BioCPPNet: Automatic Bioacoustic Source Separation with Deep Neural Networks

Peter C Bermant^{1,*}

¹Earth Species Project, Berkeley, CA, 94709, USA

*peter@earthspecies.org

ABSTRACT

We introduce the Bioacoustic Cocktail Party Problem Network (BioCPPNet), a lightweight, modular, and robust UNet-based machine learning architecture optimized for bioacoustic source separation across diverse biological taxa. Employing learnable or handcrafted encoders, BioCPPNet operates directly on the raw acoustic mixture waveform containing overlapping vocalizations and separates the input waveform into estimates corresponding to the sources in the mixture. Predictions are compared to the reference ground truth waveforms by searching over the space of (output, target) source order permutations, and we train using an objective function motivated by perceptual audio quality. We apply BioCPPNet to several species with unique vocal behavior, including macaques, bottlenose dolphins, and Egyptian fruit bats, and we evaluate reconstruction quality of separated waveforms using the scale-invariant signal-to-distortion ratio (SI-SDR) and downstream identity classification accuracy. We consider mixtures with two or three concurrent conspecific vocalizers, and we examine separation performance in open and closed speaker scenarios. To our knowledge, this paper redefines the state-of-the-art in end-to-end single-channel bioacoustic source separation in a permutation-invariant regime across a heterogeneous set of non-human species. This study serves as a major step toward the deployment of bioacoustic source separation systems for processing substantial volumes of previously unusable data containing overlapping bioacoustic signals.

Introduction

Bioacoustic source separation, informally referred to as the "cocktail party problem" (CPP), encompasses the problem of detecting, recognizing, and extracting meaningful information from conspecific signaling interactions in the presence of concurrent vocalizers in a noisy social environment¹. Automatic source separation functions to isolate individual speaker vocalizations from a recording with overlapping signals. While human speech separation is an active and well-studied area of work involving the recent emergence of supervised deep neural networks (DNNs) for speaker separation²⁻⁷, the bioacoustic CPP remains comparatively understudied⁸⁻¹¹. Occurrence of overlapping acoustic mixtures in recordings is especially common and problematic in the non-human bioacoustic domain. For instance, 58% of recordings of African elephant (*Loxodonta africana*) vocalizations contained at least two concurrent signalers¹²; 16% of sperm whale (*Physeter macrocephalus*) codas were overlapped by codas generated by another whale, and 22% were followed by another coda within 2s¹³, which is similar to how geladas (*Theropithecus gelada*) synchronize the onsets of their vocalizations¹⁴; in the NIPS4Bplus richly annotated birdsong dataset¹⁵, of the duration of recordings with vocalizations, nearly 20% contain simultaneously active classes. With no formal method for isolating individual-specific calls, biologists are often forced to omit data containing overlapping calls produced by simultaneously signaling vocalizers. In this paper, we offer a lightweight and modular neural network (NN) that acts on raw waveforms directly to reconstruct sources from recordings containing signal mixtures. Our machine learning (ML) model, the Bioacoustic Cocktail Party Problem Network (BioCPPNet), serves as an effective end-to-end source separation system to extract source estimations from composite mixtures, and we optimize the architecture specifically to accommodate bioacoustic vocal diversity across a range of taxonomic groups, even in limited data regimes.

Throughout the entire pipeline from data preprocessing to separator model evaluation, non-human bioacoustic source separation presents numerous challenges that complicate the problem relative to human speech separation. With non-human animals, technical and logistical difficulties associated with procuring bioacoustic data render it challenging to compile comparably-sized annotated datasets, and unique non-human vocal behavior demands differential treatment according to the species of interest. In the human speech domain, it is common practice to use large datasets of high-quality recordings downsampled to 8 kHz² to minimize computational costs and to train models on segments several seconds in duration. However, this approach to data preprocessing is not always feasible for non-human animals. For instance, while the fundamental frequency (F0) of macaque vocalizations is on the order of 0.1-1 kHz¹⁶, which suggests that they could be treated similarly as human speech in terms of downsampling to reduce computational costs, the mean maximum F0 of bottlenose dolphin (*Tursiops truncatus*) signature whistles can range from 9.3-27.3 kHz¹⁷, which means that with sampling rates of 96 kHz, even the lowest

order overtones approach or exceed the native Nyquist frequency. This is further exacerbated in animals such as bats, which can emit broadband calls with dominant frequencies ranging from 11-212 kHz¹⁸. Given the wide-ranging vocal behavior of non-human species, it is often infeasible to downsample recordings, resulting in long input sequences that pose a challenge to modern ML models. Though music source separation models such as Demucs¹⁹ have addressed higher sampling rates up to 44.1 kHz, the spectrotemporal parameters of non-human acoustic signals demand that we consider source separation with sampling rates up to 250kHz while still employing sufficiently long input sequences to reflect timescales related to the relevant biological behavior. The combination of relatively small datasets recorded with high sampling rates has significant implications for automatic bioacoustic source separation. In particular, time-domain separator NNs may struggle to capture long term temporal dependencies in the input time series, and recurrent-based approaches are often prohibitively expensive in terms of memory and computational costs. Further, heavyweight models such as those used to address human speech separation may overfit small datasets, impairing the models' capacities to generalize. With BioCPPNet, we construct a lightweight model both to efficiently separate long mixture waveform sequences and to minimize the risk of overfitting small bioacoustic datasets.

In general, two families of algorithms have been implemented to solve the human speech CPP. These include frequency masking approaches^{2,20} that operate on handcrafted time-frequency representation (TFR) inputs generated by defining a mapping $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}^{T \times F}$ that projects the N -sample acoustic waveform to a spectral representation with T features of dimension F ; and time-domain methods^{3,5} that operate on raw acoustic waveforms with learnable encoders. While existing research²¹ has compared various options for the TFR input, frequency masking-based methods for human speech separation commonly implement the mel-scaled short-time Fourier transform (STFT)-based spectrogram²² given its correspondence to human auditory processing and perception. In the bioacoustic regime, it remains unclear which TFRs may be optimal for ML applications. Though spectrograms or log-magnitude spectrograms are conventionally employed regardless of vocal characteristics, there exist numerous options for TFRs according to the vocal properties of the particular species of interest. For example, the Hilbert Huang transform (HHT) may provide advantages over Fourier-based analysis for transient signals such as those produced by sperm whales²³; wavelet transforms have been used for automated birdsong detection²⁴; though not extensively implemented in the bioacoustic literature, invariant scattering representations^{25,26} provide yet another option for representing animals sounds. Further, there also exist learnable spectral feature representations generated using convolutional neural networks (CNNs) acting on raw waveforms,^{3,27,28} but these fully learnable transforms remain unexplored in the bioacoustic domain. To address the fundamental bioacoustic representation problem, BioCPPNet provides a modular architecture to enable rapid experimentation using handcrafted or fully learnable encoders in combination with various fixed or learnable inverse transform decoders to recover the separated raw acoustic waveforms. In this study, we employ various STFT-based or learnable representations, but the architecture is compatible with other TFRs, which we defer to future studies.

