

# UNSUPERVISED DISCRIMINATIVE LEARNING OF SOUNDS FOR AUDIO EVENT CLASSIFICATION

Sascha Hornauer<sup>1</sup>, Ke Li<sup>1</sup>, Stella X. Yu<sup>1</sup>, Shabnam Ghaffarzadegan<sup>2</sup>, Liu Ren<sup>2</sup>

UC Berkeley<sup>1</sup>, Robert Bosch LLC<sup>2</sup>

## ABSTRACT

Recent progress in network-based audio event classification has shown the benefit of pre-training models on visual data such as ImageNet. While this process allows knowledge transfer across different domains, training a model on large-scale visual datasets is time consuming. On several audio event classification benchmarks, we show a fast and effective alternative that pre-trains the model unsupervised, only on audio data and yet delivers on-par performance with ImageNet pre-training. Furthermore, we show that our discriminative audio learning can be used to transfer knowledge across audio datasets and optionally include ImageNet pre-training.

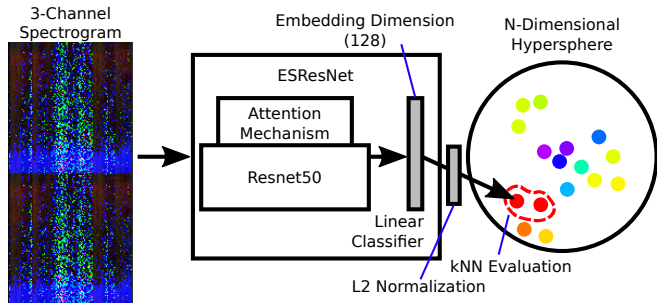
## 1. INTRODUCTION

Deep learning for audio event detection and classification benefits from large datasets. Despite unlimited access to audio data (e.g. YouTube, Freesound, etc), labeling audio events is labor intensive and noisy due to ambiguity in start and end times and short duration of some audio events.

On the other hand, due to similarities of the most commonly used audio features (i.e. spectrograms) to images, it is possible to benefit from advances in the image and video domain. Recent work shows improved performance when pre-training models on pretext tasks such as image classification or video based prediction [1, 2].

However, using image/video pre-training is very time demanding due to the size of visual data. Also, for every architectural change, this time-demanding pre-training needs to be repeated. Finally, it limits network design since the feature extractor has to be able to process image data which might not be necessarily suitable for an audio task.

Furthermore, for audio applications on embedded devices, such as voice assistants, it is desirable to be able to improve the model performance on the edge over time through fine-tuning on new recorded data. The need of on the edge computation might be due to privacy concerns, to avoid sending users' audio data to the Internet, or due to missing network availability. As a result, task performance needs to improve fast with few epochs and little data on the device. Large network models, as used for image data, may be too computationally expensive for many devices.




**Fig. 1.** Embedding of spectrograms into the feature space with ESResNet. Stereo channels are stacked for illustration. First the network learns to embed similar spectrograms close together on a hypersphere. Second, sound event classification is trained, starting with the pre-trained network.

We present a pre-training method that accelerates network fine-tuning on the task of sound classification while being itself fast, efficient and versatile. Compared with the state-of-the-art approach ESResNet [3], we show how our method achieves competitive results and significantly outperforms training from scratch. On one benchmark we even outperform state-of-the-art pre-training in early epochs.

We are faster at pre-training because we need only three audio datasets which combined have a fraction of the size of ImageNet. We focus on achieving fast results in very few epochs for edge computation and do not aim to go beyond state of the art performance on sound classification after extensive training. For a fair comparison we apply our method to the ESResNet codebase.

With unsupervised training on a pretext task, using only audio data, we also avoid the need for labels. By using Non-Parametric Instance-level Discrimination (NPID) [4] to train ESResNet on audio datasets we learn features beneficial for downstream audio classification tasks, illustrated in fig. 1. This allows us to integrate all data seamlessly across datasets and after deployment to train on novel unlabeled audio data. We call this approach Unsupervised Discriminative Learning of Sounds (DLS) and give an overview in fig. 2.



