



LAE: Language-Aware Encoder for Monolingual and Multilingual ASR

Jinchuan Tian^{1,2}, Jianwei Yu^{2,3}, Chunlei Zhang², Chao Weng^{2,3}, Yuexian Zou¹, Dong Yu²

¹ADSPLAB, School of ECE, Peking University, Shenzhen, China

²Tencent AI Lab, ³Tencent ASR Oteam

tomasyu@tencent.com; zouyx@pku.edu.cn

Abstract

Despite the rapid progress in automatic speech recognition (ASR) research, recognizing multilingual speech using a unified ASR system remains highly challenging. Previous works on multilingual speech recognition mainly focus on two directions: recognizing multiple monolingual speech or recognizing code-switched speech that uses different languages interchangeably within a single utterance. However, a pragmatic multilingual recognizer is expected to be compatible with both directions. In this work, a novel language-aware encoder (LAE) architecture is proposed to handle both situations by disentangling language-specific information and generating frame-level language-aware representations during encoding. In the LAE, the primary encoding is implemented by the shared block while the language-specific blocks are used to extract specific representations for each language. To learn language-specific information discriminatively, a language-aware training method is proposed to optimize the language-specific blocks in LAE. Experiments conducted on Mandarin-English code-switched speech suggest that the proposed LAE is capable of discriminating different languages in frame-level and shows superior performance on both monolingual and multilingual ASR tasks. With either a real-recorded or simulated code-switched dataset, the proposed LAE achieves statistically significant improvements on both CTC and neural transducer systems. Code is released¹.

Index Terms: code-switch, multilingual, bilingual, speech recognition

1. Introduction

As more than 60% of the population in the world are multilingualists[1], the automatic speech recognition (ASR) system for multilingual speech is in urgent need. Although rapid progress has been made in monolingual speech recognition, the recognition of multilingual speech still remains a challenging task[2]. In real applications, a pragmatic recognizer is commonly expected to address two scenarios simultaneously. One is the monolingual speech from multiple languages, in which only one language is used within each utterance but the languages used in different utterances vary. The other is the code-switched speech, in which multiple languages are used interchangeably within a single utterance.

Most of the previous works mainly focus on monolingual speech from various languages or the code-switched speech one or the other. For the former situation, the monolingual speech from multiple languages is handled by mixing multiple monolingual corpora during training[3, 4]. The latter situation

is mainly addressed by approaches like predicting language-switching points[5] or identifying language-IDs[6, 7], building unified acoustic[8, 9] or textual[10] representations for all considered languages, integrating language models customized for code-switch ASR[11, 12, 13], incorporating external monolingual speech and text resources[14, 15, 11], or enriching the multilingual corpora by simulated training utterances[16, 17, 18, 19]. To our knowledge, the works that indiscriminately consider both scenarios are still limited[1, 20, 21, 22, 23, 24, 25]. Our previous work [22] addresses this problem by factorizing the multilingual ASR task into monolingual sub-tasks and fine-tuning on encoders that are inter-independently pre-trained over the monolingual corpora. Although considerable improvement is achieved in [22], we argue the pre-training process is time-consuming and its performance is compromised due to the lack of interaction among the separated encoders.

To recognize both monolingual and code-switched speech consistently, a desired property of the system is to disentangle the language-specific information and generate frame-level language-aware representations during the encoding stage. Then these representations can be partially used in the monolingual situations and fully explored during the recognition of code-switched speech. In this work, a novel language-aware encoder (LAE) architecture is proposed, in which the utterance is firstly encoded by the shared block and the language-specific representations for each language are then extracted by several language-specific blocks. Additionally, a new language-aware training method enables the model to discriminate different language-specific information in both utterance-level and frame-level. For the recognition of each mono-language utterance, only one of the language-specific blocks is informative while the others stay implicitly idle. For the code-switched speech, the whole LAE is activated and the utterance-level encoder output is interlaced from the frame-level language-aware representations of all languages. The proposed LAE is applicable to two of the widely used ASR frameworks: connectionist temporal classification (CTC) [26] and neural transducers (NTs) [27]. Experiments conducted on ASRU 2019 Mandarin-English Code-Switch Challenge dataset suggest that the adoption of LAE architecture outperforms our previous work [22] and achieves statistically significant improvements: for models trained with either real-recorded or simulated code-switched data, up to 23.3% and 18.8% mixed error rate (MER) reductions are obtained.

