COMPARISON OF SELF-SUPERVISED SPEECH PRE-TRAINING METHODS ON FLEMISH DUTCH

Jakob Poncelet, Hugo Van hamme

KU Leuven

Department Electrical Engineering ESAT-PSI
Kasteelpark Arenberg 10, Bus 2441, B-3001 Leuven Belgium
{*jakob.poncelet, hugo.vanhamme*}@esat.kuleuven.be

ABSTRACT

Recent research in speech processing exhibits a growing interest in unsupervised and self-supervised representation learning from unlabelled data to alleviate the need for large amounts of annotated data. We investigate several popular pre-training methods and apply them to Flemish Dutch. We compare off-the-shelf English pre-trained models to models trained on an increasing amount of Flemish data. We find that the most important factors for positive transfer to downstream speech recognition tasks include a substantial amount of data and a matching pre-training domain. Ideally, we also finetune on an annotated subset in the target language. All pre-trained models improve linear phone separability in Flemish, but not all methods improve Automatic Speech Recognition. We experience superior performance with wav2vec 2.0 and we obtain a 30% WER improvement by finetuning the multilingually pre-trained XLSR-53 model on Flemish Dutch, after integration into an HMM-DNN acoustic model.

Index Terms— speech recognition, self-supervised learning, pre-training, cross-lingual

1. INTRODUCTION

Building a good speech recogniser typically requires a large amount of annotated data from a specific language. Obtaining high-quality labelled data is a costly and time-intensive process, and for many languages this remains a big issue. However, even in a highly resourced language like English, recent work has shown impressive results in Automatic Speech Recognition (ASR) by pre-training on unlabelled data and transferring that knowledge to regular speech recognition models [1, 2, 3] or even completely unsupervised speech recognition [4]. This paradigm shift towards unlabelled data is of great significance as untranscribed recordings of speech are much easier to acquire.

Self-supervised learning is a clever way to learn general information from data without requiring any labels. Recently many successful methods have emerged for self-supervised representation learning from speech. The general idea is to

implicitly learn the global structure and local characteristics that are inherently present in speech. Depending on the task, both local information, such as the pronunciation of a specific phoneme, and more global information, such as speaker traits and recording properties, can be useful. By pre-training a network with a well-chosen objective function, these relevant attributes about the input speech can be captured and summarised in rich feature vectors. This improves several downstream tasks like speech recognition and typically reduces the amount of required data, since the principal characteristics are already extracted and more easily accessible. Moreover, the structure of a speech waveform is to some extent general and language-independent, which explains the improvements with these features in low-resource languages [5, 6].

The objective function used in self-supervised learning techniques is the driving force behind the extraction of powerful speech representations. In fact, the self-supervised objective has more impact on the learned representation than architectural differences between methods [7]. In Autoregressive Predictive Coding (APC) [8, 9, 10, 11], the objective is to predict a frame a few steps ahead, given the information up to that point. Another branch of research focuses on predicting the current frame given past and future context, by reconstructing several masked frames [12, 13], similar to the Masked Language Modeling (MLM) approach in Natural Language Processing [14]. Finally, Contrastive Predictive Coding (CPC) [15] is a popular technique in representation learning, where the objective is to predict the future in the latent space and a contrastive loss is applied to maximise mutual information. CPC has been successfully applied to speech recognition [1, 16, 17, 18] and has shown to be able to learn robust and cross-lingual speech representations [5, 6, 19]. We refer to the literature for other related work in self-supervised and unsupervised representation learning [20, 21, 22, 23, 24, 25, 26, 27].

Following the widespread improvements in ASR as a result of self-supervised pre-training, this paper will focus on Flemish Dutch, a medium-resourced language. Flemish is the language spoken in Flanders, the Dutch-speaking part of Bel-

gium. It is closely related to the Dutch variant spoken in The Netherlands, but there are still many noticeable differences [28]. A few seconds of speech suffice to distinguish the two variants. Although geographically relatively small, Flemish Dutch is diverse with several dialects, roughly corresponding to the five provinces in Flanders, though natives will observe even finer detail. Furthermore, Dutch belongs to the family of West-Germanic languages, like English and German, which makes it a very interesting language to examine whether pre-training on English leads to strong improvements in Dutch. While there is some overlap in phones, there are also several vowels and diphthongs that do not occur in English.