**Fig. 2.** Discriminative Learning of Sounds (DLS) for Audio Event Classification. We compare the same network, pre-trained either on our proxy-task, ImageNet image classification or not at all. ImageNet pre-training takes several days on common GPUs. DLS can train unsupervised, only on sound data and within a few hours. When fine-tuning on Audio Event Detection, DLS stays on par with ImageNet pre-training over many epochs. It constitutes an efficient alternative, especially useful for edge computing or to accelerate the design phase of novel network architectures.

## 2. RELATED WORK

**Audio Event Detection.** Audio event detection was largely improved in the last decade by leveraging new available datasets [5]. We use: ESC-50, its subset ESC-10, Urban-Sound8K and the DCASE 2013 scene classification dataset (SCD) (e.g. evaluated on in [1]). Others are: AudioSet [6] which holds predominantly music and speech examples, and the newer DCASE datasets [7]. These were not evaluated since we focused on comparison with a specific recent state-of-the-art method that does not use these datasets.

**Unsupervised representation learning.** Unsupervised approaches in the vision domain improve in great strides towards their supervised counterparts. MoCo [8], SimCLR [9] and NPID [10] have shown to produce valuable features for downstream tasks. We base our contribution on NPID and extend it to the audio domain.

**Audio representation learning.** Compared to traditional hand-crafted features unsupervised training can lead to more robust and compact audio representations. [11] applied deep belief networks to learn audio representations for speech and music. Generative methods have been explored in [12, 13, 14] using variants of autoencoders. Audio representation learning has also been studied for speech [15] and music [16].

**Knowledge transfer from vision.** Transfer learning from visual tasks and exploiting audio-visual correspondences has been explored in the past. [17] was able to learn associations between free form audio, i.e. spoken sentences, and related images. By predicting simply if parts of an audio and video correspond, [18] was able to learn good audio and visual representations. They were also able to extend this to localize

sounding objects in a scene [19] and their approach was extended to be used in audio classification [1].

Concurrent work which we use as basis for our evaluation improves downstream task performance by directly fine-tuning ImageNet pre-trained visual models. By mapping spectrograms to the format of a color image and by using a pre-trained Resnet50 they achieve state of the art results on audio event detection [3]. However, these methods rely on a large network to leverage visual input for training. We present a method which can be used to pre-train networks optimized for the audio domain and low resources.

## 3. DISCRIMINATIVE LEARNING OF SOUNDS (DLS)


Here we show how to pre-train ESResNet using DLS on four datasets and fine-tune on each dataset. We use the best performing ESResNet with attention for all experiments and pre-training steps.

**Audio Datasets.** Most sounds in the datasets contain audio events recorded in natural environments, such as glass breaking or dog barking. DCASE2013 sounds also include longer recordings, such as riding in a bus. Across datasets, sounds differ in length and number of classes, summarized in table 1. Some contain events only on a fraction of their length, such as a single *dog bark*, compared to other files filling the entire standardized length, such as *kids playing* (see fig. 3).

**Spectrogram Network Input.** Power spectrograms, created with Short-Time Fourier Transform (STFT), are the input for all stages of training, following the method from [3]. We follow this method exactly to compare on equal grounds. Magnitude and phase are squared separately and the results

	US8K	ESC 50	ESC 10	DCASE 2013
Events	8732	2000	400	200
Classes	10	50	10	10
Length	$\leq 4s$	5s	5s	30s
Fold	1 of 10	1 of 5	1 of 5	1 of 2

**Table 1.** Dataset setup. Classification is hardest on ESC50 because it has the least data per class. We tune hyperparameters on one fold and use all data for unsupervised pre-training.



**Fig. 3.** Spectrograms after pre-processing. Frequency ranges are split into three and concatenated as color channels. For illustration stereo channels are separated into rows. The leftmost two columns show *children playing*, followed by *car horns*. Large appearance changes exist within some classes.

are added. Spectrograms are divided into three equal sized parts to separate the higher, middle and lower frequencies. Finally the three parts are concatenated along a new channel dimension to create color images from the spectrograms to be processed by the network.

