

# BIOACOUSTIC EVENT DETECTION WITH PROTOTYPICAL NETWORKS AND DATA AUGMENTATION

## Technical Report

*Mark Anderson, Naomi Harte*

SIGMEDIA

School of Engineering, Trinity College Dublin  
Dublin 2, Ireland  
anderm3, nharte [@tcd.ie]

### ABSTRACT

This report presents deep learning and data augmentation techniques used by a system entered into the Few-Shot Bioacoustic Event Detection for the DCASE2021 Challenge. The remit was to develop a few-shot learning system for animal (mammal and bird) vocalisations. Participants were tasked with developing a method that can extract information from five exemplar vocalisations, or shots, of mammals or birds and detect and classify sounds in field recordings. In the system described in this report, prototypical networks are used to learn a metric space, from which classification is performed by computing the distance of a query point to class prototypes, classifying based on shortest distance. We describe the architecture of this network, feature extraction methods, and data augmentation performed on the given dataset and compare our work to the challenge’s baseline networks.

**Index Terms**— Few-shot Learning, Prototypical Networks, Data Augmentation, SED, Activity Detection

## 1. INTRODUCTION

Few-shot learning has emerged as a promising tool for sound event detection, and is particularly relevant for bioacoustics applications. Few-shot learning systems must generalise to new classes unseen during training, given only a few labelled instances of each class. This makes them highly suitable to tasks such as the monitoring of animal populations through their vocalisations, where the data for a dedicated classifier may be costly to annotate, and may lose some generalisation.

In the DCASE 2021 Few-shot Bioacoustic Event Detection challenge, participants are asked to provide the onset and offset times of events in a given audio file, given 5 labelled examples of the class of interest. Four separate data sets are provided for training, with a further two for validation. The training set comprises of 14 hours, 20 minutes and contains 19 separate classes, and 4686 events. These classes include bird species, hyenas, and meerkats. The validation set contains 5 hours of audio, 4 classes and 310 events. There is no overlap between these two development sets. Further details are provided in [1] and [2]. Evaluation Data contains only the first 5 positive events for each file, with the aim to use the system to predict all other positive events [3].

Prototypical networks can be applied to a few-shot learning problem such as the task proposed in the challenge. They are considered a model based approach to the problem of few-shot learning

by Wang et al [4], which aims to reduce to constrain the hypothesis space using prior knowledge. This is achieved by learning a non-linear mapping between inputs and a metric/embedding space. Class support points should then be sufficiently separated from each other and a prototype representation of the classes constructed [5]. This prototype is constructed as the mean of the support set (which are the labelled instances of the new unseen data to be classified). Classification of each query point is done by finding its nearest class prototype according to some distance function  $d_\phi(\mathbf{x}, \mathbf{x}')$ , typically the distance function employed is Euclidean distance.

In this regard, prototypical networks are similar to clustering algorithms and nearest neighbour classification in particular, where the prototype representations of each class are the neighbours to the query point. The efficiency and simplicity in these networks makes them an appealing approach to the problem posed in this challenge.

## 2. IMPLEMENTATION DETAILS

As mentioned in both Section 1 and the challenge description, few-shot learning is a novel task. Several methods have emerged as solutions to the task, such as matching networks, however we felt prototypical networks offered the most promising approach to the problem. As such, we propose in this section a modified version of the baseline Prototypical Network, implementing a modified network architecture, as well as data augmentation to supplement training data and validation data.

### 2.1. Data Preparation

All audio files in both the development and evaluation sets are first resampled to a sampling rate of 22050Hz and normalised prior to their transformation into Mel spectrograms. We do not apply any band pass filtering to the audio signals due to the variety of species in the available datasets. Future work could apply a custom filter to each dataset based on the typical frequency ranges present, but as this network is designed to generalise to new, unseen classes, we decided against implementing this for this submission. Each audio file is transformed to a Mel spectrogram, with 128 Mel bins, an FFT size of 1024 samples and a hop of 256 samples. This is the same configuration present in the baseline system.