In this work, we compare several popular self-supervised pre-training methods when applied to Flemish. First of all, we look at the applicability of off-the-shelf models that are pre-trained on English and assess the transferability to Flemish. This would be convenient for several research domains and technological applications. Additionally it would eliminate the need for large computational resources necessary to pre-train these models, which scales with the model size (e.g. the high-capacity wav2vec 2.0 model [1]) and for many models also with the amount of data. Second, we examine the importance of matching the pre-training language to the target language as opposed to the amount of data used in pre-training. To this end, we compare pre-trained models in English and Netherlands Dutch to models trained on Flemish Dutch. Recent work [5] has shown that low-resource languages can greatly benefit from higher-resource languages when they are more similar due to positive transfer, but cross-lingual representation learning degrades the performance on high-resource languages due to interference. Furthermore, self-supervised pre-training has shown to improve robustness and reduce the degradation on out-of-domain data [29, 30]. We show that simply augmenting the finetuning data leads to a strong speech recognition improvement in noisy and reverberated environments. Finally, we investigate ASR improvements with the recent wav2vec 2.0 model [1, 5] and study several pre-training and finetuning scenarios with an increasing amount of data, yielding substantial reduction of Word Error Rates (WER) compared to the baselines.

We evaluate the models in terms of linear phone separability by reporting the classification accuracy of an external linear classifier. The classifier is trained to predict Flemish Dutch phones from the features extracted from each model. For the ASR experiments, we report the results of an HMM-DNN hybrid model [31] where the DNN is trained with the learned features.

2. MODELS

The procedure consists of three separate phases: 1) pre-training a model on data without labels, 2) optionally finetuning the model on a labelled set with transcripts, 3) extracting the learned features to perform a downstream evaluation task.

2.1. Self-Supervised Pre-training

We start with a short description of the investigated pre-training techniques and refer to the corresponding papers for more details. Table 1 gives an overview of all models.

2.1.1. APC

In APC, autoregressive models encode the temporal information in the past sequence of frames, for example with Gated Recurrent Units (GRU). A future frame, n steps ahead of the current frame, is linearly predicted from the autoregressive outputs. The model is then trained with an L1 reconstruction loss on the predicted frame. We use a model with 3 GRU layers and predict 5 steps ahead [8, 10]. The outputs of the last GRU layer are extracted as features for the downstream task.

2.1.2. Mockingjay

While APC conditions its prediction on past context only, Mockingjay leverages both past and future context to predict a frame that has been masked out. The encoder is a deep bidirectional Transformer [32] that learns contextualised representations, which are extracted from the last layer. These representations are linearly mapped to predict the masked frames, and the model is trained with a reconstruction loss between the predicted and true frames. We use the base model with 3 Transformer blocks in the encoder [12, 33].

2.1.3. CPC

CPC directly applies a stack of strided convolutional layers to the raw waveform to encode the sequence in the latent space. An autoregressive model (the aggregator) then looks at the representations of the past sequence and its output is mapped to predict the latent representations for several steps in the future. The loss is not reconstructive, but contrastive: given the aggregator output, the model has to distinguish the correct sample out of a bunch with distractors from windows more distant in time or from different sequences. We use the modified CPC approach [6] where the encoder consists of 5 CNN layers, the autoregressive model is an LSTM and the prediction network is a 1-layer Transformer network. The model predicts 1 to 12 steps in the future, with a separate projection layer for every step, and is trained with 10 distractors. The outputs of the autoregressive model are the extracted features.

2.1.4. wav2vec

Wav2vec is built on CPC but uses a fully convolutional model. The autoregressive model is replaced by a context network consisting of 12 convolutional layers. Two additional linear transformations increase the capacity of the encoder (this architecture is called *wav2vec large* in the corresponding paper [16]). The outputs of the context network are the feature vectors [34].

Table 1: Shallow overview and comparison of all pre-training techniques.

