

The Cocktail Fork Problem: Three-Stem Audio Separation for Real-World Soundtracks

Petermann, Darius; Wichern, Gordon; Wang, Zhong-Qiu; Le Roux, Jonathan

TR2022-022 March 05, 2022

Abstract

The cocktail party problem aims at isolating any source of interest within a complex acoustic scene, and has long inspired audio source separation research. Recent efforts have mainly focused on separating speech from noise, speech from speech, musical instruments from each other, or sound events from each other. However, separating an audio mixture (e.g., movie soundtrack) into the three broad categories of speech, music, and sound effects (understood to include ambient noise and natural sound events) has been left largely unexplored, despite a wide range of potential applications. This paper formalizes this task as the cocktail fork problem, and presents the Divide and Remaster (DnR) dataset to foster research on this topic. DnR is built from three well-established audio datasets (LibriSpeech, FMA, FSD50k), taking care to reproduce conditions similar to professionally produced content in terms of source overlap and relative loudness, and made available at CD quality. We benchmark standard source separation algorithms on DnR, and further introduce a new multi-resolution model to better address the variety of acoustic characteristics of the three source types. Our best model produces SI-SDR improvements over the mixture of 11.0 dB for music, 11.2 dB for speech, and 10.8 dB for sound effects.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
2022*

THE COCKTAIL FORK PROBLEM: THREE-STEM AUDIO SEPARATION FOR REAL-WORLD SOUNDTRACKS

Darius Petermann^{1,2}, Gordon Wichern¹, Zhong-Qiu Wang¹, Jonathan Le Roux¹

¹Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

²Indiana University, Department of Intelligent Systems Engineering, Bloomington, IN, USA

ABSTRACT

The cocktail party problem aims at isolating any source of interest within a complex acoustic scene, and has long inspired audio source separation research. Recent efforts have mainly focused on separating speech from noise, speech from speech, musical instruments from each other, or sound events from each other. However, separating an audio mixture (e.g., movie soundtrack) into the three broad categories of speech, music, and sound effects (understood to include ambient noise and natural sound events) has been left largely unexplored, despite a wide range of potential applications. This paper formalizes this task as the *cocktail fork problem*, and presents the Divide and Remaster (DnR) dataset to foster research on this topic. DnR is built from three well-established audio datasets (LibriSpeech, FMA, FSD50k), taking care to reproduce conditions similar to professionally produced content in terms of source overlap and relative loudness, and made available at CD quality. We benchmark standard source separation algorithms on DnR, and further introduce a new multi-resolution model to better address the variety of acoustic characteristics of the three source types. Our best model produces SI-SDR improvements over the mixture of 11.0 dB for music, 11.2 dB for speech, and 10.8 dB for sound effects.

Index Terms— audio source separation, speech, music, sound effects, soundtrack

1. INTRODUCTION

Humans are able to focus on a source of interest within a complex acoustic scene, a task referred to as the cocktail party problem [1, 2]. Research in audio source separation has been dedicated to enabling machines to solve this task, with many studies taking a stab at various slices of the problem, such as the separation of speech from non-speech in speech enhancement [3, 4], speech from other speech in speech separation [5–7], or separation of individual musical instruments [8–10] or non-speech sound events (or sound effects) [11–14]. However, separation of sound mixtures involving speech, music, and sound effects/events has been left largely unexplored, despite its relevance to most produced audio content, such as podcasts, radio broadcasts, and video soundtracks. We here intend to bite into this smaller chunk of the cocktail party problem by proposing to separate such soundtracks into these three broad categories. We refer to this task as the *cocktail fork problem*, as illustrated in Fig. 1.

While there has been much work on labeling recordings based on these three categories [15–17], the ability to separate audio signals into these streams has the potential to support a wide range of novel applications. For example, an end-user could take over the final mixing process by applying independent gains to the separated speech, music, and sound effects signals to support their specific listening environment and preferences. Furthermore, this three-stream




Fig. 1: Illustration of the cocktail fork problem: given a soundtrack consisting of an audio mixture of speech, music, and sound effects, the goal is to separate it into the three corresponding stems.

separation could be a front-end for total transcription [18] or audio-visual video description [19] where we want to not only transcribe speech but also semantically describe in great detail the non-speech sounds present in an auditory scene. A recent concurrent work [20] also explores the task of speech, music, and sound effects (therein referred to as noise) separation, but only considers the unrealistic case of fully-overlapped mixtures of the three streams, and a low sampling rate of 16 kHz. This sampling rate is not conducive to applications where humans may listen to the separated signals, and it is often difficult or impractical to transition systems trained only on fully-overlapped mixtures to real-world scenarios [21].