Our main contributions are summarized as follows: Firstly, we propose a novel language-aware encoder (LAE) architecture, which is compatible with both multiple monolingual and code-switched speech. In contrast, most of the previous works mainly focus on one of these two scenarios. Secondly, a language-aware training method is provided and the LAE is allowed to discriminate language-specific information in both utterance-level and frame-level. Such methods can significantly

Work done when Jinchuan Tian was an intern in Tencent AI Lab. Jianwei Yu and Yuexian Zou are the corresponding authors. This paper was partially supported by Shenzhen Science & Technology Research Program (No:GXWD20201231165807007-20200814115301001) and NSFC (No: 62176008).

¹https://github.com/jctian98/e2e_lfmmi/egs/asrucs

boost the performance of multilingual ASR systems with or without real-recorded code-switched data.

2. Language-Aware Encoder

In this section, the proposed language-aware encoder (LAE) architecture, the language-aware training method and the unified decoding workflow are introduced in the following three parts respectively.

2.1. Language-Aware Encoder Architecture

Fig.1 demonstrates a Mandarin-English bilingual speech recognizer with the proposed LAE architecture. The input frames are firstly fed into the shared block followed by the parallel Mandarin and English blocks. The hidden representations from the two language-specific blocks are then combined as the LAE hidden outputs h_{bil} and then sent to the global decoder to facilitate speech recognition. The shared encoder is adopted to encode the input frames globally and provide more interaction in primary stage of encoding. In each of the language-specific blocks, information from the target language is captured while the non-target information is expected to be filtered out.

2.2. Language-Aware Training

Although the language-specific blocks are designed with the goal of disentangling language-specific representations, there is no explicit guarantee for this discriminative property as the encoding process of these blocks is symmetric. To allow the language-specific blocks to discriminatively extract the information from the target language and filter out the undesired information from other languages, the language-aware training method is proposed.

During training, a monolingual Mandarin token sequence Y_{man} is generated from the bilingual target token sequence Y by masking every English token with a special token $\langle eng \rangle$. Similarly, a monolingual English token sequence Y_{eng} is generated with the special token $\langle man \rangle$. Next, the monolingual Mandarin token sequence is considered as the target to optimize the output of Mandarin-specific decoder, which is an auxiliary criterion during training:

$$J_{man} = \text{Aux.Criterion}(Y_{man}, \text{Decoder}_{man}(h_{man})) \quad (1)$$

A similar auxiliary criterion is also used in the English side:

$$J_{eng} = \text{Aux.Criterion}(Y_{eng}, \text{Decoder}_{eng}(h_{eng})) \quad (2)$$

Finally, the global training objective is the interpolation between the original training objective J_{ori} and the mean of these auxiliary criteria. the weights of these two objectives are empirically set equal.

$$J = J_{ori} + (J_{man} + J_{eng})/2 \quad (3)$$

where the original training criterion J_{ori} can be either CTC criterion or transducer criterion according to the choice of the global decoder architecture.

In this work, the CTC criterion is consistently adopted as the auxiliary criterion for simplicity and the language-specific decoders are linear classifiers. We share the parameters of these language-specific decoders. As there is no overlap between the sequences Y_{man} and Y_{eng} , sharing these parameters will not result in a leak of supervision information and can achieve better alignment between the Mandarin and English blocks in the prediction of CTC $\langle blank \rangle$. Also, we make Y , Y_{man} and Y_{eng} in the same lengths in the hope that the language-specific blocks are able to predict the number of tokens when being inactive.