Model	Feature encoder	Aggregator	Objective	Output dimension	# Parameters
APC	Filterbank	GRU	Reconstruct future frame	512	4.1M
Mockingjay	Filterbank	Bidirectional Transformer	Reconstruct masked frame	768	21.3M
CPC	CNN	LSTM	Identify future feature	256	1.8M
wav2vec	CNN	CNN	Identify future feature	512	32.5M
wav2vec 2.0	CNN	Transformer	Identify quantised future feature	768 (base), 1024 (large)	95.0M (base), 317.3M (large)

2.1.5. wav2vec 2.0

Wav2vec 2.0 combines ideas from wav2vec [16], vq-wav2vec [17] and MLM. The encoder computes latent speech representations from the raw waveform with 7 temporal convolution blocks. A certain proportion of the latent features is masked before feeding to the aggregator, which is a Transformer network. At the same time, a quantisation module maps the latent feature vectors to discretised versions. The final training objective is then to distinguish the true quantised representation for a masked time step, given the aggregator output [1]. We differentiate between the base and large architecture of the model, which contain respectively 12 and 24 Transformer blocks in the aggregator. The contextual features at the output of the aggregator are extracted for downstream tasks [34]. We duplicate them in time to mimic a stride of 10ms instead of 20ms.

The wav2vec 2.0 model can be finetuned on a labelled set. To this end, an extra linear layer is added on top of the context network and a CTC loss is applied with the transcription characters as targets. The encoder is frozen during finetuning. Finetuning is done after the pre-training is completed.

Finally, XLSR-53 is a large wav2vec 2.0 model pre-trained on 53 languages simultaneously [5]. The authors have shown that the quantised speech representations can express connections between languages when trained in a multilingual setup.

Due to limited resources, we pre-train wav2vec 2.0 base models for 100k updates and finetune for 500k updates, and we don't pre-train our own wav2vec 2.0 large models but only finetune existing pre-trained models.

2.2. Downstream Feature Evaluation

2.2.1. Phone Classification

We train an external phone classifier consisting of just one linear layer and a softmax layer [33], with as input the features extracted from the pre-trained models. All pre-trained features are compared to the baseline of 80-dimensional log-mel filterbank features, including second order delta features and mean-variance normalisation. For every utterance, there

is a phone label every 10ms, corresponding to the stride of the input features. The classifier is trained with a cross-entropy loss. We report the accuracy of the classifier of predicting the correct phone label for every 10ms window, instead of using the most voted phone during its entire duration, because the learned representations should contain phonetic information even at the start of a phone.

For English experiments we use the phone labels from [15], which have been generated by forced alignment with Kaldi [31] using pre-trained models on LibriSpeech, and mapped to 41 classes. For Flemish experiments we use the phone labels provided in the Corpus Gesproken Nederlands (Section 3.1.1). The phone sequences have been computed by forced alignment on the manually checked orthographic transcripts with SPRAAK [35], and have been partly manually checked as well. There are 49 distinct phone classes [36].

2.2.2. ASR

We train a baseline HMM-DNN model with Kaldi [31] on MFCC features. The HMM-GMM models triphones and includes LDA, MLLT and fMLLR transformations. It is trained on MFCC features to compute alignments and build a phonetic tree with one state per phone. For the pre-trained models, we reuse the alignments and tree from the MFCC model and only train the DNN model with the extracted features as input. We make a distinction between a large DNN model containing 14 TDNN-F layers [37] (similar to the Switchboard recipe) and a small DNN model with only 3 TDNN-F layers. We leave out iVector extraction and speed perturbation, and remove the delta layers for pre-trained features. We decode with a pruned trigram language model and use a lexicon of 100k words. We report Word Error Rates based on the Levenshtein distance, but make a correction for inconsistencies in compounding (which occur frequently in Dutch).

3. DATA

3.1. Flemish Dutch datasets

3.1.1. Labelled data

Corpus Gesproken Nederlands (CGN) [36] - also called Spoken Dutch Corpus - is a manually annotated speech database of around 900 hours of Dutch, of which 270 hours correspond to Flemish Dutch. CGN contains both phonetical and word-level transcriptions and segmentations. The labelled data can be used for finetuning, for ASR model training and for the proposed evaluation procedures. We make the distinction between three training sets of data, based on the type of speech.

VL-train-clean This set contains 35h of prepared, read speech by professional readers. This corresponds to component O of CGN.

VL-train-other This set contains several types of speech, including read speech (*VL-train-clean*), news reports, interviews, lectures, sports commentary, etc. This set holds 145h of data from components B,F,G,H,I,J,K,L,M,N,O of CGN.