To provide a realistic high-quality dataset for the cocktail fork problem, we introduce the Divide and Remaster (DnR) dataset, which is built upon LibriSpeech [22] for speech, Free Music Archive (FMA) [23] for music, and Freesound Dataset 50k (FSD50K) [24] for sound effects. DnR pays particular attention to the mixing process, specifically the relative level of each of the sources and the amount of inter-class overlap, both of which we hope will ease the transition of models trained with DnR to real-world applications. Furthermore, DnR includes comprehensive speech, music genre, and sound event annotations, making it potentially useful for research in speech transcription, music classification, sound event detection, and audio segmentation in addition to source separation.

In this paper, we provide a detailed description of the DnR dataset, and benchmark various source separation models. We find the CrossNet unmix (XUMX) architecture [25], originally proposed for music source separation, also works well for DnR. We further propose a multi-resolution extension of XUMX, to better handle the wide variety of audio characteristics in the sound sources we are trying to separate. We also address several important practical questions often ignored in the source separation literature, such as the impact of sampling rate on model performance, predicted energy in regions where a source should be silent [26], and performance in various overlapping conditions. While we only show here objective evaluations based on synthetic data due to the lack of realistic data with stems, we confirmed via informal listening tests that the trained models perform well on real-world soundtracks from YouTube. Our dataset and real-world examples are available online.¹

¹This work was performed while D. Petermann was an intern at MERL.

¹[cocktail-fork.github.io](https://github.com/cocktail-fork)

2. THE COCKTAIL FORK PROBLEM

We consider an audio soundtrack y such that

$$y = \sum_{j=1}^3 x_j, \quad (1)$$

where x_1 is the submix containing all music signals, x_2 that of all speech signals, and x_3 that of all sound effects. We use the term sound effects (SFX) to broadly cover all sources not categorized as speech or music, and choose it over alternatives such as sound events or noise, as the term is especially relevant to our target application where y is a soundtrack. We here define the cocktail fork problem as that of recovering, from the audio soundtrack y , its music, speech, and sound effect submixes, as opposed to extracting individual musical instruments, speakers, or sound effects.

Our goal is to train a machine learning model to obtain estimates \hat{x}_1 , \hat{x}_2 , and \hat{x}_3 of these submixes. We explore two general classes of models for estimating \hat{x}_j . The first one, exemplified by ConvTasNet [27], takes as input the time-domain mixture y , and outputs time-domain estimates \hat{x}_j . The second one operates on the time-frequency (TF) domain mixture, i.e., $Y = \text{STFT}(y)$, and estimates a real-valued mask \hat{M}_j for each source, obtaining time-domain estimates via inverse STFT as $\hat{x}_j = \text{iSTFT}(\hat{M}_j \odot Y)$.

3. MULTI-RESOLUTION CROSSNET (MRX)

In our benchmark of various network architectures in Section 5.2, we find consistently strong performance from CrossNet unmix (XUMX) [25], which uses multiple parameter-less averaging operations when simultaneously extracting multiple stems (musical instruments in [25]). XUMX is an STFT masking-based architecture, and choosing appropriate transform parameters is a key design choice. Longer STFT windows provide better frequency resolution at the cost of poorer time resolution, and vice versa for shorter windows. Mixtures of signals with diverse acoustic characteristics could thus benefit from multiple STFT resolutions in their TF encoding. Previous research has proven the efficacy of multi-resolution systems for audio-related tasks, such as in the context of speech enhancement [28], music separation [29], speech recognition [30], and sound event detection [31]. We thus introduce a multi-resolution extension of XUMX which addresses the typical limitations brought by a single-resolution architecture. In [25], the authors show that using multiple parallel branches to process the input can help in the separation task. We here apply this reasoning further towards multiple STFT resolutions.

