



Improving Hypernasality Estimation with Automatic Speech Recognition in Cleft Palate Speech

Kaitao Song¹, Teng Wan², Bixia Wang², Huiqiang Jiang¹, Luna Qiu¹, Jiahang Xu¹,
Liping Jiang², Qun Lou², Yuqing Yang¹, Dongsheng Li¹, Xudong Wang^{2,*}, Lili Qiu¹

¹Microsoft Research

² Department of Oral and Craniomaxillofacial Surgery, Shanghai Ninth People's Hospital,
Shanghai Jiao Tong University School of Medicine

{kaitaosong, hjiang, lunaqiu, jiahangxu, Yuqing.Yang, dongshengli, lililiqiu}@microsoft.com,
wonton1984@gmail.com, {bixial987, wangxudong70}@hotmail.com, jlping64@126.com,
wonderful_0322@163.com

Abstract

Hypernasality is an abnormal resonance in human speech production, especially in patients with craniofacial anomalies such as cleft palate. In clinical application, hypernasality estimation is crucial in cleft palate diagnosis, as its results determine the subsequent surgery and additional speech therapy. Therefore, designing an automatic hypernasality assessment method will facilitate speech-language pathologists to make precise diagnoses. Existing methods for hypernasality estimation only conduct acoustic analysis based on low-resource cleft palate dataset, by using statistical or neural network-based features. In this paper, we propose a novel approach that uses automatic speech recognition model to improve hypernasality estimation. Specifically, we first pre-train an encoder-decoder framework in an automatic speech recognition (ASR) objective by using speech-to-text dataset, and then fine-tune ASR encoder on the cleft palate dataset for hypernasality estimation. Benefiting from such design, our model for hypernasality estimation can enjoy the advantages of ASR model: 1) compared with low-resource cleft palate dataset, the ASR task usually includes large-scale speech data in the general domain, which enables better model generalization; 2) the text annotations in ASR dataset guide model to extract better acoustic features. Experimental results on two cleft palate datasets demonstrate that our method achieves superior performance compared with previous approaches.

Index Terms: Cleft Palate, Hypernasality, Automatic Speech Recognition.

1. Introduction

Cleft lip and palate (CLP) [1], is the most common congenital deformity in the oral and maxillofacial region, with the incidence of 1.4 in 1000 live births in human. In CLP patients, they usually have inability to pronounce the normal sound due to the incomplete closure of their soft palate (*i.e.*, velopharyngeal dysfunction), and thus results in the production of hypernasality [2, 3]. To treat this illness, CLP palates are usually required to conduct a series of palate surgery and subsequent speech therapy. Therefore, how to estimate the CLP severity is a critical element to determine the final treatment for patients.

Hypernasality is regarded as one of the primary symptoms in cleft palate and it is usually characterized as excessive nasal resonance in pronunciation. As a result, the CLP patients with

hypernasality cannot pronounce vowel and consonant as clear as normal people. Therefore, to a certain degree, the severity of hypernasality reflects the degree of the opening and closing of a velopharyngeal passageway between the oral cavity and nasal cavity. How to estimate the severity of hypernasality has been considered as an important metric to evaluate the outcome of primary cleft palate repair, and to determine the need for the further treatment, such as pharyngoplasty and speech therapy. In clinical examinations, hypernasality rating is usually evaluated by speech-language pathologists (SLPs). However, availability of expert SLPs is usually limited, and the diversity of intra-rater and inter-rater reliabilities will also affect the result of subjective hypernasality evaluation. Accordingly, some scientists attempted to investigate the possibility of detecting hypernasality by using machines or instruments. For example, nasometer [4] is a widely used instrument to measure nasalance with a rating of 0-100. Nevertheless, these instrumental methods do not demonstrate high correlation with clinical perception of hypernasality and these instruments also require experienced clinicians to operate [5].

