

PYCHAIN: A Fully Parallelized PyTorch Implementation of LF-MMI for End-to-End ASR

Yiwen Shao¹, Yiming Wang¹, Daniel Povey³, Sanjeev Khudanpur^{1,2}

¹Center for Language and Speech Processing, Johns Hopkins University, Baltimore, USA

²Human Language Technology Center of Excellence, Johns Hopkins University, Baltimore, USA

³Xiaomi Corporation, Beijing, China

{yshao18, yiming.wang, khudanpur}@jhu.edu, dpovey@gmail.com

Abstract

We present PYCHAIN, a fully parallelized PyTorch implementation of end-to-end lattice-free maximum mutual information (LF-MMI) training for the so-called *chain models* in the Kaldi automatic speech recognition (ASR) toolkit. Unlike other PyTorch and Kaldi based ASR toolkits, PYCHAIN is designed to be as flexible and light-weight as possible so that it can be easily plugged into new ASR projects, or other existing PyTorch-based ASR tools, as exemplified respectively by a new project PYCHAIN-EXAMPLE, and ESPRESSO, an existing end-to-end ASR toolkit. PYCHAIN's efficiency and flexibility is demonstrated through such novel features as full GPU training on numerator/denominator graphs, and support for unequal length sequences. Experiments on the WSJ dataset show that with simple neural networks and commonly used machine learning techniques, PYCHAIN can achieve competitive results that are comparable to Kaldi and better than other end-to-end ASR systems.

Index Terms: end-to-end speech recognition, lattice-free MMI, PyTorch, Kaldi

1. Introduction

In the past few years, end-to-end or pure neural approaches to automatic speech recognition (ASR) have received a lot of attention. Among them, connectionist temporal classification (CTC) [1], RNN-Transducer [2] and sequence-to-sequence models with attention [3, 4] are of high interests. This trend is largely caused by two main reasons: (i) the increasing demand on a simpler pipeline without several stages as in traditional hidden Markov model (HMM) based methods, and (ii) the easy access to the latest advances in deep learning, supported by powerful deep learning platforms like PyTorch [5] and TensorFlow [6]. As a result, many end-to-end ASR toolkits have been developed and have achieved impressive results, such as Deep Speech [7], ESPNET [8] and ESPRESSO [9].

On the other hand, in most scenarios (with small amount of data available especially) traditional hybrid (HMM-DNN) systems exemplified by Kaldi [10] perform better. In particular, chain models, with lattice-free MMI (LF-MMI) [11] training, are the state-of-the-art (SOTA) model in Kaldi [12]. Driven by the need for single stage training, an end-to-end version of LF-MMI (E2E LF-MMI) was proposed by Hadian et. al [13, 14] which removes any dependencies on HMM-GMM alignments, and the context-dependency trees customarily used in chain model training. But E2E LF-MMI is still implemented in Kaldi, and not compatible with PyTorch, TensorFlow etc.

This work was partially supported by the DARPA GARD program and an unrestricted gift from Applications Technology (AppTek).




Figure 1: The pipeline of doing end-to-end LF-MMI training with PYCHAIN.

To bridge the gap between Kaldi and other mainstream deep learning platforms, a lot of excellent work has been done recently, such as PyTorch-Kaldi [15], PyKaldi [16] and PyKaldi2 [17]. However, none of them have managed to implement fully parallel LF-MMI training¹, which tends to be the most effective and widely used loss function for training Kaldi ASR systems.

To fill in this gap, we present PYCHAIN, a light-weight yet powerful PyTorch implementation of the E2E LF-MMI criterion written in C++/CUDA but wrapped with PyTorch.² Also, we present examples of using PYCHAIN in two different scenarios: (a) PYCHAIN-EXAMPLE, a toy example written from scratch with only necessary utilities³, and (b) ESPRESSO, a more integrated end-to-end ASR toolkit originally built for sequence-to-sequence models⁴. A high-level overview of the pipeline is shown in Figure 1. By doing so, we are also able to make a more direct comparison of the E2E LF-MMI method and other end-to-end approaches without using Kaldi-specific functionality, such as natural gradient SGD, and parameter averaging across parallel jobs [18].

