

AUTOMATIC SOUND EVENT DETECTION AND CLASSIFICATION OF GREAT APE CALLS USING NEURAL NETWORKS

Zifan Jiang^{1,2}, Adrian Soldati^{1,3}, Isaac Schamberg², Adriano R. Lameira⁴, Steven Moran^{1,5}

¹University of Neuchâtel, ²University of Zurich, ³University of St Andrews,

⁴University of Warwick, ⁵University of Miami

jiang@cl.uzh.ch, {adrian.soldati, steven.moran}@unine.ch, isaac.schamberg@uzh.ch, adriano.lameira@warwick.ac.uk

ABSTRACT

We present a novel approach to automatically detect and classify great ape calls from continuous raw audio recordings collected during field research. Our method leverages deep pretrained and sequential neural networks, including wav2vec 2.0 and LSTM, and is validated on three data sets from three different great ape lineages (orangutans, chimpanzees, and bonobos). The recordings were collected by different researchers and include different annotation schemes, which our pipeline preprocesses and trains in a uniform fashion. Our results for call detection and classification attain high accuracy. Our method is aimed to be generalizable to other animal species, and more generally, sound event detection tasks. To foster future research, we make our pipeline and methods publicly available.¹

Keywords: sound event detection, neural networks, primatology, phonetics, computational linguistics

1. INTRODUCTION

In primatology, as in documentary linguistics, the collection, annotation, and analysis of primary field data are time-consuming and expensive. The recordings of primate calls are also often undertaken in not ideal environmental conditions (e.g., in a dense and noisy forest where other vocal species are present), making the process even more challenging. Therefore it is typical that primatologists first manually annotate their recordings and then conduct acoustic analyses on their species' specific calls.

Our goal in this paper is to automatically and accurately detect and classify primate calls from raw audio data. After discussing related work (§2), we describe our data sets (§3) and present our method that leverages pretrained and sequential neural networks (§4). We carry out reliable and reproducible experiments (§5) to support the effectiveness of our approach and compare our results between different architectural choices on

the data sets of three great ape lineages. Overall, our models achieve more than 80% frame-level classification accuracy and weighted F1-score. Interestingly, we find that the wav2vec 2.0 [1] model – even though pretrained on a large corpus of human speech [2] – generalizes surprisingly well as an audio feature representation layer for great ape calls without additional fine-tuning. This finding perhaps bides well with the similar vocal tract morphology between extant great apes and humans, and the larger prediction that ancestral ape-like calls evolved to become the building blocks of human speech.

2. RELATED WORK

To the best of our knowledge, there is no precedent research on the automatic detection and classification of great ape calls from raw audio. Therefore, we find our work lies in between general-purpose sound event detection and human speech recognition. We briefly discuss these also in light of existing research on animal sound classification.

Sound Event Detection (SED) aims at automatically recognizing *what* is happening in an audio signal and *when* it is happening [3]. SED tasks are usually general-purpose (e.g., birds singing or footsteps) as opposed to domain-specific tasks like human speech or music analysis, and thus encounter particular challenges, such as the *cocktail party effect* [4]. Notably, the DCASE challenge that took place in recent years involves relevant SED tasks [5]. A variant of SED is the audio tagging task [6], in which only *what* is happening in an audio signal is annotated and recognized, but not *when*.

Prominent work on both tasks [7, 8] tends to consist of common components, including: (1) an acoustic feature representation layer, either (a) traditional spectrogram-based approaches or (b) convolutional neural networks (CNN) on spectrogram features or on raw waveforms [9]; (2) a sequence modeling layer that captures temporal interaction between frame-level features, which takes the forms of mean/max pooling strategies,

data set	# audio clips	mean / call / total duration	# indiv. (♂ / ♀)	# call types (duration ratio) / units
chimpanzee	235	~ 8s / 1,955s / 1,964s	11 (11 / 0)	4 (6:3:3:1) / 686
orangutan	65	~ 74s / 2,793s / 4,817s	10 (10 / 0)	7 (700:40:200:100:20:1:1) / 9016
bonobo	28	~ 24s / 62s / 677s	7 (3 / 4)	18 (20:40:10:20:....:200:....:5:1) / 356

Table 1: Data set overview and stats. Shown in the table are the number of audio clips, mean duration of each clip, total duration of calls of all clips, total duration of all clips, number of individuals (male and female), number of call types (approximate total duration ratio of them, partially omitted for bonobo due to space limit), and number of individually annotated call units.

recurrent neural networks (RNN) or Transformers [10]; and (3) an objective function, usually cross-entropy classification on frame-level (for SED) or clip-level (for audio tagging), or occasionally CTC [11] for “sequentially labeled data” [12].

