



Pretext Tasks selection for multitask self-supervised speech representation learning

Salah Zaiem, Titouan Parcollet, Slim Essid, Abdelwahab Heba

► To cite this version:

Salah Zaiem, Titouan Parcollet, Slim Essid, Abdelwahab Heba. Pretext Tasks selection for multitask self-supervised speech representation learning. IEEE Journal of Selected Topics in Signal Processing, 2022, 16 (6), pp.1439-1453. 10.1109/JSTSP.2022.3195430 . hal-03601330

HAL Id: hal-03601330

<https://hal.science/hal-03601330>

Submitted on 8 Mar 2022

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.

PRETEXT TASKS SELECTION FOR MULTITASK SELF-SUPERVISED SPEECH REPRESENTATION LEARNING

Salah Zaiem

Télécom Paris
Palaiseau, France
zaiemsalah@gmail.com

Titouan Parcollet

Avignon Université, LIA
Avignon, France

Slim Essid

Télécom Paris
Palaiseau, France

ABSTRACT

Through solving pretext tasks, self-supervised learning leverages unlabeled data to extract useful latent representations replacing traditional input features in the downstream task. In audio/speech signal processing, a wide range of features where engineered through decades of research efforts. As it turns out, learning to predict such features (a.k.a pseudo-labels) has proven to be a particularly relevant pretext task, leading to useful self-supervised representations which prove to be effective for downstream tasks. However, methods and common practices for combining such pretext tasks for better performance on the downstream task have not been explored and understood properly. In fact, the process relies almost exclusively on a computationally heavy experimental procedure, which becomes intractable with the increase of the number of pretext tasks. This paper introduces a method to select a group of pretext tasks among a set of candidates. The method we propose estimates calibrated weights for the partial losses corresponding to the considered pretext tasks during the self-supervised training process. The experiments conducted on automatic speech recognition, speaker and emotion recognition validate our approach, as the groups selected and weighted with our method perform better than classic baselines, thus facilitating the selection and combination of relevant pseudo-labels for self-supervised representation learning.

1 INTRODUCTION

Self-supervised learning (SSL) methods usually rely on a supervision obtained from the data itself through solving specific pretext tasks leveraging the underlying structure of the considered data (Doersch et al., 2016; Arandjelovic & Zisserman, 2018). This technique is used in various domains including image processing (Misra & Maaten, 2020; Jing & Tian, 2020; Grill et al., 2020), natural language understanding (Chen et al., 2020b; Du et al., 2020; Lan et al., 2019) or speech and audio processing (Baevski et al., 2020; Liu et al., 2020; Jiang et al., 2020). It offers numerous advantages, such as the independence from labeled data, stronger performance on downstream tasks, more robust models and an easier transfer to low-resource setups (e.g., low-resource languages) (Baevski et al., 2020; Jing & Tian, 2020).

The numerous existing SSL approaches are characterized by the nature of the pretext tasks they solve. For instance, common techniques include predictive coding (Baevski et al., 2020; Liu et al., 2020; Song et al., 2020; Zhang et al., 2020; Hsu et al., 2021), pseudo-label learning (Pascual et al., 2019; Ravanelli et al., 2020), auto-encoding (Renshaw et al., 2015; Algayres et al., 2020), generative modelling (Khurana et al., 2020) or contrastive learning (Saeed et al., 2020; Jiang et al., 2020). More precisely, these pretext tasks may be defined through the choice of pretext labels, hereafter referred to as *pseudo-labels*. The automatic extraction of pseudo-labels for SSL (i.e. from the data itself) is common in many application domains, such as computer vision (Noroozi & Favaro, 2017; Gidaris et al., 2018), music processing (Hung et al., 2019; Wu et al., 2021) and speech processing

(Pascual et al., 2019; Shukla et al., 2020), and is commonly referred to as *multitask self supervised learning*. In the specific context of speech processing, the process of designing pseudo-labels may benefit from decades of research in signal processing. For instance, potential candidates are pitch estimators, energy-based features, voicing state and many more.

As demonstrated by Pascual et al. (2019), multitask speech representation learning is a powerful tool to build representations that are beneficial for a wide range of distinct downstream tasks, by combining different pseudo-labels which “intuitively” correspond to these tasks. Unfortunately, there is no clear understanding of how these pseudo-labels may interact when optimised together, and therefore, no common practice of how to select groups of pseudo-labels to obtain better performance on a known downstream task. As a matter of fact, this design process has been essentially driven by empirical validation and there is therefore no evidence that the obtained model is even the best one. This empirical approach can rapidly become intractable with modern SSL architectures which may contain hundreds of millions or billions of parameters trained on thousands of hours of speech, not to mention the carbon footprint of such pseudo-label searches. For instance, the self-supervised training of a single state-of-the-art large wav2vec 2.0 model (Baevski et al., 2020) on 53.2k hours of speech requires 128 GPUs for 5.2 days.

This work aims at providing a clear, efficient and theoretically motivated procedure for pseudo-label group selection and weighting based on conditional independence (CI). The method presented allows one to design ahead of training the most adapted multitask self-supervised speech representation learning model which perfectly suits the considered downstream tasks. Such an approach may also enable researchers to save a substantial amount of time and computation usually devoted to pseudo-label search. Hence, the contributions of this work are fourfold:

1. Introduce a theoretically motivated and computationally efficient method for the selection of pseudo-label *groups* among a set of candidates and with respect to the considered downstream tasks (Sections 3 and 4).
2. Validate empirically the proposed approach with a first model SSL model relying on different sets of pseudo-labels corresponding to the ones obtained for three considered speech tasks. (Sections 5).
3. Extend our method to the SOTA wav2vec 2.0 to enhance its performance (Section 6).
4. Release the code base developed with SpeechBrain (Ravanelli et al., 2021) for replication and to encourage further investigations.¹

The conducted experiments demonstrate that the proposed method allows for a more intelligent, *i.e.* better informed, pseudo-label group selection for multitask SSL settings. Indeed, we find that the models built with the proposed method obtain a word error rate and an equal error rate, respectively, 31.6% and 27.4% lower than the baseline, without the need for any empirical search.

2 RELATED WORKS AND MOTIVATIONS

SSL recently became a key component to achieve good performance on downstream tasks especially with low-resource setups, either in speech (Baevski et al., 2020; Conneau et al., 2020), natural language processing (Lan et al., 2019; Chen et al., 2020b) or computer vision (Gidaris et al., 2019; Misra & Maaten, 2020; Jing & Tian, 2020). Due to its very nature, SSL relies on large amounts of unlabeled data used to train large deep neural networks over long periods of time. It is thus crucial to understand properly what makes a good SSL model to lower the amount of computation and time needed to obtain the best downstream performance.

SSL for Speech. Self-supervised learning for speech has recently enabled researchers to reach state-of-the-art results on various speech processing tasks (Fan et al., 2021). The most successful models rely on predictive and contrastive objectives (Baevski et al., 2020; Chung et al., 2019; Saeed et al., 2020) performing well across the different tasks even in low-resource settings. This led to the design of different benchmarks evaluating the self-supervised representations in different languages (Yang et al., 2021; Evain et al., 2021). However, in contrast to this proposition, these works have not tried to motivate beforehand the different choices made in the self-supervision pipeline.

¹<https://github.com/salah-zaiem/Multitask-pretext-task-selection>

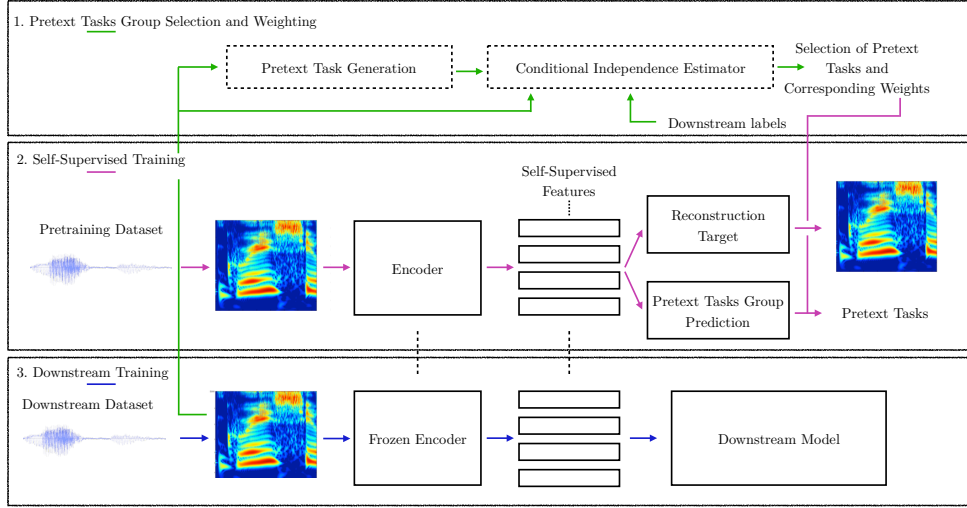


Figure 1: Illustration of the training pipeline. The three steps are depicted: 1. Selecting the group of pseudo-labels and their corresponding weights; 2. SSL training with the selected pretext task; 3. Training on the downstream task with the pretrained SSL model.