We perform ASR benchmark experiments on WSJ dataset [19]. Without bells and whistles, PYCHAIN is able to achieve SOTA results comparable to Kaldi, and better than other end-to-end ASR systems, but using simpler models and techniques.

The rest of the paper is organized as follows. We briefly discuss the LF-MMI loss function in Section 2. And then in Section 3, we describe PYCHAIN's architecture and implementation in detail. The experimental setup and results are shown in Section 4, and we present our conclusions in Section 5.

2. End-to-End LF-MMI

Maximum mutual information (MMI) [20], is one of the most popular criteria for discriminative sequence training in ASR. It

¹PyKaldi2 only implemented LF-MMI loss with minibatch size 1, which is unable to achieve competitive results

²<https://github.com/YiwenShaoStephen/pychain>

³https://github.com/YiwenShaoStephen/pychain_example

⁴<https://github.com/freewym/espesso>

takes into account the entire word-sequence (utterance) holistically in the objective instead of only considering individual frames, as in frame-level functions. It aims to maximize the *ratio* of the probability of the acoustics and reference transcription to that of the acoustics and all other possible transcriptions.

When MMI was first proposed, the marginal probability over all possible transcriptions was approximated using N-best lists, and later using lattices [21, 22]. Subsequently, Povey et. al [11] proposed lattice-free MMI (LF-MMI) training by utilizing an n-gram phone language model for the *denominator* computation. But they still require an alignment for the *numerator* computation. Hadian et. al [13, 14] extended LF-MMI training to not rely on any alignment, or even context-dependency trees, from a bootstrap model, such as an HMM-GMM system. In this work, we will mainly focus on this end-to-end (alignment-free) version of LF-MMI training.

2.1. LF-MMI Loss

For one utterance, MMI objective can be formulated as:

$$F_{MMI} = \log \frac{P(\mathbf{X}|\mathbf{W}_r)P(\mathbf{W}_r)}{\sum_{\hat{\mathbf{W}}} P(\mathbf{X}|\hat{\mathbf{W}})P(\hat{\mathbf{W}})} \quad (1)$$

where \mathbf{X} is the input frames sequence, \mathbf{W}_r is the gold transcription of \mathbf{X} , and $\hat{\mathbf{W}}$ is any possible transcription.

In LF-MMI [11], a n-gram phone language model (LM) is integrated with the acoustic part to encode all possible word sequences into one single HMM graph called *denominator* graph \mathbb{G}_{den} . Thus, we can replace the denominator part in Eq. (1) with $P(\mathbf{X}|\mathbb{G}_{den})$.

Similarly, the numerator part in Eq. (1) can be replaced with $P(\mathbf{X}|\mathbb{G}_{num})$ where \mathbb{G}_{num} is the numerator graph generated by composing the true transcript to the denominator graph \mathbb{G}_{den} . We will use *numerator* and *denominator* for short in this paper.

Extending to the corpus or batch level, we can get the final LF-MMI loss function as:

$$F_{MMI} = \sum_{u=1}^U \log \frac{P(\mathbf{X}^{(u)}|\mathbb{G}_{num}^{(u)})}{P(\mathbf{X}^{(u)}|\mathbb{G}_{den}^{(u)})} \quad (2)$$

2.2. LF-MMI Derivatives

Another big advantage of MMI loss function is that, although seemingly complicated, its derivatives can be finally reduced to the difference of occupation probability between numerator and denominator graphs[]:

$$\frac{\partial F_{MMI}}{\partial \mathbf{y}^{(u)}(s)} = \gamma_{num}^{(u)}(s) - \gamma_{den}^{(u)}(s) \quad (3)$$

where $\mathbf{y}^{(u)}$ is the network output of the u -th utterance, which we interpret as the log-likelihood for each state. And $\gamma_{num}^{(u)}(s)$ and $\gamma_{den}^{(u)}(s)$ are the state occupation probability of state s in the numerator and denominator graphs respectively, calculated by the Forward-Backward algorithm in HMM.

