

## FEW-SHOT BIOACOUSTIC EVENT DETECTION AT THE DCASE 2023 CHALLENGE

*I. Nolasco*<sup>1</sup>, *B. Ghani*<sup>2</sup>, *S. Singh*<sup>1</sup>, *E. Vidaña-Vila*<sup>3</sup>, *H. Whitehead*<sup>10</sup>, *E. Grout*<sup>4,5</sup>, *M.G. Emmerson*<sup>7</sup>, *I. Kiskin*<sup>9</sup>,  
*F. H. Jensen*<sup>8</sup>, *J. Morford*<sup>6</sup>, *A. Strandburg-Peshkin*<sup>4,5</sup>, *L. Gill*<sup>11</sup>, *H. Pamulā*<sup>12</sup>, *V. Lostanlen*<sup>13</sup>, *D. Stowell*<sup>2,14</sup>

<sup>1</sup> Centre for Digital Music (C4DM), Queen Mary University of London, London, UK

<sup>2</sup> Naturalis Biodiversity Centre, Leiden, NL

<sup>3</sup> La Salle Campus Barcelona, Ramon Llull University, Barcelona, ES

<sup>4</sup> Dept. of Biology & Centre for the Advanced Study of Collective Behaviour, University of Konstanz, DE

<sup>5</sup> Dept. for the Ecology of Animal Societies, Max Planck Institute of Animal Behavior, DE

<sup>6</sup> The Oxford Navigation group, Dept. of Zoology, Oxford University, Oxford, UK

<sup>7</sup> School of Biological and Behavioural Sciences, Queen Mary University of London, London, UK

<sup>8</sup> Biology Dept, Syracuse University, NY, USA

<sup>9</sup> Institute for People-Centred AI, FHMS, University of Surrey, Surrey, UK

<sup>10</sup> School of Science, Engineering and Environment, University of Salford, Manchester, UK

<sup>11</sup> Landesbund für Vogel- und Naturschutz; Naturkundemuseum Bayern/BIOtopia Lab, DE

<sup>12</sup> AGH University of Science and Technology, Kraków, PL

<sup>13</sup> Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004, F-44000 Nantes, FR

<sup>14</sup> Tilburg University, Tilburg, NL

### ABSTRACT

Few-shot bioacoustic event detection consists in detecting sound events of specified types, in varying soundscapes, while having access to only a few examples of the class of interest. This task ran as part of the DCASE challenge for the third time this year with an evaluation set expanded to include new animal species, and a new rule: ensemble models were no longer allowed. The 2023 few-shot task received submissions from 6 different teams with F-scores reaching as high as 63% on the evaluation set. Here we describe the task, focusing on describing the elements that differed from previous years. We also take a look back at past editions to describe how the task has evolved. Not only have the F-score results steadily improved (40% to 60% to 63%), but the type of systems proposed have also become more complex. Sound event detection systems are no longer simple variations of the baselines provided: multiple few-shot learning methodologies are still strong contenders for the task.

**Index Terms**— Few-shot learning, bioacoustics, sound event detection

### 1. INTRODUCTION

Bioacoustic event detection, the identification of animal vocalizations within specific timeframes, shares many similarities with sound event detection (SED) in varying contexts like urban settings [1] or secured spaces [2, 3]. Nonetheless, bioacoustics poses a unique set of challenges due to the varied recording conditions and diverse animal vocalizations [4]. This makes it an exciting and complex domain within machine learning, with several specialized sub-disciplines focused on different animals. Recent advances in supervised deep convolutional neural networks (CNN) have potential for enhancing feature detection.

However, their supervised nature necessitates extensive, well-categorized acoustic event data and hundreds of annotated examples per class. Gathering this data can be an uphill battle, considering the uneven distribution of species, the labor-intensive nature of audio annotation, and the variable taxonomy based on the use case [5]. The limitations of a supervised sound event detection system become more prominent when extrapolating techniques used in speech to other animal sounds. This complexity arises from the differences in sound duration, interest units, and the context in which the sounds are made. Crucially, understanding the commencement and termination times of animal sounds is vital to community ecology, shedding light on various patterns of communication and influence among species [6]. Unlike speech science with its relatively limited granularity, bioacoustic studies operate at multiple levels, from coarse classification of species to fine distinction of individual call types. Moreover, the diversity in recording equipment used for animal sounds, from far-field to underwater, adds another layer of complexity, transforming bioacoustic event detection into a collection of small-data problems, each requiring specialized systems. This fragmentation, although useful for species classification tasks, impedes the practical application of deep learning in bioacoustics and life sciences more broadly [5].

