

# LEARNING AUDIO REPRESENTATIONS WITH SELF-SUPERVISION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We explore self-supervision as a way to learn general purpose audio representations. Specifically, we propose self-supervised tasks that exploit the temporal context in the spectrogram domain. The `TemporalGap` task estimates the distance between two short audio segments extracted at random from the same audio clip. The `Audio2Vec` task is inspired by `Word2Vec`, a popular technique used to learn word embeddings, and aim at reconstructing a spectrogram slice from past and future slices or, alternatively, at reconstructing the context of surrounding slices from the current slice. We evaluate the quality of the embeddings produced by the self-supervised learning models, measuring the accuracy of linear classifiers, which receive the embeddings as input and aim at addressing a variety of downstream audio tasks. Our results show that the learned representations partially bridge the performance gap with fully supervised models of similar size, and for some tasks even approach their performance.

## 1 INTRODUCTION

Thanks to advances in supervised audio learning, it is now possible to train models that are able to successfully perform different tasks, including audio annotation (Hershey et al., 2017), music recognition (Arcas et al., 2017), automatic speech recognition (Chan et al., 2016), speaker identification (Matejka et al., 2016), etc. Despite the indisputable success, this approach suffers from three main shortcomings. First, it requires collecting large annotated datasets specific to each task to be solved. Second, separate models are typically trained for each task, making it difficult to reuse computational resources when multiple such models are deployed on a mobile device. Third, inference is performed on device, but model training is still done on the server side using datasets representing surrogate distributions, which might potentially differ from the true data distribution.

Unsupervised learning attempts to overcome these limitations, by making it possible to learn from widely available unlabelled datasets and by learning general purpose representations that can be reused for different downstream tasks. In addition, unsupervised learning lends itself to be deployed on device, where no explicit labeling of the data is available. Therefore, by leveraging the recent advances in federated learning (Bonawitz et al., 2019), it might be possible to distribute the training process across numerous devices, thus training models directly on the true data distribution, while fully preserving users' privacy.

In the area of unsupervised learning, self-supervised learning has emerged as an attractive approach. In a nutshell, an auxiliary task is formulated based on the available unlabelled data and a fully supervised model is trained to solve such a task. The key idea is that, by solving the auxiliary task, the model is also learning some general purpose representations in a lower dimensional embedding space. Therefore, the embedding encoder, e.g., the portion of the model architecture mapping the input data to the embedding space, can be reused as a feature extractor for different downstream tasks.

One of the earliest successes of self-supervised learning was obtained in the context of language models, where `Word2Vec` is used to map one-hot-encoded words to word embeddings (Mikolov et al., 2013). `Word2Vec` can be formulated in two variants: i) continuous bag-of-words (CBoW), or ii) skip-gram. In the former, the model predicts the current word based on the context of surrounding words. In the latter, the model predicts surrounding words given the current word. Recently, a similar approach has been proposed to map speech to fixed-dimensional embeddings (Chung & Glass, 2018;




Figure 1: Overview of the proposed self-supervised learning tasks.

Chung et al., 2018b). The `Speech2Vec` architecture consists of a RNN encoder-decoder which can handle variable-length inputs and outputs.

In this paper we explore self-supervised learning of audio representations. We posit that contextual temporal information can be exploited in the case of general audio signals without resorting to any form of explicit supervision. We argue that solving properly designed tasks that involve the temporal context requires extracting some sort of high level semantic information from the underlying raw data, thus leading to reusable embeddings. In this respect, this paper makes the following main contributions:

- We propose `Audio2Vec`, a self-supervised learning task that is inspired by `Word2Vec`, but applied to audio spectrograms. In the CBoW formulation (Figure 1a) the auxiliary task consists of reconstructing a temporal slice of pre-determined duration from a number of past and future slices. In the skip-gram formulation (Figure 1b) the roles of the target and surrounding slices are reversed.
- We propose `TemporalGap`, a self-supervised learning task that consists of estimating the distance in time between any two pairs of audio segments extracted at random from a longer audio clip (Figure 1c).
- We quantitatively evaluate the quality of the embeddings produced by the feature encoders obtained by solving the aforementioned self-supervised tasks. To this end we consider a wide variety of downstream tasks, including speech, music detection, speaker identification and language identification, among others. Our results show that all self-supervised models are able to partially bridge the accuracy gap with fully supervised models. This is true despite the fact that labels are only used to train simple linear classifiers that receive the embeddings as input.

While self-supervised learning has been widely explored in the context of visual and multi-modal representations, its potential to learn audio representations has been largely left unexplored, with the exception of just a few works (further discussed in Section 2). Therefore, our paper contributes to filling this gap, proposing in Section 3 two self-supervised tasks specifically tailored to the temporal nature of audio data. The proposed methods are compared to unsupervised baselines and fully supervised benchmarks in Section 4. Conclusions and future work are given in Section 5.

## 2 RELATED WORK

The work presented in this paper is related to several different areas that have received attention in the recent literature. In particular, learning representations has been explored for different modalities.

