

THE ACCENTED ENGLISH SPEECH RECOGNITION CHALLENGE 2020: OPEN DATASETS, TRACKS, BASELINES, RESULTS AND METHODS

Xian Shi^{1*}, Fan Yu^{1*}, Yizhou Lu², Yuhao Liang¹, Qiangze Feng³, Daliang Wang³, Yanmin Qian², Lei Xie^{1†}

¹Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, Xi'an, China

²SpeechLab, Department of Computer Science and Engineering, Shanghai Jiao Tong University, China

³Datatang (Beijing) Technology Co., LTD, Beijing, China

ABSTRACT

The variety of accents has posed a big challenge to speech recognition. The Accented English Speech Recognition Challenge (AESRC2020) is designed for providing a common testbed and promoting accent-related research. Two tracks are set in the challenge – English accent recognition (track 1) and accented English speech recognition (track 2). A set of 160 hours of accented English speech collected from 8 countries is released with labels as the training set. Another 20 hours of speech without labels is later released as the test set, including two unseen accents from another two countries used to test the model generalization ability in track 2. We also provide baseline systems for the participants. This paper first reviews the released dataset, track setups, baselines and then summarizes the challenge results and major techniques used in the submissions.

Index Terms— Accented speech recognition, accent recognition, acoustic modeling, end-to-end ASR

1. INTRODUCTION

Accent is one of the major variable factors in human speech, which poses a great challenge to the robustness of automatic speech recognition (ASR) systems. English is one of the most common languages speaking around the world. It is inevitable to produce varieties of English accents in different areas. The difference between accents is mainly reflected in three aspects of pronunciation: stress, tone and duration, which brings difficulties to ASR models. There has been much interest in accent recognition to distinguish different English accents [1, 2, 3], and it is also valuable to improve the generalization capability of ASR models on varieties of English accents.

The Interspeech2020 Accented English Speech Recognition Challenge (AESRC)¹ is specifically designed to provide a common testbed and a sizable dataset for both English accent recognition (set as track 1) and accented English speech recognition (set as track 2). A 180-hour speech dataset is opened to participants, which contains 10 types of English accents – Chinese, American, British, Korean, Japanese, Russian, Indian, Portuguese, Spanish and Canadian. The two tracks run on the dataset to compare the submissions fairly.

The rest of this paper is organized as follows. Section 2 is a summary of related works on accent recognition, robust speech recognition on accented speech and current datasets available for related research. Section 3 describes the dataset released by the challenge. In Section 4, the baseline experiments are introduced. Section 5 mainly summarizes the outcome of the challenge, specifically discussing on

the major techniques used in the submitted systems. Section 6 concludes the challenge with important take-home messages.

2. RELATED WORK

Accent recognition is similar to language identification [4, 5, 6] and speaker identification [7, 8, 9]. They all classify variable-length speech sequences to utterance-level posteriors to obtain accent, speaker or language ID. In order to distinguish different accents in English, Teixeira et al. [10] proposed to use context-dependent HMM units to optimize parallel networks and Deshpande et al. [11] introduced format frequency features into GMM models. Ahmed et al. [12] presented VFNet (Variable Filter Net), a convolutional neural network (CNN) based architecture which applies filters with variable size along the frequency band to capture a hierarchy of features, aiming at improving the accuracy of accent recognition in dialogues. Winata et al. [13] proposed an accent-agnostic approach that extends the model-agnostic meta-learning (MAML) algorithm for fast adaptation to unseen accents. Transfer learning and multi-task learning were also found useful for spoken accent recognition tasks [14, 15].

As for speech recognition on accented speech, adaptation methods and adversarial training related techniques were proved effective. Assuming that the labelled data for specific accent is limited, the adaptation method is first to train a base model with standard speech data that is usually with a large volume, and then adapt the model to the specific accent with the respective data [16, 17]. Domain adversarial training (DAT) was used by Sun et al. to obtain accent-independent hidden representation in order to achieve a high-performance ASR system for accented Chinese [18]. A generative adversarial network (GAN) based pre-training framework named AIPNet was proposed by Chen et al. [19]. They pre-trained an attention-based encoder-decoder model to disentangle accent-invariant and accent-specific characteristics from acoustic features by adversarial training. Accent-dependent acoustic modeling approaches take accent-related information into network architecture by accent embedding, accent-specific bottleneck features or i-vectors [20, 21]. In a closed set of known accents, accent-dependent models usually outperform the accent-independent universal models, while the latter ones usually achieve a better average model under the situations where accent labels are unavailable.