Our proposed architecture takes a time-domain input mixture and encodes it into I complex spectrograms Y_{L_i} with different STFT resolutions, where L_i denotes the i -th window length in milliseconds. Figure 2 shows an example with $I = 3$ and $\{L_i\}_i = \{32, 64, 256\}$. We use the same hop size (e.g., 8 ms in the example of Fig. 2) for all resolutions, so they remain synchronized in time, and N denotes the number of STFT frames for all resolutions. In practice, we set the window size in samples to the nearest power of 2, and the number of unique frequency bins is denoted as F_{L_i} . Each resolution is then passed to a fully connected block to convert the magnitude spectrograms of dimension $N \times F_{L_i}$ into a consistent dimension of 512 across the resolution branches. This allows us to average them together prior to the bidirectional long short-term memory (BLSTM) stacks, whose outputs are averaged once again. While the averaging operators in XUMX were originally intended to efficiently bridge independent architectures for multiple sources, in our case, the input averaging allows the network to efficiently combine inputs with multiple resolutions.

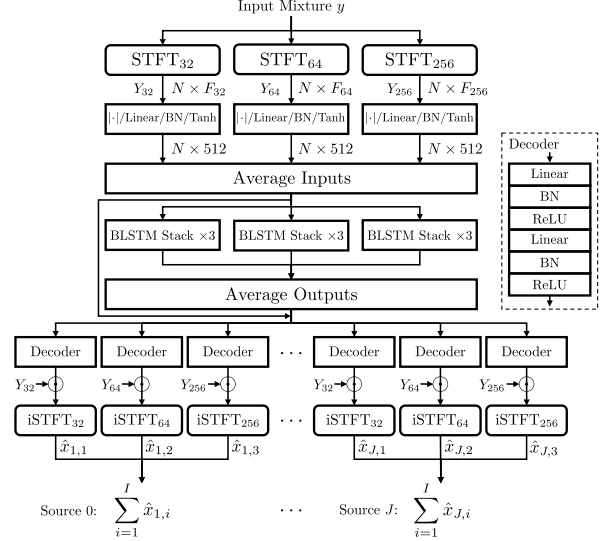


Fig. 2: Multi-resolution CrossNet (MRX) architecture.

The average inputs and outputs of the BLSTM stacks are concatenated and decoded back into magnitude soft masks $\hat{M}_{j,i}$, one for each of the three sources j and each of the I original input resolutions i . The decoder consists of two stacks of fully-connected layers, each followed by batch normalization (BN) and rectified linear units (ReLU). For a given source j , each magnitude mask $\hat{M}_{j,i}$ is multiplied element-wise with the original complex mixture spectrogram Y_{L_i} for the corresponding resolution, a corresponding time-domain signal $\hat{x}_{j,i}$ is obtained via inverse STFT, and the estimated time-domain signal \hat{x}_j is obtained by summing the time-domain signals:

$$\hat{x}_j = \sum_{i=1}^I \hat{x}_{j,i} = \sum_{i=1}^I \text{iSTFT}(\hat{M}_{j,i} \odot Y_{L_i}). \quad (2)$$

For the cocktail fork problem, the network has to estimate a total of $3I$ masks (9 in the example of Fig. 2). Since ReLU is used as the final mask decoder nonlinearity, the network can freely learn weights for each resolution that best reconstruct the time-domain signal.

4. DNR DATASET

4.1. Dataset Building Blocks

In selecting existing speech, music, and sound effects audio datasets for the cocktail fork problem, we had three primary objectives: (1) the data should be freely available under a Creative Commons license; (2) the sampling rate of the audio should be high enough to cover the full range of human hearing (e.g., 44.1 kHz) to support listening applications (one can always downsample as needed); and (3) the audio should contain metadata labels such that it can also be used to explore the impact of separation on downstream tasks, such as transcribing speech and/or providing time-stamped labels for sound effects and music. We selected the following three datasets.

FSD50K - Sound effects: The Freesound Dataset 50k (FSD50K) [24] contains 44.1 kHz mono audio, and clips are tagged using a vocabulary of 200 class labels from the AudioSet ontology [32]. For mixing purposes, we manually classify each of the 200 class labels in FSD50K into one of 3 groups: foreground sounds (e.g., dog bark), background sounds (e.g., traffic noise), and speech/musical instruments (e.g., guitar, speech). Speech and musical instrument clips are filtered out to avoid confusion with our speech and music datasets, and we use different mixing rules for foreground and background

events as described in Section 4.2. We also remove any leading or trailing silence from each sound event prior to mixing.

