Finding Strength in Weakness: Learning to Separate Sounds with Weak Supervision

Pishdadian, Fatemeh; Wichern, Gordon; Le Roux, Jonathan

Abstract

While there has been much recent progress using deep learning techniques to separate speech and music audio signals, these systems typically require large collections of isolated sources during the training process. When extending audio source separation algorithms to more general domains such as environmental monitoring, it may not be possible to obtain isolated signals for training. Here, we propose objective functions and network architectures that enable training a source separation system with weak labels. In this scenario, weak labels are defined in contrast with strong time-frequency (TF) labels such as those obtained from isolated sources, and refer either to frame-level weak labels where one only has access to the time periods when different sources are active in an audio mixture, or to cliplevel weak labels that only indicate the presence or absence of sounds in an entire audio clip. We train a separator that estimates a TF mask for each type of sound event, using a sound event classifier as an assessor of the separator's performance to bridge the gap between the TF-level separation and the ground truth weak labels only available at the frame or clip level. Our objective function requires the separator to estimate a source such that the classifier applied to it will assign high probability to the class corresponding to that source and low probability to all other classes. The objective function also enforces that the separated sources sum up to the mixture. We benchmark the performance of our algorithm using synthetic mixtures of overlapping events created from a database of sounds recorded in urban environments, and show that the method can also be applied to other tasks such as music source separation. Compared to training a network using isolated sources, our model achieves somewhat lower but still significant SI-SDR improvement, even in scenarios with significant sound event overlap.

Finding Strength in Weakness:
Learning to Separate Sounds with Weak Supervision

Fatemeh Pishdadian, *Student Member, IEEE,* Gordon Wichern, *Member, IEEE,*
and Jonathan Le Roux, *Senior Member, IEEE*

Abstract—While there has been much recent progress using deep learning techniques to separate speech and music audio signals, these systems typically require large collections of isolated sources during the training process. When extending audio source separation algorithms to more general domains such as environmental monitoring, it may not be possible to obtain isolated signals for training. Here, we propose objective functions and network architectures that enable training a source separation system with weak labels. In this scenario, weak labels are defined in contrast with strong time-frequency (TF) labels such as those obtained from isolated sources, and refer either to frame-level weak labels where one only has access to the time periods when different sources are active in an audio mixture, or to clip-level weak labels that only indicate the presence or absence of sounds in an entire audio clip. We train a separator that estimates a TF mask for each type of sound event, using a sound event classifier as an assessor of the separator's performance to bridge the gap between the TF-level separation and the ground truth weak labels only available at the frame or clip level. Our objective function requires the separator to estimate a source such that the classifier applied to it will assign high probability to the class corresponding to that source and low probability to all other classes. The objective function also enforces that the separated sources sum up to the mixture. We benchmark the performance of our algorithm using synthetic mixtures of overlapping events created from a database of sounds recorded in urban environments, and show that the method can also be applied to other tasks such as music source separation. Compared to training a network using isolated sources, our model achieves somewhat lower but still significant SI-SDR improvement, even in scenarios with significant sound event overlap.

Index Terms—source separation, weak supervision, deep learning, mask inference, sound event detection

I. INTRODUCTION

AUDIO source separation aims to isolate individual sound sources in a complex auditory scene. This process plays an essential role in a variety of applications, including speech recognition in noisy environments [1], speaker identification in a multi-speaker scenario [2], and music remixing [3].

Time-frequency (TF) domain mask inference is a common approach to solving the under-determined source separation problem, in which the number of audio sources exceeds the number of recorded channels [4]–[6]. In such an approach, a raw audio mixture is first transformed into an intermediary representation, e.g., the short-time Fourier transform (STFT).

F. Pishdadian is with Interactive Audio Lab, Northwestern University, Evanston, IL, USA (e-mail: fpishdadian@u.northwestern.edu). G. Wichern and J. Le Roux are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA (e-mail: {wichern,leroux}@merl.com).
This work was performed while F. Pishdadian was an intern at MERL.

Each source is then estimated by applying a weighting function with values typically in $[0, 1]$, referred to as a *mask*, to the mixture in the transform domain before converting back to the time domain.