VL-train-all This set contains all components from the CGN database and corresponds to 270 hours of speech. The difference with *VL-train-other* is the inclusion of narrowband telephone speech (8kHz resampled to 16kHz) and spontaneous conversational speech, which correspond to respectively components C,D and component A of CGN.

In a similar way, we make a distinction between **VL-test-clean** (4h) and **VL-test-other** (15h, including the 4h from *VL-test-clean*). There is no overlap in speakers with the train sets.

For phone classification experiments in English, we use the *train-clean-100* set of LibriSpeech [38] and use the train-test split and phone labels from [15].

3.1.2. Unlabelled data

We have created a dataset of 450h of unlabelled data for unsupervised experiments in Flemish Dutch, by extracting audio from online available resources. We refer to this set as **VL-unsup**. This set consists of 200h of data from recordings in the Flemish parliament, 100h of audio from broadcast TV news and 150h of audio from TV talkshows. For pre-training, we use this set and the labelled sets without the transcriptions.

3.2. Pre-trained models

For some experiments, we use off-the-shelf available pre-trained models for APC, Mockingjay, CPC, wav2vec and wav2vec 2.0 [33, 34]. These models have been pre-trained on English audiobooks from LibriSpeech (*LS-960*) [38], LibriLight (*LL-60k*) [39] or both, i.e. LibriVox (*LV-60k*) [40]. The XLSR-53 model is trained on 56k hours of data from 53 different languages. The XLSR data originates from CommonVoice [41], Multilingual LibriSpeech [40], and BABEL [42]. It includes around 1.6k hours of Dutch [5] of which

we recon only a very small part is Flemish Dutch (a few hours in CommonVoice). We also use a wav2vec 2.0 model pre-trained on the Dutch part of VoxPopuli (*VP-NL-4.5k*) [43], which contains 4.5k hours of Netherlands Dutch speech recordings from the European parliament.

4. DISCUSSION

4.1. Phone classification

4.1.1. Applicability of off-the-shelf models to Flemish Dutch

First, we perform phone classification as explained in Section 2.2.1. For experiments in English, we train and test the classifier on a train-test split of LibriSpeech *train-clean-100*. For experiments in Flemish, we either train a classifier on *VL-train-clean* and test on *VL-test-clean*, or train on *VL-train-other* and evaluate on *VL-test-other*. Table 2 shows the phone classification accuracies with features extracted from English pre-trained models that are online available (see Section 3.2).

Table 2: Linear phone classification accuracy (%) with features extracted from off-the-shelf models pre-trained on English. We evaluate classification on English and Flemish.

Model	English	Flemish	
	<i>LS-tc100</i>	<i>VL-test-clean</i>	<i>VL-test-other</i>
Baseline	48.0	48.5	39.3
APC	72.7	71.4	60.1
Mockingjay	68.1	71.4	59.1
CPC	71.3	71.7	60.5
wav2vec	78.4	73.3	62.4
wav2vec 2.0 (base)	75.1	71.7	58.8

The relative improvements with respect to the baseline as a result of pre-training are consistent across both languages. The accuracy on *VL-test-clean* is of a similar magnitude as the accuracy on *LS-tc100*, which can be explained by the fact that both sets contain rather easy, clean speech. On *VL-test-clean* and *VL-test-other*, we see absolute accuracy improvements of more than 20%. This shows that the pre-training techniques improve linear phone separability, even when the target language differs from the pre-training language.

4.1.2. Language Matching

Second, we examine the effect of matching the domain (i.e. the language, but also the type of speech) of the pre-training speech to the target speech. We pre-train models on Flemish Dutch data, compare them to other pre-trained models, and investigate the effect of finetuning wav2vec 2.0 on a Flemish subset. Table 3 reports phone classification accuracies for pre-training and finetuning on several datasets.

For APC, we notice an improvement over the English pre-trained model when we match the training and target language

Table 3: Phone classification accuracy (PCA) percentage when training a classifier on *VL-train-clean* and testing on *VL-test-clean* ('clean'), and when training a classifier on *VL-train-other* and testing on *VL-test-other* ('other').