We also preform Per-Channel Energy Normalisation (PCEN) on the spectrograms (proposed as method of improving robustness to channel distortion in keyword spotting tasks [6]), in order to both reduce noise in the spectrogram and provide normalisation, gain

control and compression the file. PCEN has been proposed as an alternative logarithmic scaling of the Mel spectrograms [7]. PCEN does reduce the noise present in the spectrogram, but non-stationary sources of noise still pose an issue and can contaminate features. In our experiments, we tested our systems using both log-Mel spectrograms and PCEN Mel spectrograms.

### 2.2. Data Augmentation

To increase the model performance and create an embedding space that can generalise to loss of information, temporal/frequency shift and new classes, we apply various data augmentation techniques to the resulting spectrogram representations. These data augmentations are based upon work in [8] and [9]. These augmentations are achieved using pyTorch’s Audio toolset, Torchaudio, which uses the methods specified in [8].

- **Time stretching** - We apply a time stretch to each spectrogram, both shortening and lengthening of the spectrogram without any change in pitch info. This translates to a warping along the horizontal axis of the spectrogram according to a specified factor. In our system we have chosen a stretching factor of 5%, corresponding to two augmented spectrograms at effectively 0.95x and 1.05x the playback speed of the original audio file.
- **Time masking** - Time masking is applied such that  $T$  consecutive time steps are masked, with all frequency information set to 0 within that interval. The interval is randomly chosen from a uniform distribution, as per Torchaudio’s implementation. Due to the considerable length of each file in the data set, this is applied to each 10s chunk of audio, before being reconstructed prior to dataset construction.
- **Frequency masking** - Frequency masking performs a similar augmentation on the vertical, or frequency axis of the spectrogram. Again this is applied in chunks of 10s so as to increase the variability in augmented data.

Examples of these augmentations, and the PCEN front-end, are seen in Figure 1. Note, that while we show both time masking and frequency masking in the one spectrogram to save space, we did not perform this augmentation on the data itself.

These augmentations create more data to train and validate the model during development. After augmentation, the validation and evaluation data is transformed into the appropriate dataset form, comprising of positive data points to construct the positive class prototype, negative data points to construct the negative class prototype, and all remaining data to be query points. To prevent class imbalance when training, Random Over Sampling is employed. This is necessary as there is a severe skew in the class distribution of the data, with the majority of occurrences belonging to the negative class.

### 2.3. Model Architecture

The submitted model architecture is similar to that proposed in the baseline, we have modified the encoder to a less deep model, only using 3 ConvBlocks in place of the 4 in the baseline. Each layer grouping is made up of a Conv2D layer, Batch Normalisation, a ReLU activation layer, and followed by a 2D Max Pooling layer, with pool-size (2, 2). Table 1 details the resulting architecture and the ConvBlock sub-block.

Table 1: Architecture of ConvBlock layer and Encoder Architecture

ConvBlock Architecture		
Sub-Layer 1	Conv2D	
Sub-Layer 2	BatchNorm	
Sub-Layer 3	ReLU	
Sub-Layer 4	MaxPool2D((2,2))	
Encoder Architecture		
Layer 1	ConvBlock	128
Layer 2	ConvBlock	128
Layer 3	ConvBlock	128
Layer 4	Flatten	

### 2.4. Training

Training of the prototypical network involves minimisation of the negative log-probability of a point to its true class, achieved using the Stochastic Gradient Descent (SGD) algorithm with an initial learning rate of 0.01 and a momentum factor of 0.85. Learning rate is scheduled to halve when a plateau is reached, with a patience of 5 epochs and a threshold of 0.01. In our experiments this provided the best environment for training. All our models are trained using 150 epochs, however the best performing model had minimum loss at epoch 126. As we are using additional augmented data, training time can be increased without risk of over-fitting to the training data.