Numerous additional factors further complicate the bioacoustic CPP. For instance, native environmental conditions often include acoustic interferences and energetic maskings that impair the ability of animals to recognize and discriminate signals. These interferences can consist of spectrotemporally overlapping calls produced by heterospecific signalers, other sources of biotic and abiotic noise, and anthropogenic sounds¹. With these challenges, it remains difficult to construct sufficiently large and suitably annotated datasets to enable the training of supervised learning algorithms with access to ground truth signals. Given this limited and noisy data regime, we design BioCPPNet as a less complex model than many existing human speech separation networks³⁻⁵, and we integrate fixed or learnable noise reduction into the model architecture. Finally, compared to human speech separation applications, bioacoustic source separation models are more difficult to evaluate. While subjective assessment of the perceptual quality of human speech separation models can rely on panels of human listeners³, the human auditory system is often not well-suited for evaluating non-human vocalizations. This, in combination with the obscure interpretability of objective metrics such as the scale-invariant signal-to-distortion ratio (SI-SDR)²⁹ in the context of different species with different vocal behaviors, implies that the evaluation of bioacoustic source separation models could benefit from other metrics. Qualitatively, we assess the separation performance of BioCPPNet using visual representations of the acoustic data in the form of spectrograms, and we quantitatively address the performance of BioCPPNet by considering downstream tasks. In particular, following the separation of the mixture into estimated source outputs, we feed the predictions into trained individual identity classifier models under the assumption that classification accuracy on downstream ML-based tasks serves as an objective proxy for the quality of reconstructed separated waveforms.

While the human speech separation problem is a competitive area of work, the bioacoustic CPP has received comparatively less attention, as current bioacoustic research often emphasizes other ML-based tasks such as automated detection and classification of bioacoustic sounds³⁰⁻³². However, recent work has implemented both semi-classical and deep ML-based approaches to address bioacoustic source separation, employing time domain and TFR-based algorithms. Fast fixed-point independent component analysis (FastICA), principal component analysis (PCA), and non-negative matrix factorization (NMF) have been used to separate mixtures of overlapping frog sounds⁹. Mixed signals of overlapping dolphin signature whistles were separated with the joint approximate diagonalization of eigenmatrix (JADE) algorithm⁸. Time-frequency masking was used to address humpback whale (*Megaptera novaeangliae*) song separation³³. More contemporary approaches have made

use of DNNs, but they remain limited in their scope. For instance, bi-direction long short-term memory (BLSTM) networks were implemented to separate overlapping bat echolocation and communication calls¹⁰. In this application, the separation of composite mixtures containing signals of disparate classes or natures is more analogous to human music instrument separation into predefined classes³⁴; it treats the separation problem as a multi-class regression problem and thus avoids the fundamental permutation problem² of the CPP in which the source estimate channel order is arbitrary, potentially yielding conflicting gradients during training³⁵. Other DNN-based approaches have adopted multi-step schemes to predict source number and to extract mask estimates using a segmentation-based approach to bat call separation optimized for fundamental frequency contours¹¹; however, the omission of higher-order harmonics presents a dilemma since in biological systems, the inharmonic mistuning of spectral components may serve as an important acoustic cue for overcoming the CPP³⁶; further, this bounding box-based approach performs best on signals containing low-to-moderate spectrotemporal overlap, which means they may be limited in their generalizability or practical implementation.

BioCPPNet represents a complete state-of-the-art permutation-invariant bioacoustic source separation pipeline across multiple species characterized by differential vocal behavior. The model is specifically designed to address key challenges of bioacoustic data, and we apply it to a diverse set of non-human species, including rhesus macaques (*Macaca mulatta*), bottlenose dolphins, and Egyptian fruit bats (*Rousettus aegyptiacus*). The multidisciplinary approach requires that we integrate a range of fields of study, which means the scope of our treatment is necessarily broad. We advocate for future studies to expand on our work.

Methods

Our novel approach to bioacoustic source separation involves an end-to-end pipeline consisting of multiple discrete steps, including (1) synthesizing a dataset, (2) developing and training a separator network to disentangle the input mixture, and (3) constructing and training a classifier model to employ as a downstream evaluation task. This workflow requires few species-level hyperparameter modifications to account for unique vocal behavior across different biological taxa but is otherwise general and makes no assumptions about the spectrotemporal structure of the overlapping calls. We develop a complete framework for bioacoustic source separation in a permutation-invariant mode using overlapping waveforms drawn from the same class of signals. We apply BioCPPNet to macaques, dolphins, and Egyptian fruit bats, and we consider two or three concurrent "speakers". Note that we henceforth refer to non-human animal signalers as "speakers" for consistency with the human speech separation literature². We address both the closed speaker regime in which the training and evaluation data subsets contain calls produced by individuals drawn from the same distribution as well the open speaker regime in which the model is tested on calls generated by individuals not present in the training dataset.

Bioacoustic Data

We investigate a set of species with dissimilar vocal behaviors in terms of spectral and temporal properties. We apply BioCPPNet to a macaque coo call dataset³⁷ consisting of 7285 coos produced by 8 unique individuals; a bottlenose dolphin signature whistle dataset³⁸ comprised of 400 signature whistles generated by 20 individuals, of which we randomly select 8 for the purposes of this study; and an Egyptian fruit bat vocalization dataset³⁹ containing a heterogeneous distribution of individuals, call types, and call contexts. In the case of the bat dataset, we extract the data (31399 calls) corresponding to the 15 most heavily represented individual bats, reserving 12 individuals (27586 calls) to address the closed speaker regime and the remaining 3 individuals (3813 calls) to evaluate model performance in the open speaker scenario.

Datasets

Motivated by WSJ0-2mix², a preeminent reference dataset used for human single-channel acoustic source separation, we adopt a similar approach of constructing bioacoustic datasets by mixing multiple nonoverlapping ground truth vocalizations produced by individual animals to enable supervised training of our model. The dataset consists of an input array of the composite mixture waveforms, a target array containing the separated ground truth waveforms corresponding to the respective mixtures, and a class label array denoting the individual identities of the vocalizing animals responsible for generating the signals. For each species, the mixture dataset is generated from a corpus of bioacoustic recordings annotated according to the timestamps of the signal and the identity of the signaller. In the case of macaques, we here consider closed speaker set mixtures of two and three simultaneous speakers, but our method is functionally not limited in the number of sources (N) it can handle. For dolphins, we consider the closed speaker regime with two overlapping calls, and for bats, we consider both the closed and open speaker scenarios with two overlapping sources.