Model	Pre-training	Finetuning	PCA	
			clean	other
Baseline (Filletbank)	–	–	48.5	39.3
APC	LS-960	–	71.4	60.1
	VL-train-clean	–	66.9	54.8
	VL-train-other	–	67.6	57.1
	VL-unsup	–	73.3	63.3
	VL-train-all + VL-unsup	–	73.3	63.0
Mockingjay	LS-960	–	71.4	59.1
	VL-train-clean	–	65.2	53.5
CPC	LL-60k	–	71.7	60.5
	VL-train-clean	–	72.6	55.6
	VL-train-other	–	69.3	59.5
	VL-unsup	–	66.8	57.4
	VL-train-all + VL-unsup	–	67.5	58.1
wav2vec	LS-960	–	73.3	62.4
wav2vec 2.0 base	LS-960	–	71.7	58.8
	VP-NL-4.5k	–	64.5	50.0
	VL-train-other	–	47.4	36.1
	VL-train-all + VL-unsup	–	54.7	43.9
	LS-960	VL-train-other	83.6	76.2
	VP-NL-4.5k	VL-train-other	83.6	76.1
	VL-train-other	VL-train-other	81.3	74.1
wav2vec 2.0 large	VL-train-all + VL-unsup	VL-train-other	82.2	75.0
	LS-960	–	55.9	45.2
	LV-60k	–	24.6	14.3
	XLSR-53	–	34.4	21.8
	VP-NL-4.5k	–	58.2	45.7
	LS-960	VL-train-other	81.2	73.4
	LV-60k	VL-train-other	85.0	76.6
	XLSR-53	VL-train-other	86.4	79.1
	VP-NL-4.5k	VL-train-other	84.12	76.3

and use a sufficient amount of data (but still less than LibriSpeech). For CPC, we experienced converging difficulties and a high sensitivity to the number of training cases. We see an improvement over the pre-trained model on *VL-test-clean* when only training on *VL-train-clean*. This might suggest that domain matching is important for CPC. The LibriLight pre-trained model is trained on much more data (60k hours), which can explain the strong performance on *VL-test-other*.

For wav2vec 2.0, the base models trained on Flemish data and the pre-trained model on VoxPopuli perform worse than the LibriSpeech model, despite matching language (VL) or using more data in a related language (VP). The former is most likely explained by sub-optimal training, the latter could be explained by the fact that the VoxPopuli parliament recordings reflect different acoustic conditions to certain components with clean speech in CGN. Finetuning on Flemish leads to very high phone classification accuracies for all models.

For the large high-capacity wav2vec 2.0 models, we note low accuracies without finetuning. Other works corroborate this finding and ascribe it to the problem-agnostic pre-training, and have shown that certain audio features are more easily accessible from the middle layers in very deep transformer models than the output layer [4, 30]. Finetuning (with graphemic transcripts) alleviates this discrepancy and gives an improved accuracy over the base models with the large models. It seems that the large model also benefits more from a large amount of pre-training data, as the LibriVox and XLSR model show.

The XLSR-53 model reports the highest score. It is trained on a similar amount of data as the LibriVox model, but the training data contains Dutch, German and English. We also note that the (Flemish) Dutch stops, which are pre-voiced voiced stops, differ from the aspirated voiceless stops in English. This is also better covered in the XLSR-53 set.

4.2. ASR Results

4.2.1. Clean ASR

Table 4 reports WER results on *VL-test-other* of HMM-DNN ASR experiments with a large DNN and features from different models, as described in Section 2.2.2. Every ASR model (including the baseline model) is trained on *VL-train-other*.

Table 4: ASR experiments with large DNN ASR model, reporting WER on *VL-test-other*.

Model	Pre-training	Finetuning	WER
Baseline (MFCC)	–	–	15.10
APC	LS-960	–	16.02
	VL-train-all + VL-unsup	–	16.20
CPC	LS-960	–	15.03
wav2vec	LS-960	–	14.89
wav2vec 2.0 base	LS-960	–	13.84
	VL-train-other	–	14.44
	VL-train-all + VL-unsup	–	13.52
	LS-960	VL-train-other	11.42
	VL-train-other	VL-train-other	13.41
	VL-train-all + VL-unsup	VL-train-other	11.76
wav2vec 2.0 large	LS-960	–	14.33
	LV-60k	–	14.72
	VP-NL-4.5k	–	16.32
	XLSR-53	–	13.40
	LS-960	VL-train-other	10.87
	LV-60k	VL-train-other	12.65
	XLSR-53	VL-train-other	10.61

The improvements compared to the baseline with MFCC are small, if any, except for wav2vec 2.0. In contrast to the phone classification experiments, we see an improvement over the LibriSpeech model (960h) when we pre-train the base model on a comparable amount of Flemish (720h) without finetuning, which supports the idea that the combination of a matching pre-training language and a large amount of data is key.