Free Music Archive - Music: The Free Music Archive (FMA) [23] is a music dataset including over 100,000 stereo songs across 161 musical genres at 44.1 kHz sampling rate. FMA was originally proposed to address music information retrieval (MIR) tasks and thus includes a wide variety of musical metadata. In the context of DnR, we only use track genre as the music metadata. We use the medium subset of FMA which contains 30 second clips from 25,000 songs in 16 unbalanced genres, and is of a comparable size to FSD50K.

LibriSpeech - Speech: DnR’s speech class is drawn from the LibriSpeech dataset [22], an automatic speech recognition corpus based on public-domain audio books. We use the 100 h TRAIN-CLEAN-100 subset for training, chosen over TRAIN-CLEAN-360 because it is closer in size to FSD50K and FMA-medium. For validation and test, we use the clean subsets DEV-CLEAN and TEST-CLEAN to avoid noisy speech being confused with music or sound effects. We incorporate the provided speech transcription for each utterance as part of the DnR metadata. LibriSpeech provides its data as clips containing a single speech utterance at 16 kHz. Fortunately, the original 44.1 kHz mp3 audio files containing the unsegmented audiobook recordings harvested from the LibriVox project are also available along with the metadata mapping each LibriSpeech utterance to the original LibriVox filename and corresponding time-stamp, which we use to create a high sampling rate version of LibriSpeech.

4.2. Mixing procedure

In order to create realistic mixtures of synthetic soundtracks, we focused our effort in two main areas, class overlap and relative level between the different sources in the mixture. Multi-channel spatialization is another important aspect of the mixing process, however, we were unable to find widely agreed upon rules for this process, and therefore focus exclusively on the single-channel case. We also note that trained single-channel models can be applied independently to each channel of a multi-channel recording, and the outputs combined with a multi-channel Wiener filter for post-processing [33]. For the purposes of the mixing procedure described in this section, there are four classes: speech, music, foreground effects, and background effects, but the foreground and background sounds are combined into a single submix in the final version of the DnR dataset.

In order to ensure that a mixture could contain multiple full speech utterances and feature a sufficient number of onsets and offsets between the different classes, we decided to make each mixture 60 seconds long. We do not allow within-class overlap between clips, i.e., two music files will not overlap, but foreground and background sound effects can overlap. The number of files for each class is sampled from a zero-truncated Poisson distribution with expected value λ . The values of λ are chosen based on the average file length of each class, e.g., music and background effects tend to be longer (see Table 1). For speech files, we always include the entire utterance so that the corresponding transcription remains relevant, while for other classes, we randomly sample the amount of silence between clips of the same class, the clip length, and the internal start time of each clip. Using this mixing procedure, the “all sources active” frames account for $\approx 55\%$ of the DnR test set, the “two sources” frames for $\approx 32\%$, and “one source” frames for $\approx 10\%$, leaving silent frames at $\approx 3\%$ (See Table 4 for more details).

Regarding the relative amplitude levels across the three classes, after analyzing studies such as [34] and informal mixing rules from industries such as motion pictures, video games, and podcasting, we found that levels remain fairly consistent across classes, where speech is generally found at the forefront of the mix, followed by

Table 1: Parameters used in the DnR creation procedure.

	Music	Speech	SFX-FG	SFX-BG
λ	7	8	12	6
Target LUFS	-24.0	-17.0	-21.0	-29.0

foreground sound effects, then music, and finally background ambiances. Table 1 depicts the levels used in the DnR dataset in loudness units full-scale (LUFS) [35]. To add variability while keeping a realistic consistency over an entire mixture, we first sample an average LUFS value for each class in each mixture, uniformly from a range of ± 2.0 around the corresponding Target LUFS. Then each sound file added to the mix has its individual gain further adjusted by uniformly sampling from a range of ± 1.0 .

We base our training, validation, and test splits off of those provided by each of the dataset building blocks. The number of test set mixtures is determined such that we exhaust all utterances from the LibriSpeech TEST-CLEAN set twice. We then choose the number of training and validation set mixtures to correspond to a .7/.1/.2 split between training/validation/test, which is roughly in line with the split percentages for FMA (.8/.1/.1) and FSD50k (.7/.1/.2). In the end, DnR consists of 3,406 mixtures (≈ 57 h) for the training set, 487 mixtures (≈ 8 h) for the validation set, and 973 mixtures (≈ 16 h) for the test set, along with their isolated ground-truth stems.