In general, we construct the dataset as follows. We first extract the labeled waveforms either by truncating or zero-padding the waveforms to ensure that all the samples are of fixed duration. We select the number of frames either by computing the mean plus three-sigma of the durations of the calls contained in the corpus from which we draw the signals, by selecting the maximum duration of all calls, or by choosing a fixed value. For macaques, dolphins, and bats, we use 23156 frames

(0.95s), 290680 frames (3.03s), and 250000 frames (1.0s), respectively. We then randomly select vocalizations from N different speakers drawn from the distribution of individuals used in the study (8 macaques, 8 dolphins, 12 bats for the closed speaker regime, 3 bats for the open speaker regime) and mix them additively, ensuring to randomly shift the overlaps to simulate a more plausible scenario and to provide for asynchronicity of start times, an important acoustic cue that has been suggested as a mechanism with which the animal brain can solve the CPP¹. Despite higher computational and memory costs, we opt to use native sampling rates, since certain animal vocalizations may reach frequencies near the native Nyquist frequency. With this in mind, however, our method does provide for resampling when the vocalizations of the particular species of interest are amenable to downsampling. Explicitly, for the three species we consider including macaques, dolphins, and bats, we use sampling rates of 24414 Hz, 96 kHz, and 250 kHz, respectively. For the closed speaker regime, the training and evaluation subsets contain calls produced by the same distribution of individuals to ensure a closed speaker set. We segment the original nonoverlapping vocalizations into 80/20 training/validation subsets. We generate the mixture training waveforms using 80% of the calls, and we construct the mixture validation subset using the remaining 20% of calls held out from the training data. In the case of overlapping bat calls (for which the corpus of bioacoustic recordings contains $\mathcal{O}(10$ hours) of data as opposed to $\mathcal{O}(10^{-1})$ hours) for macaques and dolphins), we also address the open speaker source separation problem by constructing a further testing data subset of mixtures of calls of additional vocalizers not contained in the training distribution. For macaques, we construct a training data subset comprised of 12k samples and a validation subset with 3k samples, all of which contain calls drawn from 8 animals. For dolphins, we randomly select 8 individuals and construct training/validations subsets with 8k and 2k samples, respectively. For bats, we select 15 individuals, randomly reserving 12 for the closed speaker problem and the remaining 3 for the open speaker situation. We train the bat separator model on 24k mixtures. We evaluate performance in both the closed and open speaker scenarios using data subsets consisting of 6k mixtures containing unseen vocalizations produced by the appropriate distribution of individuals according to the regime under consideration. We repeat the bat training using a larger mixture dataset (denoted by +) containing 72k samples. We here report validation metrics to ensure that we are evaluating model performance on unseen mixtures of unseen calls in the closed speaker regime and on unseen mixtures of unseen calls of unseen individuals in the open speaker regime.

For the downstream classification task, we extract vocalizations annotated according to the individual identity, and we segment the calls into an 80/20 training/testing split to ensure that we are evaluating model performance on unseen calls. For both the training and evaluation data subsets, we employ an augmentation scheme in which we apply random temporal shifts to call onsets to better reflect more plausible real-world scenarios.

Classification Models

In an effort to provide a more physically interpretable evaluation metric to supplement the commonly-implemented SI-SDR used in human speech separation studies, we develop CNN-based classifier models to label the individual identity of the separated vocalizations as a downstream task. This requires training classification networks to predict the speaker class label of the original unmixed waveforms. For each species we consider, we design and train custom simple and lightweight CNN-based architectures motivated by previous work³⁰, tailored to accommodate the unique vocal behavior of the given species.

The first layer in the model is an optional high pass filter constructed using a nontrainable 1D convolution (Conv1D) layer with weights determined using a windowed sinc function to eliminate low-frequency background noise. We omit this computationally intensive layer for macaques and Egyptian fruit bats, but we implement a high pass filter for the dolphin dataset, selecting an arbitrary cutoff frequency of 4.7 kHz to remove background without impinging on the region of support for dolphin whistles. After the optional filter is an encoder layer to compute on-the-fly feature extraction. We experimented with a fully learnable free Conv1D filterbank, a spectrogram, and a log-magnitude spectrogram and observed optimal performance using a non-decibel (dB)-scaled STFT layer computed with a $nfft$ window width, a hop window shift, and a Hann window where $nfft$ and hop are species-dependent variables. For macaques, we select $nfft=1024$ and $hop=64$ corresponding to temporal scales on the order of 40ms and frequency resolutions on the order of 20 Hz. We choose $nfft=1024$ and $hop=256$ for dolphins and $nfft=2048$ and $hop=512$ for bats, corresponding to temporal resolutions of 10ms and 8ms and frequency resolutions of 90 Hz and 120 Hz, respectively.

Following the built-in feature engineering, the architecture includes 4 convolutional blocks, which consist of two sequential 2D convolution (Conv2D) layers with leaky ReLU activation and a max pooling layer with pool size 4. Next is a dense fully connected layer with leaky ReLU activation followed by another linear layer with log softmax activation to output the V log probabilities (i.e. confidences) where V is the number of individual vocalizers used in the study (8, 8, 12 for macaques, dolphins, and bats, respectively). We also include dropout regularization with $p=0.25$ for the macaque classifier and $p=0.5$ for the dolphin and bat classifiers to address potential overfitting. With these architectures, the macaque, dolphin, and bat classifier models have 230k, 279k, and 247k trainable parameters, respectively.

For all species, we minimize the negative log-likelihood objective loss function using the Adam optimizer⁴⁰ with learning rate $lr=3e-4$. For macaques, dolphins, and bats, respectively, we train for 100, 50, and 100 epochs with batch sizes 32, 8, and 8.

We save the model after each epoch and select the top-performing models. We opt not to carry out hyperparameter optimization since the classification task is of secondary importance and is used solely as a downstream task.

Figure 1 – Source vocalization waveforms are mixed before being transformed to a learnable or handcrafted representation (Rep), which then passes through a 2-dimensional UNet⁴¹ composed of downsampling convolutional (DC) and upsampling convolutional (UC) blocks. The contracting blocks consist of two convolutional layers (C2D) with leaky ReLU activation and batch normalization (BN) followed by a max pooling (P). The expanding path consists of an upsampling (U) and skip connection concatenation (Conc) followed again by the C2D BN layers. The UNet predicts masks (M1 and M2), the number of which is determined by the number of sources (N), that are multiplicatively applied to the original mixture representation. The predicted representations for the separated waveforms are inverted (iRep) to output the raw waveforms, which are compared to the source ground truths, up to a permutation.

