



**HAL**  
open science

# Some voices are too common: Building fair speech recognition systems using the Common Voice dataset

Lucas Maison, Yannick Estève

## ► To cite this version:

Lucas Maison, Yannick Estève. Some voices are too common: Building fair speech recognition systems using the Common Voice dataset. Interspeech 2023, ISCA, Aug 2023, Dublin, Ireland. hal-04163615

**HAL Id: hal-04163615**

**<https://hal.science/hal-04163615>**

Submitted on 17 Jul 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

# Some voices are too common: Building fair speech recognition systems using the Common Voice dataset

Lucas Maison<sup>1,2</sup>, Yannick Estève<sup>1</sup>

<sup>1</sup>Laboratoire Informatique d’Avignon, Avignon, France

<sup>2</sup>Thales SIX, Multimedia Lab, Gennevilliers, France

lucas.maison@alumni.univ-avignon.fr

## Abstract

Automatic speech recognition (ASR) systems become increasingly efficient thanks to new advances in neural network training like self-supervised learning. However, they are known to be unfair toward certain groups, for instance, people speaking with an accent. In this work, we use the French Common Voice dataset to quantify the biases of a pre-trained wav2vec 2.0 model toward several demographic groups. By fine-tuning the pre-trained model on a variety of fixed-size, carefully crafted training sets, we demonstrate the importance of speaker diversity. We also run an in-depth analysis of the Common Voice corpus and identify important shortcomings that should be taken into account by users of this dataset.

**Index Terms:** automatic speech recognition, self-supervised learning, common voice, fine-tuning, biases

## 1. Introduction

Recent years have seen great advances in automatic speech processing, thanks to the generalization of new techniques for building speech representations using neural networks, including self-supervised learning (SSL) [1, 2, 3, 4], self-training (PT) [5], or a combination of both [6, 7, 8, 9]. SSL and PT are used to leverage huge amounts of unannotated data to build powerful speech representations that can be used to tackle various downstream tasks like automatic speech recognition (ASR) or speaker identification. To this end, models pre-trained with a SSL approach are then fine-tuned on a target domain. Although it has been shown that speech recognition models can be trained effectively in a fully-supervised manner [10, 11], to this day self-supervision remains popular for building powerful ASR systems [12].

There is a long history of bias in speech recognition [13, 14]. Despite their state-of-the-art performance, there is growing evidence that even the most recent models are not robust to domain shifts [15, 16] or data bias [17, 18]. Examples of domain shifts include acoustic conditions [19], vocabulary and grammar [16], and speech style (read speech, spontaneous speech) [15, 20], whereas models can be biased towards specific demographic attributes, like gender [18, 21], age [22], accent [18, 14, 23], or prosody (speech rate) [17].

Various methods for dealing with domain shifts have been referenced in the literature, like using data augmentation [24], adding adapters to the backbone model [22], exploiting both spectral and SSL features [25], or increasing the amount of pre-training data [15, 20].

In this work, we show how pre-trained models’ performance is biased toward certain demographics, and investigate whether fine-tuning on carefully designed training sets can counterbalance these biases. Our work is closely related to [21],

where they pre-train gender-specific models to investigate gender bias. In [15, 20], authors pre-train models using various datasets and data sizes but do not investigate the role of fine-tuning, nor do they break down results by demographics. Fairness in ASR has been investigated [26, 27], however, these works do not analyze the impact of biased fine-tuning sets. In [18], authors fine-tune SSL models on different domains and evaluate gender and accent biases, but they do not control for the size of the datasets. In this work, we investigate the impact of fine-tuning a SSL model on a single domain, controlling for data size and demographic biases. We also provide a critical analysis of the French Common Voice dataset. To the best of our knowledge, this is the first work investigating in-depth the impact of the fine-tuning step on the demographic biases of a wav2vec 2.0 model.