Supervised mask inference methods, especially those using deep neural networks, have gained much popularity over the past decade, due to their successful performance in speech enhancement [7]–[10], speech separation [11]–[15], music separation [16]–[21], and sound effect separation [22]. These approaches typically require a large dataset of isolated sound sources to construct training targets for estimating the TF masks from the corresponding sound mixtures. However, obtaining the isolated sources that compose a mixture may be expensive, require complicated recording setups, or necessitate the creation of synthetic mixtures that lack a certain amount of realism. In the extreme case, some sounds may never be recorded in isolation, such as the sound of a specific machine part that only occurs when a machine is running and other parts might also be making some sound.

In cases where isolated sources are not available for training the separation system, it is also unrealistic for humans to use signal processing tools to manually label the audio at the granularity level of TF bins, especially to do so accurately and at scale. However, it is reasonable to assume that they can produce limited labels for the activity of sound sources within some time range. Even non-expert users have successfully provided labels for musical instrument detection [23] and sound event detection (SED) [24], where the labels consisted of the type of audio events as well as the precise time of their occurrence in a given recording. The annotation burden can be further reduced, as has been considered in the SED task, by replacing the fine resolution labels on the precise sound event onsets and offsets (typically defined at a resolution of a few dozens milliseconds) by a coarse temporal label indicating the presence or absence of a sound event within a particular audio clip (e.g., on the order of 10 s). Since the fine resolution labels are typically defined at the level of an STFT frame, we hereafter refer to them as frame-level labels, while we refer to the coarse labels as clip-level labels.

In this paper, our goal is to investigate whether separation methods using deep learning, which are typically trained in a fully supervised setup using TF-bin-level labels, can be trained using weaker (frame-level or clip-level) labels. We thus attempt to perform, similarly to what has been done in the context of SED, a transition from strong to weak labels. We shall however point out an important difference regarding the notion of strength of a label in the context of SED and

separation. In SED, the goal is to estimate the type of an audio event together with its precise onset and offset, with the corresponding ground truth referred to as a strong label. In contrast, ground truth limited to presence or absence of a sound within a coarser time range is referred to as a weak label. We refer the interested reader to the description of a weakly labeled SED task in the DCASE 2017 challenge [25]. In the context of source separation, complete ground truth consists in each source's isolated signal, which amounts to having information on each source at the granularity of a TF bin. Strong labels for SED are thus only weak labels for source separation. To our knowledge, no deep-learning-based source separation system has so far been presented that can be trained under the assumption of sole availability of frame-level labels (let alone clip-level ones) and is able to separate mixtures at test time without side information.

Weakly labeled SED approaches typically leverage multiple instance learning, where an instance-level (i.e., fine temporal resolution) predictor is trained by aggregating or pooling the instance-level predictions to match the labels at the "bag" level (i.e., a chunk of audio on the order of several seconds, and its associated coarse ground truth label). We would like to use a similar concept for source separation, where the instance-level prediction is now at the level of the TF bin, and the bag level is either that of a frame or that of a whole clip. A different approach to pooling is however needed in SED and source separation systems. In weakly-labeled SED, consecutive time frames will often share the same class labels. In weakly-labeled separation, on the other hand, the structure is much more intricate, as frequency bins sharing the same label may be far from each other, often harmonically spaced in a highly variable manner even among the same types of sounds.

To overcome these difficulties in pooling over the frequency and time dimensions, we propose a form of discriminative pooling, where an SED classifier is employed as the principal metric for loss calculation when training the separator. When applied to a separated source, the classifier is expected to detect that only a single class is present, while all other sources are inactive. Furthermore, we propose a multi-task learning approach in training the separator, combining the audio event classification objective with an additional separation-specific objective that enforces the separated sources to sum up to the mixture. Our model learns to separate based solely on weak labels, either at the frame level or at the clip level. Clip-level labels are equivalents of SED weak labels, which do not require the sound to be active throughout the entire time period for which the label applies. In this work, we extend [26] to further investigate the contribution of the classification and separation objective function terms to the quality of learned masks, as well as the correlation between classifier and separator performance. We also explore different training strategies, where the classifier and separator are trained jointly, or we first train the classifier on the mixtures, and then fix or fine-tune its weights while training the separator. Empirical comparison of weakly-labeled separation performance to the strongly-labeled case (when isolated sources are available) is carried out using synthetic mixtures created from the *UrbanSound8K* [27] and *MUSDB18* [28] datasets.