Separation Models

BioCPPNet (Fig. 1) is a lightweight and modular architecture with a modifiable representation encoder, a 2D UNet core, and an inverse transform decoder, which acts directly on raw audio via on-the-fly learnable or handcrafted transforms. The structure of the network is designed to provide for extensive experimentation, optimization, and enhancement across a range of species with variable vocal behavior.

Model Architecture

As with the classifier model, the network's encoder consists of a feature engineering block, the initial layer of which is an optional high pass filter. This is followed by the representation transform, which includes several options including the Conv1D free encoder, the STFT filterbank, and the log-magnitude (dB) STFT filterbank. We choose the same kernel size ($nfft$) and stride (hop) parameters defined in the classifier model. Sequentially following the feature extraction encoder is a 2D UNet core. This architecture consists of B (4 for macaques, 3 for dolphins, and 4 for bats) downsampling convolutional blocks, a middle convolutional block, and B upsampling convolutional blocks. The downsampling blocks consist of two 2D convolutional layers with filter number that increases with model depth with leaky ReLU activation followed by a max pooling with pool size 2, 6, and 3 for macaques, dolphins, and bats. The middle block contains two 2D convolutional layers with leaky ReLU activation. The upsampling blocks include an upsampling using the bilinear algorithm and a scale factor corresponding to the pool size used during downsampling, followed by skip connections in which each corresponding level of the contracting and expanding paths are concatenated before passing through two 2D convolutional layers with leaky ReLU activation. All convolutional layers in the downsampling, middle, and upsampling blocks include batch normalization after the activation function to stabilize and expedite training and to promote regularization. Though our default implementation is phase-unaware, we also offer the option for a parallel UNet pathway working directly on phase information, which has been shown to improve performance in other applications^{42–44}. The final layer in the UNet core is a 2D convolutional layer with N channels, which are then split

Figure 2 – A schematic diagram demonstrating the application of BioCPPNet to dolphin signature whistles using handcrafted STFT-based encoders and decoders. The source waveforms are overlapped and transformed to time-frequency space using an STFT layer. The UNet predicts masks that are applied to the mixture representation. The estimated spectrogram separations are inverted using an iSTFT layer, and the raw output waveforms are compared with the ground truth waveforms.

prior to entering the inverse transform decoder. For the inverse transform, we again provide numerous choices including a free filterbank decoder based on a 1D convolutional transpose layer, an iSTFT layer, an iSTFT layer accepting dB-scaled inputs, and a multi-head convolutional neural network (MCNN) for fast spectrogram inversion⁴⁵. In detail, the UNet returns N masks that are then multiplied by the original encoded representation of the mixture waveform. The separated representations are then passed into the inverse transform layer in order to yield the raw waveforms corresponding to the model's predictions for the separated vocalizations. We initialize all trainable weights using the Xavier uniform initialization. In the case of macaques, we experiment across all combinations of representation encoders and inverse transform decoders, and we find optimal performance using the handcrafted non-dB STFT/iSTFT layers operating in the time-frequency domain. Since the model with the fully learnable Conv1D-based encoder/decoder uniquely operates in the time domain, we report evaluation metrics for this model, as well. For dolphins and bats, we here report metrics using exclusively the non-dB STFT/iSTFT technique.

BioCPPNet (Fig. 1) is designed as a lightweight model in order to efficiently process large amounts of bioacoustic data sampled at high sampling rates while simultaneously minimizing computational costs and limitations and the likelihood of overfitting. For the macaque separators, the networks consist of 1.2M parameters (for the STFT, iSTFT combination), 2.5M parameters (for the STFT, iSTFT combination with parallel phase pathway), or 2.8M parameters (for the Conv1D free filterbanks). For the dolphin separator, the model has 304k parameters, while the bat separator model has 1.2M parameters. This is to be contrasted with the comparatively heavyweight default implementations of models commonly used in human speech separation problems, such as Conv-TasNet³, which has 5.1M parameters; DPTNet⁴ with 2.7M parameters; or Wavesplit⁵ with 29M parameters. Regardless of the lower complexity of BioCPPNet, the model achieves comparable performance or even outperforms reference human speech separator models while still being lightweight enough to train on a single NVIDIA P100 GPU.

Model Training Objective

The model training objective aims to optimize the reconstruction of separated waveforms from the aggregated composite input signal. We adopt a permutation-invariant training (PIT)⁴⁶ scheme in which the model's predicted outputs are compared with the ground truth sources by searching over the space of permutations of source orderings. This fundamental property of our training objective reflects that the order of estimations and their corresponding labels from a mixture waveform is not expressly germane to the task of acoustic source separation, i.e. separation is a set prediction problem independent of speaker identity ordering.

Source separation involves training a separator model f to reconstruct the source single-channel waveforms given a mixture $x = \sum_{i=1}^N s^i$ of N sources, where each source signal s^i for $i \in [1, N]$ is a real-valued continuous vector of fixed dimension T , i.e., $s^i \in \mathbb{R}^{1 \times T}$. The model outputs the predicted waveforms $\{\hat{s}^i\}_{i=1}^N$ where $\forall i, \hat{s}^i = f^i(x)$, and a loss function is evaluated by comparing the predictions to the ground truth sources $\{s^i\}_{i=1}^N$ up to a permutation. Explicitly, we consider a permutation-invariant objective function⁵,

$$\mathcal{L}(\hat{s}, s) = \min_{\sigma \in S_N} \frac{1}{N} \sum_{i=1}^N \ell(\hat{s}^{\sigma(i)}, s^i) \quad \text{where } \forall i, \hat{s}^i = f^i(x)$$

Here, $\ell(\cdot, \cdot)$ represents the loss function computed on an (output, target) pair, σ indicates a permutation, and S_N is the space of permutations. In certain scenarios, we include the L2 regularization term,

$$\mathcal{L} \mapsto \mathcal{L} + \lambda \sum_{j=1}^P \beta_j^2$$

where β_j represent the model parameters, P denotes the model complexity, and λ is a hyperparameter empirically selected to minimize overfitting (i.e. enhance convergence of training and evaluation losses and metrics).