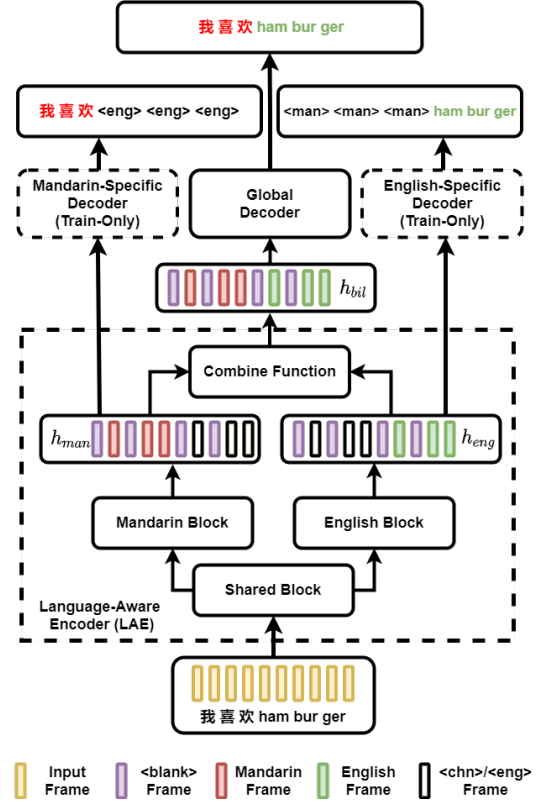


Figure 1: Mandarin-English Bilingual Speech Recognizer with Proposed Language-Aware Encoder (LAE) Architecture. (English translation of input utterance: I like hamburger)

2.3. Unified decoding with global decoder

As shown in Fig.2, for code-switched input, the proposed method is capable to disentangle the Mandarin and English segments in both monolingual and code-switched speech. Frames belonging to English and Mandarin can be extracted in English block outputs h_{man} and Mandarin block outputs h_{eng} respectively. For the monolingual input, the corresponding language-specific block is activated to predict ordinary token sequence while the other language-specific blocks are implicitly idle: they only predict the number of tokens by predicting repetitive masking symbols. Since there will be limited overlapping between the Mandarin and English frames, it is reasonable to simply combine h_{man} and h_{eng} with the frame-level addition:

$$h_{bil} = h_{man} \oplus h_{eng} \quad (4)$$

h_{bil} keeps the original order of the English and Mandarin segments and keeps the highly representative nature of h_{man} and h_{eng} for decoding. With h_{bil} , the global decoder can be used to recognize both the Mandarin and English monolingual speech and the code-switch speech in a unified manner.

3. Experiments

In this section, the experimental setup is introduced in section 3.1. The performance of the proposed method is firstly verified on the real-recorded code-switch speech dataset in section 3.2. The model's ability to discriminate languages is demonstrated in section 3.3. As the scarcity of multilingual data is widely concerned and data simulation methods are generally adopted for alleviation, it is of great importance to confirm that the proposed method can benefit from the simulated multilingual data. To this end, experiments on the simulated code-switch speech are further conducted in section 3.4.




Figure 2: CTC posterior probabilities along time-axis for different input speech. The upper row and the downer row are the output probabilities from Mandarin-specific decoder and English-specific decoder respectively. <Man-Mask> and <Eng-Mask> are the special Mandarin mask token and English mask token respectively. Left: code-switch input; middle: monolingual Mandarin input; right: monolingual English input.

3.1. Experimental Setup

All experiments are conducted on the ASRU 2019 Mandarin-English Code-Switch dataset[28], in which 200h code-switched speech and 500h Mandarin speech are provided. Following [22], 500h accented English data² is also used as the mono-English data. By default, all monolingual utterances are used in training. A simulated code-switch speech training set called *Simu-CS* is additionally generated by simply splicing the raw waves and the corresponding transcriptions from the monolingual speech. This simulated data also has the size of 200h and the length of each utterance is below 12 seconds. The code-switch evaluation set is officially provided and the monolingual evaluation sets are generated following [22]. For both global decoder and language-specific decoders, the predictions are based on a fixed vocabulary, which contains 4294 Mandarin characters, 5000 English BPE [29] tokens and three special tokens: <blank>, <man> and <eng>.