English accent recognition and accented English speech recognition are also hindered by data insufficiency. Existing open-source accented English datasets are limited in data volume and accent varieties [22, 23]. This motivates us to provide a sizable dataset and a common testbed to advance the research in the related areas.

* The first two authors contributed equally to this work.

† Lei Xie is the corresponding author.

¹ <https://www.datatang.ai/interspeech2020>

Table 1. Results of baseline systems on the separated cv set.

Track 1: English Accent Recognition												
		Network	Accuracy (%)									
			Total	RU	KR	US	PT	JPN	UK	CHN	IND	
Self-Attention Classification Network		1a: Transformer-3L	54.1	30.0	45.0	45.7	57.2	48.5	70.0	56.2	83.5	
		1b: Transformer-6L	52.2	34.0	43.7	30.6	65.7	44.0	74.5	50.9	75.2	
		1c: Transformer-12L	47.8	49.6	26.0	21.2	51.8	42.7	85.0	38.2	66.1	
		1d: ASR init-Transformer-12L	76.1	75.7	55.6	60.2	85.5	73.2	93.9	67.0	97.0	
Track 2: Accented English Speech Recognition												
Training Strategy		LM and Decoding	WER (%)									
			Ave.	RU	KR	US	PT	JPN	UK	CHN	IND	
Chain Model		2a: Accent160	13.68	6.67	11.46	15.95	10.27	9.78	16.88	20.97	17.48	
		2b: Accent160 (Libri960 Base)	13.13	6.61	10.95	15.33	9.79	9.75	16.03	19.68	16.93	
		2c: Accent160+Libri160	13.14	6.95	11.76	13.05	9.96	10.15	14.21	20.76	18.26	
Transformer		2d: Accent160	+0.3 RNNLM	8.63	5.26	7.69	9.96	7.45	6.79	10.06	11.77	10.05
		2e: Accent160 (Libri960 Base)	+0.3 RNNLM	6.92	4.6	6.4	7.42	5.9	5.71	7.64	9.87	7.85
		2f: Accent160+Libri160	-	8.38	5.35	9.07	8.52	7.13	7.29	8.6	12.03	9.05
			+0.3 RNNLM	7.48	4.68	7.59	7.7	6.42	6.37	7.76	10.88	8.41
		+0.3 RNNLM + 0.3 CTC	7.38	4.76	7.81	7.71	6.36	6.4	7.23	10.77	8.01	

3. OPEN DATASET

An accented English speech dataset was released to participants in the challenge. It was collected from both native speakers in UK and US, and also English speakers in China, Japan, Russia, India, Portugal, Korea, Spain and Canada. We suppose that the data collected in each country is belong to one type of accent of English, and in total we have 10 ‘accents’. The speakers, aged between 20 to 60, were asked to read sentences covering common conversation and human-computer speech interaction commands. All the speech recordings were collected in relatively quiet indoor environment with Android phones or iPhones. The training set, named Accent160, for both challenge tracks (introduced in Section 4) has 160 hours of speech including 8 accents (20 hours/accent). Spanish and Canadian accents are not included in the training set. The test set for track 1 includes 16 hours of data (2 hours for each accent) and the test set for track 2 has 20 hours of data including extra 4 hours of data from Spanish and Canadian accents (as unseen accent data). Speech recordings were provided in Microsoft wav format (16KHz, 16bit and mono) with manual transcriptions.

Training speech data and the corresponding transcriptions were first released to participants together with metadata, in which detailed information about speakers and recording environment including speaker gender, age, region, recording device and others are provided. In order to make a fair comparison with the provided baseline experiments, we also release a speaker list for participants to divide the CV set from the training set. The test set was released later to the participants with only audio recordings.

4. TRACKS AND BASELINES

4.1. Track 1: English Accent Recognition

Track 1 aims to study the English accent recognition problem. The rules are as follows. 1) Data used for accent classification training is limited to the 160 hours Accent English data and 960 hours of Librispeech [24] data, and data augmentation based on the above data is permitted. 2) Multi-system fusion techniques are prohibited; there is no other restriction to the model and training techniques. 3) The final ranking is based on the recognition accuracy on the whole test set and the accuracy for each accent is for reference only.

