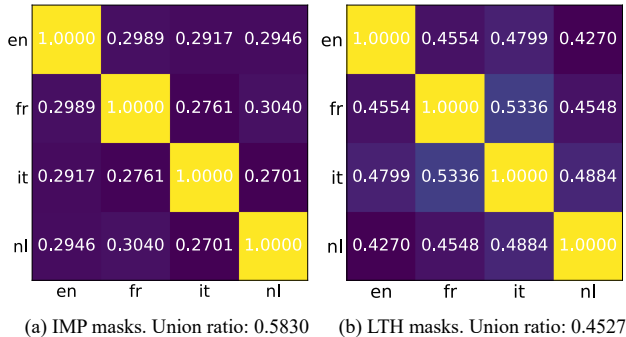


Fig. 3. IOUs and union ratios of (a) IMP masks and (b) LTH masks.

Third, we compare the proposed ASR pathways trained with different language-specific masks. Both *ASR Pathways - IMP mask* and *ASR Pathways - LTH mask* outperform *ASR Pathways - Random mask*, which suggests that IMP and LTH find better language-specific masks. Note that even with random masks, *ASR Pathways* outperforms *LAP*. This again confirms the importance of having language-specific masks in multilingual ASR scenario. In addition, *ASR Pathways - LTH mask* achieves comparable or better performance compared to the monolingual *LSP - LTH* on most languages (except for *en*), meaning that our proposed multilingual sparse ASR pathways can match the performance of monolingual sparse models. The improvements on *fr*, *it* and *nl* indicate the benefit of multilingual joint training on low-resource languages. On the flip side, the performance gap on *en* might be an outcome of the language interference problem [15]: high-resource languages can often suffer from negative transfer during multilingual joint training. The high sparsity makes it an even more challenging learning task. Finally, we observe that using LTH masks lead to better performance than IMP masks, consistent with the finding in [5].

## 5.2. Analysis on the language-specific masks

We conduct analyses on the learned language-specific masks to understand their distributions across different languages. We compute Intersection Over Union (IOU) between a mask pair ( $m_i, m_j$ ):

$$IOU(m_i, m_j) = \frac{|m_i = 1 \cap m_j = 1|}{|m_i = 1 \cup m_j = 1|} \quad (2)$$

A higher IOU suggests more overlapping between the surviving parameters of two language-specific masks. In addition, we define union ratio (UR) of the four masks ( $m_1, \dots, m_4$ ) as

$$UR(m_1, \dots, m_4) = \frac{|\bigcup_{n=1}^4 m_n = 1|}{|m_1|} \quad (3)$$

which measures how the union of the surviving parameters of the sub-networks span the whole parameters space. A higher UR suggests a higher total effective parameters usage.

We show the IOUs and union ratios in Figure 3. Comparing LTH masks with IMP masks, we find significantly higher IOUs in LTH masks. This indicates that LTH masks have more overlaps between language pairs, resulting in more shared parameters across different languages. As a result, the union ratio is also lower than IMP masks, meaning that LTH masks are using less total effective model parameters. One possible reason is that, in each iteration LTH starts from the same model parameters ( $\theta_0$  in Section 3.1.1), while

Model	WER (%)			
	en	fr	it	nl
Dense (100M)	<b>13.48</b>	<b>11.27</b>	21.14	19.55
+ group lasso	14.03	11.74	<b>18.90</b>	<b>17.69</b>
LSP - IMP	13.01	13.28	20.94	20.93
+ group lasso	<b>11.34</b>	<b>11.88</b>	<b>18.85</b>	<b>18.34</b>
LSP - LTH	14.80	18.50	28.87	24.92
+ group lasso	<b>11.80</b>	<b>11.85</b>	<b>18.38</b>	<b>18.16</b>
ASR Pathways - IMP mask	14.05	13.20	20.15	18.87
+ group lasso	<b>12.94</b>	<b>11.98</b>	<b>19.91</b>	<b>17.65</b>
ASR Pathways - LTH mask	14.65	14.04	21.67	20.77
+ group lasso	<b>12.74</b>	<b>11.59</b>	<b>17.79</b>	<b>17.12</b>

Table 2. Test set WERs with and without group lasso (Section 3.3).

IMP starts from different parameters ( $\theta_T$  in Section 3.1.1), making LTH better at keeping the universal part of parameters in  $\theta_0$  that are useful for all languages. In contrast, IMP will gradually lose this information and learn more language-specific parameters, producing less mask overlaps. Note that even with a lower union ratio (smaller total parameters usage), LTH masks can outperform IMP masks.