Understanding SSL. A few works have tried to shed some theoretical light on the mainly empirical field of self-supervised learning. Following the different paradigms in SSL, various tracks have been followed to understand what makes for a good self-supervised representation, exploring different approaches (Lee et al., 2020; Arora et al., 2019; Wei et al., 2020). On the one hand, contrastive learning (Oord et al., 2018; Chen et al., 2020a) has been advocated both theoretically and empirically to achieve a balance in the mutual information between alternative representations of the data, keeping just enough shared information to keep the class-related content (Tschannen et al., 2020; Tian et al., 2020; Bachman et al., 2019). In a recent work from Li et al. (2021), independence testing has been used to produce better transformations in contrastive learning settings for image representations. Predictive learning, on the other hand, requires the model to predict masked elements in sequential data. This technique is powerful on downstream tasks that can be reduced to a masking problem, as suggested by research on language modeling (Saunshi et al., 2020). However, all these works have been focusing on computer vision or text-related applications, and none of them addressed the multi-task self supervision problem.

Multi-task self-supervised learning. While the literature on multi-tasking in self-supervised learning remains scarce, it has been shown in classic supervised learning settings, that through estimates of similarity between tasks or thorough empirical testing, several tasks can take advantage of being solved with a common encoder (Zamir et al., 2018; Dwivedi & Roig, 2019; Shafey et al., 2019; Chen et al., 2015). More specifically, combining pretext tasks with SSL has been mainly explored in computer vision and speech (Pascual et al., 2019; Ravanelli et al., 2020). Pretext tasks such as Jigsaw (Doersch et al., 2016), colourisation and rotation (Gidaris et al., 2018) have been combined successfully to improve downstream performance (Kim et al., 2018; Shin’ya Yamaguchi et al.). The two closest works to our line of research are from Lee et al. (2020) and Doersch et al. (2017). The former shows that a theoretical link can be established between conditional independence and an improvement of the performance on the downstream task, while the latter proposes to select layers from a multitask self-supervised encoder according to the pretext task to be solved. However, in both cases, the studies do not offer practical and theoretical solutions to select groups of pseudo-labels to build an adapted SSL model that will perform well on the considered downstream tasks.

Group feature selection. Finally, feature selection, and especially group feature selection is another close and inspiring field given the problem we consider. The relationship and interactions between features have been largely investigated in the supervised learning literature (Guyon & Elisseeff, 2003). This led to multiple solutions to the feature group selection problem, including LASSO-based techniques (Yuan & Lin, 2006), or multiple kernel formulations (Sonnenburg et al., 2006;

Rakotomamonjy et al., 2007). However, these works do not involve any self-supervision, and links between feature selection and self-supervision design and pretext-task selection are yet to be proved. We will further consider these lines of works for concurrent baselines.

With this work, we aim at shortening the process of designing SSL models while giving insights on the pseudo-label importance and the underlying mechanisms between pretext and downstream tasks at the same time. We decided to experiment with speech due to the lack of literature on this domain for multitask SSL, and for the various pseudo-labels available, which are based on decades of signal processing research. The whole pipeline starting from the acoustic feature extraction to the downstream task scoring follows three major steps summarized in Figure 1. First, for every downstream task, our method produces a pretext task selection and weighting. Then, a SSL model is trained, before being used as a feature extractor front-end to one or many downstream tasks.

3 CONDITIONAL INDEPENDENCE FOR UTILITY ESTIMATION

As a first step, we require a function that estimates the utility of learning to solve a pretext-task to improve the performance on the downstream task. We use an estimation of the conditional independence between the pseudo-label values and the downstream data points given the downstream labels. Hereafter, we explain the theoretical motivations and describe the computation steps.

3.1 PROBLEM DEFINITION AND INTUITION

Let X , Y and Z be, respectively, the downstream data points, their downstream labels and their pseudo-labels. Let also \mathcal{C} be the set of possible downstream classes. As an example, if one considers speaker recognition as a downstream task, X would be the speech samples, Y the speaker IDs, \mathcal{C} the set of unique speaker IDs, and Z a computed signal feature, such as the fundamental frequency.

As stated in Section 2, Lee et al. (2020) linked the utility of a pseudo-label (Z) to the conditional independence (CI) between Z and X given Y . The approach prescribes that, given the labels Y , one may seek to *quantify how much it is possible to predict the pseudo-labels Z without knowing much about X* . The authors bounded, under certain assumptions, the downstream classifier error with a function of the downstream training set size, and a measure of the CI. More precisely, the main theorem shows that the bounding function decreases linearly with the downstream-task dataset size (M) and quadratically with the CI, thus making it a potential estimator for pseudo-label utility.

The proposed function depends on the final downstream task to be solved. This is motivated by two main reasons. First, it can be seen through the large literature on feature selection for various speech or computer vision tasks (Liu et al., 2020; Serizel et al., 2017; Schuller et al., 2007; Wang et al., 2019), that different tasks require the description of different aspects of the data. This suggests that different downstream tasks may perform better after different pre-trainings. A second argument is the difficulty to evaluate representations quality intrinsically, *i.e.* independently from the choice of a particular downstream task. A few metrics and tests (Schatz et al., 2013; Carlin et al., 2011; Lakhota et al., 2021) have been proposed for speech, but the correlation between these and downstream-task performance has not been clearly identified (Algayres et al., 2020; Gump et al., 2020).

The principal issue with CI is the difficulty of computing good estimates of how much two variables are independent given a third one on realistic data (Shah & Peters, 2018). In a previous work (Zaiem et al., 2021), we proposed a simple way to get an estimation of the conditional independence. This method has proven effective for individual pretext task selection, as the utility estimator correlated highly with the final downstream performances. In our case, the considered pseudo-labels are not independent of the speech samples, as they are signal features. The approach resorts to performing classic independence testing on data sliced by the downstream classes, to check whether this dependence remains given the downstream labels information.

3.2 CONDITIONAL INDEPENDENCE ESTIMATOR COMPUTATION

This section details the computation of the conditional independence estimate that is used as a measure of pseudo-label utility. Let $X = \{x_i\}_{i \in \{0, \dots, M\}}$ with M being the cardinal of X and x_i data samples (*e.g.* Mel-band spectrogram for audio and speech). Every sample x_i has a corresponding downstream label y_i and an automatically generated pseudo-label z_i . We assume that y_i is discrete

reducing the task to a classification problem such as with speaker ID for speaker recognition. We also assume that for every pretext-task Z , a single z_i value corresponds to each x_i . In our case, z_i values are the mean of the frame-wise pseudo-label values.

For independence testing, we decided to rely on the kernel-based Hilbert Schmidt Independence Criterion (HSIC) (Gretton et al., 2007) for two reasons. First, HSIC has already proven successful for textual data in testing statistical dependence between translated sentences (Gretton et al., 2007). Second, kernel-based techniques facilitate the handling of multivariate and varying-length data such as speech, as the estimation then boils down to the computation of a similarity measure between the considered variables.