The poor performance of the large VoxPopuli model can possibly be explained by a poor transfer from Netherlands Dutch to Flemish and different acoustic conditions compared to the test set. Also, contrarily to phone classification, the large wav2vec 2.0 model pre-trained on LibriSpeech outperforms the model pre-trained on LibriVox.

Analogous to the phone classification experiments, the XLSR-53 model - which has a large variability in its training data - yields a significant WER improvement, manifesting a positive cross-lingual transfer to Flemish, and the best results are obtained when finetuning a model on an annotated Flemish subset. We obtain almost 30% relative WER improvement when finetuning XLSR-53 on Flemish, compared to the baseline. The difference with the LibriSpeech model is however small. We postulate that a large wav2vec 2.0 model trained on more Flemish data would equal or improve this result.

4.2.2. Effect of amount of pre-training and finetuning data

We quantitatively examine the effect of an increasing amount of data used for pre-training wav2vec 2.0 base models (unlabelled) or finetuning XLSR-53 (labelled). We shuffle all available data from all sets. We also evaluate finetuning on different types of speech (from different sets). For the unlabelled dataset, this distinction is not trivial. Table 5 shows the results. The ASR model is always trained on *VL-train-other*.

Table 5: WER with DNN ASR model in function of the amount of Flemish data used for pre-training or finetuning.

10h	30h	50h	100h	150h	250h	350h	500h	700h
31.87	20.85	16.76	15.55	15.74	14.76	14.34	14.73	13.52

(a) Unlabelled data for pre-training a base wav2vec 2.0 model (no finetuning), large ASR DNN.

0h	1h	10h	20h	30h	50h	90h	150h	250h
27.75	13.84	12.08	11.32	11.19	10.71	10.61	10.53	10.50

(b) Labelled data for finetuning XLSR-53, small ASR DNN.

VL set	No FT (0h)	Clean (29h)	Other (128h)	All (248h)
WER	13.40	12.35	10.61	10.58

(c) Different sets of CGN for finetuning XLSR-53, large ASR DNN.

For pre-training, it is necessary to have a considerable amount of data to improve upon the baseline, and more Flemish data gives improvements. It seems that the learned audio representations include acoustic details aside from more abstract phoneme qualities, as the WER on 150h of shuffled data is higher than when pre-training on an equal amount of matched data (*VL-train-other* in Table 4). This might suggest a high dependency on acoustic conditions. For finetuning, the data should match the type or conditions of the test set for optimal results, and the improvements saturate with more data. Note that a small DNN suffices after finetuning, but a large DNN is required when the XLSR model is not finetuned (first column

of Table 5b). This is in line with the poor phone classification results of the large models without finetuning.

4.2.3. ASR in noisy environments

We investigate the robustness of wav2vec 2.0 to noisy and reverberated speech by replicating the *VL-test-other* set in 4 different scenario's: filtered with RIRs (*rev*), with added noise at a certain SNR (*noise1*: 5-20dB, *noise2*: 0-15dB) and both (*rev + noise3*: 5-15dB). We use noises from multiple sources (NTT Noise-DB, CHIME2, NoiseX, DEMAND, Humming) and RIRs from the Aachen Impulse Response Database. We compare wav2vec 2.0 large models pre-trained on LibriVox: without finetuning, finetuned on *VL-train-other* ('clean') and finetuned on a fourfold augmented *VL-train-other* ('aug') by adding noise and reverberation in Table 6.

Table 6: WER with large DNN ASR on augmented *VL-test-other* with LibriVox pre-trained large wav2vec 2.0 models.