With a union ratio at 0.4527, *ASR Pathways - LTH mask* is essentially having more effective parameters for the 4 languages compared to the 70.6%-sparse *LAP*. In order to have a fairer comparison to *LAP*, we also train a *LAP* model with 56.5% sparsity and obtain test WERs of 13.93 (*en*), 12.58 (*fr*), 26.57 (*it*), 26.09 (*nl*). Corroborating with the results in Table 1, the overall trend still holds, confirming the superiority of language-specific masks.

## 5.3. Effect of group lasso regularization

In this section, we discuss the effect of the group lasso regularization (Section 3.3). We pre-train two dense models (i.e.  $\theta_0$  in Section 3.1.1) with and without group lasso regularization. Then we compare the performance of the sparse models using the two dense models as the pruning starting point, respectively.<sup>2</sup> Table 2 shows test set WERs, where the “+ group lasso” rows denote models based on the regularization-enabled dense model (same as in Table 1). With group lasso, *LSP - IMP* and *LSP - LTH* gain significant improvements. Correspondingly, the ASR pathways also show significant boosts. Interestingly, without group lasso, the performance of LTH largely falls behind IMP, and group lasso appears to have a larger improvement for LTH than IMP. As we analyzed in Section 5.2, due to the rewinding step, the initialization  $\theta_0$  has a major impact on LTH. Hence, it is essential to have structured-pruning-friendly pre-trained dense weights.

## 6. CONCLUSION

We have presented multilingual ASR pathways - a design that comes with language-specific activation. We show that such a sparse multilingual ASR model outperforms both dense model and language-agnostic pruning, and has better performance on low-resource languages compared to the monolingual sparse models. Further, via our analyses on the language-specific masks, we reveal that with the help of the group lasso regularization, LTH masks are superior to IMP masks, even with less total effective parameters. In the future, we plan to extend ASR pathways to more languages and different tasks.

<sup>2</sup>Note that for best pruning performance we enable the regularization for both during pruning.