3. PYCHAIN

Our work consists of two separate parts. The first part is PYCHAIN, namely, the loss function itself written in C++/CUDA and wrapped in Python for PyTorch. It has following key components as shown in Figure 2:

- *openfst_binding*: Functions to read Kaldi numerator/denominator graphs stored as Finite State Transducers (FSTs) format and transform them into tensors.




Figure 2: The layout of PYCHAIN modules.

- *pytorch_binding*: Functions for Forward-Backward computation.
- *graph.py*: Classes for HMM graphs (ChainGraph and ChainGraphBatch)
- *loss.py*: Loss function ChainLoss (nn.Module) and ChainFunction (autograd.Function)

The second part is about the examples of using PYCHAIN. In order to illustrate the easy-to-use property of PYCHAIN, we give examples on either writing a speech recognition project from scratch (i.e. PYCHAIN-EXAMPLE, which is similar to PyTorch/examples⁵), or plugging PYCHAIN into an integrated ASR toolkit like ESPRESSO. It only takes to write basic PyTorch utilities such as dataloaders, models and train/test scripts with minimal codes.

3.1. Pipeline

As shown in Figure 1, we take advantage of both Kaldi and PyTorch to form a full recipe for an end-to-end ASR task. We do data preparations and final decoding in Kaldi for efficiency and consistency, but all other parts in PyTorch. Data preparation includes both feature extraction (e.g. MFCC) and numerator/denominator graph (FSTs) generation. Please note that, because there is no need for alignment in the E2E LF-MMI, we will not do any HMM-GMM pre-training in this stage. After all data is prepared, we go to PyTorch for dataloading and network training, which we will show in details below. After the model is trained and saved, we will load the data and model in PyTorch and then do a forward pass to get the output (posterior). They will be either piped to Kaldi for decoding on the fly, or dumped to the disk and then decoded afterwards.

3.2. Input Data

As in many other ASR toolkits [8, 9], we use Kaldi for data preparation. Both acoustic features and numerator/denominator graphs (FSTs) will be saved in scp/ark format by Kaldi. We use kaldi_io⁶ to read matrix data like input features as numpy arrays and then transform them into PyTorch Tensors. For FSTs, we write our own functions in C++ based on OpenFST [23] and then bind it with Python using pybind [24], so that all these functions can be called in Python seamlessly.

3.3. Numerator & Denominator Graphs

As shown in Eq. (2), HMM graphs are used as supervision in LF-MMI. We follow Kaldi's practice of using probability distribution function (pdf) to estimate the likelihood [10] of an HMM emission. As a result, both the network output and the occupation probability in Eq. (3) will be computed with respect to a

⁵<https://github.com/pytorch/examples>

⁶<https://github.com/vesis84/kaldi-io-for-python>

pdf-index (pdf-id) instead of an HMM state. And a typical transition in a HMM graph will be from a from-state s_f to a to-state s_t emitting a pdf-id d with a probability p . In this way, a dense transition matrix of an HMM graph will be of size (S, S, D) with transition probability p in each cell, where S is the total number of HMM states and D is the total number of pdf-ids.

However, due to the heavy minimization and pruning done in LF-MMI, both numerator and denominator graphs are very sparse. In PYCHAIN, we manage to represent numerator and denominator in a uniform way with a `ChainGraph` object \mathbb{G} that is similar to the COO (Coordinate list) format for a sparse matrix:

- A transition $\mathbf{T}^i = [s_f^i, s_t^i, d^i]$ where i is the index of the transition, s_f is the from-state and s_t is the to-state. d is the pdf-id on this transition. It basically denotes the coordinate of a transition.
- The forward transition matrix $F_T = [\mathbf{T}^i]$ of size $(I, 3)$ and forward transition probability vector $F_p = [p^i]$ of length I where I is the total number of transitions in a graph. They act like the coordinate list and value vector for a COO sparse matrix. They are denoted as F (meaning the "forward") as they are sorted in ascending order of the from-state s_f .
- The forward indices matrix $F_I = [(I_{start}^s, I_{end}^s)]$ of size $(S, 2)$ where I_{start}^s and I_{end}^s denote the start and end row index in F_T , between which are all transitions from s_f .

