

# Domain Specific Wav2vec 2.0 Fine-tuning For The SE&R 2022 Challenge

Alef Iury Siqueira Ferreira<sup>1</sup>, Gustavo dos Reis Oliveira<sup>1</sup>

<sup>1</sup>Federal University of Goiás, Goiânia, Brazil

## Abstract

This paper presents our efforts to build a robust ASR model for the shared task *Automatic Speech Recognition for spontaneous and prepared speech & Speech Emotion Recognition in Portuguese (SE&R 2022)*. The goal of the challenge is to advance the ASR research for the Portuguese language, considering prepared and spontaneous speech in different dialects. Our method consist on fine-tuning an ASR model in a domain-specific approach, applying gain normalization and selective noise insertion. The proposed method improved over the strong baseline provided on the test set in 3 of the 4 tracks available.

## Keywords

speech recognition, Portuguese, prepared speech, spontaneous speech, wild data

## 1. Introduction

The performance of Automatic Speech Recognition systems (ASRs) has increased significantly with the development of modern neural network topologies and the use of massive amount of data to train the models [1]. Although the accuracy of recent models improved for high-resource languages, such as English, the development of ASR models in other languages is still a difficult task using the same technologies [2, 3]. In this scenario, Self-Supervised Learning (SSL), a method in which representations with semantic information are learned by using unlabelled data, emerged as an important advance, allowing the training of deeper models using less labelled data [4, 5]. In this line of work, this paper explores the use of the Wav2vec 2.0 [6], a framework for self-supervised learning of discrete representations from raw audio data.

Wav2vec 2.0 (Figure 1) is inspired by previous works in unsupervised pre-training for speech recognition, that is, Wav2vec [7] and Vq-Wav2vec [4]. During pre-training the model learns speech representations solving a contrastive task which requires identifying the correct quantized latent speech representations of a masked time step among a set of distractors. After the self-supervised pre-training, the model can be fine-tuned on labeled data in a supervised task, like ASR, adding a randomly initialized linear projection on top of the context network with  $N$  classes and a loss function specific to the task at hand, like CTC, for instance.

The model shows important results among low re-

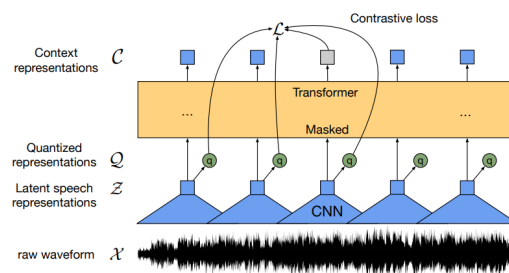


Figure 1: Illustration of the Wav2vec 2.0 framework [6].

source languages. In Portuguese, for example, [8] and [9] demonstrated that the fine-tuning of the Wav2vec 2.0 model achieves state-of-the-art (SOTA) results only using publicly available datasets.

An important aspect to consider when training an ASR model is the quality and the domain of the data [10, 11, 12]. While most of the available public datasets are composed of prepared speech [9], mostly read sentences [13, 14], the domain of real ASRs are far more complex, mainly because it is formed by spontaneous speech and different speech dialects. Quality is another issue: most of ASR use cases involve high noise environments or low recording equipment, which is not addressed in most of the public datasets available [9].

To stimulate research that can advance the present SOTA in ASR in Portuguese, for both prepared and spontaneous speech, the shared-task *Automatic Speech Recognition for spontaneous and prepared speech & Speech Emotion Recognition in Portuguese (SE&R 2022)* introduces a new baseline for ASR and a new dataset in Portuguese [9]. The Corpus of Annotated Audios (CORAA ASR), a large corpus of spontaneous and prepared speech, is composed by various subsets in Portuguese with differ-

*Proceedings of the First Workshop on Automatic Speech Recognition for Spontaneous and Prepared Speech & Speech Emotion Recognition in Portuguese (SE&R 2022), co-located with PROPOR 2022, March 21st, 2022 (Online).*

✉ alef\_iury\_c.c@discente.ufg.br (A. I. S. Ferreira);

gustavofreis2\_c.c@discente.ufg.br (G. d. R. Oliveira)

© 2022 Copyright for this paper by its authors. Use permitted under Creative

Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)



ent characteristics. The baseline achieves a Word Error Rate (WER) of 24.18% on CORAA ASR test set, a difficult dataset containing samples with low quality, noise, and a variety of domains and dialects.