Computation steps. The estimation of the CI of a pseudo-label Z for a downstream task (X, Y) consists of three steps. We start by splitting the data samples X according to the downstream (discrete) classes. Then, we compute for every downstream class $c \in \mathcal{C}$, the kernel matrices K_c and L_c representing the similarity measures for the data samples, and the pseudo-labels, respectively. Finally, we perform the independence test for every split group using K_c and L_c and aggregate the estimates with a weighted mean taking into account the number of samples per downstream class. Thus, for two speech samples x_i and x_j , holding two pseudo-label values z_i and z_j , the coefficients of the similarity matrices K_c and L_c are computed as follows:

$$K_{ij} = K(x_i, x_j) = \cos(GD(x_i), GD(x_j)), \quad L_{ij} = RBF(z_i, z_j), \quad (1)$$

with $GD(\cdot)$ the Gaussian Downsampling function (more details in the appendix A.6), $\cos(\cdot, \cdot)$ the cosine similarity, and $RBF(\cdot, \cdot)$ the Radial Basis Function kernel, defined as:

$$\cos(x, x') = \frac{\text{trace}(x^T x')}{\|x\| \cdot \|x'\|}, \quad RBF(z, z') = \exp\left(-\frac{\|z - z'\|^2}{2\sigma^2}\right), \quad (2)$$

where σ is the width of the RBF kernel and $\text{trace}(\cdot)$ the sum of elements of the main diagonal. Note that we compute the matrices K_c and L_c , for each group of samples sharing the same downstream class $c \in \mathcal{C}$. Hence, K_c and L_c correspond to the definitions above, but restricted to the points with c as a downstream label. For each downstream class c , and as in Gretton et al. (2007), the HSIC value is given by:

$$HSIC_c(X, Z) = \frac{1}{n_c^2} \text{trace}(K_c H_c L_c H_c), \quad (3)$$

with $H_c = I_{n_c} - \frac{1}{n_c} \mathbf{1}_{n_c} \mathbf{1}_{n_c}^T$, n_c being the number of points with downstream label c , and $\mathbf{1}_{n_c}$ a vector of ones of size $n_c \times 1$.

The HSIC value is non-negative and corresponds to the Hilbert norm of their cross-covariance matrix. It is used to characterize the independence of the two considered quantities. Intuitively, the HSIC value is high if samples similar in K_c are similar in L_c . Therefore, the lower this value is, the more independent the two arguments of HSIC are and the better the pseudo-label should be for self-supervision before fine-tuning on the downstream class. The final value for a given pseudo-label and a downstream task is expressed as:

$$HSIC(X, Z|Y) = \frac{1}{M} \sum_{c \in \mathcal{C}} HSIC_c(X, Z) \times n_c. \quad (4)$$

4 PRETEXT TASK GROUP SELECTION AND WEIGHTING

While we now are able to predict the utility of every considered pretext task independently, the next step remains to define a clever way to combine them optimally within the same pre-training process. Hence, we introduce a method to select a group of pseudo-labels and weight their respective losses to increase or decrease their importance in the self-supervised representation. Such an optimisation of the latent representation is expected to provide better downstream performance. Our method minimises the conditional dependence between the resulting group pretext task, entailing the prediction of a selected pseudo-label group and the downstream samples given the downstream labels.

Given a set of k possible pseudo-labels $(Z_i)_{i \in [0, k]}$, we seek to estimate a set of parameters $(\lambda_i)_{i \in [0, k]}$ weighting the loss of every pseudo-label k during the pre-training phase. Hence, we define a grouping pseudo-label Z_λ as an orthogonal concatenation of $(Z_i)_{i \in [0, k]}$ weighted with $(\lambda_i)_{i \in [0, k]}$:

$$Z_\lambda = (\lambda_1 Z_1, \dots, \lambda_k Z_k).$$

The custom conditional HSIC computation pipeline described above is fully differentiable with respect to $(\lambda_i)_{i \in [0, k]}$ as proved in appendix A.1. In the HSIC computation, the data similarity matrices $\{K_c\}_{c \in C}$ are independent of Z and therefore of λ . Only the pseudo-label similarity matrices $\{L_c\}_{c \in C}$ are changed. For every downstream class c , L_c is defined as:

$$[L_c]_{i,j} = RBF((Z_\lambda)_i, (Z_\lambda)_j) = \exp\left(\frac{-1}{2\sigma^2} \sum_{h=1}^k \lambda_h \|z_{h,i} - z_{h,j}\|_2^2\right), \quad (5)$$

with $z_{h,i}$ denotes the mean value of the h -th pseudo-label for the i -th data point in the dataset.

4.1 CONSTRAINTS ON THE WEIGHTS

The conditional-independence based utility estimator must be optimized with respect to the weighting parameters $(\lambda_i)_{i \in [0, k]}$ and three constraints.

First, the parameters $(\lambda_i)_{i \in [0, k]}$ must be positive, as they are used as weights for the corresponding losses. A negative weighting loss would lack interpretability as it could imply that the self-supervised model should “unlearn” the corresponding pretext task. This may be the case for adversarial learning methods, but we are not considering this case in the present work.

Second, the value of the weights must be high enough. Indeed, the presented method for estimating the conditional independence assumes that the considered pseudo-label Z is not independent of X . It is fortunately true for speech features, as Z is a function of X , but not always valid. For instance, with $(\lambda_i)_{i \in [0, k]} = 0$, the utility estimator would be null and thus the lowest (*i.e.* the best), but it would break the assumption of non independence between Z and X . Furthermore, the *HSIC* value decreases with positive decreasing values of $(\lambda_i)_{i \in [0, k]}$ and we thus need to ensure that the sum of the weights is significantly higher than zero, or it would mean that we are not really doing multi-task learning as the losses of the pseudo-labels would be barely considered.

Finally, for a fair comparison between the weighting choices during the optimization, the sum of the weights should remain constant. In the following, the sum of the $(\lambda_i)_{i \in [0, k]}$ is fixed to one and the problem is summarized as follows:

$$\min_{\lambda \in \mathbb{R}^k} HSIC(Z_\lambda, X, Y), \text{ s.t. } Z_\lambda = (\lambda_1 Z_1, \dots, \lambda_k Z_k), \lambda_i \geq 0, \forall i \in [0, k], \sum_i \lambda_i = 1. \quad (6)$$

To minimize the estimator quantity while respecting the constraints, the weights used in the computation of the CI value are the softmax output of free learnable parameters $(W_i)_{i \in [0, k]}$. Softmax enforces the conditions while being differentiable and the problem becomes:

$$\min_{W \in \mathbb{R}^k} HSIC(Z_\lambda, X, Y), \text{ s.t. } \lambda = \text{Softmax}(W), Z_\lambda = (\lambda_1 Z_1, \dots, \lambda_k Z_k). \quad (7)$$

4.2 WEIGHTS SPARSITY

Another trait that is desirable for the weighting vector is sparsity. If a few pseudo-labels are not needed for the given downstream task, they would rather be discarded than given a low weight. First, this would save computation time including the extraction of the pseudo-labels, and their extraction and prediction during the self-supervised training process. Second, it would help with transparency to understand what features are included or not in the latent space. This sparsity property is also related to feature selection such as with LASSO (Yuan & Lin, 2006). To ensure the sparsity of the output weighting vector, while maintaining the desired property of differentiability, we choose the sparsemax function (Martins & Astudillo, 2016) to replace the softmax in eq. 7.

5 EXPERIMENTAL STUDY

This section details the experiments validating the introduced estimator. It describes the selection and weighting processes (Section 5.1), the SSL models (Section 5.2), the downstream tasks (Section 5.3), and the obtained results (Section 5.4).

5.1 GROUP SELECTION AND WEIGHTING

Individual pseudo-labels of interest are obtained with the OpenSmile library (Eyben et al., 2010). We decided to focus on markers mostly related to prosody and spectral descriptors as the signal processing literature commonly associates them to the three considered downstream tasks (*i.e.* speech, speaker and emotion recognition). Selected pseudo-labels include: *Loudness*, *F0*, *Voicing*, *α Ratio* (Sundberg & Nordenberg, 2006), *Zero Crossing Rate*, *L1 Norm of Rasta Spectrum* (Hermansky et al., 1992), *log of Harmonicity to Noise Ratio* (Murphy & Akande, 2005). Then, and according to Figure 1 (*step 1*), we group these pseudo-labels based on their weights, *i.e.* by optimising equation 7 to obtain the different λ values associated to each one of them.

Comparative baselines follow common feature group selection strategies or natural intuitions. The first one simply bundles all the pseudo-labels together without any weighting (*i.e.* $\lambda = 1$ for all pseudo-labels) as proposed for PASE (Pascual et al., 2019). As SSL group pretext-task selection is yet to be fully explored, the two other baselines come from the feature selection literature as it represents the closest field with well established techniques. The CI-based pseudo-label selection is thus compared to Recursive Feature Elimination (RFE, Guyon et al. (2002)) and Maximum Relevance Minimum Redundancy (MRMR, Peng et al. (2005)). More details about these baseline algorithms are given in Appendix A.7, while Appendix A.9 shows the workers selected and their corresponding weights for every experiment. Noise (HNR) seems to be the most important information to learn to predict for speaker recognition while fundamental frequency is privileged for ASR.