Related work: As previously mentioned, there has been a resurgence in multiple instance learning approaches for audio following the DCASE 2017 challenge [25], where several studies examine deep network architectures [29]–[31] and/or pooling functions [32]–[34] for the weakly-labeled SED task. There have also been several applications of multiple instance learning for music, including detecting instruments in mixtures [35], applying artist-/album-level labels to individual tracks [36], and saliency-based singing voice detection [37].

Deep learning based techniques are currently dominant for fully supervised source separation, and typically trained to separate a single source class of interest, such as vocals or a particular instrument type from music mixtures [17], [19], [38], or speech from noise [9], [33]. An alternative class of techniques such as deep clustering [11] and permutation invariant training [11], [13] is required when the source types to be separated are very similar, e.g., separating speech from speech. The fully supervised approaches most relevant to the current study are those that train a single network to separate multiple classes of musical instruments [18], [21], [39].

Semi-supervised separation methods based on generative adversarial learning were proposed in [40], [41]. The key assumption of these methods is that estimated sources produced by an optimal separator should be indistinguishable from real sound sources, i.e., they should be samples drawn from the same distribution. These adversarial approaches are semi-supervised, since one-to-one correspondence between the mixtures and the real isolated sources used for training is not required. Nevertheless, their training is indeed dependent on the existence of some dataset of isolated sources. However, the need for isolated data can be relaxed when separating a single type of source while only observing isolated background and the target source in the background [42], [43]. Another class of source separation techniques based on weak labels assumes the availability of weak labels at both training and inference time, such as the score-informed approach in [44], the variational auto-encoder in [45], and the audio-visual approach in [46], where the video provides (weak) class labels to guide the audio separation. Our approach can separate multiple source classes, does not require seeing any sources in isolation, and requires only the audio mixture (no labels) during inference.

Another line of research performs source separation implicitly when training SED systems using either NMF [47] or deep networks [48]. The method in [48] is composed of two stages: first, a segmentation mapping is applied to the TF representation of an audio recording to obtain TF segmentation masks, then a classification mapping is applied to the segmentation masks to estimate the presence probability of sound events. The authors suggest that the separation task can be performed as a byproduct of event detection using the learned segmentation masks. However, their objective function is only event detection cross-entropy and does not include any terms modeling the separation problem explicitly, such as enforcing each separated mask to belong only to a single source, or enforcing estimated sources to sum up to the mixture as in our approach. Furthermore, they test their method only on isolated sources in background noise, whereas our experiments deal with multiple overlapping sound events.

II. JOINT SEPARATION-CLASSIFICATION APPROACH

We take a joint separation-classification approach to audio source separation through weak labels. In this section, we first provide basic definitions for the mask-based single-channel source separation problem and briefly review the fully supervised separation setup. We then present our weakly supervised separation model, formulate the objective function, and discuss the training setup in detail.

A. Mask-based Single-channel Audio Source Separation

Throughout this work, we assume a monaural source separation scenario, where only one recording channel of the mixture is available. We observe a mixture

$$x(t) = \sum_{i=1}^{n} s_i(t), \tag{1}$$

where $x(t)$ and $s_i(t)$ respectively denote the mixture signal and the i-th sound source signal in the time domain, and n is the total number of sound sources in the mixture. Note that each *sound source* is here assumed to belong to a distinct *sound class* (e.g., musical instrument, human speech, dog bark, etc.), in other words all instances of the same sound class are considered as a single sound source. We thus use these two terms interchangeably hereafter.

As mentioned earlier, a common approach to solving the under-determined separation problem is to perform masking on the mixture in some time-frequency (TF) domain, where there is less overlap between sources than in the time domain. We denote the magnitude TF representation (e.g., magnitude STFT) of the mixture by $X_{\omega,\tau}$, where ω and τ are frequency and time-frame indices, respectively.