## 2. Common Voice dataset

Common Voice [28] is a massive crowd-sourced multilingual corpus of read speech. It is maintained by Mozilla and is freely available online<sup>1</sup>. We use the French subset of the Common Voice 12.0 corpus, which is the latest available version at the time of writing. It consists of more than 1,000 hours of speech ( $\approx 760,000$  utterances) spoken by more than 17,000 different speakers. Anyone can contribute to expanding the dataset by recording its voice or validating other users’ recordings. Interested readers can refer to the article [28] for more details on the dataset creation process.

Utterances are partitioned into three splits (validated, invalidated, and other) depending on their validation status. The validated split itself contains three disjoint subsets (train, dev, test), each with its unique speakers and sentences. Statistics for each split are reported in Table 1.

Table 1: *Statistics for the different splits of the dataset. Duration is hours:minutes.*

Split	# Speakers	# Utterances	Duration
validated	16140	666754	915:42
→ train	6756	499535	718:38
→ dev	3140	16104	25:46
→ test	4774	16104	26:14
invalidated	9365	58418	91:56
other	845	34255	47:25

We can see that more than 92,000 utterances are yet to be validated or have been deemed *invalid* by the community. While a lot of these rejects are justified (because of inaudible

<sup>1</sup>commonvoice.mozilla.org/datasets

audio, wrong sentence, multiple voices), manual inspections of the data show that some samples were wrongfully rejected. Conversely, some samples are deemed *valid* when they are not. We explore the influence of samples' quality on ASR performance in the following sections.

## 2.1. Demographic bias

One important feature of Common Voice is that users can self-report their demographic information on their profile; this information is then shipped alongside the dataset. Demographic data comprises age category, gender, and accent. We report the distributions of each of these attributes in Figure 1.

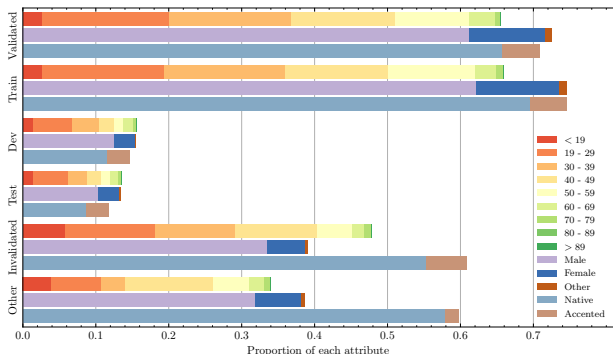


Figure 1: Distribution of each demographic attribute for the different data splits. Contributors may opt out of sharing demographic data; hence the proportions not summing to 1.

We can clearly see the data imbalances for each attribute. Taking the train set as an example, we observe that 83% of all labeled utterances (62% of the total) come from male speakers, whereas only 15% (11% of the total) come from female speakers. The discrepancy is even higher for the accent attribute, and we observe strong differences between age classes. Although these can be partially explained by the French population pyramid, the gender gap is more difficult to justify and could motivate further investigations.

We expect that these data biases will have repercussions on models fine-tuned on this trainset (or on random subsets of it). We test this hypothesis in the following sections.

## 2.2. Speaker bias

There are no limits on a user's contributions to the dataset ; hence, it is expected that some users will record more sentences than others. To get a better sense of whether this is the case, we plot the number of utterances by speaker (Figure 2, left). As we can see, the dev and test sets are balanced, with no more than 8 utterances per speaker. On the other hand, the train set is heavily unbalanced, with some speakers having several orders of magnitude more utterances than others. To better visualize this distribution, we plot the cumulative distribution of utterances on the right panel of Figure 2 (note the logarithmic axis). We observe that 7% of all utterances were spoken by a single user, while the top ten speakers represent 25% of all utterances. One-half and three-quarters of all utterances are spoken by only 1% and 10% of all the speakers, respectively. The training set is therefore clearly biased toward certain speakers. This could be a concern when taking sub-samples of this set, and indeed we find that random samples of size 50,000 contain 5,275 speakers

on average, a 22% decrease. We explore this problem further in the following sections.