**Learning audio representations:** Unsupervised feature learning can lead to more compact and descriptive representations than traditional handcrafted features, e.g., MFCCs. For example, (Lee et al., 2009) adopt convolutional deep belief networks to learn audio representations, applicable to both speech and music related tasks. More recently, different autoencoder architectures have been explored, e.g., denoising (Xu et al., 2017), convolutional LSTM autoencoders (Meyer et al., 2017) and sequence-to-sequence autoencoders (Chung et al., 2016). A self-supervised version of the triplet loss is proposed in (Jansen et al., 2018). In the absence of labels, the authors create anchor-positive pairs by adding noise, shifting in time and/or frequency, and sampling temporal neighbors. When

tested on AudioSet (Gemmeke et al., 2017), the self-supervised embeddings partially bridge the gap between a simple log spectrogram baseline and a fully supervised classifier.

**Learning visual representations:** Several auxiliary tasks have been explored to learn image representations, e.g., predicting the relative position of a pair of patches extracted from the same image (Doersch et al., 2015), re-ordering image patches and solving jigsaw puzzles (Noroozi & Favaro, 2016), or asking the model to discriminate between a patch and transformed version of it (Dosovitskiy et al., 2016). In some cases solving seemingly simple tasks can lead to very powerful representations such as, for example, detecting image rotations (Gidaris et al., 2018). In other cases, representations can be learned as a by-product of solving useful tasks, e.g., in the case of image colorization (Zhang et al., 2016) and image inpainting (Pathak et al., 2016). The latter is to some extent similar to our work, since the CBoW version of `Audio2Vec` can be seen as a form of inpainting in the spectrogram domain. The representations learned by different self-supervised learning tasks can also be combined to obtain a single representation, as presented in (Doersch & Zisserman, 2017). In the case of video, it is possible to exploit the temporal dimension to learn visual representations by asking a model to learn whether frames are in the correct temporal order (Misra et al., 2016; Fernando et al., 2017), to infer motion by observing a static image (Pathak et al., 2017), or detect whether a video is playing forwards or backwards (Wei et al., 2018).

**Learning multimodal representations:** Several papers have recently investigated learning audio representations exploiting the correlation with other modalities, e.g., text (Chung et al., 2018a), images (Owens et al., 2018) and videos (Owens & Efros, 2018; Gao et al., 2018; Arandjelović & Zisserman, 2018; Korbar et al., 2018; Cramer et al., 2019).

**Contextual predictions:** After the seminal work on `Word2Vec`, contextual prediction has been successfully explored as a means for learning representations in other modalities, e.g., in the case of image inpainting (Pathak et al., 2016), symbolic music prediction (Bretan et al., 2017), and speech (Chung & Glass, 2018). Recently (van den Oord, Yazhe Li, 2019) proposed to use contrastive predictive coding, i.e., predicting future samples directly in the embedding space, reporting promising results also in the case of audio-based tasks. Our work is mostly related to this strand of research, in that we evaluate contextual prediction for general audio-based tasks, but we put particular emphasis on learning models that can be deployed on device.

### 3 METHODS

Let  $x = \{x_1, x_2, \dots, x_n\}$  denote an audio clip of  $n$  samples in the time domain and  $X \in \mathbb{R}^{T \times F}$  the corresponding mel spectrogram, which consists of  $T$  temporal frames and  $F$  frequency bins. Note that we compute the logarithm of the modulus of the spectrogram to compress the dynamic range of the amplitudes. Let  $X_i$  denote a  $N \times F$  slice of the spectrogram  $X$ , starting at frame  $i$  with  $N < T$  temporal frames and  $z_i = Enc(X_i)$  a  $d$ -dimensional embedding computed by processing the input spectrogram  $X_i$  with an encoder  $Enc()$ , whose architecture is detailed in Section 4. Using this notation, in the following we describe the proposed self-supervised learning models.

**Audio2Vec (CBoW):** The first self-supervised learning task that we propose, `Audio2Vec`, comes in two variants. In the CBoW variant, we first select a target slice at random, together with a set of surrounding slices used for prediction. Each of the predictor slices is processed by the same encoder, which maps its input into a fixed-dimensional embedding. These embeddings are then concatenated and fed into a decoder, mimicking the same architecture as the encoder, which computes a reconstruction of the target slice. More specifically, let  $X_{(0)} = X_i$  be a slice selected at random from  $X$ . Then, a set of past ( $X_{(-P)}, \dots, X_{(-1)}$ ) and future slices ( $X_{(1)}, \dots, X_{(P)}$ ) are extracted from the same audio clip. The temporal location of the slice  $X_{(p)}$  is equal to  $X_{i+p(N+G)}$ , i.e., we consider non-overlapping slices of size  $N$ , with an extra gap of  $G$  temporal frames between any two consecutive slices. The gap is introduced to avoid that the self-supervised model exploits the leakage between adjacent STFT temporal frames as a *shortcut* to solve the task. Each slice is processed by the same encoder to obtain  $z_{(p)} = Enc(X_{(p)})$ . Then, a vector  $z_{(0)} = [z_{(-P)}, \dots, z_{(-1)}, z_{(1)}, \dots, z_{(P)}]$  is obtained by concatenating the embeddings of each of the predictor slices and fed into a convolutional decoder to obtain a reconstruction  $\hat{X}_{(0)} = Dec(z_{(0)})$ . Note that the architecture of the decoder is obtained by reversing the order of the layers in the encoder and replacing max-pooling with nearest-neighbor upsampling. The overall encoder-decoder architecture is trained end-to-end by minimizing the mean-square error loss function  $\|X_{(0)} - \hat{X}_{(0)}\|$ .