5.2 SELF-SUPERVISED TRAINING

In the second step of Figure 1, the SSL model learns to predict the selected pseudo-labels. For every one of those, the loss is multiplied by the corresponding assigned weight. Based on previous work conclusions (Ravanelli et al., 2020; Jiang et al., 2020) and apart from the considered pretext task, the network learns to reconstruct the input Mel spectrograms, and to compute 40-dimensional Mel-Frequency Cepstral Coefficients (MFCC) feature vectors. These targets are usually kept to avoid information loss harming heavily downstream performance and are used in all our experiments. For a given weighting vector $(\lambda_i)_{i \in [0, k]}$, the self-supervised loss is defined as:

$$L_{SSL} = MSE_{mel} + MSE_{mfcc} + \sum_{i=1}^k \lambda_i \ell_1(Z_i), \quad (8)$$

with MSE the classic mean squared error computed for Mel spectra (MSE_{mel}) and MFCC (MSE_{mfcc}), and $\ell_1(Z)$ the ℓ_1 -loss of the pretext task related to pseudo-label Z .

Prior to extending our method to state-of-the-art architectures such as wav2vec 2.0 that are particularly costly to train, we propose to first employ a PASE-like model to empirically validate the approach. Hence, the encoder is composed of three distinct parts: a VGG-like feature extractor (Simonyan & Zisserman, 2015), a bidirectional LSTM, and a two-layered dense neural network. All the details of the architecture are given in the appendix A.4. Then, and inspired by Ravanelli et al. (2020), the encoder is followed by simple one-layered predictors voluntarily limited in capacity.

SSL dataset. The SSL model is optimised on the training set of the English Common Voice dataset (version 5.1, 700 hours of training, Ardila et al. (2020)). Common Voice is a collection of speech utterances from worldwide users recording themselves from their own devices. Hence, the closeness to natural settings makes it a suitable choice for self-supervised learning. 700 hours of speech is a relatively small amount compared to what is generally used for state-of-the-art SSL models. However, we believe it is a sound choice as this is generally greater than what is typically available in SSL use-cases like low-resource languages. We decided to not use the LibriSpeech dataset for pre-training as it is part of our downstream evaluation protocol hence alleviating a strong bias.

5.3 DOWNSTREAM TASKS

Our proposed pseudo-label selection strategy is compared with the two baselines on two different downstream tasks leading to different groups of pseudo-labels: automatic speech recognition (ASR, with LibriSpeech 100 hours) and speaker recognition (SR, with VoxCeleb 1). Datasets and downstream architectures are inspired from the SUPERB benchmark (Yang et al., 2021) for self-supervised learning representations and carefully described in Appendix A.4.3. Prior to downstream

Table 1: Results observed with the proposed selection strategies on the two considered downstream tasks. Word Error Rate (WER) and Equal Error Rate (EER) are expressed in percentage and used for LibriSpeech 100 hours and VoxCeleb1 respectively (*i.e.* lower is better). ASR results are given with and without Language Modeling (LM). All SSL models contain 16.3M neural parameters.

Selections	LibriSpeech (WER % ↓)		VoxCeleb1 (EER % ↓)
	No LM	LM	
All	21.98	11.70	11.90
MRMR	18.94	10.36	10.56
RFE	20.02	11.42	11.91
Softmax	13.17	8.00	9.24
Sparsemax	17.18	10.41	8.63

training, the SSL model are frozen to be used as a feature extractor with the new pipeline that is task-dependent. We do not use any data augmentation for a pristine comparison of the learned models.

5.4 RESULTS

Baselines detailed in Section 4 are respectively referred to as “*All*”, “*RFE*” and “*MRMR*”. All the details about the selection and weights are available in Appendix A.9. First, it is clear from the results reported in Table 1 that, for the considered downstream tasks, the two introduced strategies (*Sparsemax* and *Softmax*) perform better than the group selection baselines with a gain of 3.28 of EER for *Sparsemax* against the *RFE* approach on VoxCeleb, and 8.81 of WER with *Softmax* compared to the *All* baseline. Interestingly, simply bundling all the pseudo-labels together may lead to poor performance as observed on LibriSpeech with a very high 21.98% of WER obtained. Hence, *intuitively* building sets of labels could be harmful for the final representation. This motivates the need for a better pseudo-label selection strategy such as the one introduced in this work, as the WER dropped to 13.17%. As a comparison, the exact same architecture trained with Mel spectra only (*i.e.* no SSL) obtains a WER of 17.3% without LM. Hence, our method even further decreases the WER while being only pretrained with a reasonable amount of data (*i.e.* only 700 hours compared to a few thousands for common SSL techniques (Baevski et al., 2020)). As expected, introducing the joint decoding with a language model strongly decreases the WER but also introduces a bias in our comparison as probabilities are smoothed with a third-party neural model. Nevertheless, and even in this scenario, our weighting strategy outperforms all the baselines. In the context of speaker recognition, *Sparsemax* beats *Softmax* with an EER 0.61 lower.

6 EXTENDING WAV2VEC 2.0 TO MULTITASK SSL

To the best of our knowledge, multi-task speech representation learning has not been scaled to a point where it could represent a state-of-the-art alternative. Contrastive predictive coding (Oord et al., 2018) based techniques like wav2vec 2.0 (Baevski et al., 2020), on the other hand, currently trust most of the leader-boards for speech-related tasks. Recently, Sadhu et al. (2021) showed that combining a consistency loss and contrastive predictive coding improves the results of the wav2vec 2.0 architectures in noisy conditions. Following this idea, we propose to further validate our selection method with an extension of wav2vec 2.0 to multitask SSL to demonstrate its scaling capabilities. Hence, the training loss is extended to:

$$L_{SSL} = L_{W2V} + \sum_{i=1}^k \lambda_i \ell_1(Z_i). \quad (9)$$

We use the standard *BASE* wav2vec 2.0 first described in (Baevski et al., 2020) as a SSL model and train it with the same Common Voice dataset. The pre-training pipeline is implemented within SpeechBrain. The trained *BASE* model has been compared to one obtained with the official Fairseq implementation from Baevski et al. (2020), and results are strictly equivalent. The entire recipe alongside with the large set of hyperparameters needed to properly train a wav2vec 2.0 model are released under our anonymous repository and will be made available with SpeechBrain afterwards.

Table 2: Results observed retraining the Wav2vec2 model with and without weighted pretext tasks using the sparsemax method. “Fr.” and “Fine.” also respectively refer to Frozen and Finetuned settings. Adding selected pretext tasks improves the downstream performance on all three considered tasks. All models contain $100M$ neural parameters.

Selections	LibriSpeech (WER % ↓)		VoxCeleb1 (EER % ↓)		IEMOCAP (Acc % ↑)	
	Fr.	Fine.	Fr.	Fine.	Fr.	Fine.
wav2vec 2.0 <i>BASE</i>	17.93	10.21	7.20	5.35	56.6	74.0
wav2vec 2.0 <i>BASE</i> + multitask SSL	16.70	9.18	6.57	5.30	59.5	74.0

To further investigate our methodology, we extend this analysis to emotion recognition (ER) with the IEMOCAP (Busso et al., 2008) dataset, as prosody has shown important in ER literature (Luengo et al., 2005). We follow the SUPERB benchmark conventions (Yang et al., 2021) both at the data and downstream architecture levels. Hence, and conversely to the previous experiments, the ASR system solely optimises the CTC criterion over characters. For each of the three tasks (*i.e.* ASR, SV, ER) we compare the standard *BASE* wav2vec 2.0 model with one trained following the sparsemax selection of multitask SSL. Sparsemax is chosen over softmax because it enforces the sparsity criterion and removes completely a few pseudo-labels from the training, which is one of the objectives of this work. As for the other experiments, the exact weights of each pseudo-label are reported in Appendix A.9. Each wav2vec 2.0 model required 24 NVIDIA Tesla V100 GPUs to train for 150 epochs (40 hours). Finally, we also propose to compare frozen and unfrozen (*i.e.* where the wav2vec 2.0 encoder is fine-tuned with the downstream task) SSL models.