Similarly, we have another three tensors for backward transitions, namely, B_T, B_P and B_I . They are equivalent to these forward ones except that they are sorted by s_t . They are arranged in this way for a quick indexing of any transition by its from/to state.

In the end, there is a unique initial state s_{init} for each graph and a final probability vector $\mathbf{p}_{final} = [p^s]$ of length S where p^s is the probability of state s being the final state in this HMM graph.

As in most cases, LF-MMI will be trained with batches of utterance, we extend `ChainGraph` to its batch version, `ChainGraphBatch`. As its name suggests, it contains a batch of graphs and can be initialized by either a list of different `ChainGraph` objects (for numerators), or by a single `ChainGraph` object (for denominator). It has exact the same type of tensors as `ChainGraph` does, except that there is one more dimension for each of these tensors representing the batch dimension. And the default zero-padding is employed here when the sizes of these tensors differ.

3.4. Forward-Backward Computations

We follow the basic routine suggested by PyTorch custom C++ and CUDA extensions guide to write kernels for `ChainFunction`.

Our implementation has following key features:

- Both numerator and denominator use the same piece of codes for forward-backward computation with full support for CPU/GPU computation.
- The computation is done in probability space instead of log-probability space by utilizing leaky HMM [11] to solve underflow issues for both numerator and denominator.
- It supports variable lengths of sequences without doing mandatory silence padding or speed perturbation [25] that was required by [13, 14].

A detailed description of our implementation is shown in Algorithm 1. We only present Forward algorithm as Backward algorithm is similar to the forward pass and we will not repeat here. Also, for simplicity and clarity, we omit the leaky HMM and readers are referred to [11] for details.

In Algorithm 1, L denotes the log-likelihood of pdf-ids along the sequences as the output from a neural network. It has size of (B, T, D) , where B is the batch size, T is the length and D is the total number of pdf-ids. We get a batch of variable lengths sequences by firstly sorting these sequences by their lengths in descending order and then do the zero-padding to the right. We use $B_v(t)$ to denote the valid batch size at each sequence step t before any padding. For example, if there are 3 sequences of lengths (100, 99, 98) in a batch, B_v would be $[3, 3, \dots, 2, 1]$. Forward trellis is represented by α of size $(B, T + 1, S)$. Finally, we sum over the probability of each state in the final step of α to get the output O of the forward algorithm, which is $\log P(\mathbf{X}|\mathbb{G})$ in Eq. (2).

Algorithm 1 The Forward Algorithm. Loops over sequences and states in line 5-6 can be parallelized due to no dependency.

Input: $L, B_v, B_T, B_p, B_I, p_{final}$.
network output L of size (B, T, D) ;
valid batch size at each time step B_v of length T ;
backward transition matrix B_T of size $(B, I, 3)$;
backward transition probability B_p of size (B, I) ;
backward transition index B_I of size $(B, S, 2)$;
final probability of each state p_{final} of size (B, S) ;

Output: O .
total log-probability O of size (B) ;

```

1: procedure FORWARD( $L, B_v, B_T, B_p, B_I, p_{final}$ )
2:    $\alpha[:, :, :] := 0; \alpha[:, 0, s_{init}] := 1$   $\triangleright$  initialize  $\alpha$ 
3:   for  $t \leftarrow 1$  to  $T$  do  $\triangleright$  loop over time steps
4:      $bs := B_v[t]$ 
5:     for  $b \leftarrow 0$  to  $bs - 1$  do  $\triangleright$  loop over sequences
6:       for  $s \leftarrow 0$  to  $S - 1$  do  $\triangleright$  loop over states
7:          $(I_{start}, I_{end}) := B_I[b, s, :]$ 
8:         for  $i \leftarrow I_{start}$  to  $I_{end} - 1$  do
9:            $(s_f, s_t, d) := B_T[b, i, :]$ 
10:           $p := B_p[b, i]$ 
11:           $\alpha[b, t, s_t] += p \cdot \alpha[b, t-1, s_f] \cdot L[b, t, d]$ 
12:          if  $t = T$  then  $\triangleright$  multiply final prob
13:             $\alpha[b, t, s] *= p_{final}[b, s]$ 
14:    $O := \alpha[:, T, :].\text{sum}(1)$   $\triangleright$  sum over all states