5. EXPERIMENTAL VALIDATION

5.1. Setup

We benchmark the performance of several source separation models in terms of scale-invariant signal-to-distortion ratio (SI-SDR) [36] for the cocktail fork problem on the DnR dataset, both in the original 44.1 kHz version and in a downsampled 16 kHz version. Unless otherwise noted, we compute the SI-SDR on each 60 second mixture, and average over all tracks in the test set.

XUMX and MRX models: We consider single-resolution XUMX baselines with various STFT resolutions. We opt to cover a wide range of window lengths L (between 32 and 256 ms) to assess the impact of resolution on performance. For our proposed MRX model, we use three STFT resolutions of 32, 64, and 256 ms, which we found to work best on the validation set. We use XUMX_L to denote a model with an L ms window. We set the hop size to a quarter of the window size. For the MRX model, we determine hop size based on the shortest window. To parse the contributions of the multi-resolution and multi-decoder features of MRX, we also evaluate an architecture adding MRX’s multi-decoder to the best single-resolution model (XUMX₆₄), referred to as XUMX_{64,multi-dec}. This results in an architecture of the same size (i.e., same number of parameters) as our proposed MRX model. In all architectures, each BLSTM layer has 256 hidden units and input/output dimension of 512, and the hidden layer in the decoder has dimension 512.

Other benchmarks: We also evaluate our own implementations of Conv-TasNet [27] and a temporal convolution network (TCN) with mask inference (MaskTCN). MaskTCN uses an identical TCN to the one used internally by Conv-Tasnet, but the learned encoder/decoder are replaced with STFT/iSTFT operations. For MaskTCN, we use an STFT window/hop of 64/16 ms, and for the learned encoder/decoder of Conv-TasNet, we use 500 filters with a window size of 32 samples and a stride of 16 at 16 kHz, and a window size of 80 samples and a stride of 40 at 44.1 kHz. All TCN parameters in both Conv-TasNet and MaskTCN follow the best configuration of [27]. Additionally, we evaluate Open-Unmix (UMX) [9], the predecessor to XUMX, by training a separate model for each source, but without the parallel branches and averaging operations introduced by XUMX. We also

Table 2: SI-SDR [dB] results of proposed models and baselines on DnR.

Model	16 kHz			44.1 kHz		
	Music	Speech	SFX	Music	Speech	SFX
No processing	-6.8	1.1	-5.2	-6.8	1.0	-5.0
Oracle PSF [37]	11.9	18.0	13.6	11.6	17.8	13.7
Conv-TasNet [27]	1.1	9.4	2.7	0.3	8.5	2.0
MaskTCN ₆₄ [27]	1.9	10.2	3.1	1.7	9.7	3.8
UMX ₆₄ [9]	3.2	11.8	4.1	3.1	11.7	4.4
XUMX ₃₂ [25]	3.2	11.5	4.5	2.9	11.2	4.7
XUMX ₆₄ [25]	3.6	11.8	4.7	3.5	11.7	5.1
XUMX ₁₂₈ [25]	3.5	11.3	4.5	3.7	11.6	5.1
XUMX ₂₅₆ [25]	2.9	10.2	3.8	2.9	10.5	4.4
XUMX _{64,multi-dec}	3.7	11.7	4.8	3.5	11.8	5.0
MRU _{single-stack} (proposed)	3.7	12.1	4.1	3.5	11.8	4.4
MRU _{multi-stack} (proposed)	3.7	12.2	4.2	3.8	12.0	4.6
MRX (proposed)	4.1	12.2	5.1	4.2	12.3	5.7

explore a new multi-resolution UMX (MRU), which uses the same settings as the MRX model in Fig. 2, but features a single BLSTM stack and a single decoder and is trained separately for each source.

Training setup: The Conv-TasNet, XUMX_L, XUMX_{64,multi-dec}, UMX, MRU, and MRX models all use SI-SDR [27, 36] as loss function, while MaskTCN uses the waveform domain L_1 loss. All models are trained on 9 s chunks, except MaskTCN, trained on 6 s chunks, and Conv-TasNet, trained on 4 s chunks at 16 kHz and 2 s chunks at 44.1 kHz; we found these values to lead to best performance under our GPU memory constraints. All models are trained for 300 epochs using ADAM. The learning rate is initialized to 10^{-3} , and halved if the validation loss is not improved over 3 epochs.