## 7. REFERENCES

- [1] Ladislav Mošner, Minhua Wu, Anirudh Raju, Sree Hari Krishnan Parthasarathi, Kenichi Kumatani, et al., “Improving noise robustness of automatic speech recognition via parallel data and teacher-student learning,” in *Proc. ICASSP*, 2019.
- [2] Sankaran Panchapagesan, Daniel S Park, Chung-Cheng Chiu, et al., “Efficient knowledge distillation for rnn-transducer models,” in *Proc. ICASSP*, 2021.
- [3] Yuan Shangguan, Jian Li, Qiao Liang, et al., “Optimizing speech recognition for the edge,” *arXiv preprint arXiv:1909.12408*, 2019.
- [4] Haichuan Yang, Yuan Shangguan, Dilin Wang, et al., “Omni-sparsity dnn: Fast sparsity optimization for on-device streaming e2e asr via supernet,” in *Proc. ICASSP*, 2022.
- [5] Shaojin Ding, Tianlong Chen, and Zhangyang Wang, “Audio lottery: Speech recognition made ultra-lightweight, noise-robust, and transferable,” in *Proc. ICLR*, 2022.
- [6] Chunxi Liu, Yuan Shangguan, Haichuan Yang, Yangyang Shi, Raghuraman Krishnamoorthi, and Ozlem Kalinli, “,” in *Proc. SLT*, 2023.
- [7] Yanzhang He, Tara N. Sainath, Rohit Prabhavalkar, et al., “Streaming end-to-end speech recognition for mobile devices,” in *Proc. ICASSP*, 2019.
- [8] Tara N. Sainath, Yanzhang He, Bo Li, et al., “A streaming on-device end-to-end model surpassing server-side conventional model quality and latency,” in *Proc. ICASSP*, 2020.
- [9] Shaojin Ding, Phoenix Meadowlark, Yanzhang He, et al., “4-bit conformer with native quantization aware training for speech recognition,” *arXiv preprint arXiv:2203.15952*, 2022.
- [10] Shubham Toshniwal, Tara N. Sainath, Ron J. Weiss, et al., “Multilingual speech recognition with a single end-to-end model,” in *Proc. ICASSP*, 2018.
- [11] Vineel Pratap, Anuroop Sriram, Paden Tomasello, et al., “Massively Multilingual ASR: 50 Languages, 1 Model, 1 Billion Parameters,” in *Proc. Interspeech*, 2020.
- [12] Bo Li, Ruoming Pang, Tara N. Sainath, et al., “Scaling end-to-end models for large-scale multilingual asr,” in *Proc. ASRU*, 2021.
- [13] Anjali Kannan, Arindrima Datta, Tara N. Sainath, et al., “Large-Scale Multilingual Speech Recognition with a Streaming End-to-End Model,” in *Proc. Interspeech*, 2019.
- [14] Junwen Bai, Bo Li, Yu Zhang, et al., “Joint unsupervised and supervised training for multilingual asr,” in *Proc. ICASSP*, 2022.
- [15] Yizhou Lu, Mingkun Huang, Xinghua Qu, et al., “Language adaptive cross-lingual speech representation learning with sparse sharing sub-networks,” in *Proc. ICASSP*, 2022.
- [16] Neeraj Gaur, Brian Farris, Parisa Haghani, et al., “Mixture of informed experts for multilingual speech recognition,” in *Proc. ICASSP*, 2021.
- [17] Jeff Dean, “Introducing pathways: A next-generation ai architecture,” <https://blog.google/technology/ai/introducing-pathways-next-generation-ai-architecture/>, accessed: Sep-2022.
- [18] Michael Zhu and Suyog Gupta, “To prune, or not to prune: exploring the efficacy of pruning for model compression,” *arXiv preprint arXiv:1710.01878*, 2017.
- [19] Jonathan Frankle and Michael Carbin, “The lottery ticket hypothesis: Finding sparse, trainable neural networks,” in *Proc. ICLR*, 2019.
- [20] Yangyang Shi, Yongqiang Wang, Chunyang Wu, Ching-Feng Yeh, Julian Chan, et al., “Emformer: Efficient memory transformer based acoustic model for low latency streaming speech recognition,” in *Proc. ICASSP*, 2021.
- [21] Yangyang Shi, Chunyang Wu, Dilin Wang, Alex Xiao, et al., “Streaming transformer transducer based speech recognition using non-causal convolution,” in *Proc. ICASSP*, 2022.
- [22] Alex Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [23] Sharan Narang, Eric Undersander, and Gregory Diamos, “Block-sparse recurrent neural networks,” *arXiv preprint arXiv:1711.02782*, 2017.
- [24] Jinchuan Tian, Jianwei Yu, Chunlei Zhang, Chao Weng, Yuexian Zou, and Dong Yu, “Lae: Language-aware encoder for monolingual and multilingual asr,” *arXiv preprint arXiv:2206.02093*, 2022.
- [25] Yun Zhu, Parisa Haghani, Anshuman Tripathi, et al., “Multilingual Speech Recognition with Self-Attention Structured Parameterization,” in *Proc. Interspeech*, 2020.
- [26] Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli, “Unsupervised cross-lingual representation learning for speech recognition,” *arXiv preprint arXiv:2006.13979*, 2020.
- [27] Zhaofeng Wu, Ding Zhao, Qiao Liang, Jiahui Yu, Anmol Gulati, and Ruoming Pang, “Dynamic sparsity neural networks for automatic speech recognition,” in *Proc. ICASSP*, 2021.
- [28] Tianxiang Sun, Yunfan Shao, Xiaonan Li, Pengfei Liu, Hang Yan, Xipeng Qiu, and Xuanjing Huang, “Learning sparse sharing architectures for multiple tasks,” in *Proc. AAAI*, 2020.
- [29] Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, et al., “The lottery ticket hypothesis for pre-trained bert networks,” in *Proc. NeurIPS*, 2020.
- [30] Tianlong Chen, Xuxi Chen, Xiaolong Ma, Yanzhi Wang, and Zhangyang Wang, “Coarsening the granularity: Towards structurally sparse lottery tickets,” *arXiv preprint arXiv:2202.04736*, 2022.
- [31] Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, et al., “MLS: A large-scale multilingual dataset for speech research,” *arXiv preprint arXiv:2012.03411*, 2020.
- [32] Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhotia, Qiantong Xu, et al., “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [33] Daniel S. Park, William Chan, Yu Zhang, et al., “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech*, 2019.
- [34] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, et al., “Conformer: Convolution-augmented transformer for speech recognition,” *Proc. Interspeech*, 2020.
- [35] Chunxi Liu, Frank Zhang, Duc Le, Suyoun Kim, Yatharth Saraf, and Geoffrey Zweig, “Improving rnn transducer based asr with auxiliary tasks,” in *Proc. SLT*, 2021.