In this work, we investigate the fine-tuning of the baseline model [9] proposed by the shared task, a fine-tuned model based on the Wav2vec 2.0 XLSR-53 [15], using only public available Portuguese datasets, including the CORAA ASR dataset. We conducted several experiments in different domains for the challenge and explored the use of selective noise insertion and audio normalization during training. This work is organized as follows: Section 2 discuss the proposed methods and Section 3 shows and discusses the obtained results. Finally, Section 4 presents the conclusions of this work.

## 2. Methods

### 2.1. Datasets

We used several publicly available datasets in Portuguese. Besides CORAA ASR, most of them is composed by prepared speech. In general, we opted to use all the data in the gathered datasets for training, except the dev part of the CORAA ASR, as presented in Table 1. The datasets used in this work are:

- CETUC [13]: contains approximately 145 hours of Brazilian Portuguese speech distributed among 50 male and 50 female speakers, each pronouncing approximately 1,000 phonetically balanced sentences selected from the CETEN-Folha<sup>1</sup> corpus;
- Common Voice (CV) 7.0 [16]: is a project proposed by Mozilla Foundation with the goal to create a wide open dataset in different languages. In this project, volunteers donate and validate speech using the official site<sup>2</sup>;
- Multilingual LibriSpeech (MLS) [14]: a massive dataset available in many languages. The MLS is based on audiobook recordings in public domain like LibriVox<sup>3</sup>. The dataset contains a total of 6k hours of transcribed data in many languages. The set in Portuguese used in this work<sup>4</sup> (mostly Brazilian variant) has approximately 284 hours of speech, obtained from 55 audiobooks read by 62 speakers;
- Multilingual TEDx [17]: a collection of audio recordings from TEDx talks in 8 source languages. The Portuguese set (mostly Brazilian Portuguese variant) contains 164 hours of transcribed speech.

<sup>1</sup><https://www.linguatca.pt/cetenfolha/>

<sup>2</sup><https://commonvoice.mozilla.org/pt>

<sup>3</sup><https://librivox.org/>

<sup>4</sup><http://www.openslr.org/94/>

- Corpus of Annotated Audios (CORAA ASR) v1 [9]: is a public available dataset that contains 290.77 hours of validated pairs (audio-transcription) in Portuguese (mostly Brazilian Portuguese variant) and is comprised by five other corpora: ALIP [18], C-ORAL Brasil I [19], NURC-Recife [20], SP2010 [21] and TEDx Portuguese talks.

### 2.2. Experiments

Our experiments consists on the fine-tuning of the baseline model of [9]. For each experiment, we trained the model for 5 epochs, using a batch size of 192, and using Adam [22] where the learning rate is  $3e^{-05}$  that is warmed up for the first 400 updates, then linearly decayed for the remained. For the experiments, we used a NVIDIA TESLA V100 32GB, a NVIDIA TESLA Tesla P100 16GB and a NVIDIA A100 80GB, depending on the type of audio pre-processing used. The code to replicate the results is available at [https://github.com/alefiury/SE-R\\_2022\\_Challenge\\_Wav2vec2](https://github.com/alefiury/SE-R_2022_Challenge_Wav2vec2).

In total, we conducted five main experiments to test our methods:

- Experiment 1: Wav2vec 2.0 XLSR-53 - Base: For this experiment, the model was fine-tuned considering the whole train set, but did not receive neither normalization nor noise addition;
- Experiment 2: Wav2vec 2.0 XLSR-53 - Norm: The model was fine-tuned with the whole train set with normalization. For the normalization, the mean gain of all the audios in the train set was considered;
- Experiment 3: Wav2vec 2.0 XLSR-53 - Norm and SNA: The model was fine-tuned with gain normalization and selective noise addition. The audios were normalized considering the mean gain of all the audios in the train set, and those audios pertaining to datasets that were considered to have a low presence of noise, namely MLS and CETUC, received randomly one of the following 5 possible types of noises: additive noise, being music or nonspeech noises from the MUSAN Corpus [23], Room impulse responses [24], Addition or reduction of gain, Pitch shift and Gaussian noise;
- Experiment 4: Wav2vec 2.0 XLSR-53 - Norm + Prepared Speech: Model fine-tuned based on the final trained model of Experiment 2, but considering just the prepared speech data from the CORAA ASR dataset, trained for more 5 epochs;
- Experiment 5: Wav2vec 2.0 XLSR-53 - Norm + Spontaneous Speech: Model fine-tuned based on the final trained model of Experiment 2, but considering just the spontaneous speech data