We consider the reconstruction performance by computing evaluation metrics using a similar expression⁵,

$$\mathcal{M}(\hat{s}, s) = \max_{\sigma \in S_N} \frac{1}{N} \sum_{i=1}^N m(\hat{s}^{\sigma(i)}, s^i) \quad \text{where } \forall i, \hat{s}^i = f^i(x)$$

where $m(\cdot, \cdot)$ is the single-channel evaluation metric computed on permutations of (output, target) pairs.

For the single-channel loss function ℓ , we consider a linear combination of several loss terms that compute the error in estimated waveform reconstructions $\{\hat{s}^i\}_{i=1}^N$ relative to the ground truth waveforms $\{s^i\}_{i=1}^N$.

- L1 Loss

$$|\hat{s}^{\sigma(i)} - s^i|$$

This represents the absolute error on raw time domain waveforms.

- STFT L1 Loss

$$|\text{STFT}(\hat{s}^{\sigma(i)}) - \text{STFT}(s^i)|$$

This term functions to minimize absolute error on time-frequency space representations. Empirically, the inclusion of this contribution enhances the reconstruction of signal harmonicity.

- Spectral Convergence Loss

$$\|\text{STFT}(\hat{s}^{\sigma(i)}) - \text{STFT}(s^i)\|_F / \|\text{STFT}(s^i)\|_F$$

where $\|\cdot\|_F$ denotes the Frobenius norm over time and frequency. This term emphasizes high-magnitude spectral components.

We also experimented with additional terms including L1 loss on log-magnitude spectrograms to address spectral valleys and negative SI-SDR (nSI-SDR), but the inclusion of these contributions did not yield empirical improvements in results.

To assess the quality of reconstruction, we compute two evaluation metrics. As is commonly implemented in the human speech separation literature, we compute the signal-to-distortion ratio (SDR)², defined as the negative log squared error normalized by reference signal energy, $\text{SDR}(\hat{s}, s) = -10 \log_{10}(|\hat{s} - s|^2) + 10 \log_{10}(|s|^2)$. Further, we consider the scale-invariant SDR (SI-SDR), which disregards prediction scale and searches over gains²⁹.

To provide a physically interpretable metric, we evaluate the performance of the trained classifier models in labeling separated waveforms according to the predicted identity of the vocalizer. This metric assumes that the classification accuracy on a downstream task reflects the fidelity of the estimated signal relative to the ground truth source and thus serves as a proxy for reconstruction quality.

For macaques, we modify the training algorithm according to the representation transform and inverse transform built into the model. For the model with the fully learnable Conv1D encoder and decoder, we train using the AdamW⁴⁷ optimizer with a learning rate 3e-4 and batch size 16 for 100 epochs. In order to stabilize training and avoid local minima when using handcrafted STFT and iSTFT filterbanks, we initially begin training the models for 3 epochs with batch size 16 using stochastic gradient descent (SGD) with Nesterov momentum 0.6 and learning rate 1e-3 before switching to the AdamW optimizer until reaching 100 epochs.

For dolphins, we provide the model with the original mixture as input, but we use high pass-filtered source waveforms as the target, which means the separation model must additionally learn to denoise the input. We again initialize training with 3 epochs and batch size 8 using SGD with Nesterov momentum 0.6 and learning rate 1e-3 before switching to the AdamW optimizer with learning rate 3e-4 for the remaining 97 epochs. We use a similar training scheme for bats, initially training with SGD for 3 epochs before employing the optimizer switcher callback to switch to AdamW and to complete 100 epochs.

Objective		Model	SI-SDR	Δ SI-SDR	Accuracy
Macaque	2SpeakerClosed	STFT	26.1	26.1	93.7%
		Conv1D	24.5	24.5	86.4%
		Conv-TasNet	23.8	23.8	85.1%
	3SpeakerClosed	STFT	16.9	22.0	86.5%
Dolphin	2SpeakerClosed	STFT	15.5	17.3	99.9%
Bat	2SpeakerClosed	STFT	9.7	9.7	57.5%
	2SpeakerOpen		10.0	10.0	—
	2SpeakerClosed+		10.3	10.3	58.6%
	2SpeakerOpen+		10.4	10.4	—

Table 1 – Summary of experiments implemented in this study organized according to species, number of speakers, open/closed speaker regime, and model architecture. STFT represents the model with STFT encoder and iSTFT decoder, and Conv1D represents the fully time domain model with Conv1D encoder and ConvTranspose1D decoder. The + denotes the expanded training dataset. We report SI-SDR with (Δ) and without improvement over input SI-SDR averaged over mixtures.

Results

We here report the validation metrics for the classifier and separator models. For the closed speaker separation task, we report both the maximal SI-SDR and classification accuracies since it remains unclear which metric reflects optimal performance of the separator model. Finally, we evaluate an open speaker bat separator model using a dataset containing vocalizations produced by individuals not contained in the training distribution. The results are summarized in Table 1. Supplementary visualizations and audio samples are included in our GitHub repository <https://github.com/earthspecies/cocktail-party-problem>.

Classification Models

The macaque classifier model attains an accuracy of 99.3% (where chance level was 12.5%, i.e. 1 out of 8) for the 8 vocalizing animals used in this study, which represents a state-of-the-art improvement over the existing literature¹⁶. The dolphin classifier model achieves 99.4% accuracy (where chance level was 12.5%, i.e. 1 out of 8) in classifying signature whistles produced by 8 individuals, again reinforcing the individually distinctive characteristics of dolphin signature whistles³⁸. The Egyptian fruit bat classifier network yields 79.7% accuracy (where chance level was 8.3%, i.e. 1 out of 12) in labeling the identity of 12 bats, supporting previous studies demonstrating individual-level specificity of everyday bat vocalizations⁴⁸.