**Audio2Vec (skip-gram):** The skip-gram variant of `Audio2Vec` uses a similar architecture. In this case we compute the embeddings of the middle slice  $z_{(0)} = \text{Enc}(X_{(0)})$ , and then let the decoder reconstruct the surrounding slices, i.e.,  $[\hat{X}_{(-P)}, \dots, \hat{X}_{(-1)}, \hat{X}_{(1)}, \dots, \hat{X}_{(P)}] = \text{Dec}(z_{(0)})$ . The decoder is identical to the one used by the CBoW variant, except for one important difference: the last convolutional layer has  $2P$  output channels, one for each of the slices to be reconstructed. The loss function minimizes the average mean-square error computed across the  $2P$  reconstructed slices.

**Temporal gap:** For the `TemporalGap` task, we ask the model to estimate the absolute value of the distance in time between two slices sampled at random from the same audio clip. More specifically, we sample the ground truth temporal gap from a uniform distribution, i.e.,  $\Delta \sim \mathcal{U}(0, N_{max} - N)$ , where  $N$  and  $N_{max}$  are the lengths (in time frames) of the slices and the original sample, respectively, and define the normalized temporal gap as  $\delta = \Delta / (N_{max} - N) \in [0, 1]$ . Then, we extract two slices  $X_i$  and  $X_j$  such that  $\Delta = |i - j|$ . Note that we do not impose a temporal order between the two slices. We concatenate the embedding representations in a single  $2d$ -dimensional vector  $z = [\text{Enc}(X_i), \text{Enc}(X_j)]$  and we feed this vector into a fully connected feed forward network with a single hidden layer of size 64 that produces the scalar output  $\hat{\delta}$ . We train the model end-to-end so as to minimize a cross-entropy loss  $\mathcal{L}_{CE}(\delta, \hat{\delta})$  between the ground-truth and the predicted gap. In our experiments, we found that this loss is to be preferred to the mean-square error  $\|\delta - \hat{\delta}\|$ , presumably because it gives more weight to errors when the ground truth gap  $\delta$  is small.

## 4 EXPERIMENTS

To evaluate the quality of the embeddings produced by different self-supervised learning methods, it is interesting to observe how the learned representations transfer to several, potentially heterogeneous, downstream tasks. Therefore, for each downstream task, we train a simple fully supervised head. The head consists of a linear classifier that receives the embeddings as input, and produces as output one of the task classes. This is compared with a fully supervised model of similar size, in which the parameters of both the encoder and the head can be trained having access to labelled data. In the following we describe the datasets used in our evaluation campaign and the baselines to which we compare our results.


**Datasets:** We use *AudioSet* (Gemmeke et al., 2017) to train all the self-supervised learning tasks, unless stated otherwise. *AudioSet* contains excerpts of 10 seconds from the soundtracks of YouTube videos. Although the dataset is annotated with labels of more than 500 classes, we discard them in our study. Note that each *AudioSet* sample can be potentially reused multiple times during training, each time extracting a different target slice (together with surrounding slices) uniformly at random.

We use six publicly available datasets to evaluate a variety of downstream tasks, covering both speech and non-speech related tasks. We use the *Speech Commands* dataset (Warden, 2018) to evaluate keyword spotting on 35 distinct keywords. *LibriSpeech* (Panayotov et al., 2015) contains audio books read by 251 different speakers. We use the 100 hours training set to evaluate a speaker identification task. The *Spoken Language Identification* dataset (Oponowicz, 2018) contains samples that belong to three different languages: English, Spanish and German, while the *MUSAN* dataset (Snyder et al., 2015) distinguishes across three classes, namely music, speech and noise. Finally, we use two datasets released in the context of the recent DCASE2018 Challenge, *Bird Audio Detection* (Stowell et al., 2018) and *TUT Urban Acoustic Scenes 2018* (Mesaros et al., 2018), which contains labeled audio samples from 10 different urban environments. To the best of the authors' knowledge, this is the first paper that comprehensively evaluates the quality of self-supervised learning for audio on such a wide variety of tasks. Note that we deliberately avoid exploiting during training any of the datasets used for evaluating the downstream tasks.

Since each dataset is characterized by samples having different durations, during training we preprocess the downstream datasets extracting equal-length slices uniformly at random from the original sample and assign the corresponding label to all of the extracted slices. We consider input samples having the duration of  $T = 975$  ms, so as to match the size of the temporal slices used when training the self-supervised tasks. During evaluation, we apply a sliding window of size  $T$  and a hop size of  $T/2$ , so as to one or more predictions for each input samples, depending on its length. In order to aggregate such predictions and produce a single output for each sample, we apply a simple naive-Bayes classifier. Note that the choice of the tasks used for the evaluation is consistent with the

Table 1: Encoder architecture. Size of activations, number of parameters and FLOPs.