**Table 1**

Dataset splits used in this work.

Dataset	Subset	Type	Train	Dev (validation)	Test
Others	CETUC	Prepared Speech	144.65h	-	-
	Common Voice	Prepared Speech	112.08h	-	-
	MLS (Portuguese)	Prepared Speech	168.34h	-	-
	Multilingual TEDx (Portuguese)	Prepared Speech	152.17h	-	-
CORAA ASR[9]	ALIP	Spontaneous Speech	33.40h	0.99h	1.57h
	C-ORAL Brasil I	Spontaneous Speech	6.54h	1.13h	1.97h
	NURC-Recife	Spontaneous Speech	137.08h	1.29h	2.94h
	SP2010	Spontaneous Speech	27.83h	1.13h	2.18h
	TEDx Portuguese	Prepared Speech	68.67h	1.37h	2.70h
<b>Total</b>			850.76h	5.91h	11.36h

from the CORAA ASR dataset, trained for more 5 epochs.

### 3. Results and Discussion

The shared-task consists of 4 tracks. Each track has a domain specific scenario, that includes prepared speech and spontaneous speech. In this regard, we conducted a prior analysis (Section 3.1) using the dev set to select the best approaches based on 3 of the 4 tracks available: Mixed, Prepared Speech PT\_BR and Spontaneous Speech. The best models were selected and then submitted for evaluation. Our final results are presented in Section 3.2.

#### 3.1. dev set Analysis

Overall, our models did not show a huge improvement in performance when compared to the baseline model. Even though we fine-tuned a model that is considered the state-of-the-art in Brazilian Portuguese, we suspect that the number of training epochs might have been insufficient to obtain an increase in performance, or that the baseline model might have already reached a local optima.

Furthermore, as presented in Table 2, the model that was fine-tuned using prepared speech clearly improved the results in the Prepared Speech subset (and consequently the Mixed subset). The same phenomenon could not be seen in the Spontaneous Speech subset. A possible explanation to this fact is that most of the data of the datasets that were added to the train set are comprised of prepared speech, which might have contributed to the increase in performance in this particular domain. Another possible explanation is the low number of training epochs used to train the models.

Additionally, the noise insertion did not gave further improvement in performance. Nevertheless, the results of the SDA model in some more noisy subsets of the CORAA ASR dataset, like ALIP and NURC, for instance, showed some interesting and promising results when compared to the baseline. These results are shown in Table 3.

#### 3.2. Final Results

Table 4 compares the baseline with our selected models in the test set. The model Wav2vec 2.0 XLSR-53 - Norm + Prepared Speech surpassed the strong baseline in the Prepared Speech PT\_BR, Prepared Speech PT\_PT and the Mixed tracks. As seen in the results based on the dev set, the fact that most of the data of the datasets that were added to the train set are comprised of prepared speech, might have contributed to the increase in performance in this domain in both Portuguese variants. Lastly, even though we were not able to surpass the baseline model in the Spontaneous Speech track, we achieved competitive results with both submitted models.

#### 3.3. Additional Experiments

After selecting and submitting the best results, we performed some additional experiments to further explore our proposed methods using the dev set. We tried to use text correction in the outputs of the ASR models, and train the normalized models for a longer period of time using early stopping considering the prepared speech and spontaneous speech data from the CORAA ASR dataset.