```

4. Experiments

We do most of our experiments on WSJ (Wall Street Journal) [19] dataset, which is a database with 80 hours of transcribed newspaper speech. We consider the standard subsets *si284*, *eval92* for training and test, respectively. We use exactly the same subset from *si284* for validation as in Kaldi, i.e., randomly selected 300 utterances plus their corresponding speed-perturbed versions (if there are any).

4.1. Data Preparation

The 40-dimensional MFCC extracted from 25 ms frames every 10 ms are used as input features. They are then normalized on a per-speaker basis to have zero mean. For numerator and denominator graphs generation, we adopt the best setting in [13, 14]. The phone language model for the denominator graph

is estimated using the training transcriptions (choosing a random pronunciation for words with alternative pronunciations in the phoneme-based setting), after inserting silence phones with probability 0.2 between the words and with probability 0.8 at the beginning and end of the sentences. We use a trivial **mono-phone** tree without CD (context-dependency) modeling and a 2-state-skip HMM topology.

4.2. Model Architecture

As one of our motivations for this work is to show the generality of the LF-MMI training outside Kaldi, we use common components to build our model, which only includes 1D dilated convolution (CNN or TDNN) [26], batch normalization [27], ReLU [28] and dropout [29]. They are stacked 6 times in the sequence of conv-BN-ReLU-dropout with residual connections and finally followed by a fully connected layer. The network is of input dimension 40 (MFCC) and output dimension 84 (number of pdf-ids). The hidden dimension is 640 for all other layers. The dropout rate is set to 0.2, and the convolutional layers are of kernel sizes of (3, 3, 3, 3, 3, 3), strides of (1, 1, 1, 1, 1, 3) (equivalent to a subsampling factor of 3 in the original LF-MMI), and dilations of (1, 1, 1, 3, 3, 3).⁷

4.3. Training Schedule

For training schedule, we use Adam [30] as the optimization method. The learning rate is set to start from 10^{-3} and will be halved if no improvement on the validation set is seen at the end of each epoch, and finally fixed to 10^{-5} . Similar to the findings in [14], we also find out that curriculum training [31] (i.e. training utterance from short to long) is very helpful to E2E LF-MMI training in terms of robustness and ease to converge. However, we only do curriculum training for the first 1 or 2 epochs to achieve the best randomness in training. Otherwise, we sort all sequences in the corpus by its length so that sequences with similar lengths would form a minibatch. In other words, we only do shuffle on the batch level. Finally the model with the best validation loss is selected for decoding.

4.4. Empirical Results

We compare our results on WSJ *eval92* with the original E2E LF-MMI and other SOTA end-to-end ASR systems. As shown in Table 1, we not only achieve competitive results but also use a smaller number of parameters and much simpler structure. Note that we only use n-gram LMs for decoding, while the others (except Hadian et al. [14]) use more powerful neural LMs.

Our implementation is also more efficient than the original one where extra time is wasted on the padded silence parts. Some might argue that since the Forward-Backward algorithm has a temporal dependency on its previous step, the longest sequence inside a batch will decide the computational time and thus our unequal length version will not make a difference on the efficiency. The explanation is, as shown in Algorithm 1, because we are doing parallel computation state-wise and sequence-wise, there would be at most $B \cdot S$ (128×7 , $398 = 946, 944$ in our case) concurrent jobs on a GPU at once. However, there are much fewer CUDA cores in a modern GPU (e.g. 3584 for a NVIDIA GTX 1080 Ti in this experiment) which becomes the actual bottleneck for speed. As a result, the computational time will be proportional to the number of jobs we

⁷The detailed configuration can be found in https://github.com/freewym/espresso/blob/master/examples/asr_wsj/run_chain_e2e.sh.

have in total.