5.2. Results and Discussion

Model comparisons: Table 2 presents the SI-SDR of various models trained and tested on DnR, in addition to the no processing condition (lower bound, using the mixture as estimate) and oracle phase sensitive mask [37] (upper bound). For each model, SI-SDR improvements are fairly consistent across source types, despite the differences in their relative levels in the mix, which can be seen in the “No Processing” SI-SDR. For both sampling rates, we observe that our proposed MRX model outperforms all single resolution baselines on all source types. This implies that the network learns to effectively combine information from different STFT resolutions to more accurately reconstruct the target sources. The performance of XUMX_{64,multi-dec} further confirms this hypothesis by performing nearly identically to XUMX₆₄, showing that the use of multiple decoders alone does not improve performance. We also observe that the single-source models (UMX, MRU) tend to perform comparably to the cross-source models (XUMX, MRX) for speech, but perform worse for music and SFX. We speculate that because music and SFX are quieter in the mix, it is harder for the network to isolate them effectively without the support of the other sources, while louder sources (here, speech) do not benefit from joint estimation.

Sampling rate comparisons: Table 3 compares the average SI-SDR performance of the MRX model across sampling rates. In the 44.1 kHz column, we observe a reduction in SI-SDR for all source classes when upsampling the 16 kHz model output to 44.1 kHz (“Resampled”). This is to be expected as all frequencies above 8 kHz are zero in the upsampled 44.1 kHz signal, but we see that these frequencies only contribute a small amount to the difference in SI-SDR scores, as there is comparatively little energy there. In the 16 kHz column, we observe a minor performance gain in the “Resampled” row where the 44.1 kHz model output is downsampled to 16 kHz, showing that the model can make use of information above 8 kHz to improve separation under 8 kHz. This result could be beneficial

Table 3: SI-SDR [dB] of the MRX model across sample rates. In Original row, processing is done at the evaluation sample rate indicated in each column. In Resampled row, processing is done at the other sample rate (44.1 kHz for 16 kHz and vice versa) and the output resampled for evaluation.

Model	16 kHz			44.1 kHz		
	Music	Speech	SFX	Music	Speech	SFX
Original	4.1	12.2	5.1	4.2	12.3	5.7
Resampled	4.2	12.3	5.4	3.9	11.3	4.1

Table 4: Performance of MRX at 44.1 kHz for each source on the seven overlapping use-cases. M, S, X indicate presence within a frame of music, speech, and SFX respectively. Scores indicate SI-SDR [dB] improvement when the source is present with another, final SI-SDR [dB] when the source is the only one present (indicated by *), and Predicted Energy at Silence [dB] (PES) when a source is not present (indicated by †).

Source	{M,S,X}	{ \emptyset ,S,X}	{M, \emptyset ,X}	{M,S, \emptyset }	{M, \emptyset , \emptyset }	{ \emptyset ,S, \emptyset }	{ \emptyset , \emptyset ,X}
Frames	31910	6351	8702	3761	1330	1761	2737
Music	10.2	-13.6 [†]	2.9	9.0	10.7*	-23.3 [†]	-21.6 [†]
Speech	11.8	9.0	-22.0 [†]	10.4	-28.7 [†]	23.4*	-29.0 [†]
SFX	12.0	10.8	7.2	-6.8 [†]	-11.1 [†]	-13.4 [†]	11.6*

for transcription applications where ASR or sound event detection models are pre-trained at 16 kHz, but a front-end source separation model can obtain better separated signals using 44.1 kHz input.

Overlap scenario comparisons: We here compute metrics over 1 s segments to evaluate performance independently in regions where only certain source classes overlap. Metrics such as SI-SDR or SDR [38] are undefined for signals containing silent target and/or estimated sources. This limitation is usually circumvented by disregarding the problematic frames in the evaluation process [9] (i.e., frames containing one or more silent target(s) and/or estimated source(s)). For example, in the MUSDB test set [8], it is reported that at least 45 minutes out of a total of 210 minutes of test data are systematically ignored for that reason [26]. Although these regions with fewer active sources may be seen as less challenging, we believe it is important to also evaluate performance when not all sources are present, and here consider all types of overlap.