Training is performed in episodic batches using a custom episodic batch sampling function, easily added to the pyTorch framework. Each episode is randomly generated from the available data, creating a randomised support and query set. Prototypes are computed from the support samples and loss on query points is calculated as per Equation 1.

$$J \leftarrow J + \frac{1}{N_C N_Q} \left[ d(f_\phi(\mathbf{x}), \mathbf{c}_k) + \log \sum_{k'} e^{-d(f_\phi(\mathbf{x}), \mathbf{c}_{k'})} \right] \quad (1)$$

Where,

- $d(\cdot)$  → Euclidean distance function
- $f_\phi$  → embedding of query set
- $\mathbf{c}_k$  → class prototype for class  $k$
- $N_C, N_Q$  → number of classes and query points

In Figure 2 we have provided the average validation loss during training of a model trained on Mel spectrograms using log-scaling and no augmentation, Mel spectrograms using log-scaling and data augmentation, and finally Mel spectrograms using PCEN and data augmentation. Models using augmented features take longer to train than those which do not, and the usage of PCEN further increases the time necessary for the model to converge on a minimum. Although models using augmented data and PCEN do not achieve lower loss than the original data using log-Mel spectrograms, as seen in Section 3 this does not lead to worse results, in fact they do perform better on this task.

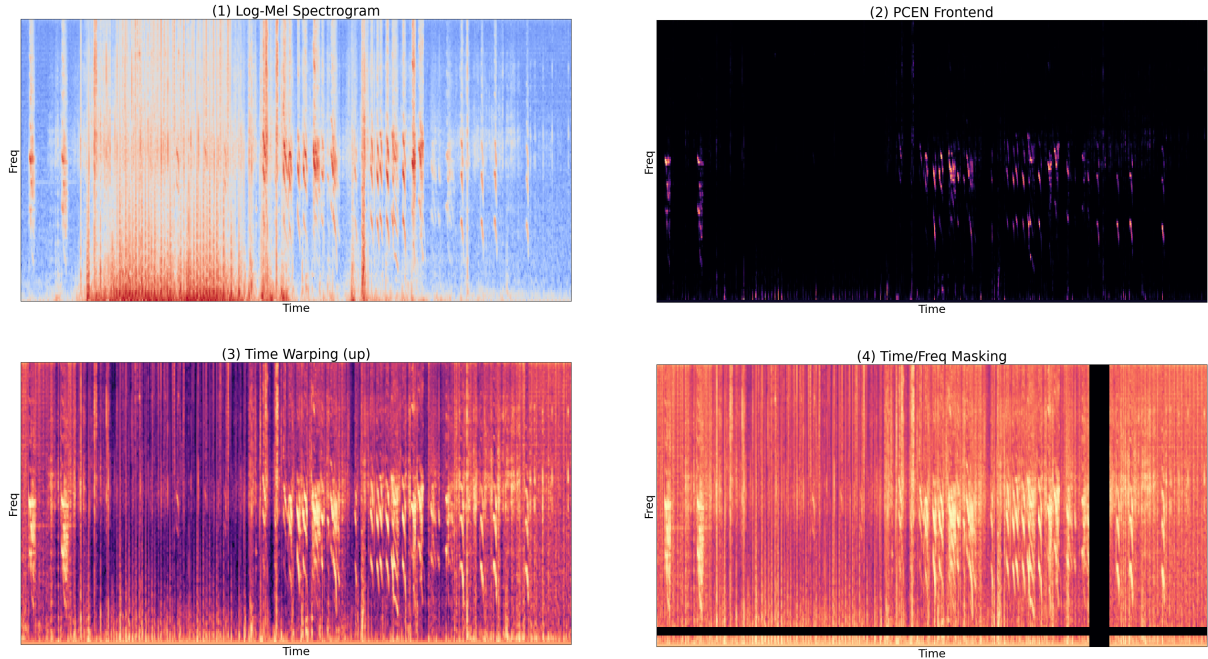


Figure 1: Example spectrograms on a 10s clip from the JD subset, (1) Log-Mel Spectrogram, (2) PCEN Frontend, (3) Time warping by 20%, (4) Mixture of time and frequency masking. PCEN representations when converted to an image, are more difficult to see than log scaled spectrograms, therefore we have decided to present the augmentations using log-scaled spectrograms.