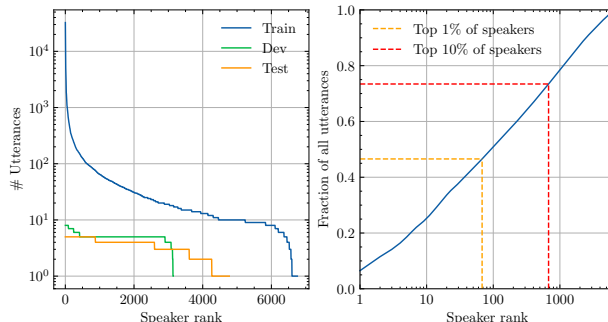


Figure 2: Left: Distribution of utterances between speakers for train/dev/test splits. Right: Cumulative distribution of utterances for the train split. Speaker rank is the rank in the list of speakers sorted by decreasing number of utterances.

## 3. Training sets

In order to answer the questions raised in section 2, we build various training sets using different strategies for sub-sampling the original train set. To ensure our comparisons are fair, each training set contains exactly 50,000 utterances ( $\approx 71 \pm 3$  hours of audio data). Choosing such a reduced data size has several benefits. It makes the training faster and consumes less energy. It also allows us to create fully biased train sets (for instance, 100% female speech, see section 3.1) even for low-resource demographic subgroups.

We begin by creating *reference* train sets, which are simply random sub-samples of the train split. These sets will provide us with baseline performance. Next, we create a *"bad quality"* set using invalidated samples, and two *"high quality"* sets containing only samples with no downvotes or samples with the most upvotes, respectively. These sets will be useful to determine the relevance of the samples' quality.

### 3.1. Demographic train sets

To measure the impact of demographic biases, we build a series of train sets whose characteristics are summarized in Table 2. For each of the three demographic attributes, we create sets voluntarily biased towards certain values of that attribute. Whenever possible, we try to affect the same number of speakers to the different train sets related to the same attribute (see for instance how "male only", "female only" and "mixed genders" all have 598 speakers). We do so in order to eliminate a possible confusion factor. Finally, we also create a *"high diversity"* set which is optimal in terms of demographic distribution: each attribute has its values equally represented in the train set.

### 3.2. Varying the number of speakers

We suspect that the number of speakers in a training set has a strong influence on the final recognition performance. To test this hypothesis independently of the data size, we create several sets of size 50,000 with  $2^2$ , 10, 100, 1,000, and  $6,756^3$  speakers.

<sup>2</sup>We use all the data from the top speaker (32,504) and complete with the second top speaker.

<sup>3</sup>The train split contains 6,756 unique speakers

Table 2: Description of the content of the demographic train sets, along with the number of speakers. †This train set contains repeated utterances to compensate for a lack of data.

Trainset	Description	Speakers
Reference	A random subsample of the train split	$\approx 5,275$
Female only	100% female	598
Male only	100% male	598
Mixed genders	50% female, 50% male	598
Youngs only	25% teens, 75% twenties	1,303
Middle-age only	50% thirties, 50% forties	1,280
Seniors only	60% fifties, 30% sixties, 10% seventies+	510
Mixed ages	10% seventies+, 15% for each other age	1,285
Native only	100% "French of France"	432
Accented only	100% other accents†	432
Mixed accents	50% native, 50% accented speech	432
High diversity	Equal distribution among attributes	1,079

We also want to test whether it is more important to include more samples or to diversify the speakers. To do so, we create two additional sets, *Small* and *Medium*, with 6,756 speakers each, and respectively one and three utterances per speaker.