The first step in a typical TF-masking-based method is to estimate source magnitudes by performing element-wise multiplication of the mixture magnitude with a set of estimated masking functions. Let $\hat{M}_{i,\omega,\tau}$ denote a TF mask estimate for the i-th source, taking on values in $[0, 1]$, with $\hat{M}_{i,\omega,\tau}$ being ideally very close or equal to 0 where the source is inactive and very close or equal to 1 where the source is dominant in the mixture. The masking operation can be formulated as

$$\hat{S}_{i,\omega,\tau} = \hat{M}_{i,\omega,\tau} X_{\omega,\tau}, \tag{2}$$

where $\hat{S}_{i,\omega,\tau}$ is the estimated magnitude of the i-th source. The estimated source magnitudes are then typically combined with the mixture phase and converted back to the time domain through an inverse transform (e.g., iSTFT). We leave extensions of our method that consider estimation of the phase or the complex spectrogram of the sources to future work.

B. Fully Supervised Separation

The supervised mask inference task aims at training a model to generate estimates of the sources present in a given audio mixture via the estimation of masks to be applied to a TF representation of the mixture.

In the fully supervised separation scenario, the time-domain signals of the isolated sources, their TF-domain representations, or TF masks built from them (e.g., the ideal binary mask or the ideal ratio mask [9]) are used as targets in model training. We refer to such targets as "strong labels," as they provide information about sound classes at the TF-bin level.

Various loss functions have been used in fully supervised mask inference, such as mask approximation (MA), magnitude spectrum approximation (MSA), phase spectrum approximation (PSA), and waveform approximation (WA) [9], [49]. We here focus on the MSA objective with L^1 norm for simplicity:

$$
\begin{aligned}
\mathcal{L}_{\mathrm{mi}} &= \sum_{i,\omega,\tau} |\hat{S}_{i,\omega,\tau} - S_{i,\omega,\tau}| \\
&= \sum_{i,\omega,\tau} |X_{\omega,\tau} \hat{M}_{i,\omega,\tau} - S_{i,\omega,\tau}|,
\end{aligned} \tag{3}
$$

where $S_{i,\omega,\tau}$ denotes the magnitude TF representation of the i-th isolated source and $|\cdot|$ indicates the modulus operator.

In the weakly supervised scenario, we no longer have access to TF-bin-level labels (target sources or masks). The target labels instead indicate only the sound class presence at the frame or clip level. The next two sections present our approach to training mask inference networks using only frame- or clip-level sound labels.

C. Weakly Supervised Separation

At a high level, our model is composed of two main blocks: a source separator and an audio event classifier. The block diagram of the entire system is shown in Fig. 1.

The separator block receives the TF representation (e.g., magnitude STFT) of a mixture and outputs estimates $\hat{S}_i$, $i = 1, \ldots, n$, for each of the sources in the TF domain, where n indicates the total number of sound classes available in a dataset. We assume the number of active sound classes in a given mixture ranges from 1 to n.

The input to the classifier block is also a TF representation. In general, the TF representation used as input to the classifier may be of a different type from the separator output (and the separator input, as we typically assume they are in the same domain), as long as we can pass gradients through the transform used to compute it. For instance, the classifier input can be a mel spectrogram while the separator input and output are a magnitude STFT. Given a TF representation $\boldsymbol{Y}$ as input, the classifier computes frame-level class probabilities $p_{i,\tau}(\boldsymbol{Y})$ for each source class i and time frame τ, representing how likely each source class is to be active at each time frame in $\boldsymbol{Y}$, or clip-level class probabilities $p_i(\boldsymbol{Y})$ for each source class i, representing how likely each source class is to be active (anywhere) within $\boldsymbol{Y}$.