Animal sound classification involves related research that addresses animal sounds [13, 14, 15]. While this line of research uses similar acoustic feature representation techniques as seen in SED tasks, the models tend to work on top of individual short audio units of animal vocalizations that are manually selected from raw audio recordings by human experts, and thus fall short on dealing with a continuous audio signal that contains many audio events as well as noise. Instead, our method detects and classifies great ape calls directly from raw recordings ranging from seconds to minutes (theoretically recurrent models also extend to hours, but it could pose challenges in model training time).

Automatic Speech Recognition (ASR) on humans – one of the extant five great apes – has seen significant progress [16] (while animal calls remain mysterious to fully decipher), including the recent work wav2vec 2.0 [1]. We are interested in whether the wav2vec 2.0 speech representation pretrained on Librispeech [2] can be applied successfully to great ape call detection and classification.

3. DATA

Our data consists of recordings of chimpanzee pant-hoots, orangutan long calls, and bonobo high-hoots. Table 1 provides an overview of these data sets.

Chimpanzee pant-hoots are vocal sequences composed of up to four acoustically distinct phases, typically produced in this order: *introduction*, *build-up*, *climax*, and *let-down* [17]. Most pant-hoots contain two or more phases, although single phases can be produced in specific contexts [18]. The successive nature and well-balanced phase duration facilitate the training of a classifier (§5.1). We use annotated recordings of isolated pant-hoots (i.e., no temporal overlap with others’ calls) produced during feeding, traveling, and resting context.

Orangutan long calls are composed by a *full pulse*, which is sub-divided into a *sub-pulse transitory element* and a *pulse body*, or into sequences of *bubble sub-pulse* or *grumble sub-pulse* [19, 20]. The complex temporal interrelationship between phases and the overlapping and class-unbalanced characteristics (some phases are extremely short) pose challenges for training a multi-class classifier. On the other hand, the duration of calls and non-calls in the data set are well balanced, which opens the gate to a binary call detection model (§5.2). Annotated recordings include spontaneous long calls and long calls produced in response to other males’ vocal presence or environmental disturbances.

Bonobo high-hoots – which make up the plurality of our bonobo call data set – are loud, tonal vocalizations [21]. Unlike the chimpanzee and orangutan data sets, call types other than the homologous loud calls are also present in the data set, but are rather minor and diverse (e.g., peep-yelp, soft barks), which could be challenging to find. The total call duration is relatively short.

4. METHOD

This section introduces our method, which is illustrated in Fig. 1. We describe each step in turn.

4.1. Audio Preprocessing

First, all audio clips are converted to *.wav* format and resampled to 16 kHz sample rate. We segment them into 20ms frames and pad with zeros if necessary.

4.2. Feature and Label Extraction

Next, we extract three types of acoustic features with the *torchaudio* Python package. For an audio clip of T frames, we get a sequence of frame-level features $I_{1:T}$, in the shape of $T \times feature_dim$:

- **Raw waveform:** the original amplitude values over time. $feature_dim = 0.02 \times 16000 = 320$.
- **Spectrogram:** calculated then from the raw waveform with $feature_dim = 201$.

- **wav2vec 2.0**: inferred from the raw waveform by the *WAV2VEC2_BASE* model on a CPU device with $feature_dim = 768$.

We then extract frame-level labels from our annotations. For each annotated unit, we mark all frames inside the annotated time span with a positive class index indicating the call type (e.g., 1 for *intro*); unmarked frames are 0s by default, indicating non-calls. The labels' shape of a clip $L_{1:T}$ is $T \times 1$.

4.3. Data Split

We shuffle the clips (features and labels) of each data set and split them into 80%, 10%, and 10% for training, validation, and test sets, respectively. We make three different splits by three different random seeds 0, 42, and 3407 [22] to allow multiple experiments and to study the effect of randomness.

4.4. Sequence Modeling

Since our goal is to learn a function f from $I_{1:T}$ to $L_{1:T}$, we first input $I_{1:T}$ to a sequence modeling function f^m that outputs a hidden sequence $H_{1:T}$ ($T \times hidden_dim$), which captures inter-frame interaction. $H_{1:T}$ then goes through a dense linear function f^d (output $D_{1:T}$) followed by a Softmax function f^s to produce the probability distribution over target classes, i.e., $P_{1:T}$ ($T \times num_class$).

