



HAL
open science

Improving Accented Speech Recognition with Multi-Domain Training

Lucas Maison, Yannick Estève

► **To cite this version:**

Lucas Maison, Yannick Estève. Improving Accented Speech Recognition with Multi-Domain Training. ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, Jun 2023, Rhodes Island, Greece. pp.1-5, 10.1109/ICASSP49357.2023.10096268 . hal-04163554

HAL Id: hal-04163554

<https://hal.science/hal-04163554>

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

IMPROVING ACCENTED SPEECH RECOGNITION WITH MULTI-DOMAIN TRAINING

Lucas Maison*[†] Yannick Esteve*

* Laboratoire Informatique d’Avignon, Avignon, France

[†] Thales SIX, Multimedia Lab, Gennevilliers, France

ABSTRACT

Thanks to the rise of self-supervised learning, automatic speech recognition (ASR) systems now achieve near human performance on a wide variety of datasets. However, they still lack generalization capability and are not robust to domain shifts like accent variations. In this work, we use speech audio representing four different French accents to create fine-tuning datasets that improve the robustness of pre-trained ASR models. By incorporating various accents in the training set, we obtain both in-domain and out-of-domain improvements. Our numerical experiments show that we can reduce error rates by up to 25% (relative) on African and Belgian accents compared to single-domain training while keeping a good performance on standard French.

Index Terms— automatic speech recognition, multi-domain training, accented speech, self-supervised learning, domain shift

1. INTRODUCTION AND RELATED WORK

In recent years, self-supervised learning (SSL) has proven useful in a variety of domains, including computer vision, NLP, and speech processing. For the latter, SSL is used to leverage huge amounts of unannotated speech to build powerful speech representations that can be used to tackle downstream tasks like automatic speech recognition (ASR). Systems pre-trained with an SSL approach and then fine-tuned on a target domain reach state-of-the-art performance. However, there is growing evidence that these systems are not robust to domain shifts [1, 2]. For example, a model fine-tuned on speech uttered with accent A is likely to have reduced performance when evaluated on accent B, although the language is the same in both cases. It is known that adding out-of-domain data to the training dataset improves robustness to domain shifts [3]; however, it is often the case that out-of-domain data is scarce, making it hard to gather amounts sufficient for training. Regarding the peculiar case of accented speech, the vast majority of corpora consist of speech spoken with the dominant accent.

In order to tackle the lack of accented data, we can use data augmentation to artificially increase the amount of training data. There exist several ways to do this, for instance,

Corpus	Duration	# Utt.	# Speakers	Accent
CV [9]	56:11	46991	2897	None
AAF [10]	13:20	11344	228	African ¹
CaFE [11]	1:09	936	12	Quebec
CFPB [12]	4:07	6132	9	Belgian

Table 1. Statistics for the datasets (duration in hours)

it is possible to add noise to the training data, modify voice speed, or transform voice by manipulating the vocal-source and vocal-tract characteristics [4]. Other approaches include applying speaker normalization or anonymization methods in a reverse manner, for example using Vocal Tract Length Perturbation [5] or voice conversion using X-vectors [6].

There exist several methods for reducing the effects of domain shifting. It has been shown that adding new domains to the pre-training dataset improves the model’s robustness [3]. In order to avoid expensive retraining of large SSL models, Fan et al. introduce DRAFT [7], a method for learning domain-specific knowledge using residual adapters. Viglino et al. use multi-task learning and accents embeddings [8] to tackle accented speech recognition. In this work, we learn domain-specific knowledge during the fine-tuning stage by carefully designing multi-domain training sets.

2. EXPERIMENTAL CONTEXT

2.1. Datasets

We use four different datasets of French speech representing different accents. All the datasets are supplied with transcripts. Their main statistics are reported in Table 1.

CV (Common Voice [9]) is a large crowd-sourced multilingual corpus of read speech. We use the French subset of the CommonVoice 3 database. We use the official splits of the dataset - train: 31h, dev: 12h, test: 13h. This is our reference corpus for (accent-free) French speech.