For baseline experiments, a self-attention (SA) based classification network is realized using ESPnet². A mean + std pooling layer is applied after encoder to pool the output on T dimension. *Transformer-3L\6L\12L* are different in the number of encoder layers. All of them are trained under the simple CE loss for 40 epochs. SpecAugment is applied to the input feature. As shown in Table 1, with the use of the released data, 6\12 layers of SA encoder result in over fitting. From the results on the CV set, we found that the accuracy of some accents varies a lot among different speakers. As there are only a few speakers in the CV set, the absolute value above is not statistically significant. However, it is worth noting that when we use an SA encoder trained by an ASR downstream task to initialize the encoder of the accent classification network, the accent recognition accuracy is significantly improved. Finally the total accuracy of ASR-init-Transformer-12L is up to 76.1% on the CV set. The code and configuration of the baseline can be found from our github³.

4.2. Track 2: Accented English ASR

Track 2 studies the robustness of ASR system on accented English speech where the word error rate on the whole test set is used as the evaluation metric. Test sets include accents beyond training data in order to evaluate the generalization performance of the model. Again, data usage is only limited to the released data and the librispeech data. All kinds of system combination methods including ROVER are strictly prohibited. Language model training should only use the transcripts of permitted speech training data. Data augmentation should only be applied on the permitted speech data only.

We prepare ASR baseline systems for track 2 with both Kaldi⁴ and ESPnet toolkits. Several training strategies and decoding related parameters are compared, and results are shown in the Table 1.

In all experiments, we use 71-dimensional mel-filterbank feature as the input of the acoustic model and frame length is 25 ms with a 10 ms shift. In our baseline chain-model system, the acoustic model consists of a single convolutional layer with the kernel of 3 to down-sample the input speech feature, 19 hidden layers of a 256-dimensional TDNN and a 1280-dimensional fully connection layer with ReLU activation function. A 3-gram language model trained

² <https://github.com/espnet/espnet>

³ <https://github.com/R1ckShi/AESRC2020>

⁴ <http://www.kaldi-asr.org/>

Table 2. Results and major techniques used in the top 8 submissions in track 1.

Team Code	Network	Data Augmentation									Total Acc (%)
		Noise Augment.	Speed perturb.	Volume Augment.	Reverb simulation	Cutting & splicing	TTS	Pitch Shift	Spec Augment.	ASR Model	
S2	TDNN	✓	✓			✓	✓		✓	✓	83.63
E2	Transformer									✓	72.39
Z2	Jasper + Transformer								✓	✓	69.63
F	Transformer	✓	✓		✓	✓		✓	✓	✓	69.59
Baseline 1d	Self-Attention Encoder									✓	64.90
D2	RESNET + CTC										62.17
C	TDNN-F	✓	✓						✓	✓	60.62
V	TDNN-LSTM-Attention	✓	✓	✓	✓						56.67
H	Transformer	✓	✓	✓		✓					50.39

with the transcripts of the 160 hours of speech is used in the decoding graph compiling. As for the transformer baseline, ESPnet joint CTC/Attention transformer which contains 12-layer encoder and 6-layer decoder is used, and the dimension of attention and feed-forward layer is set to 256 (4 heads) and 2048 respectively. The whole network is trained for 50 epochs with warmup for the first 25,000 iterations. We mainly try three training strategies: 1) only using Accent160; 2) using Accent160 with another 160 hours of selected data from librispeech (Libri160); 3) using 960 hours of librispeech data (Libri960) to train a base model and then fine-tuning the model using Accent160. Furthermore, we optimize the decoding with RNN language model and CTC posterior probability. The RNNLM is a 2-layer LSTM model trained using ESPnet on the transcriptions associated with Accent160, and both RNNLM and CTC are fused (weight=0.3 for both) with beam search scores. From the baseline results, we find that the end-to-end models outperform the hybrid chain models given the limited training data and fine-tuning the librispeech base model with accented English data achieves the best performance among the three training strategies. The whole recipe and results of the baselines can be found from our github.

5. CHALLENGE RESULTS AND ANALYSIS

5.1. Track 1: English Accent Recognition

Finally 25 teams submitted their results to track 1 and the accuracy on the test set and the major techniques used for the top 8 teams are summarized in Table 3. The winner goes to team S2 using a TDNN based classification network with phonetic posteriorgram (PPG) feature as input, and they use text-to-speech (TTS) to expand the training data [25]. The major techniques are summarized below.