Separation Models

For the macaque separation task, we here detail the results for the fully learnable Conv1D-based model as well as the top-performing non-dB STFT-based model. We also consider mixtures of 2 or 3 overlapping speakers. For the baseline model using the fully learnable Conv1D encoder/decoder, the model yields an SI-SDR of 24.5 and a downstream classification accuracy of 86.4%. We experiment across combinations of (Conv1D, STFT, dB-scaled STFT) encoders and (Conv1D, iSTFT, dB-scaled iSTFT, MCNN) decoders. We observe that the handcrafted (STFT, iSTFT) encoder/decoder combination results in the highest fidelity reconstructions, as assessed using SI-SDR and downstream classification accuracy, with respective (SI-SDR, accuracy) metrics of (26.1, 93.7%). Including the parallel phase pathway does not yield robust performance benefits, in contrast with results from other experimental studies⁴². We note that, even though our model is significantly less complex and computationally expensive, it outperforms Conv-TasNet, a pioneering human speech separation model, which attains a maximal evaluation metric pair of (23.8, 85.1%). Finally, training an STFT/iSTFT model with $N=3$ yields evaluation metrics of (16.9, 86.5%). For both the dolphin separation and bat separation tasks, we report the results using exclusively the fixed STFT encoder and iSTFT decoder rather than carrying out an ablation study across various TFRs and inverses. The dolphin separator yields metrics of (15.5, 99.9%) and the bat model achieves (9.7, 57.5%). Finally, to address the open speaker bat source separation problem, we evaluate the bat model on (1) the training subset, (2) the closed speaker testing subset, and (3) the open speaker testing subset, yielding respective SI-SDR metrics of 10.5, 9.7, 10.0. This finding demonstrates that even in the absence of individually recognizable acoustic cues, BioCPPNet generalizes to overlapping waveforms produced by unseen vocalizers not contained in the training subset. We repeat the bat source separation problem using the larger 72k-sample training dataset (+) and observe maximal metrics of (10.3, 58.6%) in the closed speaker regime. For the open speaker regime, the respective training, closed speaker testing, and open speaker testing SI-SDRs are 10.2, 10.3, 10.4, indicating that training on more samples improves bat separator model generalizability. In Table 1, we include the SI-SDR with improvement (Δ SI-SDR), i.e. the metric obtained using the model output minus the metric averaged over input mixtures.

Figure 3 – Visualizations of the mixture separations for (A,a) macaques, (B,b) bottlenose dolphins, and (C, c) Egyptian fruit bats. (D,d) demonstrates the separation of three overlapping macaque vocalizations.

Discussion

This work proposes BioCPPNet, a complete supervised end-to-end bioacoustic source separation framework designed to accommodate multiple species with diverse vocal behavior. We apply our model to a heterogeneous set of non-human organisms including macaques, bottlenose dolphins, and Egyptian fruit bats, and we demonstrate that our models yield high-quality reconstructions of sources given mixture inputs, as assessed objectively using the SI-SDR metric in combination with downstream classification accuracy and qualitatively by inspecting visual representations of the output audio. Though BioCPPNet performs best in the closed speaker regime in which testing subsets are drawn from the same distribution of individuals as the training subset, our results suggest that BioCPPNet can generalize to the open speaker problem given sufficient quantities of data; for instance, we observe that BioCPPNet and the Conv-TasNet reference model struggle in the open speaker regime when tested on macaques (for which we possess fewer than 45 minutes of highly stereotyped calls) and bottlenose dolphins (for which the dataset is comprised of fewer than 15 minutes of signature whistles). However, when evaluated on bat data containing several hours of varied bat vocalizations, our model yields comparable results in both the open and closed speaker cases. This finding highlights the need for larger datasets to facilitate the development of novel ML-based technologies for bioacoustic source separation applications.

Given the flexibility, versatility, and modularity of BioCPPNet, we suggest further experimental setups to expand on our results. While we implement our methods using both fully learnable and STFT-based handcrafted encoder and decoder filterbanks, future directions can include modifications to the TFR encoder and the inverse transform decoder to assess the performance of other learnable, handcrafted, and/or parameterized features. To our knowledge, there exists no landscape analysis of TFRs for bioacoustic data, so additional studies aiming to enhance separation performance by varying the representation encoder and the inverse transform decoder could provide important insight into optimizing the representation of bioacoustic signals. Additionally, though we implement a learnable signal filtering mechanism in the case of bottlenose dolphin whistles, our results could benefit from additional denoising or enhancement algorithms^{49–52}; as in Fig. 3 (a) and (b), noisy artifact is not always attributed to the original source, which means that targeting and removing noise in the spectral zone of support for a given set of vocalizations could improve the separation performance of BioCPPNet. We also suggest implementations of BioCPPNet across other non-human animals in an effort to further examine the generalizability of our approach.

Finally, while our implementation uses a supervised training scheme relying on synthesized mixtures and access to ground truth sources, the practical application of our methods could benefit from weakening the degree of supervision. Mixture invariant training (MixIT)⁵³ represents an entirely unsupervised approach and requires only single-channel acoustic mixtures, though this technique may be limited in the bioacoustic domain as it necessitates large quantities of well-defined mixtures. Another unsupervised technique includes a Bayesian approach employing deep generative priors^{54–56}, but this method may be limited to data within close bounds of the training sets since the distributions learned by acoustic deep generative models may not exhibit the right properties for probabilistic source separation⁵⁷. We also suggest self-supervised pre-training^{30, 58, 59} on relevant proxy tasks to enhance performance, especially in the low-data bioacoustic domain. Future studies should address unsupervised training criteria in an effort to apply deep ML-based models such as BioCPPNet to real-world bioacoustic mixtures.

The novel ML-based methods employed in this study provide an effective means for addressing the bioacoustic CPP, which can help to maximize the information extracted from bioacoustic recordings. Our framework represents an important step toward the realization of bioacoustic processing technologies capable of recovering large quantities of previously unusable data containing overlapping signals. Greater access to larger bioacoustic datasets can empower the scientific community to explore more areas of research, to confront increasingly complex research questions using deep ML approaches, and ultimately to design better-informed conservation and management strategies to protect the earth's non-human animal species.

Data Availability

The macaque coo dataset³⁷ that supports the findings of this study is publicly available in *Dryad* with the identifier <https://doi.org/10.5061/dryad.7f4p9>, and the Egyptian fruit bat dataset³⁹ used in this study is publicly available in *Scientific Data* with the identifier <https://doi.org/10.1038/sdata.2017.143>. The bottlenose dolphin whistle data that support the findings of this study are available from Laela Sayigh but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Data are however available from the author upon reasonable request and with permission of Laela Sayigh.

Code Availability

The custom-written code for generating mixture datasets and constructing and training ML models is publicly available at our GitHub repository <https://github.com/earthspecies/cocktail-party-problem>.