Specifically, f^m takes the following forms:

- **RNN**: a (bidirectional) LSTM with hidden size 1024. $hidden_dim = 1024 * 2 = 2048$.
- **Transformer encoder**: a Transformer encoder with 8 heads and 6 layers. $hidden_dim = 1024$.
- **Autoregressive model**: we optionally add autoregressive connections between time steps to encourage consistent output labels. The output of f^d at time step t , i.e., D_t is concatenated to the next time step's input I_{t+1} . By default, it indicates that the RNN layer goes unidirectionally, but one can also stack two autoregressive layers and add them up before the Softmax operation to retain bidirectionality.

Lastly, a cross-entropy loss is computed between the model output $P_{1:T}$ and the gold labels $L_{1:T}$, then backpropagated to train the models. To mitigate the class imbalance problem, we set class weights to the reciprocal of each class's occurrence times.

5. EXPERIMENTS AND RESULTS

Our experiments were done in PyTorch [23], with Python 3.8 on an Nvidia Tesla V100 GPU (32GB ram). Most models have $\sim 40k$ million parameters and finish the training of up to 200 epochs with early

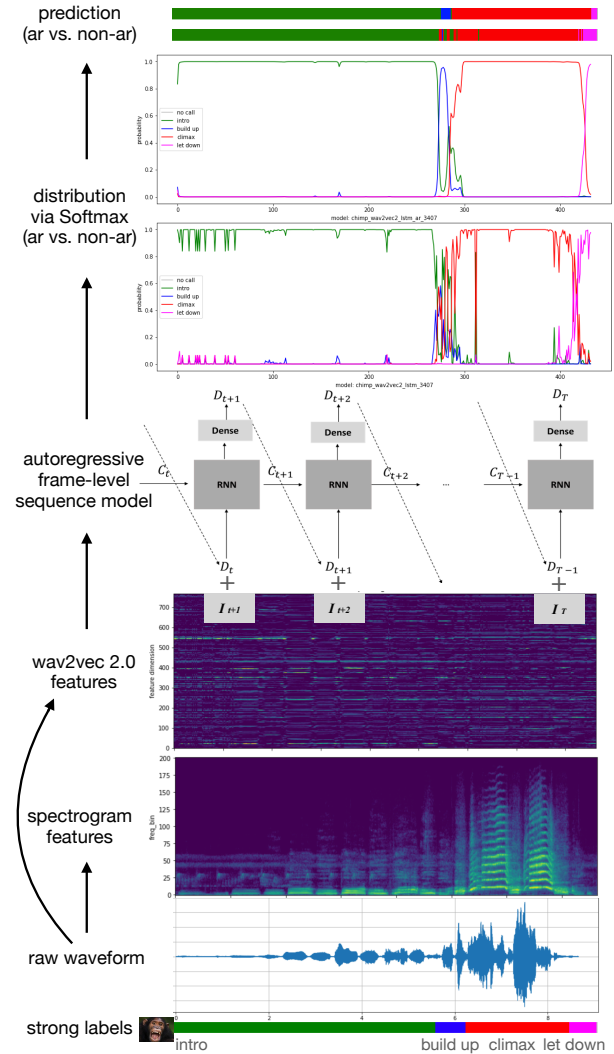


Figure 1: Our method (bottom-up). The audio clip shown in the figure is from the chimpanzee data set. (*Non-ar* stands for (non-)autoregressive.

stopping on validation F1-score within one hour (or a few hours because autoregressive recurrent models run slowly on PyTorch). Table 2 presents our results.

5.1. Initial Exploration with Chimpanzee Data

We first test the viability of our approach on the chimpanzee data in light of the simplicity of the pant-hoot annotation scheme. We start from *E1*, a simple waveform + LSTM baseline, and observe in *E2* that wav2vec 2.0 outperforms the raw waveform and spectrogram by a large margin, which demonstrates the power of transfer learning from pretraining on human speech. In *E2.1*, we find that the Transformer encoder does not outperform LSTM. Hence, we infer that Transformer's ability to capture arbitrary long-range dependency is not