AAF (African Accented French [10]) is a corpus of read speech. Speakers originate from five African countries (Cameroon, Chad, Congo, Gabon, and Niger) where French

¹While improper, we re-use the term *African accent* to refer to the mix of accents from Cameroon, Chad, Congo, Gabon, and Niger which compose the dataset.

Hours	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	95%	100%	FullCV
CV	0	0.89	2.00	3.42	5.32	7.97	11.96	18.6	30.87	30.87	30.87	30.87	56.19
AAF	7.97	7.97	7.97	7.97	7.97	7.97	7.97	7.97	7.97	3.99	1.99	0	7.97
Total	7.97	8.86	9.97	11.39	13.29	15.94	19.93	26.57	38.84	34.86	32.86	30.87	64.16
#Utt.	6780	7556	8516	9831	11481	13678	17068	22890	33580	30186	28553	26800	53771
#Spk.	140	151	166	188	210	248	300	382	574	508	474	434	3037

Table 2. Statistics for the training sets of our first experiment. Percentages indicate the proportion of CV data in the training set. *FullCV* is a dataset containing AAF together with all the splits of CV.

is (one of) the official language(s), however, their accent is clearly audible. We split this dataset as such - train: 8h, dev: 3h, test: 3h. This is our target corpus, that is, we want to obtain the best performance on this dataset.

CaFE (Canadian French Emotional [11]) is a small corpus of acted emotional speech. Speakers have a distinguishing *Québécois* accent. Due to the low amount of audio of this dataset, we do not split it and use it solely for testing.

CFPB (*Corpus de Français Parlé à Bruxelles* (Corpus of French as Spoken in Brussels) [12]) is a small corpus of interviews with Brussels speakers with a Belgian accent. We split this dataset as such - train: 3h, test: 1h.

2.2. Model

We use the following wav2vec 2.0 models from the *LeBenchmark* [13] initiative: LB-7K-base and LB-7K-large, which were pre-trained on 7,739 hours of French audio. The *base* variant refers to the standard model architecture from [14] that has 95 million parameters, while the *large* refers to their larger architecture that presents greater capacity (317 million parameters). We use the *LB-7K* variants of the models since previous work [15] has shown that for this task, models pre-trained using the greater quantity of audio performed best.

Each 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 768 or 1,024 neurons for a *base* or *large* model, respectively. Each linear layer is followed by batch normalization and a Leaky ReLU [16] 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.

2.3. Training

We use the *SpeechBrain* [17] toolkit for all our experiments. All our models are fine-tuned during 50 epochs using the CTC loss. Adam [18] and Adadelta [19] 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 [20] algorithm: it disrupts audio speed, and randomly drops chunks of audio and frequency bands.

For fine-tuning we use several different training sets, which are formed using varying amounts of audio data from one or more *speech domains* (accents). We detail the formation of these training sets in section 3. We also use a validation set (dev set) for early stopping; this set is composed of 5 hours of audio, evenly distributed between CV and AAF. We believe that using such a validation set favors the selection of a model with good performance in both domains. Note that we take care of separating by speaker when creating splits of the data; this way validation and testing are always done on unknown speakers.

2.4. Evaluation

We evaluate our trained models on four test sets, which stay identical for all the experiments: CV, AAF, CFPB (test splits), and CaFE (whole set). We use the Word Error Rate (WER) as our test metric; lower is better. Note that we do not use any language model to avoid introducing a bias during evaluation.

3. MULTI-DOMAIN DATA AUGMENTATION

3.1. Data Augmentation using McAdams coefficient

We use data augmentation to increase the amount of accented training data. Our approach consists of the alteration of the McAdams coefficient [21], and is described in [22] in the context of speaker anonymization. It is based upon simple signal processing techniques and is particularly relevant because it is simple to implement and does not require any additional training data. By applying the McAdams transformation to a speech signal with different McAdams coefficients α , we generate new signals with the same transcription but uttered by a pseudo voice with a different timbre.

We use values of α ranging from 0.7 to 1 with a step of 0.1 to generate the augmented dataset; note that using $\alpha = 1$ does not change the sample. We do not use values of $\alpha \leq 0.6$ because it deteriorates too much the intelligibility. Applying

this augmentation to AAF results in a dataset that is $4\times$ bigger than the original set, leading to a total of $\simeq 32$ hours of speech, matching the audio quantity of CV. We denote this augmented dataset **AAFaug**.