Recently, machine learning based automatic hypernasality estimation has drawn enormous attention. Existing methods are mainly based on speech signal processing as cleft palate speech exhibits abnormal nasal resonance in the spectrum features. The training pipeline of machine learning based methods is to first extract acoustic features, like Mel-frequency cepstral coefficients (MFCC) or filter banks, and then perform classification using support vector machine and Gaussian mixture models. For example, [6] applied voice intonation features, MFCC and Teager energy operator to detect hypernasality. [7] investigated the effect of energy shifts to the low-frequency bands to estimate hypernasality. [8] introduced a single frequency filter bank based long-term average spectral for hypernasality estimation. [9] analyzed multiple different acoustic features for automatic identification of hypernasality. [10, 11] also designed some novel acoustic features to better extract hypernasality-related semantics. These works mainly focused on extracting or designing advanced acoustic features for hypernasal speech detection. Inspired by the success of deep neural network in speech processing tasks, some works [12, 13, 14, 15] also tried to use recurrent neural network, convolutional network or attention model to detect hypernasality in an end-to-end manner.

Although significant progresses have been achieved in evaluating hypernasality, there still remain a few limitations in the current systems. First, existing systems for hypernasality estimation only leverage cleft palate datasets. However, differ-

* Corresponding Author: Xudong Wang

ent from common speech processing tasks, cleft palate datasets are usually extremely difficult to collect, due to privacy and its annotations that require expert SLPs labeling. In other words, CLP datasets are usually in a low-resource scenario, which results in a poor generalization for model capacity. Secondly, previous work [16, 17] have analyzed that hypernasality has obviously abnormal area in acoustic space (*e.g.*, vowel or consonant), especially in the spectrum dimension [13]. Thus, it can be concluded that learning high quality acoustic features is beneficial for predicting hypernasality. However, existing systems (especially neural network based models) for hypernasality estimation only build the connection between the CLP speech and its hypernasality rating, without modeling acoustic features explicitly or implicitly. How to address the above two items for hypernasality estimation is the target of our paper.

Different from CLP datasets, many speech processing tasks in general domain usually have rich resource. Specifically, automatic speech recognition (ASR) [18] is one of the most representative tasks in speech processing. The objective of ASR is to recognize the content of the input speech and return its corresponding text. By leveraging additional text annotations, the model is able to understand the content of speech semantic (*e.g.*, vowel or consonant) and extract high quality representations for acoustic features. In addition, many ASR datasets [19, 20] usually have enormous speech data (*e.g.*, Librispeech [19] includes 960 hrs data), and do not suffer from the low-resource issue. Considering these characteristics of ASR task, we have raised a hypothesis: is it possible to improve hypernasality estimation by leveraging the advantages of the ASR objective.

As aforementioned, hypernasality derives abnormal expressions in the acoustic space (*e.g.*, vowel and consonant). The existing advanced methods for hypernasality estimation are mainly using neural network [12, 13, 14, 15]. But all of them just simply stacked multiple neural network layers over the acoustic features (like MFCCs) to conduct classification, and cannot enable model to further extract better acoustic semantics since these methods only focused on learning hypernasality-related features and ignored to understand acoustic semantics. However, it is worthy to note that ASR models can naturally extract high-quality representations of acoustic features (especially in identifying vowel and consonant), since large-scale text annotations encourage model to thoroughly understand the content of acoustic features (*e.g.*, for text “am” and its phonation /əɪm/, the model needs to identify vowel /ə/ and consonant /m/, and then returns the answer). Inspired by success of transfer learning in other tasks [21, 22], we deem that using ASR model for initialization is beneficial for estimating hypernasality.

Therefore, in this paper, we introduce a novel approach to leverage ASR model for hypernasality estimation from the perspective of transfer learning. Specifically, we first pre-train an encoder-decoder framework by using the ASR objective on large-scale ASR corpus, and then apply the pre-trained ASR encoder to conduct hypernasality estimation on CLP corpus. Such design can take advantages of ASR model in learning hypernasality-related semantics: 1) ASR tasks usually includes enormous audio data that enables model to obtain a better generalization; 2) Benefiting from text annotations, the neural network based encoder with ASR objective is able to extract better phonetic representations for hypernasality estimation. Experimental results on two different CLP datasets also indicate that our model with ASR initialization is superior to the base model without using initialization.




Figure 1: The training pipeline of our approach. The first row is using ASR dataset for model training, and then use ASR-trained encoder for an initialization to conduct hypernasality assessment on cleft palate dataset.