	Output Size	Num. params	FLOPs
Input layer	$96 \times 64 \times 1$	-	-
Conv. layer 1	$48 \times 32 \times 8$	0.2k	2.9M
Conv. layer 2	$24 \times 16 \times 16$	1k	4M
Conv. layer 3	$12 \times 8 \times 32$	5k	4M
Conv. layer 4	$6 \times 4 \times 64$	20k	3.9M
Conv. layer 5	$3 \times 2 \times 128$	82k	3.9M
FC layer	$1 \times 1 \times 128$	16k	33k
Total	-	125k	18.7M

Figure 2: (a): Training loss for the Audio2Vec (skip-gram) model and corresponding accuracy on the *MUSAN* downstream dataset. (b) Accuracy on *MUSAN* datasets for all models under evaluation.

selected temporal granularity. As such, we do not consider automatic speech recognition tasks, which generally require processing much shorter temporal slices.

**Encoder architecture:** In our work we consistently use the same audio frontend, which processes input sequences sampled at 16 kHz, with a window size of 25 ms and a hop size equal to 10 ms to compute the short-time Fourier transform (STFT), and then computes  $F = 64$  mel-spaced frequency bins in the range 60–7800 Hz. For the encoder  $Enc()$ , we use a convolutional neural network, whose architecture is described in Table 1. We also repeated the experiments using a larger architecture, which is described together with the corresponding results at the end of this section. Each convolutional layer consists of a series of two convolutions, one along the time axis (with size  $3 \times 1 \times C_{in} \times C_{out}$ ) and one along the frequency axis (with size  $1 \times 3 \times C_{in} \times C_{out}$ ), in parallel to a pointwise  $1 \times 1$  convolution as a residual connection. All activation functions are ReLUs and batch normalization is used in all convolutional layers. Each layer is followed by max-pooling, to reduce the time-frequency dimensions by a factor of two at each layer. Finally, a global max-pooling layer produces a  $d$ -dimensional vector, which is further processed by a fully-connected layer to get the embeddings. By default, we set  $N = 96$  (corresponding to 975 ms) and  $d = 128$ , thus reducing the dimensionality of the raw audio samples by a factor of about 122.

**Models:** We compare different self-supervised models trained on *AudioSet*: Audio2Vec, in its two variants, CBoW and skip-gram, and TemporalGap. For Audio2Vec we use  $P = 2$  slices on each side of the target, and a gap of  $G = 2$  temporal frames between consecutive slices. We also include the TripletLoss methods proposed in (Jansen et al., 2018) in our evaluations. More specifically, positive/negative pairs are obtained by extracting a slice from, respectively, the same or a different original sample. In addition, we also train an AutoEncoder sharing the same encoder and decoder architectures as Audio2Vec. We tested different variants, including denoising and variational autoencoders, but we did not observe significant differences with respect to the default autoencoder. When evaluating the accuracy in downstream tasks, we extract the portion of the model corresponding to the encoder and use it to map input log-mel spectrograms to 128-dimensional embeddings.

We compare our results to two different fully supervised baselines based on a simple linear classifier model: i) the Spectrogram model receives directly the (flattened) spectrogram features as input; ii) the Untrained model computes the embeddings with the very same encoder architecture described in Section 3, but using randomly initialized weights.

Table 2: Accuracy on downstream tasks (and fraction of accuracy recovered wrt. baselines). Downstream tasks: SPC: (*Speech Commands*), LSP: (*LibriSpeech*), TUT: *TUT Urban Acoustic Scenes 2018*, MUS: *MUSAN*, BSD: *Bird Audio Detection*, LID: *Spoken Language Identification*. In **bold** the highest accuracy attained by self-supervised models for each task.

model	SPC	LID	LSP	MUS	TUT	BSD
Spectrogram	0.16 ± .01 (+0%)	0.28 ± .04 (+0%)	0.97 ± .01 (+0%)	0.74 ± .01 (+0%)	0.36 ± .03 (+0%)	0.65 ± .02 (+0%)
Untrained	0.16 ± .01 (-1%)	0.48 ± .04 (+33%)	0.54 ± .02 (-1338%)	0.93 ± .00 (+77%)	0.57 ± .03 (+35%)	0.70 ± .02 (+31%)
AutoEncoder	0.28 ± .01 (+21%)	<b>0.64</b> ± .04 (+56%)	0.99 ± .00 (+55%)	0.94 ± .00 (+81%)	0.59 ± .03 (+38%)	0.69 ± .02 (+27%)
A2V (CBoW)	<b>0.30</b> ± .01 (+23%)	0.57 ± .04 (+47%)	0.99 ± .00 (+82%)	<b>0.98</b> ± .00 (+97%)	0.66 ± .03 (+50%)	0.71 ± .01 (+40%)
A2V (SG)	0.28 ± .01 (+21%)	0.55 ± .04 (+44%)	<b>1.00</b> ± .00 (+85%)	<b>0.98</b> ± .00 (+98%)	0.67 ± .03 (+52%)	0.69 ± .02 (+28%)
TemporalGap	0.23 ± .01 (+12%)	0.45 ± .04 (+27%)	0.97 ± .01 (+11%)	0.97 ± .00 (+92%)	0.63 ± .03 (+44%)	0.71 ± .01 (+44%)
TripletLoss	0.18 ± .01 (+3%)	0.62 ± .04 (+55%)	<b>1.00</b> ± .00 (+96%)	0.97 ± .00 (+95%)	<b>0.73</b> ± .03 (+61%)	<b>0.73</b> ± .01 (+55%)
MultiHead	0.72 ± .01 (+95%)	0.82 ± .03 (+88%)	1.00 ± .00 (+99%)	0.98 ± .00 (+95%)	0.94 ± .02 (+96%)	0.78 ± .01 (+90%)
Supervised	0.75 ± .01 (+100%)	0.90 ± .03 (+100%)	1.00 ± .00 (+100%)	0.99 ± .00 (+100%)	0.97 ± .01 (+100%)	0.79 ± .01 (+100%)