The input features are 80-dim Fbank features plus 3-dim pitch features. These features are sub-sampled with a factor of 4 by a two-layer CNN before being fed into the encoders. All encoders are stacked Conformer[30] layers, in which the attention dimension, feed-forward dimension, number of attention heads, number of convolutional kernels are fixed to 512, 2048, 4, 31 respectively. Three encoder architectures with similar parameter budgets are designed for comparison: **Vanilla**: 15 stacked Conformer layers. **Bi-Encoder**: two separated encoders with 8 Conformer layers stacked each. **Language-Aware Encoder (LAE)**: the shared block contains 9 conformer layers while the language-specific blocks consist of 3 Conformer layers each. For the CTC system, the encoder output is linearly classified to make predictions. The global decoder of the NT system consists of a 2048-dim LSTM prediction network and a 256-dim MLP joint network.

All models are optimized by Adam optimizer with warm-up steps of 25k, peak learning rate of $3e-4$ and inverse square root decay schedule. The training process lasts for 100 epochs and is implemented on 8 Tesla V100 GPUs. The effective batch size is 512 with a gradient accumulation number of 8. SpecAugment[31] is consistently used with 2 time masks and 2 frequency masks. The checkpoints from the last 10 epochs are averaged before evaluation. Standard CTC prefix search and ALS algorithm [32] are adopted for CTC and NT decoding respectively with a fixed beam size of 10. A word-level 5-gram language model[33] trained from all training transcriptions (monolingual and code-switched) is optionally used

by shallow fusion[34] with a fixed weight of 0.2. We report both code-switched and monolingual results: CER / WER only counts the Mandarin characters / English words while MER (Mixed Error Rate) counts the both. Code is mainly revised from Espnet[35].

3.2. Results with real-recorded code-switch speech

The results with real-recorded code-switch and monolingual training data are reported in table 1. Our main observations are reported as follows.

Overall performance: As shown in table 1, the proposed method achieves significant performance improvement on both CTC and NT systems. For the CTC system, the relative improvements on code-switch, mono-Mandarin and mono-English evaluation sets are 23.3%, 52.9% and 16.3% respectively (exp.1 vs. exp.6) while these numbers for the NT system are 6.3%, 28.9% and 4.2% (exp.7 vs. exp.10). The significance of these improvements is further confirmed by matched pairs sentence-segment word error (MAPSSWE) based significant test³ with the significant level of $p = 0.001$. The results in exp.6 and exp.10 verify that the proposed method is applicable to both monolingual and code-switch scenarios simultaneously.

Language-Aware Training for LAE: Compared with the vanilla CTC system, considerable improvement is provided by the LAE architecture (exp.1 vs. exp.4). Similar improvement is also achieved by the Bi-Encoder architecture (exp.1 vs. exp.2), which emphasizes the necessity of handling each language by a specific block. Further study shows that the error reduction provided by language-aware training in Bi-Encoder architecture is marginal (exp.2 vs. exp.3, also observed in [22]) while the adoption of language-aware training on the proposed LAE system provides up to 0.8% absolute improvement (exp.4 vs. exp.5). This observation implies the unified encoding process in the shared block is necessary. The adoption of LAE architecture and language-aware training also provides up to 1.3% absolute improvement on NT systems (exp.8 vs. exp.9). In summary, the combination of the LAE and the language-aware training shows superior performance on both CTC and NT systems.

Comparison between CTC and NT system. With the same LAE architecture, the NT system outperforms the CTC system if no language models are integrated (exp.5 vs. exp.9). As the performance differences are mainly in the English scenarios, this phenomenon can be attributed to the well-known conditional independence assumption of the CTC system[26]: English words need more context dependency to spell words