The text correction was done with an additional post-processing step, using a KenLM [25] model. For the different tasks, we used 2 different KenLM models: one for spontaneous speech, which was built using subsets of the CORAA ASR dataset containing spontaneous speech phrases. And the other one was built considering wikipedia in portuguese texts, as proposed by [3]. Both were 4-grams. We found that this post-processing of the predictions from the ASR models did not improve the results from the dev set on the Prepared Speech PT\_BR and Spontaneous Speech tracks, as can be seen in Table 5, in fact they were worse. One possible explanation is that some of the decoder hyper-parameters did not work well with our ASR models. Another possibility is that the 4-gram trained with spontaneous text was built with a small amount of text, which might have decreased the performance of the model. However, the results on the Prepared Speech PT\_PT track were much better com-

**Table 2**  
Dev set analysis by subset

Model	Prepared Speech PT_BR		Prepared Speech PT_PT		Spontaneous Speech		Mixed	
	CER	WER	CER	WER	CER	WER	CER	WER
Wav2vec 2.0 XLSR-53 Baseline	4.53%	13.09%	17.73%	42.09%	<b>14.60%</b>	30.13%	12.87%	28.86%
Wav2vec 2.0 XLSR-53 - Base	5.65%	15.38%	16.82%	38.66%	15.13%	30.86%	13.18%	28.94%
Wav2vec 2.0 XLSR-53 - Norm	5.43%	14.85%	17.07%	38.06%	14.77%	<b>30.1%</b>	13.01%	28.28%
Wav2vec 2.0 XLSR-53 - Norm and SDA	5.26%	14.46%	17.38%	38.80%	14.77%	30.12%	13.04%	28.37%
Wav2vec 2.0 XLSR-53 - Norm + Prepared Speech	<b>4.38%</b>	<b>12.46%</b>	<b>15.63%</b>	<b>35.27%</b>	15.00%	30.83%	<b>12.50%</b>	<b>27.35%</b>
Wav2vec 2.0 XLSR-53 - Norm + Spontaneous Speech	4.71%	13.38%	16.30%	37.44%	14.93%	30.36%	12.72%	27.89%

**Table 3**  
Dev set analysis by dataset in CORAA ASR

Model	ALIP		NURC-Recife		C-ORAL-BRASIL I		SP2010	
	CER	WER	CER	WER	CER	WER	CER	WER
Wav2vec 2.0 XLSR-53 Baseline	15.08%	30.13%	14.02%	32.50%	20.09%	38.41%	<b>9.21%</b>	19.49%
Wav2vec 2.0 XLSR-53 - Base	15.42%	30.23%	14.27%	33.32%	21.23%	40.19%	9.58%	19.72%
Wav2vec 2.0 XLSR-53 - Norm	14.98%	29.68%	14.23%	32.42%	20.47%	38.98%	9.38%	19.36%
Wav2vec 2.0 XLSR-53 - Norm and SDA	<b>14.97%</b>	<b>29.24%</b>	<b>14.00%</b>	<b>32.26%</b>	20.68%	39.65%	9.42%	<b>19.31%</b>
Wav2vec 2.0 XLSR-53 - Norm + Prepared Speech	15.19%	30.36%	14.61%	33.78%	20.30%	39.30%	9.89%	19.88%
Wav2vec 2.0 XLSR-53 - Norm + Spontaneous Speech	15.31%	29.84%	14.23%	32.67%	<b>19.89%</b>	<b>37.93%</b>	10.32%	21.01%

pared to the previous experiments. This result suggests that the LM might improve results when there are few domain data used to train the Wav2vec model, since most of our training data was composed by Brazilian Portuguese audios.

Furthermore, as we had suspected earlier, the model with gain normalization that were trained considering the spontaneous speech data from the CORAA ASR corpus and for a longer period of time performed better in its respective subtrack, strengthening our hypothesis that our results did not improve in the main experiments due to a low number of training epochs.

## 4. Conclusions

In this work we presented our efforts to build a robust ASR model using multiple approaches, such as selective noise insertion and domain specific fine-tuning. In our experiments we found that fine-tuning a strong baseline

with additional public available data in multiple domains and using normalization, even for a few epochs, can improve performance. With our results we were able to improve on the test set in 3 of the 4 tracks available over the strong baseline provided.

As future works, we plan to train an ASR model using a dynamic noise insertion approach that do not depend on choosing specific datasets previously.