Since each task is characterized by a different number of target classes and intrinsic difficulty, we compare the accuracy to the level attained by task-specific fully supervised models (`Supervised`), each using the same encoder, but trained end-to-end on each of the labeled downstream datasets. In addition, we also trained a `MultiHead` model, where a single shared encoder is composed with a different fully connected layer for each downstream task. This provides an upper bound for the best performance we could expect, as it uses the same architecture as when using the self-supervised embeddings, but leverages the in-domain labeled data for end-to-end training.

All models are trained with stochastic gradient descent and Adam optimizer with default hyperparameters. The learning rate was set to  $10^{-3}$  for `Audio2Vec`, `AutoEncoder`, and all the supervised models, while it was set to  $10^{-4}$  for `TemporalGap` and `TripletLoss`. We use a mini-batch size equal to 256 and we stop training after approximately 2 days (on five Tesla V100 GPUs), thus iterating between 1.3 and 1.8 million mini-batches. In most cases, the accuracy of downstream tasks saturated after iterating over 500k mini-batches.

**Main results:** In our results we report the prediction accuracy on the eval set of each of the six datasets, averaging the values observed by repeating the training of the head five times. During training we monitor both the loss of the self-supervised task as well as the accuracy on each of the downstream tasks. As an illustration, Figure 2a shows that the accuracy of the *MUSAN* downstream task increases as the reconstruction loss of `Audio2Vec` (skip-gram) decreases, and both tend to saturate after approximately 300k iterations. For the same dataset, Figure 2b shows that all self-supervised methods attain a level of accuracy that is in-between the baselines and the fully supervised benchmarks, with `Audio2Vec` (skip-gram) outperforming the other models on this task. We repeated the evaluation on all downstream tasks and show the results in Table 2. We report the level of accuracy, with 95% confidence intervals capturing the uncertainty due to the finite size of the evaluation datasets. In brackets we also report the accuracy normalized between 0% (`Spectrogram`) and 100% (`Supervised`). We observe that the proposed self-supervised learning models are able to recover between 11% and 98% of the accuracy of the `Supervised` model. Generally, `Audio2Vec` (skip-gram) and `TripletLoss` seem to outperform other self-supervised models. The best results are obtained on *MUSAN* and *LibriSpeech*, presumably because these tasks require to capture relatively stationary spectral characteristics of the inputs. Conversely, all self-supervised models achieve relatively poor performance on the *Speech Commands* dataset. This might be explained by the fact that for this dataset it is particularly important to recognize the

Table 3: Accuracy obtained when training self-supervised models on *LibriSpeech* (and the relative difference with respect to training on *AudioSet*). Red indicates a decrease in the level of accuracy.

model	SPC	LID	LSP	MUS	TUT	BSD
AutoEncoder	0.27	0.65	0.96	0.87	0.56	0.67
	(-3%)	(+1%)	(-3%)	(-7%)	(-5%)	(-2%)
A2V (CBoW)	<b>0.26</b>	0.55	0.99	0.96	0.65	0.70
	(-13%)	(-3%)	(+0%)	(-2%)	(-1%)	(-1%)
A2V (SG)	0.23	<b>0.65</b>	0.99	<b>0.97</b>	<b>0.66</b>	<b>0.71</b>
	(-17%)	(+18%)	(-1%)	(-1%)	(-1%)	(+2%)
TemporalGap	0.18	0.55	0.93	0.94	0.59	0.64
	(-21%)	(+22%)	(-4%)	(-3%)	(-6%)	(-9%)
TripletLoss	0.10	0.34	<b>1.00</b>	0.93	0.56	0.65
	(-44%)	(-45%)	(+0%)	(-4%)	(-23%)	(-10%)

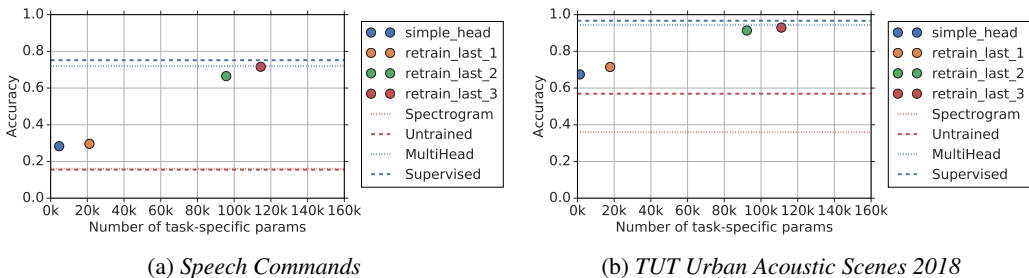


Figure 3: Accuracy obtained when retraining the last layers of the Audio2Vec (skip-gram) encoder.

non-stationary variation of the spectral features along the temporal dimension, which does not seem to be captured by the embeddings generated by the self-supervised models. Note that the different self-supervised models might be capturing different characteristics of the underlying audio data. Therefore, there might be the possibility of merging the different representations, as recently proposed in (Pascual et al., 2019) for the case of speech embeddings.