## 4. Experiments

### 4.1. Data preparation

We keep the audio files as they are, simply converting them from 48 kHz MP3 to 16 kHz WAV format. We also filter out the few audio files longer than 12 s for efficiency purposes. Regarding the text labels, we do extensive cleaning in order to obtain the most faithful transcriptions possible. First, we convert symbols like  $\alpha$  and \$ to their written forms. Then we normalize non-Latin Unicode characters into their ASCII form while taking special care of keeping French-accented characters. Finally, we filter out all the remaining special characters and punctuation. While not perfect, we believe this procedure is a good heuristic for cleaning the transcriptions and validated it empirically by listening to corner cases.

### 4.2. Model

We use the Wav2Vec 2.0 model [2] for all our experiments ; specifically, we use the LB-7K-large model from the *LeBenchmark* [29] initiative, which was pre-trained on 7,739 hours of French audio. Hyperparameters and architecture are identical to the ones first introduced in [2]. Note that we use the *large* architecture variant, which presents greater capacity (317 million parameters). We use the *LB-7K* variant of the models since previous work [30, 29] has shown that for this task, models pre-trained using the greater quantity of audio performed best.

The pre-trained Wav2Vec 2.0 model acts as a speech encoder, which is fine-tuned for the ASR task together with an additional feed-forward network. This head network consists of three linear layers with 1,024 neurons. Each linear layer is followed by batch normalization and a leaky ReLU [31] activation function. We use dropout with  $p = 0.15$  between each linear layer. At last, a final linear layer projects the output into token space, and log-softmax is applied to obtain probabilities of each

token. We use 42 tokens, which represent individual characters.

### 4.3. Training

We use the SpeechBrain [32] toolkit for all our experiments. All our models are fine-tuned during 50 epochs using the CTC loss. We decided to freeze the convolutional layers of the transformer during fine-tuning since preliminary experiments taught us that this speeds up training by 15% with little to no impact on performances. Adam [33] and Adadelta [34] optimizers with learning rates  $10^{-4}$  and 1.0 are used to update the weights of the Wav2Vec 2.0 model and the additional top layers respectively. Learning rates are reduced at each epoch in which the validation loss does not improve.

During training, we apply on-the-fly data augmentation using the Speechbrain time-domain approximation of the SpecAugment [35] algorithm: it disrupts audio speed, and randomly drops chunks of audio and frequency bands. We disable audio speed modification since we find it to destabilize training while providing little or no performance improvement.

For fine-tuning we use several different training sets, which were formed by sampling audio data following various protocols (varying speaker diversity, demographics, speech quality, etc). We detail the formation of these training sets in section 3. We use the official dev set (validation set) for early stopping; this set is composed of 26 hours of audio, see Table 1.

Each model is trained on a single V100 GPU. Fine-tuning a model for 50 epochs on 70 h of audio takes  $\approx 44$  h of compute.

### 4.4. Evaluation

We evaluate our trained models on the official test set, which stays identical for all the experiments; this set is composed of 26 hours of audio, see Table 1. We use the Word Error Rate (WER) as our test metric; lower is better. Note that we do not use any language model besides our end-to-end model.

## 5. Results

### 5.1. Does the quality of the samples matter?

We evaluate our reference models on the test set and obtain a WER of  $13.7 \pm 0.1$ . We compare this result with our two "*high quality*" models, which score 13.7 and 13.9 respectively. We deduce that the number of upvotes or downvotes of validated samples barely matters. On the other hand, our "*bad quality*" model trained using invalidated samples scores 16.7, a significant drop in performance. However, it is quite noteworthy that a model fine-tuned on samples that were deemed unusable is still able to attain a reasonable ASR performance. This suggests that the invalidated split still contains a lot of usable audio, therefore the rejection process of Common Voice could be made less strict, which would contribute to increasing the dataset size.

### 5.2. Can we reduce bias toward certain demographic groups?

For each model trained on the different train sets described in Table 2, we evaluate its performance on the test set and break down results by demographic attributes (see Table 3). Surprisingly, we observe that models trained on biased train sets perform no better than the reference on the subgroup they are supposed to target. For instance, the model trained on the "*female only*" set scores 16.9 on the female test subset, whereas the reference model scores 16.3. This trend holds true for all the subgroups for our analysis.