## Acknowledgments

This research was funded by CEIA with support by the Goiás State Foundation (FAPEG grant #201910267000527)<sup>5</sup>. We also would like to thank Cyberlabs Group<sup>6</sup> for the support for this work.

<sup>5</sup><http://centrodeia.org/>

<sup>6</sup><https://cyberlabs.ai/>

**Table 4**  
Test set analysis by subset

Model	Prepared Speech PT_BR		Prepared Speech PT_PT		Spontaneous Speech		Mixed	
	CER	WER	CER	WER	CER	WER	CER	WER
Wav2vec 2.0 XLSR-53 Baseline	3.56%	<b>11.19%</b>	17.08%	39.75%	<b>12.39%</b>	<b>26.24%</b>	11.35%	25.85%
Wav2vec 2.0 XLSR-53 - Norm	4.40%	12.71%	15.33%	35.31%	12.51%	26.50%	11.18%	25.25%
Wav2vec 2.0 XLSR-53 - Norm + Prepared Speech	<b>3.55%</b>	11.25%	<b>15.16%</b>	<b>34.68%</b>	12.61%	26.81%	<b>10.98%</b>	<b>24.89%</b>

**Table 5**

Additional Experiments. All the experiments used the Wav2vec 2.0 XLSR-53 pre-trained model for finetuning.

Model	Prepared Speech PT_BR		Prepared Speech PT_PT		Spontaneous Speech		Mixed	
	CER	WER	CER	WER	CER	WER	CER	WER
Norm + Prepared Speech + Prepared KenLM	5.59%	14.11%	<b>15.23%</b>	<b>32.36%</b>	16.26%	30.66%	13.33%	26.95%
Norm + Prepared Speech + Spontaneous KenLM	5.83%	15.58%	16.52%	36.74%	16.73%	33.78%	13.95%	29.97%
Norm + Spontaneous Speech + Prepared KenLM	5.80%	13.94%	15.80%	32.58%	16.38%	30.75%	13.59%	27.00%
Norm + Spontaneous Speech + Spontaneous KenLM	6.05%	15.52%	16.58%	36.98%	16.91%	33.45%	14.11%	29.85%
Norm + Prepared Speech + Early Stopping	<b>4.54%</b>	<b>12.96%</b>	15.73%	35.97%	14.75%	30.16%	<b>12.45%</b>	<b>27.31%</b>
Norm + Spontaneous Speech + Early Stopping	4.59%	13.14%	15.97%	36.52%	<b>14.66%</b>	<b>29.83%</b>	12.47%	27.33%