**Impact of training dataset:** All the results reported so far use the *AudioSet* dataset to train the self-supervised models. *AudioSet* contains a wide variety of audio clips, including music, speech, ambient noise, acoustic events, etc. In order to evaluate the impact of the choice of the dataset, we repeated self-supervised training using *LibriSpeech* (discarding the speaker labels). We chose *LibriSpeech* because the original samples are sufficiently long to support our self-learning tasks and because it contains audio of different content than *AudioSet* (i.e., speech only). Table 3 reports how the evaluation results shown in Table 2 change when training the self-supervised models on *LibriSpeech* instead of *AudioSet*. In most cases, we observe a decrease in the level of accuracy on downstream tasks, especially for *TemporalGap* and *TripletLoss*, suggesting that a richer content variety in the training set is preferable when learning general-purpose audio representations.

**Encoder fine-tuning:** So far we considered the case in which the encoder is shared completely across different tasks, and only the last layer is allowed to learn task-specific parameters. It is interesting to observe what happens when we relax this assumption, allowing to retrain one (or more) of the deepest layers of the encoder. Figure 3 shows the trade-off between the level of accuracy and the number of task specific parameters for two datasets, *Speech Commands* and *TUT Urban Acoustic Scenes 2018*, for which Audio2Vec (skip-gram) was able to only partially bridge the accuracy gap with respect to the Supervised model. The left-most (blue) point corresponds to the accuracy already reported in Table 2. Note that in this case the number of task-specific parameters is equal to  $128 \times C$ , where  $C$  is the number of classes (equal to, respectively, 35 and 10 for these datasets). The second (orange) point from the left corresponds to retraining the fully-connected layer, while the remaining points correspond to retraining until the fifth and fourth convolutional layers included. Generally, retraining the last two layers is needed to recover most of the accuracy of the fully supervised model. Note that, although the last two layers account for approximately 80% of the parameters, they only



Table 4: Accuracy obtained when using a larger encoder architecture (relative change wrt. Table 2).

model	SPC	LID	LSP	MUS	TUT	BSD
AutoEncoder	0.35	0.62	<b>1.00</b>	0.96	0.65	0.70
	(+24%)	(-3%)	(+1%)	(+2%)	(+10%)	(+1%)
A2V (SG)	<b>0.46</b>	<b>0.81</b>	<b>1.00</b>	<b>0.99</b>	0.78	<b>0.76</b>
	(+64%)	(+47%)	(+0%)	(+1%)	(+16%)	(+10%)
TemporalGap	0.37	0.77	<b>1.00</b>	0.98	0.73	0.74
	(+60%)	(+71%)	(+3%)	(+1%)	(+15%)	(+4%)
TripletLoss	0.30	0.73	<b>1.00</b>	<b>0.99</b>	<b>0.81</b>	0.76
	(+66%)	(+17%)	(+0%)	(+2%)	(+10%)	(+4%)

Table 5: Accuracy using a head with two fully connected layers (relative change wrt. Table 2).

model	SPC	LID	LSP	MUS	TUT	BSD
AutoEncoder	0.46	0.85	1.00	0.98	0.83	0.76
	(+58%)	(+32%)	(+2%)	(+4%)	(+45%)	(+11%)
A2V (SG)	0.33	0.60	1.00	0.99	0.89	0.76
	(+13%)	(+9%)	(+1%)	(+1%)	(+30%)	(+7%)
TemporalGap	0.28	0.60	0.99	0.99	0.83	0.76
	(+16%)	(+1%)	(+2%)	(+2%)	(+31%)	(+7%)
TripletLoss	0.22	0.37	1.00	0.99	0.92	0.76
	(+9%)	(-44%)	(+0%)	(+2%)	(+27%)	(+5%)

contribute to 20% of the FLOPs, and this could be particularly useful when deploying on mobile devices.

**Impact of encoder architecture size:** We repeated our evaluation by increasing the size of the encoder architecture described in Table 1. Namely, we increased the number of channels in each convolutional layer by a factor of 4, and we increased the number of outputs in the last fully connected layer to obtain 256-dimensional embeddings. Table 4 shows that the accuracy on downstream tasks increases, and `Audio2Vec` (skip-gram) achieves the highest accuracy on almost all datasets.

**Beyond linear separability:** All the results presented so far are obtained by training a linear classifier on top of the encoder that produces the embeddings. We repeated our experiments training a more complex head, which includes an additional fully connected layer (followed by a ReLU), to see if downstream classes can be better disentangled using a non-linear model. The results are presented in Table 5. Generally, a higher level of accuracy is achieved in this case (with one exception, which could be explained with overfitting to the relatively small dataset). Moreover, it is interesting to observe that the `AutoEncoder` model benefits more than the other self-supervised models.