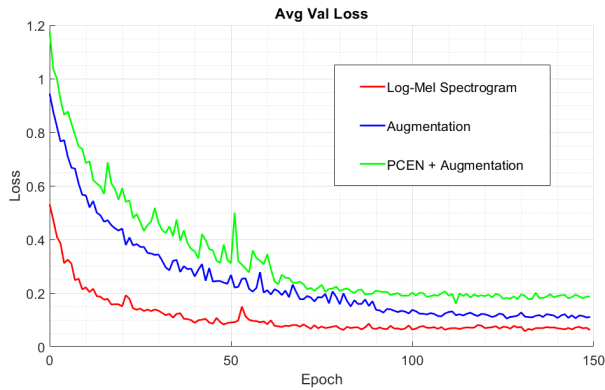


Figure 2: Average validation loss for system with log-Mel spectrograms, augmentation, and PCEN + augmentation.

## 2.5. Evaluation & Post-Processing

Evaluation of the model involves calculating the probability of whether a query point belongs to the positive class, achieved by taking the softmax of the distances between query points and the two class prototypes as in Equation 2.

$$p_i(y = c_k | \mathbf{x}) = \frac{\exp(-d(f_\phi(\mathbf{x}), c_k))}{\sum_{k'} \exp(-d(f_\phi(\mathbf{x}), c_{k'}))} \quad (2)$$

$\mathbf{p}$  is the set of probabilities for each iteration (Equation 3). Five

iterations of this operation are performed, after which the average of the probabilities is taken. The chosen class is the one with the highest mean probability over all iterations as per Equation 4. Thresholding of probabilities is set to 0.5, although this could be tuned depending on application, and whether recall or precision is valued more.

$$\mathbf{p} = \{p_1, \dots, p_i\} \quad (3)$$

$$P_\phi(y = c_k | \mathbf{x}) = E[\mathbf{p}] \quad (4)$$

For a given file, this results in a classification of each segment as either positive or negative. To prevent abrupt transitions where the classifier may have incorrectly classified a segment in the middle of a vocalisation, median filtering with a window length of 5 samples is applied. Edge detection is performed using a 1D edge detection kernel, convolved with the classification results at the output of the median filtering stage (Equations 5 and 6) to determine the points where onset the onset and offset of an event occurred. The output of this convolution contains the value 1 at the index of onset, and  $-1$  at the index of offset.

$$k[n] = [1, -1], \text{ edge detection kernel} \quad (5)$$

$$(y * k)[n] = y[n] - y[n - 1] \quad (6)$$

Onset and offset times may then be calculated and written to a CSV file. These onset and offset times are further post-processed using a method proposed by the baseline, where any positive instances that are shorter than 60% the length of the shortest positive

Table 2: Table of results on validation set.

Model	F-Meas. (%)	Pre.	Rec.
<b>Baseline Reproduction</b>	30.437	0.456	0.228
<b>Log Mel Spectrograms</b>	16.565	0.108	0.355
<b>Data Aug.</b>	20.517	0.143	0.360
<b>Data Aug. + PCEN</b>	26.243	0.200	0.381

Table 3: Results by validation data subset.

Data Sub-Set	F-Meas. (%)	Pre.	Rec.
<b>PB</b>	22.222	0.201	0.249
<b>HV</b>	32.039	0.199	0.825

labelled instance, are dropped from the final results. Evaluation of results with and without this post-processing step show that it reduces the amount of false positive events with little change in the amount of true positives registered. Results of experimentation with the minimum allowed event length vary by file, and may warrant future investigation.

### 3. RESULTS

Performance is evaluated on the validation data, provided in the development set released as part of the challenge. We utilise the evaluation method provided by the challenge organisers, which creates a confusion matrix for each file, and calculates precision, recall, and the F-measure.

Given the similarity of our system to that of the baseline, we decided to reproduce the baseline prototypical methods results for comparison with ours. Unfortunately, we could not replicate the results stated in the challenge, achieving at most an F-measure of 30.437%.

In Table 2, we summarise the results on the validation data provided in the development set on the baseline, and three progressive iterations of our system (Log-Mel spectrogram, Log-Mel Spectrogram with augmentation, Mel spectrogram with PCEN and augmentation).