We denote the frame-level label for the i-th sound source at frame τ by $l_{i,\tau}$, which indicates whether the source is active at that frame ($l_{i,\tau} = 1$) or not ($l_{i,\tau} = 0$). The clip-level label l_i indicates whether the source i is present anywhere in the clip, i.e., $l_i = 1$ would correspond to the case where $l_{i,\tau} = 1$ for at least one time-frame τ, and $l_i = 0$, otherwise, assuming that frame-level labels were available. Note that $l_{i,\tau}$ may be regarded as the output of a pooling operation across frequency applied to the TF-level labels for the i-th isolated source at frame τ, while l_i would be further pooled across time.
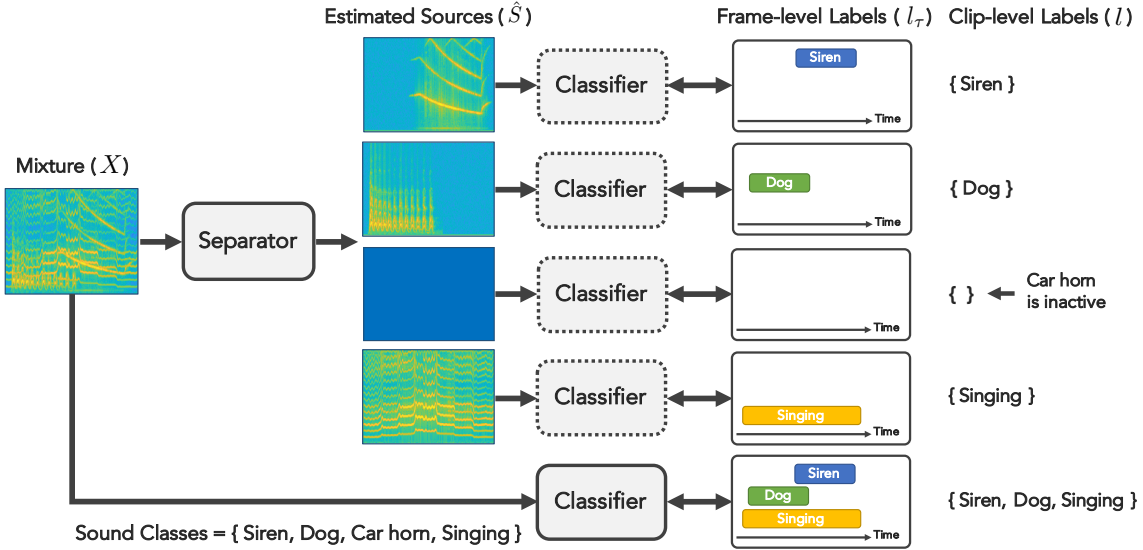
Fig. 1. The joint separation-classification model. The separator receives an audio mixture and returns source estimates (the blue square is the estimate of an inactive source). The classifier processes separately the mixture and each estimated source (dashed lines indicate shared parameters). When applied to the mixture, the classifier should output the presence probabilities for all classes. The separator is trained such that if any of the source estimates is used as input to the classifier, the classifier output is the presence probability for that source along with zeros for all other sources.

Our main idea is that, if we can train a classifier that performs well in predicting source class activities on natural mixtures, where sound classes may sometimes occur in isolation and other times overlap with other classes, we can use that classifier as a critic of the separator's performance, assessing how well the separator is able to separate each source. We can thus use weak labels, either at the frame or clip level, to train the separator through the classifier. For instance, if source i is active at frame τ, passing the estimated source $\hat{\boldsymbol{S}}_i$ as input to the classifier should result in classification outputs such that $p_{i,\tau}(\hat{\boldsymbol{S}}_i) = 1$ and $p_{j,\tau}(\hat{\boldsymbol{S}}_i) = 0$ for all $j \neq i$, because all other sources should be removed from $\hat{\boldsymbol{S}}_i$. This is illustrated in Fig. 1, where we have shown both frame-level labels, with onsets and offsets for each sound class, and clip-level labels where only presence or absence within a clip is indicated.

The classifier can be trained independently or jointly with the separator. The separator, on the other hand, requires the classification results while training, since TF-bin labels are not available and the classifier is required to pool over the TF-bin-level source activity predictions to make predictions at each time frame or for the whole clip. In this work, we consider three training strategies: i) training the separator and classifier jointly from scratch, ii) training the separator through a pre-trained classifier while the classifier is being fine-tuned, and iii) training the separator through a pre-trained and fixed classifier. It should be noted that we pre-train the classifier only on mixtures, not on isolated sources, as the latter case would violate the assumption that strong labels are unavailable.