To address these challenges, this DCASE task proposes a unified approach for bioacoustic event detection across the various sub-domains, aiming to mitigate the problems associated with data acquisition, annotation, and the fragmentation in computational bioacoustics. Hence, we compiled a unique ensemble of 14 small-scale datasets, each between 10 minutes and 10 hours long and derived from distinct sources, representing different application contexts. Breaking from the norm of training individual machine learning systems for each dataset, the idea is to develop a single, versatile system capable of identifying sound events across various datasets, with event categories specified at "query time". Additionally, dur-

	Dataset	mic type	# audio files	total duration	# labels	# events
Training set	BV: BirdVox-DCASE-10h	fixed	5	10 hours	11	9026
	HT: Hyenas	various	5	5 hours	5	611
	MT: Meerkats	animal mounted	2	70 mins	4	1294
	JD: Jackdaws	mobile	1	10 mins	1	357
	WMW: Western Mediterranean Wetlands Birds	various	161	5 hours	26	2941
Validation set	HB: Humbug mosquitoes	handheld	10	2.38 hours	1	712
	PB: Polish Baltic Sea bird flight calls	fixed	6	3 hours	2	292
	ME: Meerkats	animal mounted	2	20 mins	2	73
Evaluation Set	CHE: Transfer-Exposure-Effects birds	fixed	18	3 hours	3	2550
	DC: BIOTOPIA Dawn Chorus birds	fixed	10	95 mins	3	967
	CT: Coati	handheld	3	48 mins	3	365
	MS: Manx shearwater birds	fixed	4	40 mins	1	1087
	QU: Dolphin quacks	animal mounted	8	74 mins	1	3441
	MGE: Chick calls birds	fixed	3	32 mins	2	1195
	CHE23: Transfer-Exposure-Effects Frogs	fixed	16	40 mins	1	798
	CW: Cow moos	fixed	4	56 mins	1	293

Table 1: Summary of dataset characteristics.

ing an evaluation on an audio file, the system is provided with the initial five instances of the desired sound event. This approach employs a machine learning paradigm known as "few-shot learning" (FSL) [7, 8], where the aim is to construct precise models using less training data. In this context, FSL is explored using N-way-k-shot classification, where N and k represent the number of classes and the examples per class, respectively. Upon training with the first five occurrences of an event, the system effectively detects subsequent instances of the same event.

Our hypothesis is that bioacoustic event detectors can be trained using available bioacoustic datasets and then generalized to new targets using a few examples at the time of deployment.

## 2. DATASETS

When the DCASE challenge begins, each task releases its own development set, consisting of a training and validation sets. Participants must use this dataset to develop and validate their systems. As the challenge enters the evaluation phase, the evaluation set is released and participants apply their developed systems and output the predictions which are then used to calculate the final ranking scores. These datasets are organised in subsets that represent different acoustic sources and were gathered here with the specific purpose of broadening the targeted species. A summary of the main characteristics is presented in Table 1. Overall there are 8 sets focusing on bird species, 5 sets of mammal vocalisations (one of which underwater), 1 set of flying insect sounds (HB) and 1 set of amphibian calls (CHE23).

For the Few shot bioacoustic task, the training set is multi-label, since the provided annotations contain more than one class of interest. However, both validation and evaluation sets are single label, meaning that each audio file is annotated only for a single class of interest. While events of other classes are present these are not annotated and should not be predicted by the systems.

Also, the split between training, validation and evaluation sets does not follow the common supervised learning approach, since the classes in the validation and evaluation sets may overlap or not with the classes in the training set. This is designed so that examples of the target classes we want to detect are only provided as the 5 shots.

Given the few-shot setup of this task, each audiofile of the evaluation set is accompanied only with the annotations for the 5 initial events of the class of interest. The datasets used on the 2023 edition of the task remain the same as in previous edition (thoroughly described in [5]), but the evaluation set has been extended with two new subsets of data: Cow moos (CW) and frog croakings (CHE23).