The best performing model utilises both PCEN on the front-end, and data augmentation to give our best result of 26.243% on the F-measure metric. Breakdown of results by each subset provided for validation can be seen in Table 3.

### 4. DISCUSSION

The model architecture we have chosen is similar to that of the baseline, however we have applied various data augmentation methods to the training data to increase model robustness.

Our aim of improving on the baseline results has not been met, as our F-measure is  $\approx 4.2\%$  below that of our reproduced baseline. We have made significant improvements to recall, our system is capable of recognising 38% of calls in the validation set, but this has come at the cost of precision, as we have many false positives in our predictions. Future work to improve this may focus on the post processing applied to the data after initial inference. We may also employ deeper models in the future, as there is the possibility that the learned embedding space does not separate the two classes well enough, with support points not being sufficiently far away from each other.

Improvement was mostly gained through the usage of PCEN over the more traditional approach of log-scaling the Mel spectrogram, lending evidence to the claim in [7] that PCEN can significantly outperform log-Mel scaling when used on noisy data, with no significant increase in computational complexity. Our methods of data augmentation also improved on our initial log-Mel spectrogram model, and was further improved by the addition of PCEN. Other methods of data augmentation (such as additional noise from other files, piece-wise frequency stretching, etc.), and augmentation of support points for evaluation should also be investigated. Interestingly, although the average validation loss of models trained using data augmentation was greater than that of the log-Mel spectrograms, their performance when evaluating the validation set was greater. We also found that they could be trained for longer, reaching a minimum less quickly than the original log-based features without augmentation or PCEN.

To conclude, our results of our system show that few-shot learning for bioacoustics applications with prototypical networks is still a novel and challenging task which is in its infancy, but one which can rapidly improve going forward.

### 5. ACKNOWLEDGEMENTS

The authors would like to thank the organisers of this challenge, for proposing an interesting and novel problem which is relevant to our own research. We also want to thank the ADAPT Centre for the usage of their 'Boole' cluster when developing this system. The ADAPT Centre for Digital Content Technology is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

## 6. REFERENCES

- [1] V. Morfi, D. Stowell, V. Lostanlen, A. Strandburg-Peshkin, L. Gill, H. Pamula, D. Benvent, I. Nolasco, S. Singh, S. Sridhar, and et al., Feb 2021. [Online]. Available: <http://dcase.community/challenge2021/task-few-shot-bioacoustic-event-detection>
- [2] —, “Dcase 2021 task 5: Few-shot bioacoustic event detection development set,” Feb 2021. [Online]. Available: <https://zenodo.org/record/4543504>
- [3] —, “Dcase 2021 task 5: Few-shot bioacoustic event detection evaluation set,” May 2021. [Online]. Available: <https://zenodo.org/record/4864755>
- [4] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, “Generalizing from a few examples: A survey on few-shot learning,” *ACM Computing Surveys*, vol. 53, no. 3, pp. 1–34, 2020.
- [5] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” in *Advances in Neural Information Processing Systems*, vol. 2017-December, 2017, pp. 4078–4088.
- [6] Y. Wang, P. Getreuer, T. Hughes, R. F. Lyon, and R. A. Saurous, “Trainable frontend for robust and far-field keyword spotting,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 5670–5674.
- [7] V. Lostanlen, J. Salamon, A. Farnsworth, S. Kelling, and J. P. Bello, “Robust sound event detection in bioacoustic sensor networks,” *PLOS ONE*, vol. 14, no. 10, pp. 1–31, 10 2019. [Online]. Available: <https://doi.org/10.1371/journal.pone.0214168>
- [8] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Interspeech 2019*, Sep 2019. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2680>
- [9] M. Lasseck, “Acoustic bird detection with deep convolutional neural networks,” *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop*, 2018. [Online]. Available: [http://dcase.community/documents/challenge2018/technical\\_reports/DCASE2018\\_Lasseck\\_76.pdf](http://dcase.community/documents/challenge2018/technical_reports/DCASE2018_Lasseck_76.pdf)