3.2. Multi-domain mix of datasets

We want to study the impact of multi-accent fine-tuning on the recognition of individual accents. To do so we design several experiments which consist of varying the number of domains and their associated quantity of speech data.

In our first experiment, we create train sets using various amounts of CV and AAF. Starting with a set containing only AAF, we keep increasing its size by including increasingly larger subsets of CV. In a reverse manner, we also create train sets by starting with CV and progressively adding AAF. We denote these sets **CV - $x\%$** , where x represents the proportion of CV data in the set. In addition, we create a *FullCV* train set consisting of the whole CV dataset (train, dev, test) together with the AAF train split. See Table 2 for details on the individual training sets.

In a second experiment, we select the best train sets from experiment 1 (i.e. the sets which led to the lowest WER) and combine them with either CaFE (whole) or CFPB (train split). This leads to slightly larger train sets which have three accents instead of two.

In our third and last experiment, we repeat experiment 1 using a fixed number of hours in the training set. To do so, we make use of the augmented dataset AAFaug 3.1. Similarly to experiment 1, we define 11 train sets, starting with a 31h subset of AAFaug, and gradually replacing more and more AAFaug data by CV data with a 10% proportion increment, ending with the full CV train set. Finally, we create two larger training sets using the augmented data, namely $CV \cup \text{AAFaug}$ and $\text{FullCV} \cup \text{AAFaug}$.

4. RESULTS AND DISCUSSION

In this section, we summarize and analyze the results of our experiments. Numerical results are reported in Table 3; for the sake of brevity, we only list the best training sets. All models use the *Large* architecture unless specified otherwise.

4.1. Experiment 1: Mixing two domains

We study the impact of mixing two domains during training on performance. We observe U-shaped curves of WER with respect to CV proportion in the training set, meaning that WER first decreases when adding CV to the mix (from 0 to 80%), then slightly increases again when removing AAF (from 80 to 100%). Going from one to two domains causes both in-domain and out-of-domain improvements. When adding a new domain B to a train set containing only domain A, not only do we improve on domain B as expected, but we




Fig. 1. A representation of the performance trade-off between CV and AAF. Each point represents a model fine-tuned on a different train set. All models are *Large* unless specified otherwise. Details on the experiments can be found in section 3.2.

also improve on domain A and on unseen domains C and D as well. We obtain our best results on CV (9.4 WER) and on CaFE (36.8 WER) using CV 90% and CV 50% respectively, while also reaching competitive performance on the other test sets compared to the baselines. Remarkably, training on a CV/AAF mixed set allows us to reach WER of 34.6 and 36.8 on CFPB and CaFE respectively, which is a lot better than the scores reached when training on CV or AAF alone, see table 3.

We observe similar trends on base models although with significantly higher WER (+40% on average relative to large models). However, we see lesser improvements from the domain augmentation compared to large models. We believe the base model is less able to take advantage of multi-domain due to its reduced capacity.

It should be noted that the *LB-7K* models were pre-trained using several different corpora, including AAF and CaFE (see [13]). However, this does not prevent the catastrophic forgetting phenomenon which happens during fine-tuning. For instance, the *LB-7K-large* model fine-tuned on CV scores 18.7 on AAF, despite AAF being included in the pre-training data. This highlights the importance of mixing domains during fine-tuning.

4.2. Experiment 2: Mixing three domains

We select the five best train sets from experiment 1, which are CV 0% (i.e. AAF only), CV 80% (i.e. $CV \cup \text{AAFaug}$), CV 90%, CV 100% (i.e. CV only) and *FullCV*. When adding CaFE to those sets, we observe no significant changes on CV or AAF,

as we can see in figure 1. We observe the same trend with CFPB; however, its addition results in a massive boost on the CFPB test set (-30% WER). The best performance on CFPB is reached on the train set CV 90% \cup CFPB, with a WER of 24.8 and very good scores on CV and AAF.