5.1.1. Data augmentation

Since the size of the released training data is relatively small, most teams have done a lot of work in data augmentation. For example, noise augmentation and speed perturbation are generally used. Speed augmentation can enhance the robustness of the model modestly, while volume augmentation and reverberation simulation help a little. Simulated room impulse response (RIR) is used to convolution with the original speech to generate data with reverberation.

Moreover, several teams use random cutting and splicing to expand the data. In detail, two pieces of audio with the same accent are randomly selected from the training dataset, and then each piece is cut to two splices and the splices from the two pieces are combined as new samples. Specaugmentation is also very useful reported by many teams. Pitch shift is also an effective data augmentation method reported by team F. In addition to the above tricks, it is also worth to notice that the winner team S2 used the provided data to train a TTS system to synthesize a large number of training audio clips with the corresponding accents, and the accuracy of accent recognition was improved absolutely by 10%.

5.1.2. Training Strategy

As revealed in the baseline experiments, the training of accent classification is easy to be over fitted to the speakers as the biggest acoustic difference lies in the speaker characteristics, which leads to great difference in the accuracy of different speakers with same accent. Therefore, it is beneficial to use speaker-invariant feature input and pre-trained encoder by speaker-invariant downstream task to initialize the network. Team S2 used PPG features generated by a Kaldi ASR system as model input. Team Z2 adopted a multi-task strategy with both accent recognition and phoneme classification. The second place team E2 used the accent label together with transcripts to train a Transformer ASR model with accent classification ability at the same time. As reported, putting the accent tag at the beginning of text outperforms tagging at the end. Besides the mainstream neural networks, team H used an NN and support vector machine (SVM) combination method. An embedding layer was applied before the fully-connected layer and SVM was used to classify the embedding vector.

5.2. Track 2: Accented English ASR

Forty two teams submitted their results to track 2. Team Q2 obtained the lowest average WER of 4.06%. This was achieved by a CTC model with LAS rescoring while the CTC model was initialized by a Wav2Vec encoder trained in an unsupervised manner using Fairseq toolkit [26]. The superior performance indicates that unsupervised training is promising in improving performance when labeled data is limited. The results and major techniques used in the top 10 systems are summarized in Table 3.

Table 3. Results and major techniques used in the top 10 submissions in track 2.

Team Code	Network	Data Augmentation				LM&Decode		AM&Training				Ave. WER(%)
		Noise & Reverb. Aug.	Speed Perturbation	Volume Augmentation	Spec Augmentation	2-Pass Decoding	NN LM	Multi-level Modeling	Accent Multi-task	Unsupervised Training	Transfer Learning	
Q2	Wav2vec enc. + CTC + LAS	✓	✓	✓	✓	✓	✓	✓		✓		4.06
S2	CTC + LAS	✓	✓		✓	✓	✓	✓			✓	4.47
E2	Transformer + CTC				✓		✓				✓	4.52
A2	Conformer + CTC		✓		✓		✓					4.71
T2	Transformer + CTC		✓		✓		✓		✓		✓	4.72
F	Conformer + CTC		✓		✓		✓				✓	4.95
U2	Lite Transformer + CTC				✓		✓					5.33
M3	BLSTM-CNN-TDNNF Hybrid	✓	✓			✓	✓				✓	5.67
D	Conformer + CTC						✓				✓	5.95
M2	CNN-Multi-Stream-TDNNF-Attn		✓	✓	✓	✓	✓				✓	6.23
Baseline 2f	Transformer + CTC						✓					8.29

5.2.1. Data augmentation

Similar to track 1, various data augmentation tricks were widely adopted in the submitted system in track 2 and Table 3 shows the tricks used in the top-performing systems. According to the system descriptions provided by the teams, the relative WER reduction of 5% to 10% can be achieved by methods including volume augmentation and speed perturbation. Noise and reverberation augmentation was tried by several teams but no obvious gain was obtained. This is probably because the acoustic and channel conditions between the test set and the training set is similar.