**Cow moos (CW):** This dataset contains 4 audio files of about 15 minutes each recorded on a Cow’s farm in Catalonia, Spain. An ambient microphone connected to a Zoom H5 recorder was hung on the ceiling of a yard with multiple cows. Cow vocalizations were recorded and manually labelled by researchers from La Salle Campus Barcelona and AWEC Advisors S.L. in the framework of the projects CowTalk and CowTalk-Pro.

**Transfer-Exposure-Effects Frogs (CHE23):** This dataset is part of the same project which originated the CHE dataset, data were collected using unattended acoustic recorders (Song-meter 3) in the Chernobyl Exclusion Zone (CEZ) to capture the Chernobyl soundscape and investigate the longterm effects of the nuclear power plant accident on the local ecology. The CHE23 dataset consists in 16 audiofiles of varying lengths annotated for frog croaking events, however many other calls of other species are present throughout the recordings. The annotations were produced by Helen Whitehead using Raven Pro 1.6.

Finally, teams are allowed to use certain external datasets and pretrained models. These were selected based on two criteria: They must be open access and classes do not overlap with the evaluation target classes. The complete list of allowed datasets and pretrained models can be accessed in the DCASE challenge task page [11].

## 3. BASELINES AND EVALUATION METRICS

The benchmarks and evaluation metrics remain identical to those established in the 2022 rendition of the task [9]. The associated code can be procured from the GitHub repository<sup>1</sup>.

The few-shot bioacoustic sound event detection task adopts two baselines: 1) Template matching, and 2) Protootypical networks. Template matching represents a common practice in the bioacoustics domain. The overall approach consists in taking each of the

<sup>1</sup><https://github.com/c4dm/dc-case-few-shot-bioacoustic>

Team name	Best submission	Eval set: <i>F</i> -score %	Val set <i>F</i> -score %	Main characteristics
Du_NERCSLIP	2	63.78	75.6	Multitask learning; Frame-level embeddings; Voice activity detection
Moummad.IMT	2	42.72	63.46	Contrastive learning learns an Embedding space; fine-tuning encoder on both positive and negative events;
XuQianHu_NUdT_BIT	3	42.5	63.94	prototypical network, Delta MFCC and PCEN; Squeeze Excitation blocks
Gelderblom_SINTEF	2	31.10	36.6	Encoder based on BEATs; prototypical network.
Jung_KT	3	27.12	81.52	Prototypical network trained with a negative-based loss
Wilkinghoff_FKIE	4	16.00	62.64	Embeddings learnt with temporal dimension; template matching with Dynamic warping.
TM_baseline	-	14.89	3.37	Template matching baseline as in [9]
proto_baseline	-	2.92	28.45	Prototypical network baseline as in [9]

Table 2: F-score results per team (best scoring system) on evaluation and validation sets, and summary of system characteristics. Systems are ordered by higher scoring rank on the evaluation set. These results and technical reports for the submitted systems can be found on task 5 results’ page [10].

5 initial examples as templates and cross-correlate each template with the remaining audiofile. Events are predicted by selecting time frames where the cross-correlation values surpasses a defined threshold. Prototypical networks [7], on the other hand are trained through episodic learning and employ a 5-way-2-shot classification model in our case. Each prototype represents a coordinate in vector space, calculated as a mean of the coordinates of the 5 samples. Training comprises a *Support set* of 5 labelled samples from each class, while the remaining samples form the *Query set*. A class prototype is computed via an embedding function with learning parameters. Distances are optimised, and the network training creates a general representation where similar sounds are closer. In this way, the future data points are labelled using nearest-neighbour algorithms.

The systems are evaluated based on how well they predict events on the evaluation set. The metric used combines intersection over union and bipartite graph matching algorithms to select the best matches between the predicted events and ground truth events. After the matching phase, we count the number of true positives (TP), false positives (FP), and false negatives (FN), which are then used to calculate precision, recall and F-score. The systems are ranked using the event based F-score value. The task description and details are provided in [5].

In this year’s task, a distinctive modification has been introduced in the evaluation procedure. The use of ensemble models was no longer allowed. The objective behind this rule is to incentivise the development of truly general models, rather than a simple fusion of completely independent models.