Moreover, the best average WER on all four test sets is reached on the train set CV 80% \cup CFPB, see Table 3, last row. This shows the benefits of training on multiple domains, as we can leverage speech from other domains to improve on each domain with respect to the single-domain baseline.

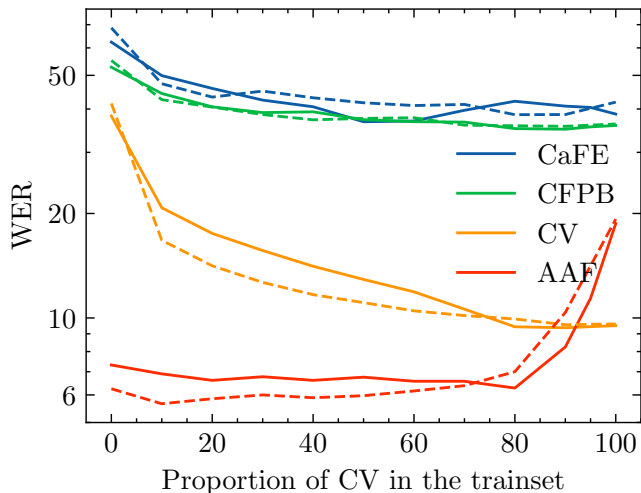


Fig. 2. Graph of the WER with respect to the CV proportion. Plain and dashed lines correspond to experiments 1 and 3 respectively (without or with augmentation).

4.3. Experiment 3: using data augmentation

We repeat experiment 1 using an augmented AAF dataset. Results are shown in figure 2. We clearly see the advantage of using the augmented train set: performance improves on CV and AAF while staying nearly identical on CaFE and CFPB. This improvement is also visible in figure 1, bottom-left corner. In addition, this figure highlights the existing trade-off between error rates on different domains. When training using a fixed number of hours, an improvement on one domain usually comes at the expense of a (possibly slight) degradation in another domain.

The best performance on AAF among these 31h sets is reached using CV 10% (which comprises 28h of African speech) with a WER of 5.7, but its performance is poor on CV (16.7 WER). However, if we train on CV \cup AAFaug (63h of speech), we obtain a WER of 5.3 and 9.7 on AAF and CV respectively (see “top result” in figure 2), which is our best score on African accent.

These results on augmentation are encouraging: thanks to it, we are able to significantly reduce the WER on African

Train set	CV	AAF	CFPB	CaFE
CV	9.5	18.71	35.86	38.68
AAF	38.2	7.32	52.78	62.34
CFPB	45.1	44.28	32.59	67.41
CV 90%	9.37	8.25	34.98	40.77
CV 100% \cup AAFaug	9.69	5.29	37.2	44.69
CV 90% \cup CFPB	9.42	8.32	24.79	41.77
CV 50%	12.91	6.75	37.27	36.76
CV 80% \cup CFPB	9.54	6.42	25.21	41.09

Table 3. Results on the different test sets. Each line corresponds to a different training set. Rows contain, in that order: single-domain train sets; best train sets for each test set; train set with the best average score. The best results in each column are shown in bold

accent (-14% WER, relative) compared to the best result without augmentation. However, the augmentation method we describe (which is based on the variation of the McAdams coefficient) may not be particularly relevant to the task. Indeed, we ran an ablation study, training using only on-the-fly SpecAugment on the same quantity of data, and obtained nearly the same results. Further investigation is needed to assess more precisely the impact of the different augmentation methods on ASR performance, and the relevance of McAdams pseudo-voices.

Finally, we discuss the impact of using a substantially larger CV domain for training, forming the *FullCV* family of train sets, which are the largest ones. Models trained on these sets tend to perform better on average on CaFE (-6.1%) and CFPB (-1.4%) compared to those trained on the second biggest sets, but also show a performance degradation on AAF (+1.7%). Besides, no model trained on *FullCV* sets has reached top performance on any of the test sets. This is likely due to the discrepancies between the accents in the training data, and indicates that blindly maximizing the amount of training data may not always be the best choice for minimizing WER, neither across domains nor for a single domain.