Table 3: Results of models trained on demographic-biased datasets. WER on the full test set is reported in the column test. We also report the mean WER across demographic attributes. †In this row we report the mean WER across our reference models. The best results in each column are shown in bold.

Train set	test	female	male	youngs	middle-aged	seniors	native	accented	mean
reference †	<b>13.7</b>	16.3	<b>14.3</b>	16.1	13.7	<b>13.4</b>	<b>12.6</b>	17.4	<b>14.8</b>
female only	14.4	16.9	15.3	16.7	14.8	14.9	13.4	18.5	15.8
male only	14.1	17.7	14.7	17.1	14.5	<b>13.4</b>	13.0	19.1	15.6
mixed genders	14.0	17.0	14.6	16.4	14.4	13.5	13.0	18.6	15.4
youngs only	<b>13.7</b>	16.3	14.4	<b>16.0</b>	13.8	14.2	12.7	17.5	15.0
middle-aged only	13.9	17.3	14.7	16.9	14.5	13.5	12.9	19.0	15.5
seniors only	14.6	18.5	15.4	17.8	14.8	14.5	13.8	19.2	16.3
mixed ages	13.9	17.2	14.8	16.6	14.3	14.4	13.2	18.5	15.6
native only	14.3	17.4	15.1	17.5	13.9	14.2	12.8	<b>19.5</b>	15.8
accented only	14.9	17.8	15.2	17.1	14.5	15.0	14.0	17.6	15.9
mixed accents	14.2	17.3	14.9	16.5	14.5	14.5	13.4	18.0	15.6
high diversity set	13.8	<b>16.1</b>	<b>14.3</b>	<b>16.0</b>	<b>13.6</b>	13.9	12.8	<b>16.9</b>	<b>14.8</b>

Overall, there are persistent biases toward some demographic subgroups: male speech is better recognized than female speech, young people’s speech is badly recognized compared to the other age categories, and native speech is far better recognized than accented speech. We believe that these biases are rooted in the underlying pre-trained model; indeed, LB-7K-large [29] has been pre-trained on a dataset that contains almost no accented data and is heavily biased toward male speech<sup>4</sup>.

Finally, we observe that even our “high diversity” model does not help to decrease subgroups biases. Although it seems to perform better than the reference on some subgroups, we find that these improvements are not statistically significant.

### 5.3. Influence of the number of speakers

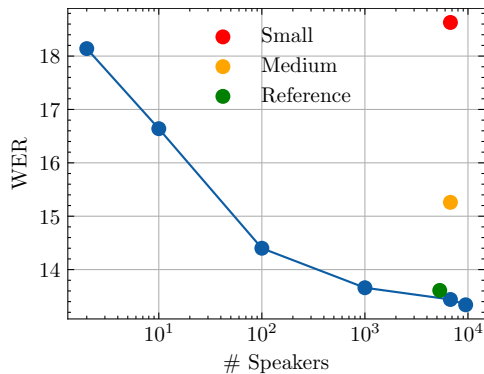


Figure 3: WER as a function of the number of speakers. Each point represents a different train set. Note that all train sets are of the same size (50,000 utterances) except Small (6,756) and Medium (19,918).

We evaluate the models fine-tuned on the train sets described in section 3.2 which have varying numbers of speakers. Results are shown in Figure 3, blue curve. We can see that the number of speakers (hence, the diversification of timbres, recording conditions, etc) in the training set has a clear impact on the WER. The higher the number of speakers, the lower the

<sup>4</sup>See [29], Table 1. Note that although data from VoxPopuli [36] is not gender-annotated, it comes from the European Parliament which used to consist of a majority of men.

WER. When fine-tuning the model on a dataset containing more speakers than the reference, it obtains a test WER of 13.4, beating the score of the reference.