D. Weakly Supervised Objective Function

1) Mixture loss: Our principal goal in training the model is to achieve high quality separation, which requires explicit optimization of mask estimates, even if ground truth TF labels are not available. To this end, a key constraint is

to enforce the output signals of the separator to add up to explain the input mixture. Indeed, this constraint is critical in preventing the separator from producing masks that solely focus on the most discriminating time-frequency components for classification without fully reconstructing the entire source. We can promote the enforcement of this constraint through a mixture loss term in the objective function that minimizes the discrepancy between the mixture and the sum of estimated source spectrograms, or between the mixture magnitude and the sum of estimated source magnitudes, assuming that all source estimates are obtained using the mixture phase. A vanilla version of such a term can be formulated using an L^1 loss as

$$\mathcal{L}_{\text{mix,vanilla}} = \sum_{\omega,\tau} |X_{\omega,\tau} - \sum_{i=1}^{n} \hat{S}_{i,\omega,\tau}|. \quad (4)$$

Thanks to the information provided by the weak labels, we can in fact further enforce that only the sum over active sources should be equal to the mixture, and all inactive sources should be silent. In the frame-level case, the vanilla loss term in (4) can therefore be replaced by a more explicitly constrained version defined in two parts:

$$\mathcal{L}_{\text{f-mix}} = \sum_{\omega,\tau} |X_{\omega,\tau} - \sum_{i \in \mathcal{A}_\tau} \hat{S}_{i,\omega,\tau}| + \sum_{\omega,\tau} \sum_{i \notin \mathcal{A}_\tau} |\hat{S}_{i,\omega,\tau}|, \quad (5)$$

where $\mathcal{A}_\tau$ is the set of active source indices in time frame τ. Moreover, given the weak labels, we can locate mixture frames where no sources are active and exclude those entirely from loss computation. We refer to $\mathcal{L}_{\text{f-mix}}$ as the frame-level mixture loss. In the clip-level case, we only have information about the set $\mathcal{A}$ of active source indices for the whole clip, which we can use to similarly modify (4) as:

$$\mathcal{L}_{\text{c-mix}} = \sum_{\omega,\tau} |X_{\omega,\tau} - \sum_{i \in \mathcal{A}} \hat{S}_{i,\omega,\tau}| + \sum_{\omega,\tau} \sum_{i \notin \mathcal{A}} |\hat{S}_{i,\omega,\tau}|, \quad (6)$$

which we refer to as the clip-level mixture loss. In our experiments, these refinements to the vanilla mixture loss in (4) proved very important for obtaining good mask estimates.

2) Frame-level loss: The sound classes identified by the classifier should match the correct labels, whether it is applied to the input mixture or any of the sources estimated by the separator. This can be achieved by including a binary cross-entropy term between the classifier output and the corresponding true labels. Let $H(l, p)$ denote the binary cross-entropy function defined as

$$H(l, p) = -l \log(p) - (1 - l) \log(1 - p), \quad (7)$$

where $l \in [0, 1]$ and $p \in [0, 1]$ respectively denote the true and estimated class probabilities. We denote by $\mathcal{L}_{\text{f-class}}(\boldsymbol{Y}, \tau)$ the frame-level classification loss at frame τ for an input spectrogram $\boldsymbol{Y}$ and its associated frame-level weak labels (where labels are left implicit for simplicity of notation). This loss is computed on the mixture $\boldsymbol{X}$ and on each separated source $\hat{\boldsymbol{S}}_i$. For the mixture $\boldsymbol{X}$, the classification loss at frame τ can be classically computed as the sum of binary cross-entropy terms over all sources,

$$\mathcal{L}_{\text{f-class}}(\boldsymbol{X}, \tau) = \sum_{i=1}^{n} H(l_{i,\tau}, p_{i,\tau}(\boldsymbol{X})), \quad (8)$$

where $l_{i,\tau} \in \{0, 1\}$ is the true frame-level label for the i-th source at frame τ. For the i-th estimated source $\hat{\boldsymbol{S}}_i$, the associated labels at each frame τ are obtained from the labels for mixture $\boldsymbol{X}$ by keeping only the label $l_{i,\tau}$ for the i-th source, whose activity should be the same as in $\boldsymbol{X}$, and replacing the labels for all other sources with zeros, as they should now be inactive. The loss is thus computed as:

$$\mathcal{L}_{\text{f-class}}(\hat{\boldsymbol{S}}_i, \tau) = H(l_{i,\tau}, p_{i,\tau}(\hat{\boldsymbol{S}}_i)) + \sum_{j \neq i} H(0, p_{j,\tau}(\hat{\boldsymbol{S}}_i)). \quad (9)$$