²<http://speechchocean.com>

³<https://github.com/talhanai/wer-sigtest>

No.	System	Code-Switch			Monolingual	
		All (MER)	Man (CER)	Eng (WER)	Man (CER)	Eng (WER)
Literature						
	Single NT [21, 1]	11.3	9.3	30.8	6.5	17.8
	Gating NT [20, 1]	11.2	8.8	34.7	5.7	34.6
	Factorized NT + Pre-train[22]	10.2	8.1	29.2	5.3	16.3
	Hybrid CTC-Attention [36]	9.5	7.6	25.0	-	-
Our Results						
1	Vanilla CTC (baseline)	11.6	9.2	38.7	5.1	20.3
2	Bi-Encoder CTC	10.3	8.1	36.1	3.6	19.4
3	+ Language-Aware Training	10.2	8.1	34.7	3.6	19.1
4	LAE CTC	10.3	8.2	34.9	3.5	18.5
5	+ Language-Aware Training (proposed)	9.5	7.5	33.1	3.2	17.7
6	+ word N-gram LM	8.9*	7.0*	30.9*	2.4*	17.0*
7	Vanilla NT (baseline)	9.5	7.9	28.6	3.8	16.6
8	Bi-Encoder NT	9.9	8.2	29.3	3.9	17.7
9	LAE NT + Language-Aware Training (proposed)	9.1	7.5	28.2	3.4	16.4
10	+ N-gram LM	8.9*	7.3*	27.7*	2.7*	15.9*

Table 1: Performance of models trained on both monolingual and real-recorded code-switch data. * means statistically significant improvement compared with the baseline systems.

Language Classification (Acc%)		
Code-Switch	Mandarin	English
99.65	99.91	99.98

Table 2: Utterance-level language classification accuracy of the proposed model over the three evaluation sets

Code-Switch		Monolingual	
Man (CER)	Eng (WER)	Man (CER)	Eng (WER)
Global Decoder (Reference)			
7.5	33.1	3.2	17.7
Mandarin-Specific Decoder			
9.3	100.0	3.5	100.0
English-Specific Decoder			
98.3	35.0	99.7	17.5

Table 3: Results of decoding using the language-specific decoders and the global decoder.

correctly. Further investigation demonstrates that the language models are more beneficial on the CTC system than on the NT system (exp.6 vs. exp.10), so we suppose the conditional-independence of CTC is alleviated by the context information in language models.

3.3. Ability to discriminate languages

In this part, we experimentally show that the proposed model is capable to discriminate the language-specific representations in both utterance-level and frame-level.

For utterance-level capability, an additional language classifier is applied to the time-dimensional mean value of LAE hidden output \mathbf{h}_{bil} during training to tell an utterance is mono-English, mono-Mandarin or code-switched. As suggested in table 2, the accuracy over the three evaluation sets is all above 99.6%, which implies the output of LAE provides strong evidence to classify the utterance.

To demonstrate that the LAE is also capable to discriminate language-specific information in frame-level, the decoding results of language-specific decoders are compared with those results of the global decoder in table 3. For both code-switch and monolingual situations, the language-specific encoders achieve comparable results with those of global encoders in the corresponding target languages but fail in other language. This observation proves that the hidden output \mathbf{h}_{man} and \mathbf{h}_{eng} are discriminative in frame-level: they keep sufficient information from the target language for independent decoding while filtering out the information from non-target languages.

Visualization is provided in Fig.2 to further exemplify the language-specific representations in frame-level. Firstly, the

language-specific encoders for both languages are well-aligned in time-axis and the CTC make spikes simultaneously. This property is desired to alleviate the *token-or-blank* conflicts in the feature combination stage. Secondly, whenever a Mandarin token is predicted by the Mandarin-specific decoder, the mask token is predicted by the English-specific decoder and vice versa. This verifies that the language-specific blocks can cooperate consistently: they only capture the information from the specific language and are activated interchangeably when the language-switching happens.

3.4. Results with simulated multilingual speech

In this part, we further demonstrate the proposed LAE method can benefit from the simulated code-switch data. Experimental results are presented in table 4.

System	Code-Switch			Monolingual	
	All MER	Man CER	Eng WER	Man CER	Eng WER
Vanilla CTC	34.0	35.5	102.5	5.8	20.6
+ Simu-CS	38.7	39.4	114.8	6.6	20.4
LAE CTC + Lang.-Aware Training	32.6	33.4	93.3	4.1	18.8
+ Simu-CS	29.5	29.3	89.1	4.0	18.5
+ N-gram LM	27.6	27.8	87.5	3.1	17.6

Table 4: Performances of models trained from monolingual data. The simulated code-switch data is optionally used.