Table 4 shows the overall results for each of the three sources in the seven possible overlapping scenarios. For regions where a source is not active in the ground truth, we report predicted energy at silence (PES) [26] to quantify the energy incorrectly assigned to a silent source. We note that speech has smaller PES values than music or SFX in Table 4, even though it is the loudest source on average. SI-SDR in the single-source cases are very high, especially for speech, showing that few artifacts are introduced. Among the two-source cases, we note that SI-SDR_i is substantially lower for music and SFX ({M, \emptyset ,X}), indicating that these two sets of varied sources are more difficult to separate from each other than from speech.

6. CONCLUSION

In this paper, we formalized the task of three-stem soundtrack separation as the cocktail fork problem, and introduced DnR, a high-quality dataset built on top of three well-established sound collections: LibriSpeech (speech), FSD50K (SFX), and FMA (music). We benchmarked several source separation algorithms on DnR and showed that our proposed multi-resolution model performed best. In the future, we plan to combine the separation models developed in this paper with speech recognition and sound classification systems for automatic caption generation of speech and non-speech sounds, and explore remixing strategies that minimize perceptual artifacts [39].

7. REFERENCES

- [1] E. C. Cherry, “Some experiments on the recognition of speech, with one and with two ears,” *J. Acoust. Soc. Am.*, vol. 25, no. 5, pp. 975–979, 1953.
- [2] J. H. McDermott, “The cocktail party problem,” *Current Biology*, vol. 19, no. 22, pp. R1024–R1027, 2009.
- [3] D. Wang and J. Chen, “Supervised speech separation based on deep learning: An overview,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [4] C. K. Reddy, V. Gopal, R. Cutler, E. Beyrami *et al.*, “The Interspeech 2020 Deep Noise Suppression challenge: Datasets, subjective testing framework, and challenge results,” in *Proc. Interspeech*, Oct. 2020.
- [5] J. R. Hershey, Z. Chen, and J. Le Roux, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proc. ICASSP*, Mar. 2016, pp. 31–35.
- [6] L. Drude, J. Heitkaemper, C. Boeddeker, and R. Haeb-Umbach, “SMS-WSJ: Database, performance measures, and baseline recipe for multi-channel source separation and recognition,” in *arXiv preprint arXiv:1910.13934*, 2019.
- [7] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu *et al.*, “WHAM!: Extending speech separation to noisy environments,” in *Proc. Interspeech*, Sep. 2019.
- [8] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimilakis *et al.*, “The MUSDB18 corpus for music separation,” Dec. 2017. [Online]. Available: <https://doi.org/10.5281/zenodo.1117372>
- [9] F.-R. Stöter, S. Uhlich, A. Liutkus, and Y. Mitsufuji, “Open-Unmix - a reference implementation for music source separation,” *Journal of Open Source Software*, 2019.
- [10] E. Manilow, G. Wichern, P. Seetharaman, and J. Le Roux, “Cutting music source separation some Slakh: A dataset to study the impact of training data quality and quantity,” in *Proc. WASPAA*, Oct. 2019.
- [11] I. Kavalerov, S. Wisdom, H. Erdogan, B. Patton *et al.*, “Universal sound separation,” in *Proc. WASPAA*, Oct. 2019.
- [12] E. Tzinis, S. Wisdom, J. R. Hershey, A. Jansen *et al.*, “Improving universal sound separation using sound classification,” in *Proc. ICASSP*, May 2020.
- [13] F. Pishdadian, G. Wichern, and J. Le Roux, “Finding strength in weakness: Learning to separate sounds with weak supervision,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 2386–2399, 2020.
- [14] T. Ochiai, M. Delcroix, Y. Koizumi, H. Ito *et al.*, “Listen to what you want: Neural network-based universal sound selector,” in *Proc. Interspeech*, Oct. 2020.
- [15] T. Theodorou, I. Mporas, and N. Fakotakis, “An overview of automatic audio segmentation,” *IJITCS*, vol. 6, no. 11, 2014.