## 5 CONCLUSION

In this paper we present self-supervised learning methods that exploit the temporal context in audio clips. Our results show that both `Audio2Vec` and `TemporalGap` are able to produce representations that can be re-used for different downstream tasks, without having access to labelled datasets during training. In our future work we will investigate training self-supervised models directly on device in a distributed fashion, by taking advantage of federated learning. Another interesting direction is merging representations learned by different self-supervised models, as recently proposed in (Pascual et al., 2019) for the case of speech embeddings.

## REFERENCES

Relja Arandjelović and Andrew Zisserman. Objects that Sound. In *European Conference on Computer Vision (ECCV)*, dec 2018. URL <http://arxiv.org/abs/1712.06651>.



- Blaise Agüera y Arcas, Beat Gfeller, Ruiqi Guo, Kevin Kilgour, Sanjiv Kumar, James Lyon, Julian Odell, Marvin Ritter, Dominik Roblek, Matthew Sharifi, and Mihajlo Velimirović. Now Playing: Continuous low-power music recognition. Technical report, nov 2017. URL <http://arxiv.org/abs/1711.10958>.
- Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, H. Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. Towards Federated Learning at Scale: System Design. Technical report, feb 2019. URL <http://arxiv.org/abs/1902.01046>.
- Mason Bretan, Sageev Oore, Doug Eck, and Larry Heck. Learning and Evaluating Musical Features with Deep Autoencoders. Technical report, jun 2017. URL <http://arxiv.org/abs/1706.04486>.
- William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals. Listen, Attend and Spell. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, aug 2016. URL <http://arxiv.org/abs/1508.01211>.
- Yu-An Chung and James Glass. Speech2Vec: A Sequence-to-Sequence Framework for Learning Word Embeddings from Speech. In *Proc. Interspeech*, pp. 811–815, mar 2018. URL <http://arxiv.org/abs/1803.08976>.
- Yu-An Chung, Chao-Chung Wu, Chia-Hao Shen, Hung-Yi Lee, and Lin-Shan Lee. Audio Word2Vec: Unsupervised Learning of Audio Segment Representations using Sequence-to-sequence Autoencoder. Technical report, mar 2016. URL <http://arxiv.org/abs/1603.00982>.
- Yu-An Chung, Wei-Hung Weng, Schrasing Tong, and James Glass. Unsupervised Cross-Modal Alignment of Speech and Text Embedding Spaces. In *Neural Information Processing Systems (NeurIPS)*, 2018a. URL <http://arxiv.org/abs/1805.07467>.
- Yu-An Chung, Wei-Hung Weng, Schrasing Tong, and James Glass. Unsupervised Cross-Modal Alignment of Speech and Text Embedding Spaces. In *Neural Information Processing Systems (NeurIPS)*, may 2018b. URL <http://arxiv.org/abs/1805.07467>.
- Jason Cramer, Ho-Hsiang Wu, Justin Salamon, and Juan Pablo Bello. Look, Listen and Learn More: Design Choices for Deep Audio Embeddings. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2019. URL <https://github.com/marl/l3embedding>.
- Carl Doersch and Andrew Zisserman. Multi-task Self-Supervised Visual Learning. In *International Conference on Computer Vision*, pp. 2051–2060, aug 2017. URL <http://arxiv.org/abs/1708.07860>.
- Carl Doersch, Abhinav Gupta, and Alexei A. Efros. Unsupervised Visual Representation Learning by Context Prediction. In *IEEE International Conference on Computer Vision (ICCV)*, pp. 1422–1430, may 2015. URL <http://arxiv.org/abs/1505.05192>.
- Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative Unsupervised Feature Learning with Exemplar Convolutional Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1734–1747, jun 2016. URL <http://arxiv.org/abs/1406.6909>.
- Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-Supervised Video Representation Learning With Odd-One-Out Networks. In *Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 3636–3645, nov 2017. URL <http://arxiv.org/abs/1611.06646>.
- Ruohan Gao, Rogerio Feris, and Kristen Grauman. Learning to Separate Object Sounds by Watching Unlabeled Video. In *European Conference on Computer Vision*, pp. 35–53, apr 2018. URL <http://arxiv.org/abs/1804.01665>.

- Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. Audio Set: An ontology and human-labeled dataset for audio events. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 776–780. IEEE, mar 2017. ISBN 978-1-5090-4117-6. doi: 10.1109/ICASSP.2017.7952261. URL <http://ieeexplore.ieee.org/document/7952261/>.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations. In *International Conference on Learning Representations (ICLR)*, mar 2018. URL <http://arxiv.org/abs/1803.07728>.
- Shawn Hershey, Sourish Chaudhuri, Daniel P.W. Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, Malcolm Slaney, Ron J Weiss, and Kevin Wilson. CNN architectures for large-scale audio classification. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 131–135, 2017. ISBN 9781509041176. doi: 10.1109/ICASSP.2017.7952132.
- Aren Jansen, Manoj Plakal, Ratheet Pandya, Daniel P. W. Ellis, Shawn Hershey, Jiayang Liu, R. Channing Moore, and Rif A. Saurous. Unsupervised Learning of Semantic Audio Representations. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 126–130, nov 2018. URL <http://arxiv.org/abs/1711.02209>.
- Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative Learning of Audio and Video Models from Self-Supervised Synchronization. In *Neural Information Processing Systems (NIPS)*, 2018.
- Honglak Lee, Yan Largman, Peter Pham, and Andrew Y Ng. Unsupervised feature learning for audio classification using convolutional deep belief networks. In *Neural Information Processing Systems (NIPS)*, 2009. URL <https://ai.stanford.edu/~ang/papers/nips09-AudioConvolutionalDBN.pdf>.
- Pavel Matejka, Ondrej Glembek, Ondrej Novotny, Oldrich Plchot, Frantisek Grezl, Lukas Burget, and Jan Honza Cernocky. Analysis of DNN approaches to speaker identification. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 5100–5104. IEEE, mar 2016. ISBN 978-1-4799-9988-0. doi: 10.1109/ICASSP.2016.7472649. URL <http://ieeexplore.ieee.org/document/7472649/>.
- Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen. Detection and Classification of Acoustic Scenes and Events. In *Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*, pp. 9–13, 2018. doi: 10.5281/zenodo.1228142. URL <https://doi.org/10.5281/zenodo.1228142>, .
- Matthias Meyer, Jan Beutel, and Lothar Thiele. Unsupervised Feature Learning for Audio Analysis. In *Workshop track - ICLR*, 2017. URL <http://people.ee.ethz.ch/matthmey/>.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. In *International Conference on Learning Representations (ICLR)*, jan 2013. URL <http://arxiv.org/abs/1301.3781>.
- Ishan Misra, C. Lawrence Zitnick, and Martial Hebert. Shuffle and Learn: Unsupervised Learning using Temporal Order Verification. In *European Conference on Computer Vision*, mar 2016. URL <http://arxiv.org/abs/1603.08561>.
- Mehdi Noroozi and Paolo Favaro. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In *European Conference on Computer Vision (ECCV)*, pp. 69–84, mar 2016. URL <http://arxiv.org/abs/1603.09246>.
- Tomasz Oponowicz. Spoken Language Identification, 2018. URL <https://www.kaggle.com/toponowicz/spoken-language-identification>.
- Andrew Owens and Alexei A Efros. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. In *European Conference on Computer Vision (ECCV)*, pp. 631–648, 2018. URL <http://andrewowens.com/multisensory>.

- Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Learning Sight from Sound: Ambient Sound Provides Supervision for Visual Learning. *International Journal of Computer Vision*, 126(10):1120–1137, 2018. URL <https://arxiv.org/pdf/1712.07271.pdf>.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. LibriSpeech: An ASR Corpus Based on Public Domain Audio Books. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2015. URL <http://www.gutenberg.org>.
- Santiago Pascual, Mirco Ravanelli, Joan Serra, Antonio Bonafonte, and Yoshua Bengio. Learning Problem-agnostic Speech Representations from Multiple Self-supervised Tasks. Technical report, 2019. URL <https://github.com/santi-pdp/pase>.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context Encoders: Feature Learning by Inpainting. In *Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 2536–2544, apr 2016. URL <http://arxiv.org/abs/1604.07379>.
- Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning Features by Watching Objects Move. In *Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 2701–2710, dec 2017. URL <http://arxiv.org/abs/1612.06370>.
- David Snyder, Guoguo Chen, and Daniel Povey. MUSAN: A Music, Speech, and Noise Corpus. Technical report, 2015. URL <http://www.itl.nist.gov/iad/mig/tests/sre/>.
- Dan Stowell, | Mike Wood, | Hanna Pamuła, Yannis Stylianou, and Hervé Glotin. Automatic acoustic detection of birds through deep learning: the first Bird Audio Detection challenge. Technical report, 2018. URL <https://arxiv.org/pdf/1807.05812.pdf>.
- Oriol Vinyals van den Oord, Yazhe Li. Representation Learning with Contrastive Predictive Coding. Technical report, 2019. URL <https://arxiv.org/pdf/1807.03748.pdf>.
- Pete Warden. Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition. Technical report, 2018. URL <https://arxiv.org/pdf/1804.03209.pdf>.
- Donglai Wei, Joseph Lim, Andrew Zisserman, and William T Freeman. Learning and Using the Arrow of Time. In *Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 8052–8060, 2018. URL <http://people.csail.mit.edu/donglai/paper/aot18.pdf>.
- Yong Xu, Qiang Huang, Wenwu Wang, Peter Foster, Siddharth Sigtia, Philip J B Jackson, and Mark D Plumbley. Unsupervised Feature Learning Based on Deep Models for Environmental Audio Tagging. *IEEE/ACM TRANSACTIONS ON AUDIO, SPEECH, AND LANGUAGE PROCESSING*, 25(6), 2017. doi: 10.1109/TASLP.2017.2690563. URL <http://ieeexplore.ieee.org>.
- Richard Zhang, Phillip Isola, and Alexei A. Efros. Colorful Image Colorization. In *European Conference on Computer Vision*, pp. 649–666, mar 2016. URL <http://arxiv.org/abs/1603.08511>.