2. Method

2.1. ASR Formulation

Automatic speech recognition (ASR) is to identify the content of human voices and then return the corresponding text. Assuming ASR corpus as $(\mathcal{S}, \mathcal{T})$, where \mathcal{S} and \mathcal{T} correspond to audios and the corresponding language descriptions, respectively. For audio data, they are first processed as a sequence of acoustic features (*e.g.*, Mel filter-bank features) at the frame level, and then feed into the model. Generally, we adopt an encoder-decoder framework [23] to handle ASR task where the decoder is formulated as an auto-regressive manner, and the parameters are defined as $\theta = \{\theta_{enc}, \theta_{dec}\}$. The objective of ASR system is to predict the most possible corresponding text from the given speech as $\arg \max P(\mathcal{T}|\mathcal{S})$. The objective function is optimized by maximum likelihood estimation as:

$$\mathcal{L}_{ASR} = - \sum_{(s,t)} \sum_{i=1}^{|t|} \log P(t_i | t_{<i}, s; \theta_{enc}, \theta_{dec}), \quad (1)$$

where (s, t) is a paired speech and text, $t_{<i}$ means tokens before position i and $\{\theta_{enc}, \theta_{dec}\}$ represent the parameters of encoder and decoder, and $P(t_i | t_{<i}, s; \theta_{enc}, \theta_{dec})$ is to predict the i -th token in the text sequence based on previous position and the given speech. Therefore, by optimizing the Eqn. 1, we can obtain a well-trained ASR model in the general domain, and its encoder can extract deep representations from acoustic features so that the decoder is able to predict the corresponding text accurately.

2.2. Hypernasality Estimation Formulation

The target of hypernasality estimation is to evaluate the level of hypernasality from the speech of CLP patients. More specifically, by adopting different evaluation criteria, hypernasality estimation can be subdivided as hypernasality detection and hypernasality assessment. The former is to detect whether the cleft palate speech includes hypernasality or not, which can be simplified as a binary classification task. The latter is to predict the hypernasality rating of the corresponding cleft palate speech to determine the hypernasality severity, which is a more challenging multi-class classification task. In this paper, the rating of hypernasality severity is formulated as a range of 0–3 (0-normal, 1-mild, 2-moderate, 3-severe), each of which represents one class in the multi-class classification task.

The CLP corpus can be denoted as a 2-tuple $(\mathcal{X}, \mathcal{Y})$, where \mathcal{X} and \mathcal{Y} represent the datasets of CLP speech and the corresponding hypernasality rating, respectively. For CLP speech, we adopt the same preprocessing step to obtain the similar acoustic features (*i.e.*, Mel filter-bank feature), like ASR task, to

Rating	# NMCPC	# CNH
Normal	25	239
Mild	11	190
Moderate	14	509
Severe	16	108

Table 1: *Statistics of NMCPC and CNH cleft palate datasets. “# NMCPC” and “# CNH” column represent the patient number of each hypernasality severity.*

guarantee the consistency. To conduct hypernasality estimation, we employ an encoder network with a classifier layer, where the parameters are defined as $\tilde{\theta} = \{\tilde{\theta}_{enc}, \tilde{\theta}_{cls}\}$. Therefore, the objective of hypernasality estimation is formulated as following:

$$\mathcal{L}_{CLS} = - \sum_{(x,y)}^{(\mathcal{X},\mathcal{Y})} \log P(y|x; \tilde{\theta}_{enc}, \tilde{\theta}_{cls}), \quad (2)$$

where $\tilde{\theta}_{enc}$ and $\tilde{\theta}_{cls}$ represent the parameters of the encoder and classifier respectively. The classifier is a linear layer to scale the dimension of the encoder outputs as the number of categories $|C|$, with a connected softmax function to predict probability. $|C|$ is respectively set as 2 and 4 for hypernasality detection and hypernasality assessment.

2.3. Transfer Learning

As mentioned above, with the help of large-scale speech data and the labeled text, the ASR encoder can be viewed as a powerful acoustic feature extractor. Consequently, we deem that the encoder with ASR training is better for hypernasality estimation due to its ability in learning high quality acoustic representation. To fulfill this target, we require the architecture of encoder used in the ASR task and hypernasality estimation to be completely identical, and same as the input acoustic features (Mel filter-bank). More specifically, we first train an encoder-decoder framework in an ASR objective, and then conduct hypernasality estimation task based on ASR encoder. Therefore, the objective function of our method is as follows:

$$\mathcal{L}_{CLS} = - \sum_{(x,y)}^{(\mathcal{X},\mathcal{Y})} \log P(y|x; \theta_{enc}, \tilde{\theta}_{cls}),$$

where θ_{enc} is the parameters of ASR encoder in Eqn. 1. Figure 1 also presents a simple description about our training pipeline.