- [16] B. Meléndez-Catalán, E. Molina, and E. Gómez, “Open broadcast media audio from tv: A dataset of tv broadcast audio with relative music loudness annotations,” *Trans. ISMIR*, vol. 2, no. 1, 2019.
- [17] S. Venkatesh, D. Moffat, A. Kirke, G. Shakeri *et al.*, “Artificially synthesising data for audio classification and segmentation to improve speech and music detection in radio broadcast,” in *Proc. ICASSP*, Jun. 2021, pp. 636–640.
- [18] N. Moritz, G. Wichern, T. Hori, and J. Le Roux, “All-in-one transformer: Unifying speech recognition, audio tagging, and event detection,” in *Proc. Interspeech*, Oct. 2020.
- [19] C. Hori, T. Hori, T.-Y. Lee, Z. Zhang *et al.*, “Attention-based multimodal fusion for video description,” in *Proc. CVPR*, Jul. 2017.
- [20] L. Zhang, C. Li, F. Deng, and X. Wang, “Multi-task audio source separation,” *arXiv preprint arXiv:2107.06467*, 2021.
- [21] Z. Chen, T. Yoshioka, L. Lu, T. Zhou *et al.*, “Continuous speech separation: Dataset and analysis,” in *Proc. ICASSP*, May 2020, pp. 7284–7288.
- [22] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *Proc. ICASSP*, Apr. 2015, pp. 5206–5210.
- [23] M. Defferrard, K. Benzi, P. Vandergheynst, and X. Bresson, “FMA: A dataset for music analysis,” in *Proc. ISMIR*, Oct. 2017.
- [24] E. Fonseca, X. Favory, J. Pons, F. Font *et al.*, “FSD50K: An open dataset of human-labeled sound events,” *arXiv preprint arXiv:2010.00475*, 2020.
- [25] R. Sawata, S. Uhlich, S. Takahashi, and Y. Mitsufuji, “All for one and one for all: Improving music separation by bridging networks,” in *Proc. ICASSP*, Jun. 2021, pp. 51–55.
- [26] K. Schulze-Forster, C. Doire, G. Richard, and R. Badeau, “Weakly informed audio source separation,” in *Proc. WASPAA*, Oct. 2019.
- [27] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [28] Y. Koizumi, N. Harada, and Y. Haneda, “Trainable adaptive window switching for speech enhancement,” in *Proc. ICASSP*, May 2019.
- [29] E. M. Grais, H. Wierstorf, D. Ward, and M. D. Plumbley, “Multi-resolution fully convolutional neural networks for monaural audio source separation,” in *Proc. LVA*, Jul. 2018.
- [30] D. T. Toledano, M. P. Fernández-Gallego, and A. Lozano-Diez, “Multi-resolution speech analysis for automatic speech recognition using deep neural networks: Experiments on TIMIT,” *PLoS ONE*, vol. 13, 2018.
- [31] D. De Benito-Gorrón, D. Ramos, and D. T. Toledano, “A multi-resolution CRNN-based approach for semi-supervised sound event detection in DCASE 2020 challenge,” *IEEE Access*, vol. 9, pp. 89 029–89 042, 2021.
- [32] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen *et al.*, “Audio set: An ontology and human-labeled dataset for audio events,” in *Proc. ICASSP*, Mar. 2017.
- [33] A. A. Nugraha, A. Liutkus, and E. Vincent, “Multichannel audio source separation with deep neural networks,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 9, 2016.
- [34] S. Chaudhuri, J. Roth, D. P. Ellis, A. Gallagher *et al.*, “Ava-speech: A densely labeled dataset of speech activity in movies,” in *Proc. Interspeech*, Sep. 2018.
- [35] E. Grimm, R. Van Everdingen, and M. J. L. C. Schöpping, “Toward a recommendation for a european standard of peak and LKFS loudness levels,” *SMPTE Motion Imaging Journal*, vol. 119, no. 3, pp. 28–34, 2010.
- [36] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR—half-baked or well done?” in *Proc. ICASSP*, May 2019.
- [37] H. Erdogan, J. Hershey, S. Watanabe, and J. Le Roux, “Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks,” in *Proc. ICASSP*, Apr. 2015.
- [38] E. Vincent, R. Gribonval, and C. Fevotte, “Performance measurement in blind audio source separation,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [39] M. Torcoli, J. Paulus, T. Kastner, and C. Uhle, “Controlling the remixing of separated dialogue with a non-intrusive quality estimate,” in *Proc. WASPAA*, Oct. 2021.