#### 4. NETWORK PRE-TRAINING.

The pre-text task of instance discrimination generates features, useful for the downstream task of audio classification. The network’s task is to assign a unique id to every spectrogram. Applied on images, this creates a feature space where *visually* similar images are grouped. However, visual similarity is not an explicitly defined criteria in the loss function, but a by-product of the pressure on the network to structure and differentiate between all images of the dataset.

We follow [4] and use spectrograms as input. We train ESResNet  $f_\theta$  with weights  $\theta$  to map a spectrogram  $x_i$  with the assigned id  $i$  in the training set to a latent output vector  $\mathbf{v}_i$ . The vector is L2 normalized, so  $\|\mathbf{v}_i\| = 1$  and updates its entry in a memory bank in each training iteration to calculate  $\{\mathbf{v}_j\}$  of all spectrograms. This allows efficient calculation of a non-parametric softmax to get the assigned instance id probabilities. The probability of a spectrogram  $x$  mapped to

vector  $\mathbf{v}$  having training set id  $i$  is defined by:

$$P(i|\mathbf{v}) = \frac{\exp(\mathbf{v}_i^T \mathbf{v} / \tau)}{\sum_{j=1}^n \exp(\mathbf{v}_j^T \mathbf{v} / \tau)} \quad (1)$$

where  $\tau$  is the Softmax temperature parameter. Stochastic gradient descent is used to minimize the log-likelihood:  $\arg \min_\theta - \sum_{i=1}^n \log P(i|f_\theta(x_i))$  during training.

The final feature mapping has the advantage that similar spectrograms are mapped close in the feature space. This enables to group sounds with similar characteristics, as sketched in fig. 1. While this enables our performance on the downstream task, larger variance within a class than between classes is an issue and limits this approach. Examples of very different car horns can be seen in fig. 3.

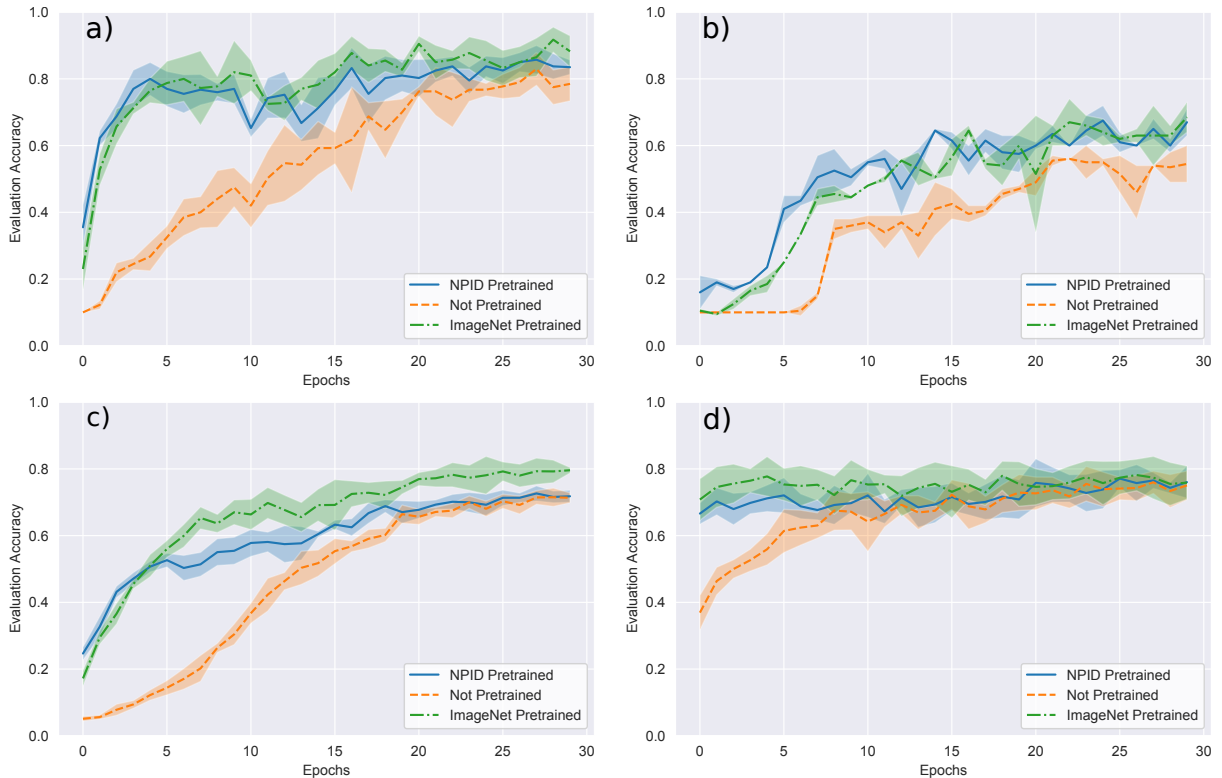
**Evaluation of Unsupervised Learning.** We first trained and evaluated on each dataset separately, according to the official training/evaluation folds (see tab. 1). Every instance of the evaluation set is mapped and we check its label alignment with its weighted  $K=5$  nearest neighbors in the training data feature space. We tune hyperparameters separately for each dataset but found the optimal configuration was almost the same for all. Therefore, we use the average parameters: *Embedding dimension*=128, *NCE-K*=64 and *NCE-T*=0.4, which are the number of negative samples and Softmax temperature of the Noise Contrastive Estimation in [4]. We trained 200 epochs with a batchsize of 64.