3. Experiments

3.1. CLP dataset

3.1.1. NMCPC-CLP

NMCPC-CLP is a dataset collected by New Mexico Cleft Palate Center [8], which is mainly composed of English speakers with cleft palate. This dataset includes 41 CLP patients and we sample 25 normal speakers as the control group. The average age of NMCPC-CLP is 9.2 ± 3.3 years. Each patient is required to record a random subset of sentences from the candidate sentences. Based on the sampled audios, each patient will be assigned a score from 0 to 3 (0 stands normal and 3 means severe). For hypernasality detection, the patients with hypernasality are defined as who has a score between 1 and 3. Refer to Table 1 for more details about the NMCPC-CLP dataset.

		ASR	CLS
CNN	stride	2, 2	
	kernel width	5, 5	
	channel	1024	
Transformer	layer	12	
	hidden size	512	
	filter size	2048	
	dropout	0.1	
	heads	8	
Optimization	batch	256	32
	learning rate	2e-3	2e-4
	Epoch	100	30

Table 2: *Hyper-parameters of experimental setup. The “ASR” and “CLS” column means the setting of automatic speech recognition, and hypernasality estimation.*

3.1.2. CNH-CLP

CNH-CLP is a dataset collected by a Chinese hospital, which includes cleft palate patients with a range from children to adults. All patients in CNH-CLP dataset are Chinese speakers. The procedure of audio collection for each patient is similar as NMCPC-CLP dataset. Table 1 introduces the detailed information about the CNH dataset.

3.2. ASR Datasets

3.2.1. Librispeech

Librispeech [19] is a large-scale speech recognition dataset in English domain. Librispeech includes 960 hrs speech data, sampled at 16,000 Hz with the corresponding text for training. During the ASR training, we select dev-clean/dev-other as the development set. We use Librispeech to train an ASR model, and then apply it to NMCPC-CLP dataset since these two datasets are both in English domain.

3.2.2. Aishell

Aishell-1 [20] is a widely used speech recognition dataset in Chinese domain. Aishell-1 dataset includes 150/10 hrs audios for the training/dev set. Each audio is sampled at 16,000 Hz, with the corresponding text. We use Aishell-1 for ASR training and then use it for CNH-CLP dataset as these two tasks belong to the same language.

3.3. Setup

For audio inputs, we first resample audios at 16,000 Hz and then extract 80-channel log mel filter-bank features (25 ms window size and 10 ms shift) for all ASR datasets and CLP datasets. We choose Fairseq-S2T [24] as the development toolkit. For ASR task, the encoder is composed of two convolutional layers for sub-sampling and a stack of transformer layers [25], and the decoder is a stack of transformer layers with cross-attention modules. During the ASR training stage, we use SpecAugment [26] for data augmentation. We adopt Adam [27] as the default optimizer. The detailed hyperparameters are reported in Table 2. In addition, to reduce variance, we use five-fold cross-validation to measure our classification accuracy. All evaluations are tested at the speaker level.

4. Results

4.1. Hypernasality Estimation Accuracy

We adopt the precision as the metric to evaluate the performance of our model in hypernasality assessment. The results are shown in Table 3. Our baseline is the base model without using ASR encoder. We also list some baselines [8], which use statistical features for reference. From Table 3, we have the following observations: 1) when configured with ASR encoder, our method can outperform the baseline method by a large margin, including hypernasality detection and assessment, in both NMCCPC and CNH cleft palate datasets; 2) compared with the statistical methods, our method also achieves significant improvement, especially in hypernasality detection. These improvements also demonstrate the effectiveness of our method. Besides, our method is capable of generality and can be applied to any neural networks used in previous works [10, 12].