Model	FT	WER				
		<i>clean</i>	<i>rev</i>	<i>noise1</i>	<i>noise2</i>	<i>rev + noise3</i>
MFCC	–	15.10	28.36	20.58	26.26	39.21
w2v2	–	14.71	27.12	19.96	25.28	39.19
w2v2	clean	12.43	22.60	16.31	20.61	33.32
w2v2	aug	12.13	18.08	14.64	17.39	24.43

Finetuning on augmented data gives strong improvements over finetuning on clean data in the reverberated and noisy settings, with 3-9% absolute WER reduction. More so, there is even a slight improvement in the clean setting as well, probably because of more finetuning data.

5. CONCLUSION

Pre-trained features on English speech transfer well to Flemish Dutch in terms of improving linear phone separability. These self-supervised pre-trained models are readily available and easy to use. Matching the pre-training and target language further improves results, but either matching the type of speech or using a larger amount of data is necessary. The recently proposed wav2vec 2.0 model appears superior, especially when finetuned on data from the target language. Finally, we obtain the best results with the large multilingually trained XLSR-53 model and see nearly 30% improvement in WER by finetuning the XLSR-53 model on Flemish, compared to the baseline. We show the importance of matching pre-training and target language and acoustic conditions.

6. ACKNOWLEDGEMENTS

This research received funding from the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" programme.

7. REFERENCES

- [1] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” 2020.
- [2] Qiantong Xu, Alexei Baevski, Tatiana Likhomanenko, Paden Tomasello, Alexis Conneau, Ronan Collobert, Gabriel Synnaeve, and Michael Auli, “Self-training and pre-training are complementary for speech recognition,” 2020.
- [3] Yu Zhang, James Qin, Daniel S. Park, Wei Han, Chung-Cheng Chiu, Ruoming Pang, Quoc V. Le, and Yonghui Wu, “Pushing the limits of semi-supervised learning for automatic speech recognition,” 2020.
- [4] Alexei Baevski, Wei-Ning Hsu, Alexis Conneau, and Michael Auli, “Unsupervised speech recognition,” 2021.
- [5] Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli, “Unsupervised cross-lingual representation learning for speech recognition,” 2020.
- [6] Morgane Rivire, Armand Joulin, Pierre-Emmanuel Mazar, and Emmanuel Dupoux, “Unsupervised pre-training transfers well across languages,” 2020, Code: https://github.com/facebookresearch/CPC_audio.
- [7] Yu-An Chung, Yonatan Belinkov, and James Glass, “Similarity analysis of self-supervised speech representations,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 3040–3044.
- [8] Yu-An Chung, Wei-Ning Hsu, Hao Tang, and James Glass, “An unsupervised autoregressive model for speech representation learning,” 2019.
- [9] Yu-An Chung and James Glass, “Generative pre-training for speech with autoregressive predictive coding,” 2020.
- [10] Yu-An Chung, Hao Tang, and James Glass, “Vector-quantized autoregressive predictive coding,” 2020, Code: <https://github.com/iamyuanchung/VQ-APC>.
- [11] Yu-An Chung and James Glass, “Improved speech representations with multi-target autoregressive predictive coding,” 2020.
- [12] Andy T. Liu, Shu-wen Yang, Po-Han Chi, Po-chun Hsu, and Hung-yi Lee, “Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders,” *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020.
- [13] Po-Han Chi, Pei-Hung Chung, Tsung-Han Wu, Chun-Cheng Hsieh, Shang-Wen Li, and Hung yi Lee, “Audio ALBERT: A lite BERT for self-supervised learning of audio representation,” 2021.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” 2019.
- [15] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” 2019.
- [16] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli, “wav2vec: Unsupervised pre-training for speech recognition,” 2019.
- [17] Alexei Baevski, Steffen Schneider, and Michael Auli, “vq-wav2vec: Self-supervised learning of discrete speech representations,” 2020.
- [18] Alexei Baevski, Michael Auli, and Abdelrahman Mohamed, “Effectiveness of self-supervised pre-training for speech recognition,” 2020.
- [19] Kazuya Kawakami, Luyu Wang, Chris Dyer, Phil Blunsom, and Aaron Van den Oord, “Learning robust and multilingual speech representations,” 2020.
- [20] Jan Chorowski, Ron J. Weiss, Samy Bengio, and Aaron Van den Oord, “Unsupervised speech representation learning using wavenet autoencoders,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 12, pp. 2041–2053, Dec 2019.
- [21] Sameer Khurana, Antoine Laurent, Wei-Ning Hsu, Jan Chorowski, Adrian Lancucki, Ricard Marxer, and James Glass, “A convolutional deep markov model for unsupervised speech representation learning,” 2020.
- [22] Andy T. Liu, Shang-Wen Li, and Hung yi Lee, “TERA: Self-supervised learning of transformer encoder representation for speech,” 2020.
- [23] Alexander H. Liu, Yu-An Chung, and James Glass, “Non-autoregressive predictive coding for learning speech representations from local dependencies,” 2020.
- [24] Mirco Ravanelli, Jianyuan Zhong, Santiago Pascual, Pawel Swietojanski, Joao Monteiro, Jan Trmal, and Yoshua Bengio, “Multi-task self-supervised learning for robust speech recognition,” 2020.