It is clear from the results reported in Table 2 that our approach improves the performance over the standard wav2vec 2.0 framework for every considered downstream task. Here, it is worth noting that the difference in performance compared to the literature mostly comes from the training conditions. For instance, wav2vec 2.0 is commonly pre-trained with larger models on LibriSpeech to achieve lower WER on this dataset. We decided to avoid this comparison since including the downstream dataset in the pre-training phase would have introduced biases in the evaluation of our method.

7 FURTHER DISCUSSIONS

Computational efficiency. Efficiency is one of the key motivations of this work, and the gain in time observed with our approach is considerable. The K and L matrices used for the CI estimate are only computed for the downstream datasets. But since these datasets may get bigger and bigger, we can sample among the considered downstream classes to keep the computations tractable and quick. For instance, the CI testing of a considered pretext task (*i.e.* pseudo-labels selection) takes less than half an hour on 20 CPUs whether it be for LibriSpeech or VoxCeleb. This is to be compared to 40 hours of GPU training (*i.e.* 24 Nvidia Tesla V100 for wav2vec 2.0 *BASE*) for a pre-training experiment. The exhaustive search for proper pseudo-label weighting is even more dramatically computationally and energy consuming, if we considered s values per pseudo-label in a grid search involving k pseudo-labels, we would need s^k experiments.

Pseudo-label interactions. A thorough analysis of the evolution of the CI estimate as a function of the weights, as given in Appendix A.4.3, provides interesting insights on which pseudo-labels seem to be interchangeable, complementary or even harmful with respect to a downstream task. This suggests that the design of good pseudo-labels and sets of pseudo-labels is not a trivial task.

8 CONCLUSION

In this work, we introduce a method to quickly and simply combine pseudo-labels into a useful pretext task for multitask self-supervised learning settings. Our approach allows for an optimal selection of pseudo-labels following a cheap optimisation process drastically decreasing the time and compute needed to design the best performing multitask SSL model. Our method is validated on three speech-related downstream tasks and outperforms common pseudo-label selection strategies

when combined with simple and state-of-the-art SSL models. This opens a range of possibilities for finding and selecting new pretext tasks in self-supervised learning for speech or other types of data.

REFERENCES

- Robin Algayres, Mohamed Salah Zaiem, Benoît Sagot, and Emmanuel Dupoux. Evaluating the reliability of acoustic speech embeddings. In *INTERSPEECH 2020 - Annual Conference of the International Speech Communication Association*, Shanghai / Virtual, China, October 2020.
- Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- 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.
- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A Theoretical Analysis of Contrastive Unsupervised Representation Learning. *36th International Conference on Machine Learning, ICML 2019*, 2019-June:9904–9923, feb 2019.
- Philip Bachman, R Devon Hjelm, and William Buchwalter. *Learning Representations by Maximizing Mutual Information across Views*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*, 2020.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Ebrahim (Abe) Kazemzadeh, Emily Mower Provost, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. Iemocap: interactive emotional dyadic motion capture database. *Language Resources and Evaluation*, 42: 335–359, 2008.
- Michael Carlin, Samuel Thomas, Aren Jansen, and Hynek Hermansky. Rapid evaluation of speech representations for spoken term discovery. pp. 821–824, 01 2011.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1597–1607. PMLR, 13–18 Jul 2020a. URL <http://proceedings.mlr.press/v119/chen20j.html>.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. 2020b.
- Zhuo Chen, Shinji Watanabe, Hakan Erdogan, and J. Hershey. Speech enhancement and recognition using multi-task learning of long short-term memory recurrent neural networks. In *INTERSPEECH*, 2015.
- Yu-An Chung, Wei-Ning Hsu, Hao Tang, and James Glass. An unsupervised autoregressive model for speech representation learning, 2019.
- Alexis Conneau, Alexei Baevski, Ronan Collobert, Abdelrahman Mohamed, and Michael Auli. Unsupervised cross-lingual representation learning for speech recognition. *CoRR*, abs/2006.13979, 2020. URL <https://arxiv.org/abs/2006.13979>.
- Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised visual representation learning by context prediction, 2016.
- Carl Doersch, Andrew Zisserman, and Deepmind. Multi-task Self-Supervised Visual Learning. Technical report, 2017.
- Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Ves Stoyanov, and Alexis Conneau. Self-training improves pre-training for natural language understanding, 2020.

- Kshitij Dwivedi and Gemma Roig. Representation Similarity Analysis for Efficient Task taxonomy & Transfer Learning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:12379–12388, apr 2019. URL <http://arxiv.org/abs/1904.11740>.
- Solene Evain, Ha Nguyen, Hang Le, Marcely Zanon Boito, Salima Mdhaffar, Sina Alisamir, Ziyi Tong, Natalia Tomashenko, Marco Dinarelli, Titouan Parcollet, Alexandre Allauzen, Yannick Esteve, Benjamin Lecouteux, Francois Portet, Solange Rossato, Fabien Ringeval, Didier Schwab, and Laurent Besacier. Lebenchmark: A reproducible framework for assessing self-supervised representation learning from speech, 2021.
- Florian Eyben, Martin Wöllmer, and Björn Schuller. Opensmile: The munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM International Conference on Multimedia*, MM ’10, pp. 1459–1462, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605589336. doi: 10.1145/1873951.1874246. URL <https://doi.org/10.1145/1873951.1874246>.
- Zhiyun Fan, Meng Li, Shiyu Zhou, and Bo Xu. Exploring wav2vec 2.0 on speaker verification and language identification, 2021.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *CoRR*, abs/1803.07728, 2018. URL <http://arxiv.org/abs/1803.07728>.
- Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8059–8068, 2019.
- Alex Graves. Connectionist temporal classification. In *Supervised Sequence Labelling with Recurrent Neural Networks*, pp. 61–93. Springer, 2012.
- Arthur Gretton, Kenji Fukumizu, Choon Hui Teo, Le Song, Bernhard Schölkopf, and Alexander Smola. A kernel statistical test of independence. 01 2007.
- Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- Michael Gump, Wei-Ning Hsu, and James Glass. Unsupervised Methods for Evaluating Speech Representations. 2020. doi: 10.21437/Interspeech.2020-2990. URL <http://dx.doi.org/10.21437/Interspeech.2020-2990>.
- Isabelle Guyon and André Elisseeff. An introduction of variable and feature selection. *J. Machine Learning Research Special Issue on Variable and Feature Selection*, 3:1157 – 1182, 01 2003. doi: 10.1162/153244303322753616.
- Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46:389–422, 01 2002. doi: 10.1023/A:1012487302797.
- Hynek Hermansky, Nathaniel Morgan, A. Bayya, and P. Kohn. Rasta-plp speech analysis technique. volume 1, pp. 121 – 124 vol.1, 04 1992. ISBN 0-7803-0532-9. doi: 10.1109/ICASSP.1992.225957.
- Nils Holzenberger, Mingxing Du, Julien Karadayi, Rachid Riad, and Emmanuel Dupoux. Learning Word Embeddings: Unsupervised Methods for Fixed-size Representations of Variable-length Speech Segments. In *Interspeech 2018, Proceedings of Interspeech 2018, Hyderabad, India, September 2018*. ISCA. doi: 10.21437/Interspeech.2018-2364.
- Wei-Ning Hsu, Yao-Hung Hubert Tsai, Benjamin Bolte, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: How much can a bad teacher benefit asr pre-training? In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6533–6537, 2021. doi: 10.1109/ICASSP39728.2021.9414460.

- Yun-Ning Hung, Yi-An Chen, and Yi-Hsuan Yang. Multitask learning for frame-level instrument recognition, 2019.
- Sergey Ioffe. Probabilistic Linear Discriminant Analysis. In Aleš Leonardis, Horst Bischof, and Axel Pinz (eds.), *Computer Vision – ECCV 2006*, pp. 531–542, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-33839-0.
- 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.
- Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- 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.
- Dahun Kim, Donghyeon Cho, Donggeun Yoo, and In So Kweon. Learning image representations by completing damaged jigsaw puzzles. *CoRR*, abs/1802.01880, 2018. URL <http://arxiv.org/abs/1802.01880>.
- Kushal Lakhota, Evgeny Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Adelrahman Mohamed, and Emmanuel Dupoux. Generative spoken language modeling from raw audio, 2021.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- Jason D. Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning, 2020.
- Yazhe Li, Roman Pogodin, Danica J. Sutherland, and Arthur Gretton. Self-supervised learning with kernel dependence maximization, 2021.
- 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. doi: 10.1109/icassp40776.2020.9054458.
- Iker Luengo, Eva Navas, Inmaculada Hernáez, and Jon Sánchez. Automatic emotion recognition using prosodic parameters. pp. 493–496, 09 2005. doi: 10.21437/Interspeech.2005-324.
- Christoph Lüscher, Eugen Beck, Kazuki Irie, Markus Kitza, Wilfried Michel, Albert Zeyer, Ralf Schlüter, and Hermann Ney. Rwth asr systems for librispeech: Hybrid vs attention. *Interspeech 2019*, Sep 2019. doi: 10.21437/interspeech.2019-1780. URL <http://dx.doi.org/10.21437/Interspeech.2019-1780>.
- André F. T. Martins and Ramón F. Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML’16*, pp. 1614–1623. JMLR.org, 2016.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. Montreal forced aligner: Trainable text-speech alignment using kaldi. pp. 498–502, 08 2017. doi: 10.21437/Interspeech.2017-1386.
- Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6707–6717, 2020.