In an effort to gather even more speakers, we split the dev set (90-10 split) and use the largest part to augment our training set, increasing the number of different speakers to 9,500. We obtain a WER of 13.3 with a model trained on this set, which is consistent with our previous findings. Moreover, models trained on these sets obtain lower WER scores than the reference for each demographic subgroup, see Table 4. We conclude that researchers looking for the fairest fine-tuning dataset should focus on maximizing the number of speakers rather than maximizing demographic diversity.

Finally, models trained on a large number of speakers (6,756) but on smaller train sets have performance on par with models trained on more data but with a small number of speakers (see Figure 3). This highlights the importance of prioritizing speaker diversity over dataset size when collecting audio data.

Table 4: Results of models trained on a larger number of speakers. We report the mean WER across demographic attributes in the column mean.

Trainset	Speakers	test	mean
Reference	5,324	13.7	14.8
All spks from train	6,756	13.4	14.4
train + dev	9,500	<b>13.3</b>	<b>14.3</b>

## 6. Conclusion

There is still a long way to go to obtain more fair speech recognition models. Initiatives like the Common Voice corpus are a first step toward the collection of big and diverse speech datasets. Even though Common Voice has important shortcomings, it is still useful to build efficient ASR models. We demonstrate how speaker diversity can be more important than demographic diversity, and therefore encourage researchers and the general public to give their voices in a collective effort to further expand the dataset. In future work, we plan to expand our analysis using biased pre-trained models and explore the influence of prosodic attributes (e.g. speech rate) on speech recognition. We also project to dive deeper into the differences between male and female speech and test for instance the impact of the pitch.