- [25] Santiago Pascual, Mirco Ravanelli, Joan Serrà, Antonio Bonafonte, and Yoshua Bengio, “Learning problem-agnostic speech representations from multiple self-supervised tasks,” 2019.
- [26] Dongwei Jiang, Wubo Li, Miao Cao, Ruixiong Zhang, Wei Zou, Kun Han, and Xiangang Li, “Speech SIMCLR: Combining contrastive and reconstruction objective for self-supervised speech representation learning,” 2020.
- [27] Shaoshi Ling, Yuzong Liu, Julian Salazar, and Katrin Kirchhoff, “Deep contextualized acoustic representations for semi-supervised speech recognition,” *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020.
- [28] Hans Van de Velde, M. Kissine, E. Tops, S. van der Harst, and R. van Hout, “Will Dutch become Flemish? Autonomous developments in Belgian Dutch,” in *Multilingua* 29, 2010.
- [29] Wei-Ning Hsu, Anuroop Sriram, Alexei Baevski, Tatiana Likhomanenko, Qiantong Xu, Vineel Pratap, Jacob Kahn, Ann Lee, Ronan Collobert, Gabriel Synnaeve, and Michael Auli, “Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training,” 2021.
- [30] Danni Ma, Neville Ryant, and Mark Liberman, “Probing acoustic representations for phonetic properties,” 2021.
- [31] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely, “The Kaldi speech recognition toolkit,” in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. Dec. 2011, IEEE Signal Processing Society, IEEE Catalog No.: CFP11SRW-USB.
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” 2017.
- [33] Andy T. Liu and Yang Shu-wen, “S3PRL: The self-supervised speech pre-training and representation learning toolkit,” 2020.
- [34] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli, “fairseq: A fast, extensible toolkit for sequence modeling,” in *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- [35] Kris Demuynck, Tom Laureys, and Steven Gillis, “Automatic generation of phonetic transcriptions for large speech corpora,” in *Proceedings International Conference on Spoken Language Processing*, 2002, vol. 1, pp. 333–336.
- [36] Nelleke Oostdijk, “The Spoken Dutch Corpus: Overview and first evaluation,” *Proceedings of LREC-2000, Athens*, vol. 2, 01 2000.
- [37] Daniel Povey, Gaofeng Cheng, Yiming Wang, Ke Li, Hainan Xu, Mahsa Yarmohammadi, and Sanjeev Khudanpur, “Semi-orthogonal low-rank matrix factorization for deep neural networks,” in *Proc. Interspeech 2018*, 2018, pp. 3743–3747.
- [38] 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, pp. 5206–5210.
- [39] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, “Libri-Light: A benchmark for ASR with limited or no supervision,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7669–7673, <https://github.com/facebookresearch/libri-light>.
- [40] Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert, “MLS: A large-scale multilingual dataset for speech research,” *Interspeech 2020*, Oct 2020.
- [41] Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber, “Common Voice: A massively-multilingual speech corpus,” 2020.
- [42] Mark J. F. Gales, Kate M. Knill, Anton Ragni, and Shakti P. Rath, “Speech recognition and keyword spotting for low-resource languages: Babel project research at cued,” in *SLTU*, 2014, pp. 16–23.
- [43] Chaghan Wang, Morgane Rivière, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux, “VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” 2021, Code: <https://github.com/facebookresearch/voxpopuli>.