- Peter Murphy and Olatunji Akande. Cepstrum-Based Harmonics-to-Noise Ratio Measurement in Voiced Speech. In Gérard Chollet, Anna Esposito, Marcos Faundez-Zanuy, and Maria Marinaro (eds.), *Nonlinear Speech Modeling and Applications*, pp. 199–218, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. ISBN 978-3-540-31886-6.
- Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. Voxceleb: A large-scale speaker identification dataset. *Interspeech 2017*, Aug 2017. doi: 10.21437/interspeech.2017-950.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles, 2017.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An asr corpus based on public domain audio books. pp. 5206–5210, 04 2015. doi: 10.1109/ICASSP.2015.7178964.
- Santiago Pascual, Mirco Ravanelli, Joan Serra, Antonio Bonafonte, and Yoshua Bengio. Learning problem-agnostic speech representations from multiple self-supervised tasks, 2019.
- Vijayaditya Peditinti, Daniel Povey, and S. Khudanpur. A time delay neural network architecture for efficient modeling of long temporal contexts. In *INTERSPEECH*, 2015.
- Hanchuan Peng, Fuhui Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238, 2005. doi: 10.1109/TPAMI.2005.159.
- Alain Rakotomamonjy, Francis Bach, Stephane Canu, and Yves Grandvalet. More efficiency in multiple kernel learning. *Proceedings of the 24th International Conference on Machine Learning (ICML)*, 227, 01 2007. doi: 10.1145/1273496.1273594.
- 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.
- Mirco Ravanelli, Titouan Parcollet, Aku Rouhe, Peter Plantinga, Elena Rastorgueva, Loren Lugosch, Nauman Dawalatabad, Chou Ju-Chieh, Abdel Heba, Francois Grondin, William Aris, Chien-Feng Liao, Samuele Cornell, Sung-Lin Yeh, Hwidong Na, Yan Gao, Szu-Wei Fu, Cem Subakan, Renato De Mori, and Yoshua Bengio. Speechbrain. <https://github.com/speechbrain/speechbrain>, 2021.
- Daniel Renshaw, H. Kamper, A. Jansen, and S. Goldwater. A comparison of neural network methods for unsupervised representation learning on the zero resource speech challenge. In *INTERSPEECH*, 2015.
- Samik Sadhu, Di He, Che-Wei Huang, Sri Harish Mallidi, Minhua Wu, Ariya Rastrow, Andreas Stolcke, Jasha Droppo, and Roland Maas. wav2vec-C: A Self-Supervised Model for Speech Representation Learning. In *Proc. Interspeech 2021*, pp. 711–715, 2021. doi: 10.21437/Interspeech.2021-717.
- Aaqib Saeed, David Grangier, and Neil Zeghidour. Contrastive Learning of General-Purpose Audio Representations. oct 2020.
- Nikunj Saunshi, Sadhika Malladi, and Sanjeev Arora. A mathematical exploration of why language models help solve downstream tasks. *CoRR*, abs/2010.03648, 2020. URL <https://arxiv.org/abs/2010.03648>.
- Thomas Schatz, Vijayaditya Peditinti, Francis Bach, Aren Jansen, Hynek Hermansky, and Emmanuel Dupoux. Evaluating speech features with the Minimal-Pair ABX task: Analysis of the classical MFC/PLP pipeline. In *INTERSPEECH 2013 : 14th Annual Conference of the International Speech Communication Association*, pp. 1–5, Lyon, France, August 2013.

- B. Schuller, B. Vlasenko, R. Minguez, G. Rigoll, and A. Wendemuth. Comparing one and two-stage acoustic modeling in the recognition of emotion in speech. In *2007 IEEE Workshop on Automatic Speech Recognition Understanding (ASRU)*, pp. 596–600, 2007. doi: 10.1109/ASRU.2007.4430180.
- Romain Serizel, Victor Bisot, Slim Essid, and Gael Richard. Acoustic Features for Environmental Sound Analysis. In Tuomas Virtanen, Mark D. Plumbley, and Dan Ellis (eds.), *Computational Analysis of Sound Scenes and Events*, pp. 71–101. Springer International Publishing AG, 2017. doi: 10.1007/978-3-319-63450-0_4. URL <https://hal.archives-ouvertes.fr/hal-01575619>.
- Laurent El Shafey, Hagen Soltau, and Izhak Shafran. Joint speech recognition and speaker diarization via sequence transduction. *CoRR*, abs/1907.05337, 2019. URL <http://arxiv.org/abs/1907.05337>.
- Rajen Shah and Jonas Peters. The hardness of conditional independence testing and the generalised covariance measure. *Annals of Statistics*, 48, 04 2018. doi: 10.1214/19-AOS1857.
- Sekitoshi Shin’ya Yamaguchi, Tetsuya Kanai, Shoichiro Shioda, Ntt Takeda, and Japan Tokyo. Multiple Pretext-Task for Self-Supervised Learning via Mixing Multiple Image Transformations. Technical report.
- Abhinav Shukla, Stavros Petridis, and Maja Pantic. Learning speech representations from raw audio by joint audiovisual self-supervision. 07 2020.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition, 2015.
- D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur. X-vectors: Robust dnn embeddings for speaker recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333, 2018. doi: 10.1109/ICASSP.2018.8461375.
- David Snyder, Daniel Garcia-Romero, and Daniel Povey. Time Delay Deep Neural Network-Based Universal Background Models for Speaker Recognition. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pp. 92–97, 2015. doi: 10.1109/asru.2015.7404779. URL <https://app.dimensions.ai/details/publication/pub.1093368586>.
- Xingchen Song, Guangsen Wang, Zhiyong Wu, Yiheng Huang, Dan Su, Dong Yu, and Helen Meng. Speech-xlnet: Unsupervised acoustic model pretraining for self-attention networks, 2020.
- Sören Sonnenburg, Gunnar Rätsch, Christin Schäfer, and Bernhard Schölkopf. Large scale multiple kernel learning. *J. Mach. Learn. Res.*, 7:1531–1565, December 2006. ISSN 1532-4435.
- Johan Sundberg and Maria Nordenberg. Effects of vocal loudness variation on spectrum balance as reflected by the alpha measure of long-term-average spectra of speech. *The Journal of the Acoustical Society of America*, 120:453–7, 08 2006. doi: 10.1121/1.2208451.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What Makes for Good Views for Contrastive Learning? In H Larochelle, M Ranzato, R Hadsell, M F Balcan, and H Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6827–6839. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/4c2e5eaae9152079b9e95845750bb9ab-Paper.pdf>.
- M. Tschannen, J. Djolonga, P. K. Rubenstein, S. Gelly, and M. Lucic. On mutual information maximization for representation learning. In *8th International Conference on Learning Representations (ICLR)*, April 2020. URL <https://openreview.net/forum?id=rkxoh24FPH>.
- Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E. Gonzalez. Tafe-net: Task-aware feature embeddings for low shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

- Colin Wei, Kendrick Shen, Yining Chen, and Tengyu Ma. Theoretical analysis of self-training with deep networks on unlabeled data. *CoRR*, abs/2010.03622, 2020. URL <https://arxiv.org/abs/2010.03622>.
- Ho-Hsiang Wu, Chieh-Chi Kao, Qingming Tang, Ming Sun, Brian McFee, Juan Pablo Bello, and Chao Wang. Multi-task self-supervised pre-training for music classification, 2021.
- Shuwen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhota, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Ko tik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung yi Lee. Superb: Speech processing universal performance benchmark, 2021.
- Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B*, 68:49–67, 02 2006. doi: 10.1111/j.1467-9868.2005.00532.x.
- Salah Zaiem, Titouan Parcollet, and Slim Essid. Conditional independence for pretext task selection in self-supervised speech representation learning, 2021.
- Amir Roshan Zamir, Alexander Sax, William B. Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. *CoRR*, abs/1804.08328, 2018. URL <http://arxiv.org/abs/1804.08328>.
- 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. oct 2020.