5. CONCLUSION

Recognition of accented speech remains a challenging task. In this work, we showed the positive influence of using multi-domain fine-tuning datasets both on in-domain and out-of-domain accents. We experimented with various ways of combining training sets and achieved remarkable gains on our evaluation accents compared to single-domain baselines. We leave as future work the collection and evaluation of new accents, as well as the experimentation of new methods for generating accented speech for low-resource accents.

6. REFERENCES

- [1] Wei Ning 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*, 2021, vol. 3.
- [2] Juan 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.
- [3] Tatiana Likhomanenko, Qiantong Xu, Vineel Pratap, Paden Tomasello, Jacob Kahn, Gilad Avidov, Ronan Collobert, and Gabriel Synnaeve, “Rethinking Evaluation in ASR: Are Our Models Robust Enough?,” in *Proc. Interspeech 2021*, 2021, pp. 311–315.
- [4] Takashi Fukuda, Raul Fernandez, Andrew Rosenberg, Samuel Thomas, Bhuvana Ramabhadran, Alexander Sorin, and Gakuto Kurata, “Data augmentation improves recognition of foreign accented speech,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*. 2018, vol. 2018-September, pp. 2409–2413, ISCA.
- [5] Navdeep Jaitly and Geoffrey E Hinton, “Vocal tract length perturbation (vtpl) improves speech recognition,” *Proceedings of the 30th International Conference on Machine Learning*, vol. 90, 2013.
- [6] Fuming Fang, Xin Wang, Junichi Yamagishi, I. Echizen, Massimiliano Todisco, Nicholas Evans, and Jean-François Bonastre, “Speaker anonymization using x-vector and neural waveform models,” in *Proceedings of the 10th ISCA Speech Synthesis Workshop*, 09 2019, pp. 155–160.
- [7] Ruchao Fan and Abeer Alwan, “DRAFT: A novel framework to reduce domain shifting in self-supervised learning and its application to children’s ASR,” in *Proceedings of the 23th Interspeech Conference*, 09 2022, pp. 4900–4904.
- [8] Thibault Viglino, Petr Motlicek, and Milos Cernak, “End-to-End Accented Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 2140–2144.
- [9] Rosana 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.
- [10] “African accented french,” <https://www.openslr.org/57/>, Accessed : 14-10-2022.
- [11] Philippe Gournay, Olivier Lahaie, and R. Lefebvre, “A canadian french emotional speech dataset,” in *Proceedings of the 9th ACM Multimedia Systems Conference*, 06 2018, pp. 399–402.
- [12] Anne Dister and Emmanuelle Labeau, “Le corpus de français parlé à bruxelles : origines, hypothèses, développements et prédictions,” *Cahiers AFLS*, 2015.
- [13] Solène Evain *et al.*, “Task agnostic and task specific self-supervised learning from speech with *LeBenchmark*,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [14] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [15] Lucas Maison, Marcelly Zanon Boito, and Yannick Estève, “Promises and Limitations of Self-supervised Learning for Automatic Speech Processing,” in *Proceedings of the 4th Conference on Artificial Intelligence for Defense (CAID)*, Rennes, France, Nov. 2022, DGA Maîtrise de l’Information.
- [16] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng, “Rectifier nonlinearities improve neural network acoustic models,” in *Proceedings of the 30th International Conference on Machine Learning*, 2013.
- [17] Mirco Ravanelli *et al.*, “SpeechBrain: A general-purpose speech toolkit,” 2021, arXiv:2106.04624.
- [18] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *ICLR 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun, Eds., 2015.
- [19] Matthew D. Zeiler, “Adadelata: An adaptive learning rate method,” *ArXiv*, vol. abs/1212.5701, 2012.
- [20] Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proceedings of the 20th Interspeech Conference*. 9 2019, vol. 2019-September, pp. 2613–2617, ISCA.
- [21] Stephen Edward McAdams, *Spectral fusion, spectral parsing and the formation of auditory images*, Stanford university, 1984.
- [22] Jose Patino, Natalia Tomashenko, Massimiliano Todisco, Andreas Nautsch, and Nicholas Evans, “Speaker anonymisation using the mcadams coefficient,” *arXiv preprint arXiv:2011.01130*, 2020.