ID	data	feature	model	dev acc.	dev f1	test acc.	test f1	aucpr
<i>Explore the best feature and model combination</i>								
E1	chimp	waveform	lstm (baseline)	51.7 ± 2.1	35.4 ± 2.5	51.0 ± 3.6	34.7 ± 4.3	-
E1.1	chimp	spectrogram	lstm	60.3 ± 1.5	55.7 ± 2.3	58.7 ± 4.7	53.9 ± 5.4	-
E2	chimp	wav2vec2	lstm	81.0 ± 4.0	79.9 ± 4.6	79.3 ± 2.3	77.9 ± 3.6	-
E2.1	chimp	wav2vec2	transformer	71.3 ± 0.6	68.3 ± 0.2	75.3 ± 0.6	72.1 ± 0.5	-
<i>Explore the hyper-parameters</i>								
E3.1	chimp	wav2vec2	lstm (E2 + <i>batch_size</i> = 4)	69.7 ± 1.5	71.8 ± 2.6	67.7 ± 4.0	69.6 ± 4.0	-
E3.2	chimp	wav2vec2	lstm (E2 + <i>batch_size</i> = 8)	63.3 ± 0.6	62.6 ± 1.0	62.0 ± 4.4	61.5 ± 4.0	-
E3.3	chimp	wav2vec2	lstm (E2 + <i>dropout</i> = 0.2)	80.7 ± 3.5	80.0 ± 4.4	78.0 ± 1.7	76.8 ± 2.7	-
E3.4	chimp	wav2vec2	lstm (E2 + <i>dropout</i> = 0.1)	81.0 ± 4.0	80.2 ± 4.8	78.7 ± 2.9	77.3 ± 3.9	-
E3.5	chimp	wav2vec2	lstm (E2 - <i>balance_weights</i>)	81.0 ± 3.6	79.6 ± 4.4	79.3 ± 2.3	78.3 ± 3.6	-
<i>Explore autoregressive modeling</i>								
E4	chimp	wav2vec2	lstm (E2 + autoregressive)	87.7 ± 1.2	87.1 ± 1.8	85.7 ± 2.1	85.6 ± 2.5	-
<i>Extend to orangutan long calls and a binary setting</i>								
E5	orang	wav2vec2	lstm (= E4)	83.0 ± 1.0	82.7 ± 1.4	81.7 ± 3.1	82.0 ± 2.6	-
E5.1	orang	wav2vec2	lstm (E5 + binary target)	92.3 ± 2.5	92.1 ± 2.5	92.0 ± 1.0	91.9 ± 1.1	0.96
<i>Extend to bonobo calls and a binary setting</i>								
E6	bonobo	wav2vec2	lstm (= E4)	87.0 ± 4.6	85.9 ± 6.3	83.7 ± 3.8	82.3 ± 2.2	-
E6.1	bonobo	wav2vec2	lstm (E6 + binary target)	92.0 ± 3.6	91.9 ± 3.4	87.7 ± 3.5	87.8 ± 2.9	0.87
<i>Zero-shot transferring from orangutan to bonobo</i>								
E7	bonobo	wav2vec2	lstm (= E5.1)	63.0 ± 13	69.2 ± 10	72.0 ± 4.0	74.2 ± 3.1	0.55

Table 2: Experimental results. We run all experiments three times based on different random seeds and report the mean and standard deviation. *acc.* stands for frame-level accuracy, *f1* stands for the frame-level average F1-score weighted by the number of true instances per class, and *aucpr* stands for the area under the precision-recall curve for the positive class in the binary case at test time when the random seed is set to 0. For hyper-parameters, we start *E1* with *batch_size* = 1, *dropout* = 0.4 and keep them by default, if not otherwise specified in the table.

beneficial to our task and/or our small data sets. Next, we explore the hyper-parameters in the second experimental group and we find in *E3.5* that balancing class weights (§4.4) has a small impact. Lastly, we show in *E4* that the autoregressive connections are beneficial for consistent output, as illustrated by the tiny gaps in *non-ar* output in Fig. 1.

5.2. Extending to Orangutan and Bonobo Data

We successfully extend the model in *E4* to as trained on the orangutan (*E5*) and bonobo (*E6*) data sets. We note that some minority classes perform less well due to data scarcity, in contrast to the well-balanced situation in the chimpanzee data set (see Table 1).

We further reduce the task to a binary (call vs. non-call) classification that resembles voice activity detection of human speech. This is a useful tool to automatically extract calls from raw recordings for further bioacoustic analysis (e.g., studying the repertoire of a given species).

Finally, to understand the generalizability of our models, we try zero-shot transferring the model trained in *E5.1* on orangutans directly to unseen bonobo data. The results show that it is promising to build a potential general-purpose sound event

detection model for all great ape calls.

6. DISCUSSION

We have addressed a gap in the bioacoustics research of non-human great apes by developing an approach for automatically identifying sound events and classifying great ape calls using a neural network architecture. Our method successfully and accurately identifies and classifies calls in three species of non-human great apes, and provides a tool for primatologists to bootstrap call identification and analysis from raw unannotated audio recordings.