A APPENDIX

A.1 DIFFERENTIABILITY PROOF

We want to show that the utility estimate is differentiable with respect to the weighting parameters $(\lambda_i)_{i \in [0, k]}$. Since the final estimate is a weighted mean of the in-class independent tests, the problem boils down to showing that within a downstream class c , $HSIC_c(X, Z_\lambda)$ is differentiable. Let us recall the definition of the considered quantities:

$$HSIC_c(X, Z_\lambda) = \frac{1}{n_c^2} \text{trace}(K_c H_c L_c H_c) \quad (10)$$

where K_c and H_c are independent of λ and L_c coefficients are defined as:

$$[L_c]_{i,j} = \text{RBF}((Z_\lambda)_i, (Z_\lambda)_j) = \exp\left(\frac{-1}{2\sigma^2} \sum_{h=1}^k \lambda_h \|z_{h,i} - z_{h,j}\|_2^2\right) \quad (11)$$

Therefore for $p \in [0, k]$:

$$\begin{aligned} \frac{\partial HSIC_c(X, Z_\lambda)}{\partial \lambda_p} &= \frac{1}{n_c^2} \sum_{i,j} \frac{\partial(\text{trace}(K_c H_c L_c H_c))}{\partial [L_c]_{i,j}} \frac{\partial [L_c]_{i,j}}{\partial \lambda_p} \\ &= \frac{1}{n_c^2} \sum_{i,j} (H_c^T K_c^T H_c^T)_{i,j} \frac{-\|z_{p,i} - z_{p,j}\|_2 [L_c]_{i,j}}{2\sigma^2} \end{aligned} \quad (12)$$

This allowed us to minimize the conditional-independence based utility estimator according to the weighting values.

A.2 CONSIDERED SIGNAL FEATURES AND DESCRIPTIONS

Table 3 contains the descriptions of the signal features used as pseudo-labels in this work.

Table 3: Candidate speech pseudo-labels and descriptions.

Feature	Description
Loudness	Intensity & approx. loudness
F0	Fundamental Frequency
Voicing	Voicing Decision
Alpha Ratio (Sundberg & Nordenberg, 2006)	Ratio of spectrum intensity % 1000 Hz
Zero Crossing Rate	Zero crossing number per frame
RastaSpec L1Norm	L1 Norm of Rasta Spectrum (Hermansky et al., 1992)
log HNR (Murphy & Akande, 2005)	log of Harmonicity to Noise Ratio

A.3 SPARSEMAX INITIALIZATION

When initialized with random parameters W , and if one parameter is high enough compared to the other, leading with the Sparsemax function to a weighting value close to 1, we observed that the minimization process falls into local minima selecting only one pseudo-label with weight 1. To avoid this, we initialize all the free parameters W with the same unitary value to which we add some Gaussian noise. Hence, $W_{init} = (1) + N(0, \epsilon)$ with $\epsilon = 0.05$.

A.4 TRAINING AND ARCHITECTURES DETAILS

All the considered audio files are sampled at 16kHz. We feed the SSL models with 80-band Mel spectrograms, with 25ms windows and 10ms stride. To every Mel band corresponds a learned vector of size 256 obtained at the output of the SSL model. So if the input spectrogram is of size $(N, 80)$ with N the number of frames, the representation fed to the downstream pipeline is of size $(N, 256)$. All models including SSL and downstream ones are developed with SpeechBrain (Ravanelli et al., 2021).

A.4.1 PRETRAINING OF THE SSL ENCODER.

The encoder is a succession of 2D CNN layers, LSTM layers and a final dense network. This representation is then fed to one dense layer that predict the selected pretext task labels. There are 3 successive CNN blocks containing each 2 CNN layers with kernel size $(3, 3)$ and 128, 200 and 256 channels for each block respectively. No time pooling is performed in order to preserve the input sequence length. 5 bidirectional LSTM layers of size 256 are then stacked. Finally, a MLP with one hidden layer with 256 neurons. The LeakyReLU activation is used across all the layers except for the LSTM. We use a dropout rate of 0.15 during the training. The AdaDelta optimizer is used to update the weights with an initial learning rate of 1.0, $\rho = 0.8$ and $\epsilon = 10^{-8}$. For every experiment, the SSL model is trained for 10 epochs (leading to the convergence of the validation loss).

A.4.2 DOWNSTREAM TRAININGS : FIRST EXPERIMENTS

Speaker recognition details. VoxCeleb1 (Nagrani et al., 2017) is used for the speaker recognition task. The training set contains 148,642 utterances from 1,251 different speakers. To compute the conditional independence estimates while limiting the computational load, we restricted ourselves to the utterances of 50 different speakers (the detailed list is given in the released repository). A standard xvector model (Snyder et al., 2018) is trained following the available VoxCeleb SpeechBrain recipe. The extracted speaker embeddings are tested on the enrol and test splits using PLDA (Ioffe, 2006) as a similarity metric. Performance is reported in terms of equal error rate (EER). While architecture details are given in appendix A.4, it is worth noticing that the whole pipeline is fully integrated to Speechbrain and can thus easily be extended.

We train an embedding model (XVector) until the validation loss converges, on top of the self supervised representations using 5 successive layers of time-delay neural networks (TDNN) (Peddinti et al., 2015). The number of channels is (512, 512, 512, 512, 1500), with kernel sizes of (5, 3, 3, 1, 1) and dilations of (1, 2, 3, 1, 1). The architecture is inspired by successful works on embeddings for speaker recognition (Snyder et al., 2015). The learned embeddings are therefore used on a list of pairs of samples to predict whether they are from the same speaker or not. The details of the recipe

can be found in the given GitHub repository. We train every embedding model on 10 epochs with an Adam Optimizer starting with a learning rate of 0.001 decaying linearly to 0.0001.

Speech recognition details. ASR is conducted with the 100-hour clean subset of the LibriSpeech dataset (Panayotov et al., 2015) to simulate the low-resource scenario commonly encountered with SSL settings. CI estimations are obtained with word-level alignments from the *Montreal Forced Aligner* (McAuliffe et al., 2017). The ASR pipeline follows the LibriSpeech recipe of SpeechBrain (Ravanelli et al., 2021) and therefore contains a CRDNN encoder (*i.e.* CNN, RNN, DNN) trained jointly with CTC (Graves, 2012) and attention (Lüscher et al., 2019) (details in appendix A.4). The decoding process is based on beam-search with and without shallow fusion with a pretrained recurrent language model that is publicly available and obtained from SpeechBrain.² Performance is expressed in word error rate (WER).

The CRDNN starts with three CNN blocks composed each with 2 2D CNN layers, layer-normalisation and (2, 2) maxpooling along the frequency dimension. The filter dimensions for each block are 64, 100, 100. Then, maxpooling of 4 is applied on the time dimension to reduce the sequence length before being fed to the RNN. The latter is made of 5 bidirectional LSTM layers of 1,024 neurons. Finally two dense layers are connected (with batch-normalisation in between). The LeakyReLU activation function is used across all the layers except for the LSTM. A dropout rate of 0.15 is employed with the encoder. The CTC decoder is a simple dense linear layer of size equal to the vocabulary. The vocabulary is obtained with byte pair encoding or sub-words units (BPE) and is of size 1,000. The attentional decoder is a one-layered location-aware GRU (1,024 neurons). Then, a beam search of depth 60 is applied to obtain the output transcripts. The model is trained for 30 epochs. The learning rate (1.0) is multiplied with a factor of 0.8 every time the validation loss is not decreasing to ensure an optimal convergence of all the models.

A.4.3 SUPERB SETTINGS

SUPERB (Yang et al., 2021) is a recent benchmark for self-supervised representations of speech data. We use this benchmark for our experiments in combining wav2vec with our selected pretext tasks. We detail here the downstream models as detailed in the benchmark paper :

Emotion Recognition. IEMOCAP (Busso et al., 2008) is used for the Emotion Recognition (ER) task. 4 classes are considered (neutral, happy, sad, angry), and only the audio data is used. The learned representations are mean-pooled then fed to a final linear classifier to compute a cross-entropy loss. We cross-validate on five folds of the standard splits. The result shown is the average of the five attempts. The evaluation metric is accuracy (ACC).