## References

- [1] J. Li, Recent advances in end-to-end automatic speech recognition, 2021. [arXiv:2111.01690](https://arxiv.org/abs/2111.01690).
- [2] 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: International conference on machine learning, PMLR, 2016, pp. 173–182.
- [3] I. M. Quintanilha, S. L. Netto, L. W. P. Biscainho, An open-source end-to-end asr system for brazilian portuguese using dnns built from newly assembled corpora, Journal of Communication and Information Systems 35 (2020) 230–242.
- [4] A. Baevski, S. Schneider, M. Auli, vq-wav2vec: Self-supervised learning of discrete speech representations, in: International Conference on Learning Representations (ICLR), 2020. URL: <https://openreview.net/pdf?id=rylwjxrYDS>.
- [5] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, F. Makedon, A survey on contrastive self-supervised learning, Technologies 9 (2021) 2.
- [6] A. Baevski, Y. Zhou, A. Mohamed, M. Auli, wav2vec 2.0: A framework for self-supervised learning of speech representations, in: H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, H. Lin (Eds.), Advances in Neural Information Processing Systems, volume 33, Curran Associates, Inc., 2020, pp. 12449–12460. URL: <https://proceedings.neurips.cc/paper/2020/file/92d1e1eb1cd6f9b3227870bb6d7f07-Paper.pdf>.
- [7] S. Schneider, A. Baevski, R. Collobert, M. Auli, wav2vec: Unsupervised pre-training for speech recognition., in: INTERSPEECH, 2019.
- [8] L. R. S. Gris, E. Casanova, F. S. de Oliveira, A. da Silva Soares, A. C. Junior, Brazilian portuguese speech recognition using wav2vec 2.0, 2021. [arXiv:2107.11414](https://arxiv.org/abs/2107.11414).
- [9] A. C. Junior, E. Casanova, A. Soares, F. S. de Oliveira, L. Oliveira, R. C. F. Junior, D. P. P. da Silva, F. G. Fayet, B. B. Carlotto, L. R. S. Gris, S. M. Aluísio, Coraa: a large corpus of spontaneous and prepared speech manually validated for speech recognition in brazilian portuguese, 2021. [arXiv:2110.15731](https://arxiv.org/abs/2110.15731).
- [10] M. L. Seltzer, D. Yu, Y. Wang, An investigation of deep neural networks for noise robust speech recognition, in: 2013 IEEE international conference on acoustics, speech and signal processing, IEEE, 2013, pp. 7398–7402.
- [11] T. Likhomanenko, Q. Xu, V. Pratap, P. Tomasello, J. Kahn, G. Avidov, R. Collobert, G. Synnaeve, Rethinking evaluation in asr: Are our models robust enough?, [arXiv preprint arXiv:2010.11745](https://arxiv.org/abs/2010.11745) (2020).
- [12] W.-N. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve, et al., Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training, [arXiv preprint arXiv:2104.01027](https://arxiv.org/abs/2104.01027) (2021).
- [13] V. Alencar, A. Alcain, Lsf and lpc-derived features for large vocabulary distributed continuous speech recognition in brazilian portuguese, in: 2008 42nd Asilomar conference on signals, systems and computers, IEEE, 2008, pp. 1237–1241.
- [14] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, R. Collobert, Mls: A large-scale multilingual dataset for speech research, Interspeech 2020 (2020). URL: <http://dx.doi.org/10.21437/Interspeech.2020-2826>. doi:10.21437/interspeech.2020-2826.
- [15] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, É. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised cross-lingual representation learning at scale, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 8440–8451.
- [16] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, G. Weber, Common voice: A massively-multilingual speech corpus, in: Proceedings of the 12th Language Resources and Evaluation Conference, 2020, pp. 4218–4222.
- [17] E. Salesky, M. Wiesner, J. Bremerman, R. Cattoni, M. Negri, M. Turchi, D. W. Oard, M. Post, The multilingual tedx corpus for speech recognition and

- translation, arXiv preprint arXiv:2102.01757 (2021).
- [18] S. C. L. Gonçalves, Projeto alip (amostra linguística do interior paulista) e banco de dados iboruna: 10 anos de contribuição com a descrição do português brasileiro, *Estudos Linguísticos* (São Paulo. 1978) 48 (2019) 276–297. URL: <https://revistas.gel.org.br/estudos-linguisticos/article/view/2430>. doi:10.21165/el.v48i1.2430.
  - [19] T. Raso, H. Mello, C-oral - brasil i: Corpus de referência do português brasileiro falado informal, 2012, pp. 362–367. doi:10.1007/978-3-642-28885-2\_40.
  - [20] M. Oliveira, Jr, Nurec digital um protocolo para a digitalização, anotação, arquivamento e disseminação do material do projeto da norma urbana linguística culta (nurec), *CHIMERA: Romance Corpora and Linguistic Studies* 3 (2016) 149–174.
  - [21] M. R B, O. L, Mapping paulistano portuguese: the sp2010 project, Firenze, Italy: Firenze University Press, 2012, pp. 459–463.
  - [22] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings, 2015. URL: <http://arxiv.org/abs/1412.6980>.
  - [23] D. Snyder, G. Chen, D. Povey, Musan: A music, speech, and noise corpus, 2015. arXiv:1510.08484.
  - [24] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, S. Khudanpur, A study on data augmentation of reverberant speech for robust speech recognition, in: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2017, pp. 5220–5224. doi:10.1109/ICASSP.2017.7953152.
  - [25] K. Heafield, KenLM: Faster and smaller language model queries, in: Proceedings of the Sixth Workshop on Statistical Machine Translation, Association for Computational Linguistics, Edinburgh, Scotland, 2011, pp. 187–197. URL: <https://aclanthology.org/W11-2123>.