#### 4. RESULTS

The third edition of the few-shot bioacoustic event detection task received participation of 6 teams, with a total of 22 submitted systems. The overall F-scores for the best submission per team are presented in Table 2 together with the main characteristics of the respective systems, and the results on each dataset of the evaluation set are presented in Fig. 1.

The winning submission is by the same team that won the previous edition of this task, namely, **Du\_NERCSLIP**. The system improves on the last year’s submission, [12] by including their frame-level embedding system into a multi-task learning architecture. The

new system now includes Target Speaker Voice Activity Detection as one of the branches. This system achieved 63% F-score which is an increase from the best F-score from last year, that was approx. 60%. Observing Fig. 2, it is possible to observe the improved results on individual datasets for this team compared to the last year’s system (columns Du22 and Du23). This shows that the described modifications are responsible for the considerable increase in the overall F-score.

Furthermore, an intriguing observation when looking at the F-scores per dataset in Fig. 1, is that overall systems performed extremely well on the CW dataset, but not the winning submission. Indeed the performance of Du\_NERCSLIP’s system on CW dataset is similar to the performance on the QU dataset, which is considered to be an extremely difficult case due to its very short events.

**Moummad.IMT** implemented a system based on Contrastive Learning, a method to learn an embedding space that maximises the distinction between positive events and negative events. During the evaluation stage, the encoder is further fine-tuned on the 5 positive

<sup>2</sup><https://dcase.community/challenge2022/task-few-shot-bioacoustic-event-detection-results>

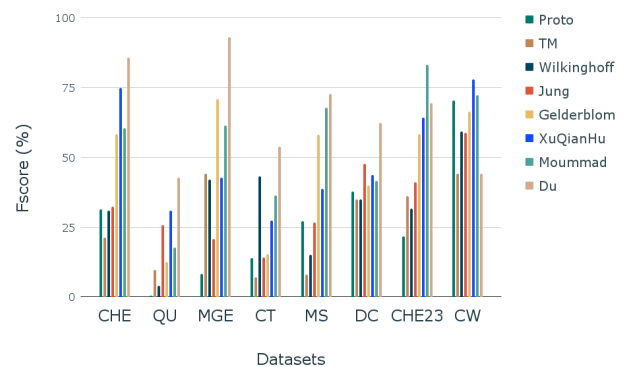


Figure 1: F-Score results of best system for each team and baselines. Systems are ordered from least to highest scoring rank on the evaluation set and indicated by the name of main author.




Figure 2: Comparison of the maximum F-score achieved by each team for the datasets used in both 2022<sup>2</sup> and 2023. The results for each year are distinguished by different colors.

examples provided and on selected negative examples for each audiofile. The final predictions are then the result of a simple binary classification in this embedding space.

**XuQianHu\_NUDT\_BIT** largely improved upon the Prototypical Network baseline by including squeeze/excitation (SE) blocks into the encoder part of the network. The idea behind this is to create an adaptive mechanism that assigns different weights to different channels of the feature map, by predicting their importance. The system also adopts the negative sample search mechanism proposed by Liu *et al.* on the previous edition of the task [13], which is designed to improve the learning of the negative prototypes. Also following from Liu *et al* submission, the input features employed are Delta MFCC and PCEN.

**Gelderblom\_SINTEF** followed the Prototypical Network approach, but use the BEATs pretrained model as encoder. BEAT stands for Bidirectional Encoder representation from Audio Transformers released by Microsoft for audio tokenisation and classification. In their submission, the authors explore how useful this model is to represent bioacoustic data and compare the embeddings produced by the pretrained model with the embeddings produced after a few epochs of fine-tuning on the ECS-50 dataset [14].

**Jung\_KT** combines Contrastive Learning and Prototypical Networks. This addresses the problem that the high imbalance between positive samples and negative samples creates in the learning of the prototypes. They propose a novel negative-based prototypical loss function that is used in a fine-tuning stage of the pipeline and drives the system to maximise the positive to negative samples distance and minimise the distance between negative samples.

**Wilkinghoff\_FKIE** adopts template matching and dynamic time warping applied to embeddings trained with temporal resolution. The embedding model is trained to predict both class and temporal position of the sound event.