References

1. Bee, M. & Micheyl, C. The “cocktail party problem”: What is it? how can it be solved? and why should animal behaviorists study it? *J. Comp. Psychol.* **122**, 235–251, DOI: <https://doi.org/10.1037/0735-7036.122.3.235> (2008).
2. Hershey, J. R., Chen, Z., Le Roux, J. & Watanabe, S. Deep clustering: Discriminative embeddings for segmentation and separation. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 31–35, DOI: <https://doi.org/10.1109/ICASSP.2016.7471631> (2016).
3. Luo, Y. & Mesgarani, N. Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM Transactions on Audio, Speech, Lang. Process.* **27**, 1256–1266, DOI: <https://doi.org/10.1109/TASLP.2019.2915167> (2019).
4. Chen, J., Mao, Q. & Liu, D. Dual-path transformer network: Direct context-aware modeling for end-to-end monaural speech separation. Preprint at <https://arxiv.org/abs/2007.13975> (2020).
5. Zeghidour, N. & Grangier, D. Wavesplit: End-to-end speech separation by speaker clustering. Preprint at <https://arxiv.org/abs/2002.08933> (2020).
6. Subakan, C., Ravanelli, M., Cornell, S., Bronzi, M. & Zhong, J. Attention is all you need in speech separation. Preprint at <https://arxiv.org/abs/2010.13154> (2021).
7. Lam, M. W. Y., Wang, J., Su, D. & Yu, D. Sandglassnet: A light multi-granularity self-attentive network for time-domain speech separation. Preprint at <https://arxiv.org/abs/2103.00819> (2021).
8. Deng, X., Tao, Y., Tu, X. & Xu, X. The separation of overlapped dolphin signature whistle based on blind source separation. In *2017 IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC)*, 1–5, DOI: <https://doi.org/10.1109/ICSPCC.2017.8242534> (2017).
9. Hassan, N. & Ramli, D. A. A comparative study of blind source separation for bioacoustics sounds based on FastICA, PCA and NMF. *Procedia Comput. Sci.* **126**, 363 – 372, DOI: <https://doi.org/10.1016/j.procs.2018.07.270> (2018).
10. Zhang, K. *et al.* Separating overlapping bat calls with a bi-directional long short-term memory network. Preprint at <https://www.biorxiv.org/content/early/2019/12/15/2019.12.15.876656> (2019).
11. Izadi, M. R., Stevenson, R. & Kloepper, L. N. Separation of overlapping sources in bioacoustic mixtures. *The J. Acoust. Soc. Am.* **147**, 1688–1696, DOI: <https://doi.org/10.1121/10.0000932> (2020).
12. Poole, J. *The Amboseli Elephants: A Long-Term Perspective on a Long-Lived Mammal*, chap. Behavioral Contexts of Elephant Acoustic Communication (University of Chicago Press, 2011).
13. Tyler M. Schulz, S. G., Hal Whitehead & Rendell, L. Overlapping and matching of codas in vocal interactions between sperm whales: insights into communication function. *Animal Behav.* **76**, 1977–1988, DOI: <https://doi.org/10.1016/j.anbehav.2008.07.032> (2008).
14. Richman, B. The synchronization of voices by gelada monkeys. *Primates* **19**, 569–581, DOI: <https://doi.org/10.1007/BF02373317> (1978).
15. Morfi, V., Bas, Y., Pamula, H., Glotin, H. & Stowell, D. NIPS4Bplus: a richly annotated birdsong audio dataset. *PeerJ Comput. Sci.* **5**, DOI: <https://doi.org/10.7717/peerj-cs.223> (2019).
16. Fukushima, M., Doyle, A., Mullarkey, M., Mishkin, M. & Averbach, B. Distributed acoustic cues for caller identity in macaque vocalization. *Royal Soc. Open Sci.* **2**, DOI: <https://doi.org/10.1098/rsos.150432> (2015).
17. Esch, H. C., Sayigh, L. S. & Wells, R. S. Quantifying parameters of bottlenose dolphin signature whistles. *Mar. Mammal Sci.* **25**, 976–986, DOI: <https://doi.org/10.1111/j.1748-7692.2009.00289.x> (2009).
18. Jones, G. & Holderied, M. W. Bat echolocation calls: adaptation and convergent evolution. *Proc. Royal Soc. B* **274**, 905–912, DOI: <https://doi.org/10.1098/rspb.2006.0200> (2007).
19. Défossez, A., Usunier, N., Bottou, L. & Bach, F. Music source separation in the waveform domain. Preprint at <https://hal.archives-ouvertes.fr/hal-02379796v2> (2021).
20. Roux, J. L., Wichern, G., Watanabe, S., Sarroff, A. & Hershey, J. R. The phasebook: Building complex masks via discrete representations for source separation. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 66–70, DOI: <https://doi.org/10.1109/ICASSP.2019.8682587> (2019).
21. Zhu, W., Wang, M., Zhang, X.-L. & Rahardja, S. A comparison of handcrafted, parameterized, and learnable features for speech separation. Preprint at <https://arxiv.org/abs/2011.14295> (2021).