Automatic Speech Recognition For ASR, the decoder is a vanilla 2-layer 1024-unit BLSTM fed with our self-supervised representations and optimized by CTC loss on characters. We use the same language model for decoding as in the first experiments. LibriSpeech Clean-100 only is used for downstream training.

Speaker Recognition The model and the dataset splits used in the first experiment correspond to the SUPERB ones, so we kept the same settings. The results are therefore comparable.

A.5 INTUITION AROUND THE USE OF CONDITIONAL INDEPENDENCE

To get an intuitive understanding of the motivations of this choice, let us consider the example of image classification as the downstream task, and image colourization as the pretext task. In this case, this pretext task would be suited to the downstream one if the final classification label can help implying the colours. For instance, if there are only two classes "Blue skies" and "Yellow deserts", then colourisation is an interesting pretext task, as knowing the final label helps a lot for the pretext task, independently of the image. However, if all the classes share the same colour palette, colourization may not be an interesting task. (In this simple example, we are ignoring the edge detection aspect of colourization, and only focusing on the colour choice part. Obviously the former aspect plays a big part on why the colourization pretext task has been successful.)

Concerning our estimation method, as the pseudo-labels considered in this work are data features, they are indeed functions of the original data samples. This ensures that the data samples are not

²<https://huggingface.co/speechbrain/asr-crdnn-rnnlm-librispeech>

independent of the pseudo-labels. The idea behind the estimator of conditional independence is that it will test whether this remains true when the considered points share the same downstream class.

A.6 KERNELS USED FOR THE SIMILARITY MATRICES

The computation of the similarity matrices used in our kernel-based independence test, requires fixed-size embeddings for the data speech samples. These embeddings allow the use of classic kernels on top. However, in the case of sequential data, as it is the case with audio/speech signals, one may want to avoid the additional burden of learning fixed-size embeddings (for possibly variable-length audio sequences). One possible solution to this, which we conveniently exploited in our application to speech data (see Section 5) is the Gaussian Downsampling method (Holzenberger et al., 2018) detailed thereafter. In this instance, after the Mel spectrogram extraction, a speech sample is a sequence of varying length input feature vectors. Therefore, to obtain fixed size embeddings aggregating the input frame-wise Mel spectrum vectors into a fixed number N of input vectors, N being a fixed hyper-parameter, we first divide the sequence into N equal length segments. Then, in each segment, a Gaussian average of the input spectra is computed around the center of the considered segment with the standard deviation σ_{gd} being another hyper-parameter. Denoting by D the dimension of the input frame-wise Mel spectrum vectors, this leads, for any speech excerpt, to a $N \times D$ tensor, without any training procedure. As in the work presenting the gaussian downsampling method (Holzenberger et al., 2018), we set $N = 20$ and $\sigma_{gd} = 0.07$. For the RBF kernel on the pseudo-labels mean value per file, we fixed the RBF kernel width to $\sigma = 0.05$.

A.7 LINKS WITH FEATURE SELECTION

We also studied the link between classic feature selection and pretext task selection through two experiments. The first one was made to check how hard it was to estimate the utility of a pseudo-label, we computed the mutual information between the pseudo-labels and the downstream labels, and checked how much it would correlate with downstream performance. It led to very low correlation values, with even changing signs between VoxCeleb and LibriSpeech. This seems to indicate that Mutual Information is not related directly to self-supervision utility.

The maximum relevance minimum redundancy (MRMR) technique (Peng et al., 2005) used as a baseline in this work relies on the Conditional Independence based estimator. It is a close to a naive selection of the best pretext tasks according to the CI based criterion, but it furthermore penalizes the mutual information between the selected pretext tasks. More precisely, we select the group of pseudo-labels $(Z)_i \in [0, p]$ maximizing :

$$Score_{MRMR}(Z) = \frac{-1}{p} \sum_{i \in [0, p]} HSIC(X, Z_i | Y) - \frac{1}{\binom{p}{2}} \sum_{i < j} I(Z_i, Z_j)$$

Recursive Feature Elimination (RFE) (Guyon et al., 2002) relies on a classifier that provides information concerning the importance of a given feature in the decision. This classifier is first trained with the whole set of pseudo-labels as features, and the least important feature is eliminated. The process is repeated until only 4 pseudo-labels are kept. We use the scikit-learn implementation with the C-Support Vector Classification as the classifier providing the feature importance values. We use the default scikit-learn hyperparameters.

These two baselines perform worse than the proposed techniques. This suggests that despite the apparent similarity, feature selection and self-supervision pretext task design do not necessarily involve the same mechanisms.

A.8 PSEUDO-LABELS' INTERACTIONS.

To understand the interactions between pseudo-labels, studying the evolution of the CI estimate as a function of the weights shows which pseudo-labels seem interchangeable, which ones are complementary and which ones seem only harmful to the considered downstream task. Figure 2 shows the CI estimates for weighted combinations of groups of three pseudo-labels. As the weights sum up to one, two pretext tasks' values are shown on the x and y axes, while the value of the remaining one, whose name is in the title, is equal to $1 - x - y$. For instance, at the origin point $(0, 0)$, only

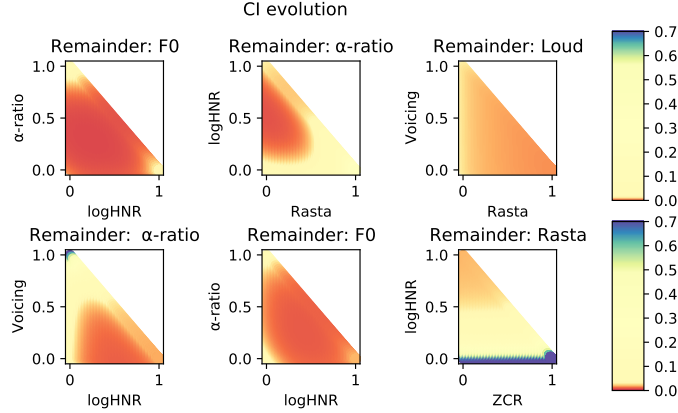


Figure 2: CI-Based utility estimator as a function of the weighting for groups of three pseudo-labels. Top line is for Librispeech, while the bottom one is for VoxCeleb. Three pseudo-labels are presented on every plot, one on the x -axis, one on the y -axis and one that is equal to $1 - x - y$ (hence being called the remainder) and whose name is on the title. Every point in the triangle corresponds to a pretext task that is the weighted combination of the three considered pseudo-labels. For instance, in the top left corner, the point $(0.5, 0.3)$ correspond to the CI value of a pretext task weighting logHNR with 0.5, α -ratio with 0.3 and F0 with 0.2.

Table 4: Weights for every pretext-tasks in every considered experiment. With techniques only leading to a selection of pretext tasks (without weights) a unitary weight is assigned for the selected tasks and zero for the non selected. We can see in this table the zeros induced by the Sparsemax function.

Selection	α -zero	F0	Loudness	Audspec Rasta	ZCR	log HNR	Voicing
All	1	1	1	1	1	1	1
VC RFE	1	1	0	0	1	0	1
VC MRMR	1	0	0	1	0	1	0
VC Sparsemax	0.28	0.26	0	0	0	0.4544	0
VC Softmax	0.27	0.11	0.18	0.04	0.06	0.31	0.03
Libri RFE	1	0	0	0	1	1	1
Libri MRMR	0	1	0	1	0	1	1
Libri Sparsemax	0.30	0.37	0	0.06	0	0.27	0
Libri Softmax	0.28	0.47	0.07	0.04	0.02	0.08	0.04
IEMOCAP Sparse.	0.16	0.22	0	0.14	0.12	0.17	0.19

the third pseudo-label is selected with a weight equal to one, while its weight is equal to zero on the hypotenuse of the right triangle. Figure 2 illustrates that the relationship leading to a lower CI-based utility estimator is not always straightforward. For instance, if we consider the second plot on the second row (*i.e.* α -ratio, F0, logHNR), we can see that selecting only one element is always worse than selecting a weighted concatenation, because the areas around the origin and the points $(1, 0)$ and $(0, 1)$ are brighter than the central area.

A.9 FULL WEIGHTS FOR THE CONSIDERED EXPERIMENTS

Table 4 shows the weights computed using the proposed selection and weighting process for every considered downstream task. It also shows the pretext tasks selected by the baseline methods.