For the LAE CTC system, the adoption of simulated code-switch training data achieves an absolute error reduction of 3.1% over the code-switch evaluation set. Further experiments with language models reduce the MER from 34.0% to 27.6%, which is a 18.8% relative improvement. By contrast, no performance improvement is observed on the vanilla CTC system when the simulated data is added. We suppose the simply stacked Conformer layers can hardly capture the dynamics of language-switching in the simulated data.

4. Conclusion

The language-aware encoder (LAE) and the language-aware training are proposed in this work to address monolingual and multilingual speech problems simultaneously. The proposed model is capable to capture the language-specific features of each language by specific encoder blocks. The proposed method is verified compatible with both CTC and neural transducer systems. Experimentally, our models achieve statistically significant improvements in both monolingual and multilingual scenarios and show their strengths on simulated bilingual corpora.

5. References

- [1] S. Dalmia, Y. Liu, S. Ronanki, and K. Kirchhoff, "Transformer-transducers for code-switched speech recognition," in *ICASSP*. IEEE, 2021, pp. 5859–5863.
- [2] S. Sitaram, K. R. Chandu, S. K. Rallabandi, and A. W. Black, "A survey of code-switched speech and language processing," *arXiv preprint arXiv:1904.00784*, 2019.
- [3] V. Pratap, A. Sriram, P. Tomasello, A. Hannun, V. Liptchinsky, G. Synnaeve, and R. Collobert, "Massively Multilingual ASR: 50 Languages, 1 Model, 1 Billion Parameters," in *Proc. Interspeech 2020*, 2020, pp. 4751–4755.
- [4] O. Adams, M. Wiesner, S. Watanabe, and D. Yarowsky, "Massively multilingual adversarial speech recognition," in *NAACL-HLT (1)*, 2019.
- [5] S. Rallabandi, S. Sitaram, and A. W. Black, "Automatic detection of code-switching style from acoustics," in *Proceedings of the Third Workshop on Computational Approaches to Linguistic Code-Switching*. Association for Computational Linguistics, Jul. 2018, pp. 76–81.
- [6] L. Zhou, J. Li, E. Sun, and S. Liu, "A configurable multilingual model is all you need to recognize all languages," *arXiv preprint arXiv:2107.05876*, 2021.
- [7] S. Punjabi, H. Arsikere, Z. Raeesy, C. Chandak, N. Bhawe, A. Bansal, M. Müller, S. Murillo, A. Rastrow, A. Stolcke, J. Droppo, S. Garimella, R. Maas, M. Hans, A. Mouchtaris, and S. Kunzmann, "Joint asr and language identification using rnn-t: An efficient approach to dynamic language switching," in *ICASSP*, 2021, pp. 7218–7222.
- [8] C. Jacobs, Y. Matuskevych, and H. Kamper, "Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation," in *SLT*. IEEE, 2021, pp. 919–926.
- [9] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised Cross-Lingual Representation Learning for Speech Recognition," in *Proc. Interspeech 2021*, 2021, pp. 2426–2430.
- [10] B. Li, Y. Zhang, T. Sainath, Y. Wu, and W. Chan, "Bytes are all you need: End-to-end multilingual speech recognition and synthesis with bytes," in *ICASSP*. IEEE, 2019, pp. 5621–5625.
- [11] G. Liu and L. Cao, "Code-switch speech rescoring with monolingual data," in *ICASSP*. IEEE, 2021, pp. 6229–6233.
- [12] H. Gonen and Y. Goldberg, "Language modeling for code-switching: Evaluation, integration of monolingual data, and discriminative training," in *EMNLP-IJCNLP*, Nov. 2019, pp. 4175–4185.
- [13] S.-P. Chuang, T.-W. Sung, and H.-y. Lee, "Training code-switching language model with monolingual data," in *ICASSP*, 2020, pp. 7949–7953.
- [14] X. Zhou, E. Yilmaz, Y. Long, Y. Li, and H. Li, "Multi-Encoder-Decoder Transformer for Code-Switching Speech Recognition," in *Proc. Interspeech 2020*, 2020, pp. 1042–1046.