22. Wang, D. & Chen, J. Supervised speech separation based on deep learning: An overview. Preprint at <https://arxiv.org/abs/1708.07524> (2018).
23. Adam, O. The use of the hilbert-huang transform to analyze transient signals emitted by sperm whales. *Appl. Acoust.* **67**, 1134–1143, DOI: <https://doi.org/10.1016/j.apacoust.2006.04.001> (2006).
24. Priyadarshani, N., Marsland, S., Juodakis, J., Castro, I. & Listanti, V. Wavelet filters for automated recognition of birdsong in long-time field recordings. *Methods Ecol. Evol.* **11**, 403–417, DOI: <https://doi.org/10.1111/2041-210X.13357> (2020).
25. Bruna, J. & Mallat, S. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis Mach. Intell.* **35**, 1872–1886, DOI: <https://doi.org/10.1109/TPAMI.2012.230> (2013).
26. Sprechmann, P., Bruna, J. & LeCun, Y. Audio source separation with discriminative scattering networks. In *Latent Variable Analysis and Signal Separation*, 259–267 (Springer International Publishing, 2015).
27. Kreuk, F., Keshet, J. & Adi, Y. Self-supervised contrastive learning for unsupervised phoneme segmentation. Preprint at <https://arxiv.org/abs/2007.13465> (2020).
28. Zeghidour, N., Teboul, O., de Chaumont Quitry, F. & Tagliasacchi, M. Leaf: A learnable frontend for audio classification. In *ICLR* (2021).
29. Roux, J. L., Wisdom, S., Erdogan, H. & Hershey, J. R. SDR – half-baked or well done? In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 626–630, DOI: <https://doi.org/10.1109/ICASSP.2019.8683855> (2019).
30. Bermant, P. C., Bronstein, M. M., Wood, R. J., Gero, S. & Gruber, D. F. Deep machine learning techniques for the detection and classification of sperm whale bioacoustics. *Sci. Reports* **9**, DOI: <https://doi.org/10.1038/s41598-019-48909-4> (2019).
31. Kirsebom, O. S. *et al.* Performance of a deep neural network at detecting north atlantic right whale upcalls. *The J. Acoust. Soc. Am.* **147**, 2636–2646, DOI: <https://doi.org/10.1121/10.0001132> (2020).
32. Dufourq, E. *et al.* Automated detection of hainan gibbon calls for passive acoustic monitoring. Preprint at <https://www.biorxiv.org/content/10.1101/2020.09.07.285502v2> (2020).
33. Zhang, Z. & White, P. R. A blind source separation approach for humpback whale song separation. *The J. Acoust. Soc. Am.* **141**, 2705–2714, DOI: <https://doi.org/10.1121/1.4980856> (2017).
34. Rafii, Z., Liutkus, A., Stöter, F.-R., Mimilakis, S. I. & Bittner, R. MUSDB18 - a corpus for music separation, DOI: <https://doi.org/10.5281/zenodo.1117371> (2017).
35. Xiao, X. *et al.* Single-channel speech extraction using speaker inventory and attention network. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 86–90, DOI: <https://doi.org/10.1109/ICASSP.2019.8682245> (2019).
36. McDermott, J., Popham, S. & Boebinger, D. The role of harmonic spectral structure in the cocktail party problem. *The J. Acoust. Soc. Am.* **139**, 2017–2017, DOI: <https://doi.org/10.1121/1.4949945> (2016).
37. Fukushima, M., Doyle, A., Mallarkey, M., Mishkin, M. & Averbek, B. Data from: Distributed acoustic cues for caller identity in macaque vocalization, DOI: <https://doi.org/10.5061/dryad.7f4p9> (2015).
38. Sayigh, L. S., Esch, H. C., Wells, R. S. & Janik, V. M. Facts about signature whistles of bottlenose dolphins, *tursiops truncatus*. *Animal Behav.* **74**, 1631–1642, DOI: <https://doi.org/10.1016/j.anbehav.2007.02.018> (2007).
39. Prat, Y., Taub, M., Pratt, E. & Yovel, Y. An annotated dataset of egyptian fruit bat vocalizations across varying contexts and during vocal ontogeny. *Sci. Data* **4**, DOI: <https://doi.org/10.1038/sdata.2017.143> (2017).
40. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *ICLR* (2015).
41. Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 234–241 (Springer International Publishing, 2015).
42. Choi, W., Kim, M., Chung, J., Lee, D. & Jung, S. Investigating U-Nets with various intermediate blocks for spectrogram-based singing voice separation. Preprint at <https://arxiv.org/abs/1912.02591> (2020).
43. Takahashi, N., Agrawal, P., Goswami, N. & Mitsufuji, Y. Phasenet: Discretized phase modeling with deep neural networks for audio source separation. In *Proc. Interspeech 2018*, 2713–2717, DOI: <https://doi.org/10.21437/Interspeech.2018-1773> (2018).
44. Yin, D., Luo, C., Xiong, Z. & Zeng, W. PHASEN: A phase-and-harmonics-aware speech enhancement network. Preprint at <https://arxiv.org/abs/1911.04697> (2019).

45. Arik, S. O., Jun, H. & Diamos, G. Fast spectrogram inversion using multi-head convolutional neural networks. *IEEE Signal Process. Lett.* **26**, 94–98, DOI: <https://doi.org/10.1109/LSP.2018.2880284> (2019).
46. Yu, D., Kolbæk, M., Tan, Z.-H. & Jensen, J. Permutation invariant training of deep models for speaker-independent multi-talker speech separation. Preprint at <https://arxiv.org/abs/1607.00325> (2016).
47. Loshchilov, I. & Hutter, F. Decoupled weight decay regularization. In *ICLR* (2019).
48. Prat, Y., Taub, M. & Yovel, Y. Everyday bat vocalizations contain information about emitter, addressee, context, and behavior. *Sci. Reports* **6**, DOI: <https://doi.org/10.1038/srep39419> (2016).
49. Park, S. R. & Lee, J. A fully convolutional neural network for speech enhancement. In *INTERSPEECH*, DOI: <https://doi.org/10.21437/INTERSPEECH.2017-1465> (2017).
50. Sonning, S., Schüldt, C., Erdogan, H. & Wisdom, S. Performance study of a convolutional time-domain audio separation network for real-time speech denoising. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 831–835, DOI: <https://doi.org/10.1109/ICASSP40776.2020.9053846> (2020).
51. Michelashvili, M. & Wolf, L. Audio denoising with deep network priors. Preprint at <https://arxiv.org/abs/1904.07612v1> (2019).
52. Stowell, D. & Turner, R. E. Denoising without access to clean data using a partitioned autoencoder. Preprint at <https://arxiv.org/abs/1509.05982> (2015).
53. Wisdom, S. *et al.* Unsupervised sound separation using mixture invariant training. In *Advances in Neural Information Processing Systems*, vol. 33, 3846–3857 (Curran Associates, Inc., 2020).
54. Jayaram, V. & Thickstun, J. Source separation with deep generative priors. In *Proceedings of the 37th International Conference on Machine Learning*, vol. 119 of *Proceedings of Machine Learning Research*, 4724–4735 (PMLR, 2020).
55. Narayanaswamy, V., Thiagarajan, J., Anirudh, R. & Spanias, A. Unsupervised audio source separation using generative priors. *Proc. Annu. Conf. Int. Speech Commun. Assoc. INTERSPEECH* 2657–2661, DOI: <https://doi.org/10.21437/Interspeech.2020-3115> (2020).
56. Jayaram, V. & Thickstun, J. Parallel and flexible sampling from autoregressive models via langevin dynamics. Preprint at <https://arxiv.org/abs/2105.08164> (2021).
57. Frank, M. & Ilse, M. Problems using deep generative models for probabilistic audio source separation. In *Proceedings on "I Can't Believe It's Not Better!" at NeurIPS Workshops*, vol. 137 of *Proceedings of Machine Learning Research*, 53–59 (PMLR, 2020).
58. Huang, S.-F. *et al.* Self-supervised pre-training reduces label permutation instability of speech separation. Preprint at <https://arxiv.org/abs/2010.15366> (2020).
59. Erhan, D., Courville, A., Bengio, Y. & Vincent, P. Why does unsupervised pre-training help deep learning? In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, vol. 9 of *Proceedings of Machine Learning Research*, 201–208 (PMLR, 2010).

Acknowledgements

The author thanks Aza Raskin, Britt Selvitelle, and Katherine Zacarian for useful discussion and Laela Sayigh, Frants Jensen, Michelle Fournet, Andrés Babino, and James Crutchfield for reviewing the manuscript.

Author Information

Affiliations

Earth Species Project, Berkeley, CA, USA

Contributions

All aspects of the work (task setting, data processing, machine learning, article writing, etc.) are performed by the author.

Corresponding author

Correspondence to P.C. Bermant

Ethics Declaration

Competing interests

The author declares no competing interests.