Method	NMCCPC		CNH	
	HD	HA	HD	HA
Baseline	90.1 ± 0.3	77.8 ± 0.6	93.1 ± 0.3	73.1 ± 0.3
+ ASR	93.4 ± 0.2	83.6 ± 0.2	96.5 ± 0.2	79.8 ± 0.5
MFCC [8]	84.07	62.64	88.1 ± 0.2	69.1 ± 0.1
CQCC [8]	84.07	70.05	-	-
SFFB [8]	89.00	82.10	-	-

Table 3: Results of hypernasality estimation on NMCCPC and CNH cleft palate datasets. The “+ ASR” row means using ASR encoder. “HD” and “HA” represent hypernasality detection and hypernasality assessment respectively.

4.2. Analysis




Figure 2: Confusion matrix of our method for hypernasality classification on CNH dataset.

4.2.1. Confusion Matrix

To demonstrate the generalization of our model in predicting hypernasality, we also visualize the confusion matrix of our model by using ASR encoder in hypernasality classification, and the results are shown in Figure 2. From Figure 2, we observe that our model is significant in predicting each label, especially in normal, mild and moderate patients. Considering that CNH is an unbalanced dataset (moderate and severe patients occupy 50% and 10% respectively), the accuracy of our model in

predicting severe cases is still acceptable. Overall, these results indicate that our model can obtain better generalization performance and avoid overfitting effectively.

4.2.2. Visualization

To better explain the advantages of our methods in identifying acoustic features, we also visualize the hidden unit of encoder output. More specifically, we assume the output of encoder as $h \in \mathbb{R}^{L \times D}$, where L and D represent the frame length and the hidden size, respectively. Therefore, we first calculate the average value of h , and then use its absolute value as the activated value of i -th frame (i.e., $\frac{1}{D} |\sum_{j=1}^D h_{i,j}|$). The result and its corresponding mel spectrum are shown in Figure 3. We can find that the outputs of our model is more sensitive to acoustic features and demonstrate higher activated value, while the model without using ASR is too smooth to distinguish the semantic of acoustic features. This phenomenon also validates our hypothesis that ASR model can indeed help to extract high quality acoustic features.




Figure 3: Comparisons between our method and baseline without using ASR in the activated value of encoder output.

5. Conclusion

In this paper, we introduce a simple and effective approach to improve hypernasality estimation from the perspective of using ASR model. More specifically, we note that existing neural network based methods only stack multiple neural network layers for classification, and cannot extract high quality representation for acoustic features, which is useful for hypernasality estimation. To address this deficiency, we propose to fine tune the encoder, which is pre-trained with the ASR objective, for hypernasality estimation. Such design allows our model to enjoy the benefits of ASR model from two aspects: 1) ASR corpus usually includes more audio data, which enables better generalization; 2) the labeling text of ASR corpus guides model to better extract acoustic features. Experimental results on two cleft palate datasets also demonstrate the effective of our methods in hypernasality assessment. In the future, we expect to focus on two research directions: 1) is it possible to utilize more powerful pre-trained speech models (like wav2vec [28]) to conduct hypernasality estimation in cleft palate speech; 2) traditional methods for hypernasality estimation usually designed some advanced statistical acoustic features, and thus we want to explore the potential of combining neural network based features and statistical features to better estimate hypernasality.