In a second phase, we use all available audio datasets and folds to combine one large training dataset, without evaluation or test set. Therefore we train the embedding with the fixed hyperparameters found in phase one and no longer evaluate the performance in this step. However, we observe that the training loss at epoch 200 is sufficiently converged. The result of this phase is the pre-trained ESResNet.

#### 5. NETWORK FINE-TUNING.

The pretrained ESResNet is fine-tuned and evaluated on each dataset individually. The final classification layer is set to the number of classes depending on the dataset we fine-tune on, which means it needs to be retrained. We found fine-tuning all layers yields faster performance gain than fixing them and only fine-tuning the classifier. Results are shown in fig. 4.

**Comparison with ImageNet Data** We compare pre-training by image classification on ImageNet with DLS on the four audio datasets: DCASE2013 (SCD), ESC-50/ESC-10 and UrbanSounds 8k which provide 9.3Gb of sound data. ImageNet consists of over 200Gb of data, with 14,197,122 images. Training a classification task on ImageNet with Resnet50, which is also the backbone of the ESResNet architecture, can take almost two weeks on a single GPU [20]. Even if we assume ImageNet training happens on a similar modern setup as ours (4x GeForce GTX 1080 Ti), the training time of ImageNet would still be at least 1-2 days. Training DLS with ESResNet on 200 epochs takes 4 hours.



**Fig. 4.** Evaluation accuracy over epochs on a) ESC10, b) DCASE2013, c) ESC50 and d) UrbanSounds8k. Shaded areas show standard deviation of average  $n$ -fold cross validation results. The average accuracy pretrained with DLS is on par with ImageNet pre-training and significantly better than training from scratch. Even though we use only sound data with a fraction of the size of ImageNet, we reach a similar performance gain. In d) pre-training shows on par results with ImageNet pretraining. In c) DLS pre-training outperforms from scratch training over the first 18 epochs. In b) DLS pre-training even outperforms ImageNet pre-training on the first 11 epochs.

## 6. EXPERIMENTAL RESULTS

We evaluated the performance of ESResNet with attention with DLS pretraining, Imagenet pretraining and training from scratch. The architectures are fully identical [3].

**DCASE2013 and ESC10.** Training on small datasets benefits largely from pre-training, as can be seen by the performance gain on ESC10 and DCASE2013 (fig. 4). On ESC10, DLS pre-training outperforms ImageNet pretraining slightly in the first three epochs, rising to 80% accuracy where training from scratch remains under 40%. On DCASE2013, DLS pre-training is best for 11 epochs, rising towards 50% accuracy and continuing on par with ImageNet pre-training.