Our method also shows the general applicability of the wav2vec 2.0 model trained on human speech for identifying vocalizations and call types in other species. Future work may apply our approach to more animals, as part of the goal to decode the communication systems of great apes and other non-human animals more broadly.²

¹ https://github.com/J22Melody/sed_great_ape

² For example: <https://www.earthspecies.org>.

7. ACKNOWLEDGEMENTS

We gratefully acknowledge funding from: UK Research & Innovation, Future Leaders Fellowship (MR/T04229X/1; ARL), Swiss National Science Foundation (PCEFP1_186841: SM, ZJ; 310030_185324: AS), St Leonard College (AS), and Swissuniversities (AS). We thank also Klaus Zuberbühler, Josep Call, and the field assistants of the Budongo Conservation Field Station.

8. REFERENCES

- [1] A. Baeviski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [2] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [3] A. Mesaros, T. Heittola, T. Virtanen, and M. D. Plumbley, “Sound event detection: A tutorial,” *IEEE Signal Processing Magazine*, vol. 38, no. 5, pp. 67–83, 2021.
- [4] B. Arons, “A review of the cocktail party effect,” *Journal of the American Voice I/O Society*, vol. 12, no. 7, pp. 35–50, 1992.
- [5] N. Turpault, R. Serizel, A. Parag 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*, New York City, United States, October 2019. [Online]. Available: <https://hal.inria.fr/hal-02160855>
- [6] A. Mesaros, T. Heittola, E. Benetos, P. Foster, M. Lagrange, T. Virtanen, and M. D. Plumbley, “Detection and classification of acoustic scenes and events: Outcome of the DCASE 2016 challenge,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 2, pp. 379–393, Feb 2018.
- [7] J. Ebberts and R. Haeb-Umbach, “Self-trained audio tagging and sound event detection in domestic environments,” in *Proceedings of the 6th Detection and Classification of Acoustic Scenes and Events 2021 Workshop (DCASE2021)*, 2021.
- [8] Y. Gong, Y.-A. Chung, and J. Glass, “Ast: Audio spectrogram transformer,” *arXiv preprint arXiv:2104.01778*, 2021.
- [9] W. Dai, C. Dai, S. Qu, J. Li, and S. Das, “Very deep convolutional neural networks for raw waveforms,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 421–425.
- [10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [11] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd International Conference on Machine Learning*, 2006, pp. 369–376.
- [12] Y. Hou, Q. Kong, J. Wang, and S. Li, “Polyphonic audio tagging with sequentially labelled data using crnn with learnable gated linear units,” *arXiv preprint arXiv:1811.07072*, 2018.
- [13] P. Fedurek, K. Zuberbühler, and C. D. Dahl, “Sequential information in a great ape utterance,” *Scientific Reports*, vol. 6, no. 1, pp. 1–11, 2016.
- [14] Y. Sun, T. M. Maeda, C. Solis-Lemus, D. Pimentel-Alarcon, and Z. Burivalova, “Classification of animal sounds in a hyperdiverse rainforest using convolutional neural networks,” *arXiv preprint arXiv:2111.14971*, 2021.
- [15] N. M. Recalde, “pykanto: a Python library to accelerate research on wild bird song,” *arXiv preprint arXiv:2302.10340*, 2023.
- [16] J. Li *et al.*, “Recent advances in end-to-end automatic speech recognition,” *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [17] P. Marler and L. Hobbett, “Individuality in a long-range vocalization of wild chimpanzees,” *Zeitschrift für Tierpsychologie*, vol. 38, no. 1, pp. 97–109, 1975.
- [18] A. Soldati, P. Fedurek, G. Dezechache, J. Call, and K. Zuberbühler, “Audience sensitivity in chimpanzee display pant hoots,” *Animal Behaviour*, vol. 190, pp. 23–40, 2022.
- [19] M. Hardus, “A description of the orangutan’s vocal and sound repertoire, with a focus on geographical variation,” *Orangutans: Geographic Variation in Behavioral Ecology and Conservation*, pp. 49–64, 2009.
- [20] A. R. Lameira and S. A. Wich, “Orangutan long call degradation and individuality over distance: a playback approach,” *International Journal of Primatology*, vol. 29, no. 3, pp. 615–625, 2008.
- [21] I. Schamberg, D. L. Cheney, Z. Clay, G. Hohmann, and R. M. Seyfarth, “Call combinations, vocal exchanges and interparty movement in wild bonobos,” *Animal Behaviour*, vol. 122, pp. 109–116, 2016.
- [22] D. Picard, “Torch. manual_seed (3407) is all you need: On the influence of random seeds in deep learning architectures for computer vision,” *arXiv preprint arXiv:2109.08203*, 2021.
- [23] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.