Observing the results spanning the two-year period (see Fig. 2), it is evident that each dataset presents unique challenges for various algorithms. Notably, the QU dataset consistently proved to be difficult for all participating teams across both years. Certain aspects, which are not fully discussed here, have had an important impact in

the performance of these systems. We highlight data augmentation and post processing techniques adapted to the task which have been fundamental to achieve a good score. A comprehensive analysis and discussion of these aspects can be read in [5].

## 5. CONCLUSION

The 2023 edition of the few-shot bioacoustic event detection task received some very innovative systems that reflect the state-of-the-art in few-shot learning. We especially underscore the introduction of a novel technique, such as contrastive learning, making its initial entry in the history of the task’s execution. Contrastive learning in the audio domain has seen increasing success and seems like a promising approach for the few-shot problem. Also of note is the quality of the evaluation set gathered this year. The dataset now extends to 3 different taxonomic groups: mammals, birds and amphibians, which is a good indicator of the variety of challenges faced in the bioacoustics domain. Moving forward we would be interested in analysing how exactly the characteristics of the different datasets impact each system and be able to understand if a single general model is indeed capable of predicting many different classes based on such few examples. The work in [5] started to tackle these questions, and while it is still not clear, the improving results on successive editions of this task indicate that the few-shot setting is a way to go.

## 6. REFERENCES

- [1] A. Mesaros, T. Heittola, A. Eronen, and T. Virtanen, “Acoustic event detection in real life recordings,” in *2010 18th European signal processing conference*. IEEE, 2010, pp. 1267–1271.
- [2] 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, pp. 1733–1746, 2015.

- [3] N. Turpault, R. Serizel, A. P. Shah, and J. Salamon, “Sound event detection in domestic environments with weakly labeled data and soundscape synthesis,” in *Workshop on Detection and Classification of Acoustic Scenes and Events*, 2019.
- [4] D. Stowell, “Computational bioacoustic scene analysis,” *Computational analysis of sound scenes and events*, pp. 303–333, 2018.
- [5] I. Nolasco, S. Singh, V. Morfi, V. Lostanlen, A. Strandburg-Peshkin, E. Vidaña-Vila, L. Gill, H. Pamuła, H. Whitehead, I. Kiskin, *et al.*, “Learning to detect an animal sound from five examples,” *arXiv preprint arXiv:2305.13210*, 2023.
- [6] D. Stowell, L. Gill, and D. Clayton, “Detailed temporal structure of communication networks in groups of songbirds,” *Journal of the Royal Society Interface*, vol. 13, no. 119, p. 20160296, 2016.
- [7] J. Snell, K. Swersky, and R. Zemel, “Prototypical networks for few-shot learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [8] Y. Wang, N. J. Bryan, M. Cartwright, J. P. Bello, and J. Salamon, “Few-shot continual learning for audio classification,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 321–325.
- [9] I. Nolasco, S. Singh, E. Vidaña-Vila, E. Grout, J. Morford, M. Emmerson, F. Jensen, I. Kiskin, H. Whitehead, A. Strandburg-Peshkin, *et al.*, “Few-shot bioacoustic event detection at the dcase 2022 challenge,” in *Detection and Classification of Acoustic Scenes and Events*, 2022.
- [10] <https://dcase.community/challenge2023/task-few-shot-bioacoustic-event-detection-results>, accessed: 10-06-2023.
- [11] <http://dcase.community/workshop2023/>.
- [12] J. Tang, Z. Xueyang, T. Gao, D. Liu, X. Fang, J. Pan, Q. Wang, J. Du, K. Xu, and Q. Pan, “Few-shot embedding learning and event filtering for bioacoustic event detection technical report,” DCASE2022 Challenge, Tech. Rep., June 2022.
- [13] H. Liu, X. Liu, X. Mei, Q. Kong, W. Wang, and M. D. Plumbley, “Surrey system for dcase 2022 task 5 : Few-shot bioacoustic event detection with segment-level metric learning technical report,” DCASE2022 Challenge, Tech. Rep., June 2022.
- [14] K. J. Piczak, “ESC: Dataset for Environmental Sound Classification,” in *Proceedings of the 23rd Annual ACM Conference on Multimedia*. ACM Press, pp. 1015–1018. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2733373.2806390>