Flexibility is another big advantages of our variable lengths supported implementation. Instead of modifying utterance lengths by either silence padding or speed perturbation, we are able to form any utterances into a minibatch to get the maximal flexibility and randomness in training. Note that simply doing zero-padding at the end of each utterance does not work easily because these padded features may not be necessarily aligned to the final states and would probably lead to a terrible alignment and thus degrade the results.

Table 1: *WERs (%) on the WSJ without data augmentation.*

System	# Params (M)	WER (%)
Zeghidour et al. [32]	17	5.6
Baskar et al. [33]	~100	3.8
Likhomanenko et al. [34]	17	3.6
Zeghidour et al. [35]	17	3.5
Wang et al. [9]	18	3.4
Hadian et al. [14]	9.1	4.3
PYCHAIN	6.3	3.5

For the completeness of comparison, we also do experiments on augmented data with a 2-fold speed perturbation (sp) as was originally required by [13] and compare the results with other end-to-end systems with data augmentation. As shown in Table 2, we again match up to the results in Kaldi but with a smaller network.

Table 2: *WERs (%) on the WSJ with data augmentation.*

System	# Params (M)	WER (%)
An et al. [36]	16	3.2
Amodei et al. ⁸ [7]	-	3.1
Kriman et al. ⁹ [37]	19	3.0
Hadian et. al [14]	9.1	3.0
PYCHAIN	6.3	3.0

5. Conclusions and Future Work

We presented PYCHAIN, a PyTorch-based implementation of the end-to-end (alignment-free) LF-MMI training method for ASR. We used tensors to represent HMM graphs, which permits seamless transformation between Kaldi and PyTorch. Examples in PYCHAIN-EXAMPLE and ESPRESSO illustrate that PYCHAIN can be easily used in a project that includes building an ASR system from scratch, or extending an existing ASR system. We would like to support the use of PYCHAIN in such projects in the future.

Experiments on WSJ with simple networks exhibit the power of PYCHAIN. Work is underway to validate its performance on other benchmark tasks with new architectures.

We hope that our experience with PYCHAIN will inspire other efforts to build next-generation hybrid ASR tools. Possible next steps include extending PYCHAIN to support regular LF-MMI, log-space computations, and context-dependent acoustic units.

⁸12k hours AM train set and common crawl LM.

⁹Data augmentation and pre-trained on LibriSpeech [38] and Mozilla Common Voice datasets.