6. References

- [1] A. Kummer, "Evaluation of speech and resonance for children with craniofacial anomalies," *Facial Plastic Surgery Clinics of North America*, vol. 24, 11 2016.
- [2] P. Vijayalakshmi, M. R. Reddy, and D. D. O'Shaughnessy, "Acoustic analysis and detection of hypernasality using a group delay function," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 4, pp. 621–629, 2007.
- [3] G. Lee, C. Wang, C. C. H. Yang, and T. B. J. Kuo, "Voice low tone to high tone ratio: a potential quantitative index for vowel [a:] and its nasalization," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 7, pp. 1437–1439, 2006.
- [4] K. Bettens, F. Wuyts, and K. Lierde, "Instrumental assessment of velopharyngeal function and resonance: A review," *Journal of communication disorders*, vol. 52, 05 2014.
- [5] T. Watterson, S. C. McFarlane, and D. S. Wright, "The relationship between nasalance and nasality in children with cleft palate," *Journal of Communication Disorders*, vol. 26, no. 1, pp. 13–28, 1993.
- [6] A. Maier, F. Hoenig, T. Bocklet, E. Noeth, F. Stelzle, E. Nkenke, and M. Schuster, "Automatic detection of articulation disorders in children with cleft lip and palate," *The Journal of the Acoustical Society of America*, vol. 126, pp. 2589–602, 11 2009.
- [7] L. He, J. Zhang, Q. Liu, H. Yin, and M. Lech, "Automatic evaluation of hypernasality and consonant misarticulation in cleft palate speech," *IEEE Signal Processing Letters*, vol. 21, no. 10, pp. 1298–1301, 2014.
- [8] M. H. Javid, K. Gurugubelli, and A. K. Vuppala, "Single frequency filter bank based long-term average spectra for hypernasality detection and assessment in cleft lip and palate speech," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, 2020, pp. 6754–6758.
- [9] M. Golabbakhsh, F. Abnavi, M. Elyaderani, F. Derakhshandeh, F. Khanlar, P. Rong, and D. Kuehn, "Automatic identification of hypernasality in normal and cleft lip and palate patients with acoustic analysis of speech," *The Journal of the Acoustical Society of America*, vol. 141, pp. 929–935, 02 2017.
- [10] A. Dubey, S. Prasanna, and S. Dandapat, "Hypernasality severity detection using constant q cepstral coefficients," 09 2019, pp. 4554–4558.
- [11] M. Saxon, J. Liss, and V. Berisha, "Objective measures of plosive nasalization in hypernasal speech," vol. 2019, 05 2019, pp. 6520–6524.
- [12] V. C. Mathad, N. Scherer, K. Chapman, J. Liss, and V. Berisha, "An attention model for hypernasality prediction in children with cleft palate," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, 2021, pp. 7248–7252.
- [13] X. Wang, S. Yang, M. Tang, H. Yin, H. Huang, and L. He, "Hypernasalitynet: Deep recurrent neural network for automatic hypernasality detection," *Int. J. Medical Informatics*, vol. 129, pp. 1–12, 2019.
- [14] V. C. Mathad, N. Scherer, K. Chapman, J. M. Liss, and V. Berisha, "A deep learning algorithm for objective assessment of hypernasality in children with cleft palate," *IEEE Trans. Biomed. Eng.*, vol. 68, no. 10, pp. 2986–2996, 2021.
- [15] X. Wang, M. Tang, S. Yang, H. Yin, H. Huang, and L. He, "Automatic hypernasality detection in cleft palate speech using cnn," *Circuits, Systems, and Signal Processing*, vol. 38, 08 2019.
- [16] M. Béchet, F. Hirsch, C. Fauth, and R. Sock, "Consonantal space area in children with a cleft palate: an acoustic study," in *Proc. Interspeech 2012*, 2012, pp. 58–61.
- [17] N. K., S. Kalita, C. Vikram, M. Pushpavathi, and S. M. Prasanna, "Hypernasality severity analysis in cleft lip and palate speech using vowel space area," in *Proc. Interspeech 2017*, 2017, pp. 1829–1833. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-1245>
- [18] J. H. Bradford, "Speech recognition by machine," *ACM SIGCHI Bull.*, vol. 22, no. 2, pp. 81–82, 1990.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*. IEEE, 2015, pp. 5206–5210.
- [20] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHELL-1: an open-source mandarin speech corpus and a speech recognition baseline," in *20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment, O-COCOSDA 2017, Seoul, South Korea, November 1-3, 2017*. IEEE, 2017, pp. 1–5.
- [21] M. Huh, P. Agrawal, and A. A. Efros, "What makes imagenet good for transfer learning?" *CoRR*, vol. abs/1608.08614, 2016.
- [22] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 2019, pp. 4171–4186.
- [23] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, 2014, pp. 3104–3112.
- [24] C. Wang, Y. Tang, X. Ma, A. Wu, D. Okhonko, and J. Pino, "Fairseq S2T: Fast speech-to-text modeling with fairseq," in *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: System Demonstrations*, dec 2020, pp. 33–39.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017, pp. 5998–6008.
- [26] D. S. Park, Y. Zhang, C. Chiu, Y. Chen, B. Li, W. Chan, Q. V. Le, and Y. Wu, "SpecAugment on large scale datasets," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*. IEEE, 2020, pp. 6879–6883.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [28] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*. ISCA, 2019, pp. 3465–3469.