## 7. References

- [1] A. Mohamed *et al.*, “Self-supervised speech representation learning: A review,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1179–1210, 2022.
- [2] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 12 449–12 460.
- [3] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [4] S. Chen *et al.*, “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal on Selected Topics in Signal Processing*, vol. 16, pp. 1505–1518, 10 2021.
- [5] D. S. Park *et al.*, “Improved noisy student training for automatic speech recognition,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 2020-October. International Speech Communication Association, 5 2020, pp. 2817–2821.
- [6] Q. Xu *et al.*, “Self-training and pre-training are complementary for speech recognition,” in *ICASSP Proceedings*, vol. 2021-June. Institute of Electrical and Electronics Engineers Inc., 2021, pp. 3030–3034.
- [7] Y. Zhang *et al.*, “Pushing the limits of semi-supervised learning for automatic speech recognition,” in *Proceedings of the NeurIPS 2020 workshop Self-Supervised Learning for Speech and Audio Processing*, 10 2020.
- [8] Y.-A. Chung *et al.*, “w2v-BERT: Combining contrastive learning and masked language modeling for self-supervised speech pre-training,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 244–250.
- [9] Y. Zhang *et al.*, “BigSSL: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition,” *IEEE Journal on Selected Topics in Signal Processing*, vol. 16, pp. 1519–1532, 9 2021.
- [10] A. Gulati *et al.*, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Proc. Interspeech 2020*, 2020, pp. 5036–5040.
- [11] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” 2022. [Online]. Available: <https://arxiv.org/abs/2212.04356>
- [12] Y. Zhang *et al.*, “Google USM: Scaling automatic speech recognition beyond 100 languages,” 2023.
- [13] R. Tatman, “Gender and dialect bias in youtube’s automatic captions,” *EACL 2017 - Ethics in Natural Language Processing, Proceedings of the 1st ACL Workshop*, pp. 53–59, 2017.
- [14] A. Koenecke *et al.*, “Racial disparities in automated speech recognition,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 117, pp. 7684–7689, 4 2020.
- [15] W. N. Hsu *et al.*, “Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 3, 2021.
- [16] J. Zuluaga-Gomez *et al.*, “How does pre-trained wav2vec 2.0 perform on domain-shifted asr? an extensive benchmark on air traffic control communications,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*, 2023, pp. 205–212.
- [17] Y. Meng, Y.-H. Chou, A. T. Liu, and H.-y. Lee, “Don’t speak too fast: The impact of data bias on self-supervised speech models,” in *ICASSP*, 2022, pp. 3258–3262.
- [18] M. Riviere, J. Copet, and G. Synnaeve, “ASR4REAL: An extended benchmark for speech models,” 2021.
- [19] Q. S. Zhu, J. Zhang, Z. Q. Zhang, M. H. Wu, X. Fang, and L. R. Dai, “A noise-robust self-supervised pre-training model based speech representation learning for automatic speech recognition,” in *ICASSP*, vol. 2022-May. Institute of Electrical and Electronics Engineers Inc., 2022, pp. 3174–3178.
- [20] T. Likhomanenko *et al.*, “Rethinking evaluation in ASR: Are our models robust enough?” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*. International Speech Communication Association, 8 2021, pp. 311–315.
- [21] M. Z. Boito, L. Besacier, N. Tomashenko, and Y. Estève, “A study of gender impact in self-supervised models for speech-to-text systems,” in *Proceedings of INTERSPEECH*, vol. 2022-September. International Speech Communication Association, 4 2022, pp. 1278–1282.
- [22] R. Fan and A. Alwan, “DRAFT: A Novel Framework to Reduce Domain Shifting in Self-supervised Learning and Its Application to Children’s ASR,” in *Proc. Interspeech 2022*, 2022, pp. 4900–4904.
- [23] A. Hinsvark *et al.*, “Accented speech recognition: A survey,” 2021.
- [24] T. Fukuda *et al.*, “Data augmentation improves recognition of foreign accented speech,” in *Proceedings of INTERSPEECH*, vol. 2018-September. International Speech Communication Association, 2018, pp. 2409–2413.
- [25] D. Berrebbi, J. Shi, B. Yan, O. López-Francisco, J. Amith, and S. Watanabe, “Combining Spectral and Self-Supervised Features for Low Resource Speech Recognition and Translation,” in *Proc. Interspeech 2022*, 2022, pp. 3533–3537.
- [26] C. Liu *et al.*, “Towards measuring fairness in speech recognition: Casual conversations dataset transcriptions,” in *ICASSP*, 2022, pp. 6162–6166.
- [27] V. A. Trinh, P. Ghahremani, B. King, J. Droppo, A. Stolcke, and R. Maas, “Reducing Geographic Disparities in Automatic Speech Recognition via Elastic Weight Consolidation,” in *Proc. Interspeech 2022*, 2022, pp. 1298–1302.
- [28] R. Ardila *et al.*, “Common voice: A massively-multilingual speech corpus,” in *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, 2020, pp. 4211–4215.
- [29] S. Evain *et al.*, “Task agnostic and task specific self-supervised learning from speech with lebenchmark,” in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, J. Vanschoren and S. Yeung, Eds., vol. 1, 2021.
- [30] L. Maison, M. Zanon Boito, and Y. Estève, “Promises and Limitations of Self-supervised Learning for Automatic Speech Processing,” in *Conference on Artificial Intelligence for Defense*, ser. Proceedings of the 4th CAID. Rennes, France: DGA Maitrise de l’Information, Nov. 2022.
- [31] A. L. Maas, A. Y. Hannun, and A. Y. Ng, “Rectifier nonlinearities improve neural network acoustic models,” in *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- [32] M. Ravanelli *et al.*, “SpeechBrain: A general-purpose speech toolkit,” 2021, arXiv:2106.04624.
- [33] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015.
- [34] M. D. Zeiler, “ADADELTA: An adaptive learning rate method,” 2012.
- [35] D. S. Park *et al.*, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proceedings of the 20th Interspeech Conference*, vol. 2019-September. ISCA, 9 2019, pp. 2613–2617.
- [36] C. Wang *et al.*, “VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” in *Proceedings of the 59th Annual Meeting of the ACL and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Aug. 2021, pp. 993–1003.