6. References

- [1] A. Graves, A.-r. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. ICASSP*, 2013.
- [2] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [3] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. NeurIPS*, 2014.
- [4] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, "End-to-end attention-based large vocabulary speech recognition," in *Proc. ICASSP*, 2016.
- [5] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Proc. NeurIPS*, 2019.
- [6] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "Tensorflow: A system for large-scale machine learning," in *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, 2016.
- [7] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," in *Proc. ICML*, 2016.
- [8] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "Espnet: End-to-end speech processing toolkit," in *Proc. INTERSPEECH*, 2018.
- [9] Y. Wang, T. Chen, H. Xu, S. Ding, H. Lv, Y. Shao, N. Peng, L. Xie, S. Watanabe, and S. Khudanpur, "Espresso: A fast end-to-end neural speech recognition toolkit," in *Proc. ASRU*, 2019.
- [10] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldı speech recognition toolkit," in *Proc. ASRU*, 2011.
- [11] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," in *Proc. INTERSPEECH*, 2016.
- [12] D. Povey, G. Cheng, Y. Wang, K. Li, H. Xu, M. Yarmohammadi, and S. Khudanpur, "Semi-orthogonal low-rank matrix factorization for deep neural networks," in *Proc. INTERSPEECH*, 2018.
- [13] H. Hadian, H. Sameti, D. Povey, and S. Khudanpur, "End-to-end speech recognition using lattice-free mmi," in *Proc. INTERSPEECH*, 2018.
- [14] —, "Flat-start single-stage discriminatively trained hmm-based models for asr," *IEEE/ACM TASLP*, 2018.
- [15] M. Ravanelli, T. Parcollet, and Y. Bengio, "The pytorch-kaldi speech recognition toolkit," in *Proc. ICASSP*, 2019.
- [16] D. Can, V. R. Martinez, P. Papadopoulos, and S. S. Narayanan, "Pykaldi: A python wrapper for kaldı," in *Proc. ICASSP*, 2018.
- [17] L. Lu, X. Xiao, Z. Chen, and Y. Gong, "Pykaldi2: Yet another speech toolkit based on kaldı and pytorch," in *Proc. ICASSP*, 2020.
- [18] D. Povey, X. Zhang, and S. Khudanpur, "Parallel training of dnns with natural gradient and parameter averaging," *Proc. ICLR workshop*, 2015.
- [19] D. B. Paul and J. M. Baker, "The design for the wall street journal-based csr corpus," in *Proceedings of the workshop on Speech and Natural Language*, 1992.
- [20] D. Povey, "Discriminative training for large vocabulary speech recognition," Ph.D. dissertation, University of Cambridge, 2005.
- [21] V. Valtchev, J. Odell, P. C. Woodland, and S. J. Young, "Lattice-based discriminative training for large vocabulary speech recognition," in *Proc. ICASSP*, 1996.
- [22] P. C. Woodland and D. Povey, "Large scale discriminative training of hidden markov models for speech recognition," *Computer Speech & Language*, 2002.
- [23] C. Allauzen, M. Riley, J. Schalkwyk, W. Skut, and M. Mohri, "Openfst: A general and efficient weighted finite-state transducer library," in *International Conference on Implementation and Application of Automata*, 2007.
- [24] W. Jakob, J. Rhinelander, and D. Moldovan, "pybind11 – seamless operability between c++11 and python," 2017, <https://github.com/pybind/pybind11>.
- [25] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Proc. INTERSPEECH*, 2015.
- [26] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proc. INTERSPEECH*, 2015.
- [27] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML*, 2015.
- [28] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proc. ICML*, 2010.
- [29] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *JMLR*, 2014.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2014.
- [31] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proc. ICML*, 2009.
- [32] N. Zeghidour, N. Usunier, G. Synnaeve, R. Collobert, and E. Dupoux, "End-to-end speech recognition from the raw waveform," in *Proc. INTERSPEECH*, 2018.
- [33] M. K. Baskar, L. Burget, S. Watanabe, M. Karafiát, T. Hori, and J. H. Černocký, "Promising accurate prefix boosting for sequence-to-sequence asr," in *Proc. ICASSP*, 2019.
- [34] T. Likhomanenko, G. Synnaeve, and R. Collobert, "Who needs words? lexicon-free speech recognition," in *Proc. INTERSPEECH*, 2019.
- [35] N. Zeghidour, Q. Xu, V. Liptchinsky, N. Usunier, G. Synnaeve, and R. Collobert, "Fully convolutional speech recognition," *arXiv preprint arXiv:1812.06864*, 2018.
- [36] K. An, H. Xiang, and Z. Ou, "CAT: crf-based ASR toolkit," *CoRR*, vol. abs/1911.08747, 2019. [Online]. Available: <http://arxiv.org/abs/1911.08747>
- [37] S. Kriman, S. Beliaev, B. Ginsburg, J. Huang, O. Kuchaiev, V. Lavrukhin, R. Leary, J. Li, and Y. Zhang, "Quartznet: Deep automatic speech recognition with 1d time-channel separable convolutions," in *Proc. ICASSP*, 2020.
- [38] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. ICASSP*, 2015.