**ESC50 and UrbanSounds8k.** On larger datasets, the accuracy gain in the first epochs is similar between DLS and ImageNet pre-training. However DLS was trained in significantly shorter time and on less data. On ESC50, it outperforms from scratch training significantly and rises to over 60% accuracy in the first 15 epochs. UrbanSounds8k is the largest dataset and both pre-training methods show similar performance gain per epoch. Each epoch contains more sound files than in any

other dataset so methods are already close after the first epoch when they are evaluated. DLS still outperforms training from scratch on 5-6 epochs.

**Densenet as backbone.** We also compared using Densenet as backbone and omit graphs for brevity. Results show slower performance gain over the first 15 epochs and then level off at similar levels. All trends are similar though on DCASE2013, all performance gains are very close before epoch 30.

## 7. DISCUSSION

We presented an unsupervised pre-training method which enables fast fine-tuning when training networks for audio event classification. While using only a comparatively small amount of sound data, we show that we can gain early performance as the same network pre-trained on ImageNet. We validate our results on three commonly used datasets for audio classification. Our approach enables network training on devices with limited computing resources, e.g. for continuous improvement by adopting novel collected data. Because pre-training with DLS is unsupervised, novel unlabeled data can be integrated with existing data seamlessly.

## 8. REFERENCES

- [1] Jason Cramer, Ho-Hsiang Wu, Justin Salamon, and Juan Pablo Bello, "Look, listen, and learn more: Design choices for deep audio embeddings," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019.
- [2] Karren Yang, Bryan Russell, and Justin Salamon, "Telling left from right: Learning spatial correspondence of sight and sound," in *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 06 2020.
- [3] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel, "Esresnet: Environmental sound classification based on visual domain models," *arXiv preprint arXiv:2004.07301*, 2020.
- [4] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [5] D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M. D. Plumbley, "Detection and classification of acoustic scenes and events," *IEEE Transactions on Multimedia*, vol. 17, no. 10, 2015.
- [6] Jort F Gemmeke, Daniel PW 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 *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017.
- [7] Archontis Politis, Sharath Adavanne, and Tuomas Virtanen, "A dataset of reverberant spatial sound scenes with moving sources for sound event localization and detection," *arXiv preprint arXiv:2006.01919*, 2020.
- [8] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton, "A simple framework for contrastive learning of visual representations," *arXiv preprint arXiv:2002.05709*, 2020.
- [10] Zhirong Wu, Alexei A Efros, and Stella X Yu, "Improving generalization via scalable neighborhood component analysis," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [11] Honglak Lee, Peter Pham, Yan Largman, and Andrew Y Ng, "Unsupervised feature learning for audio classification using convolutional deep belief networks," in *Advances in neural information processing systems*, 2009.
- [12] Matthias Meyer, Jan Beutel, and Lothar Thiele, "Unsupervised feature learning for audio analysis," *arXiv preprint arXiv:1712.03835*, 2017.
- [13] 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," *arXiv preprint arXiv:1603.00982*, 2016.
- [14] Yong Xu, Qiang Huang, Wenwu Wang, Peter Foster, Siddharth Sigtia, Philip JB 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*, vol. 25, no. 6, 2017.
- [15] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al., "Deep speech: Scaling up end-to-end speech recognition," *arXiv preprint arXiv:1412.5567*, 2014.
- [16] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li, "Yfcc100m: The new data in multimedia research," *Communications of the ACM*, vol. 59, no. 2, 2016.
- [17] David Harwath, Antonio Torralba, and James Glass, "Unsupervised learning of spoken language with visual context," in *Advances in Neural Information Processing Systems*, 2016.
- [18] Relja Arandjelovic and Andrew Zisserman, "Look, listen and learn," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- [19] Relja Arandjelovic and Andrew Zisserman, "Objects that sound," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [20] Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, and Kurt Keutzer, "Imagenet training in minutes," in *Proceedings of the 47th International Conference on Parallel Processing*, 2018.