- [15] S. Zhang, J. Yi, Z. Tian, Y. Bai, J. Tao, and Z. Wen, "Decoupling pronunciation and language for end-to-end code-switching automatic speech recognition," in *ICASSP*, 2021, pp. 6249–6253.
- [16] G. Lee, X. Yue, and H. Li, "Linguistically motivated parallel data augmentation for code-switch language modeling," in *INTERSPEECH*, 2019, pp. 3730–3734.
- [17] C.-T. Chang, S.-P. Chuang, and H.-Y. Lee, "Code-Switching Sentence Generation by Generative Adversarial Networks and its Application to Data Augmentation," in *Proc. Interspeech 2019*, 2019, pp. 554–558.
- [18] A. Pratapa, G. Bhat, M. Choudhury, S. Sitaram, S. Dandapat, and K. Bali, "Language modeling for code-mixing: The role of linguistic theory based synthetic data," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Jul. 2018, pp. 1543–1553.
- [19] M. Ma, B. Ramabhadran, J. Emond, A. Rosenberg, and F. Biadsy, "Comparison of data augmentation and adaptation strategies for code-switched automatic speech recognition," in *ICASSP*, 2019, pp. 6081–6085.
- [20] Y. Lu, M. Huang, H. Li, J. Guo, and Y. Qian, "Bi-Encoder Transformer Network for Mandarin-English Code-Switching Speech Recognition Using Mixture of Experts," in *Proc. Interspeech 2020*, 2020, pp. 4766–4770.
- [21] S. Zhang, J. Yi, Z. Tian, J. Tao, and Y. Bai, "Rnn-transducer with language bias for end-to-end mandarin-english code-switching speech recognition," in *ISCSLP*. IEEE, 2021, pp. 1–5.
- [22] B. Yan, C. Zhang, M. Yu, S.-X. Zhang, S. Dalmia, D. Berrebbi, C. Weng, S. Watanabe, and D. Yu, "Joint modeling of code-switched and monolingual asr via conditional factorization," *arXiv preprint arXiv:2111.15016*, 2021.
- [23] S. Shah, B. Abraham, S. Sitaram, V. Joshi *et al.*, "Learning to recognize code-switched speech without forgetting monolingual speech recognition," *arXiv preprint arXiv:2006.00782*, 2020.
- [24] V. M. Shetty and M. S. M. NJ, "Improving the performance of transformer based low resource speech recognition for indian languages," in *ICASSP*. IEEE, 2020, pp. 8279–8283.
- [25] V. Joshi, A. Das, E. Sun, R. R. Mehta, J. Li, and Y. Gong, "Code-Switching Sentence Generation by Generative Adversarial Networks and its Application to Data Augmentation," in *Proc. Interspeech 2019*, 2019, pp. 554–558.
- [26] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, pp. 369–376.
- [27] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [28] X. Shi, Q. Feng, and L. Xie, "The asru 2019 mandarin-english code-switching speech recognition challenge: Open datasets, tracks, methods and results," *arXiv preprint arXiv:2007.05916*, 2020.
- [29] R. Sennrich, B. Haddow, and A. Birch, "Neural machine translation of rare words with subword units," *arXiv preprint arXiv:1508.07909*, 2015.
- [30] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proc. Interspeech 2020*, 2020, pp. 5036–5040.
- [31] 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," *Interspeech 2019*, Sep 2019.
- [32] G. Saon, Z. Tüske, and K. Audhkhasi, "Alignment-length synchronous decoding for rnn transducer," in *ICASSP*, 2020, pp. 7804–7808.
- [33] J. Tian, J. Yu, C. Weng, Y. Zou, and D. Yu, "Improving mandarin end-to-end speech recognition with word n-gram language model," *IEEE Signal Processing Letters*, pp. 1–1, 2022.
- [34] C. Gulcehre, O. Firat, K. Xu, K. Cho, L. Barrault, H.-C. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On using monolingual corpora in neural machine translation," *arXiv preprint arXiv:1503.03535*, 2015.
- [35] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proceedings of Interspeech*, 2018, pp. 2207–2211.
- [36] S. Zhang, J. Yi, Z. Tian, J. Tao, Y. T. Yeung, and L. Deng, "Reducing language context confusion for end-to-end code-switching automatic speech recognition," *arXiv preprint arXiv:2201.12155*, 2022.