5.2.2. Network Structure

A variety of different models were found to be used in track 2, mainly including Transformer-based encoder-decoder models [27, 28, 29], CTC models with LAS [30] rescoring and traditional hybrid models. Attention-based end-to-end models are able to take full sentence-scale acoustic information into consideration, so they have a natural advantage over the traditional hybrid models. Team Q2 followed the work of wav2vec2.0 [31] and pre-trained a self-attention encoder using both contrastive loss and diversity loss, in order to obtain an encoder with waveform reconstruction capability and contextualized representation capability. The score of letter-level model above was combined with a word-piece level Transformer LM. The second place team S2 adopted frame-level CE loss pre-trained encoder [32] (labels are generated by a Kaldi triphone system), resulting in 5% WER reduction. Conformer and lite Transformer were used by several teams which implies the potential room for improvement of primitive Transformer, especially in enhancing the ability of local information modeling. Explicit accent related optimizations were rarely used in the submitted systems. But team T2 used accent recognition multi-task training encoder and yielded 3% relative WER reduction.

5.2.3. Language Modeling

As for language modeling, it is obvious that NN language model rescoring works well, and it brings improvement ranging from 7% to 15% on the CV set, reported from the system descriptions. Two-pass decoding was used by the top 2 teams. It is effective to rescore

the lattice generated by WFST (11.7% WER reduction by team S2) or fusion the primitive posterior probability in the process of beam search with LAS (26.4% WER reduction by Q2). Team Q2 specifically compared the performance of statistic language model and NN language model. It turns out that a well-trained Transformer LM can achieve a slightly lower perplexity, but the 4-gram model outperforms the Transformer LM in WER (3.96% to 4.01%). Team Q2 also tried a method of combining two language models using different granularity modeling units. A word-piece level Transformer LM was applied in the decoding fusion, leading to 2.4% relative WER reduction on the CV set (3.73% to 3.64%).

6. SUMMARY

According to the results of track 1, we found it necessary to fix the over-fitting problem, which means to peel off the speaker-related information from the encoder output. Therefore, a pre-trained encoder with ASR downstream task works well. The use of phonetic posteriorgram (PPG) features as network input is also effective in accent recognition. In track 2, we have seen a variety of networks including end-to-end models and traditional hybrid systems. In conclusion, CNN-based unsupervised training with contrastive loss and diversity loss can enhance the ability of waveform reconstruction and contextualized representation, leading to superior recognition performance. Strengthening local information like Conformer and Lite Transformer apparently leads to improved performance. As to language modeling, CTC with LAS 2-pass decoding performs well, which combines the time sequence modeling capability with CTC loss and full sentence scale modeling ability from self-attention. A few accent related modeling techniques are used in track 2, and most participants choose to train an average model for all accents. With limited training data, data augmentation tricks are essential for both tracks.

7. REFERENCES

- [1] Hamid Behravan, Ville Hautamäki, Sabato Marco Siniscalchi, Tomi Kinnunen, and Chin-Hui Lee, "I-vector modeling of speech attributes for automatic foreign accent recognition,"

- IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 1, pp. 29–41, 2015.
- [2] Yusnita Ma, M. P. Paulraj, Sazali Yaacob, A. B. Shahrman, and Sathees Kumar Nataraj, “Speaker accent recognition through statistical descriptors of mel-bands spectral energy and neural network model,” 2012.
 - [3] Maryam Najafian and Martin Russell, “Automatic accent identification as an analytical tool for accent robust automatic speech recognition,” *Speech Communication*, vol. 122, pp. 44–55, 2020.
 - [4] Pradeep Rangan, Sundeep Teki, and Hemant Misra, “Exploiting spectral augmentation for code-switched spoken language identification,” 2020.
 - [5] Nur Endah Safitri, Amalia Zahra, and Mirna Adriani, “Spoken language identification with phonotactics methods on minangkabau, sundanese, and javanese languages,” *Procedia Computer Science*, vol. 81, pp. 182–187, 2016.
 - [6] Chithra Madhu, Anu George, and Leena Mary, “Automatic language identification for seven indian languages using higher level features,” in *IEEE International Conference on Signal Processing*, 2017.
 - [7] Koji Okabe, Takafumi Koshinaka, and Koichi Shinoda, “Attentive statistics pooling for deep speaker embedding,” *arXiv preprint arXiv:1803.10963*, 2018.
 - [8] Weidi Xie, Arsha Nagrani, Joon Son Chung, and Andrew Senior, “Utterance-level aggregation for speaker recognition in the wild,” in *ICASSP 2019*. IEEE, pp. 5791–5795.
 - [9] Arsha Nagrani, Joon Son Chung, Weidi Xie, and Andrew Senior, “Voxceleb: Large-scale speaker verification in the wild,” *Computer Speech & Language*, vol. 60, pp. 101027, 2020.
 - [10] Carlos Teixeira, Isabel Trancoso, and António Serralheiro, “Accent identification,” in *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP’96*. IEEE, 1996, vol. 3, pp. 1784–1787.
 - [11] Shamalee Deshpande, Sharat Chikkerur, and Venu Govindaraju, “Accent classification in speech,” in *Fourth IEEE Workshop on Automatic Identification Advanced Technologies (AutoID’05)*. IEEE, 2005, pp. 139–143.
 - [12] Asad Ahmed, Pratham Tangri, Anirban Panda, Dhruv Ramani, and Samarjit Karmakar, “Vfnet: A convolutional architecture for accent classification,” in *2019 IEEE 16th India Council International Conference (INDICON)*. IEEE, 2019, pp. 1–4.
 - [13] Genta Indra Winata, Samuel Cahyawijaya, Zihan Liu, Zhaojiang Lin, Andrea Madotto, Peng Xu, and Pascale Fung, “Learning fast adaptation on cross-accented speech recognition,” *arXiv preprint arXiv:2003.01901*, 2020.
 - [14] Zhong Meng, Hu Hu, Jinyu Li, Changliang Liu, Yan Huang, Yifan Gong, and Chin-Hui Lee, “L-vector: Neural label embedding for domain adaptation,” in *ICASSP 2020*. IEEE, 2020, pp. 7389–7393.
 - [15] Thibault Viglino, Petr Motlicek, and Milos Cernak, “End-to-end accented speech recognition,” in *INTERSPEECH*, 2019, pp. 2140–2144.
 - [16] Bo Li and Khe Chai Sim, “Comparison of discriminative input and output transformations for speaker adaptation in the hybrid nn/hmm systems,” in *INTERSPEECH 2010*.
 - [17] Han Zhu, Li Wang, Pengyuan Zhang, and Yonghong Yan, “Multi-accent adaptation based on gate mechanism,” *arXiv preprint arXiv:2011.02774*, 2020.
 - [18] Sining Sun, Ching-Feng Yeh, Mei-Yuh Hwang, Mari Ostendorf, and Lei Xie, “Domain adversarial training for accented speech recognition,” *CoRR*, vol. abs/1806.02786, 2018.
 - [19] Yi Chen Chen, Zhaojun Yang, Ching Feng Yeh, Mahaveer Jain, and Michael L. Seltzer, “Aipnet: Generative adversarial pre-training of accent-invariant networks for end-to-end speech recognition,” in *ICASSP 2020*, 2020.
 - [20] Sanghyun Yoo, Inchul Song, and Yoshua Bengio, “A highly adaptive acoustic model for accurate multi-dialect speech recognition,” in *ICASSP 2019*, 2019.
 - [21] M. Chen, Zhanlei Yang, Jizhong Liang, Yanpeng Li, and Wenju Liu, “Improving deep neural networks based multi-accent mandarin speech recognition using i-vectors and accent-specific top layer,” in *INTERSPEECH*, 2015.
 - [22] Isin Demirsahin, Oddur Kjartansson, Alexander Gutkin, and Clara Rivera, “Open-source Multi-speaker Corpora of the English Accents in the British Isles,” in *Proceedings of The 12th LREC*, Marseille, France, May 2020, pp. 6532–6541, European Language Resources Association (ELRA).
 - [23] Ltd. Magic Data Technology Co., “Openslr68: Magicdata mandarin chinese read speech corpus,” .
 - [24] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *ICASSP 2015*. IEEE, 2015, pp. 5206–5210.
 - [25] Houjun Huang, Xu Xiang, Yexin Yang, Rao Ma, and Yanmin Qian, “Aispeech-sjtu accent identification system for the accented english speech recognition challenge,” in *Proc. ICASSP 2021*. IEEE.
 - [26] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli, “fairseq: A fast, extensible toolkit for sequence modeling,” in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
 - [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
 - [28] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang, “Conformer: Convolution-augmented transformer for speech recognition,” 2020.
 - [29] Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han, “Lite transformer with long-short range attention,” 2020.
 - [30] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *ICASSP 2016*. IEEE, 2016, pp. 4960–4964.
 - [31] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” 2020.
 - [32] Tian Tan, Yizhou Lu, Rao Ma, Sen Zhu, Jiaqi Guo, and Yanmin Qian, “Aispeech-sjtu asr system for the accented english speech recognition challenge,” in *Proc. ICASSP 2021*. IEEE, 2021 (to appear).