The total frame-level classification loss $\mathcal{L}_{\text{f-class}}^{\text{total}}(\tau)$ at frame τ, where the classifier is applied to the mixture and all the estimated sources, is computed as

$$\mathcal{L}_{\text{f-class}}^{\text{total}}(\tau) = \mathcal{L}_{\text{f-class}}(\boldsymbol{X}, \tau) + \sum_{i=1}^{n} \mathcal{L}_{\text{f-class}}(\hat{\boldsymbol{S}}_i, \tau). \quad (10)$$

Combining the mixture loss and the classification loss, the overall frame-level loss function to be minimized can be written as

$$\mathcal{L}_{\text{f}} = \sum_{\tau} \mathcal{L}_{\text{f-class}}^{\text{total}}(\tau) + \alpha \mathcal{L}_{\text{f-mix}}, \quad (11)$$

where $\alpha \geq 0$ is a tunable parameter controlling the contribution of the mixture loss to the total loss.

3) Clip-level loss: When only clip-level weak labels are available, we assume that the classifier outputs a single prediction at the clip level. For example, in our experiments, a time-pooling operation is applied to the output of the frame classifier to map frame labels to clip labels as is commonly done in the weakly-labeled SED literature [32], [33] (see Section II-F). The classification loss given the clip-level labels is formulated as

$$\mathcal{L}_{\text{c-class}}^{\text{total}} = \mathcal{L}_{\text{c-class}}(\boldsymbol{X}) + \sum_{i=1}^{n} \mathcal{L}_{\text{c-class}}(\hat{\boldsymbol{S}}_i), \quad (12)$$

with

$$\mathcal{L}_{\text{c-class}}(\boldsymbol{X}) = \sum_{i=1}^{n} H(l_i, p_i(\boldsymbol{X})), \quad (13)$$

$$\mathcal{L}_{\text{c-class}}(\hat{\boldsymbol{S}}_i) = H(l_i, p_i(\hat{\boldsymbol{S}}_i)) + \sum_{j \neq i} H(0, p_j(\hat{\boldsymbol{S}}_i)), \quad (14)$$

where l_i denotes the clip-level label for the i-th sound class and p_i is the clip-level class probability output by the classifier for the i-th class. Finally, the total loss in the clip-level case is computed as

$$\mathcal{L}_{\text{c}} = \mathcal{L}_{\text{c-class}}^{\text{total}} + \alpha \mathcal{L}_{\text{c-mix}}. \quad (15)$$

Table I summarizes the loss functions used with strong labels, frame-level weak labels, and clip-level weak labels. In the present work, we only consider systems trained exclusively with one of the three label types, i.e., we leave the study of combining available training data with different label-types as a topic of future work. We also note that, while the different loss functions in Table I are interrelated, the decisions they make in separating and classifying sources are not necessarily consistent with each other. For example, a clip-level classifier using max-pooling to combine frame-level decisions could make the right decision at the clip level even though all frame-level decisions are incorrect (i.e., the correct maximum score occurs at a frame that does not have the correct frame-level label).

E. Balancing Class Weights

In the preceding discussions, all sound sources contribute equally to the total loss value. This is a reasonable setup in cases where all sound sources are equally likely to be active at any given time. However, sound sources may in general occur with very different activity levels in a dataset. For instance, a dataset of urban sounds might include rare, impulsive sound events such as gun shots, as well as sounds that are active over long periods of time such as street music. Therefore, we weight each sound class during training to balance the contribution to the total loss of active and inactive frames for that class, which also equalizes the weight between classes.

Let γ_i denote the probability for the i-th source to be active at any given frame. We compute γ_i from the training data as the ratio of the number of frames in the dataset where the i-th source is active to the total number of frames in the dataset. We aim at increasing the contribution of sources occurring less frequently or for very short periods of time (e.g., $\gamma_i = 0.1$) in the total loss, while decreasing the contribution of sources that are active most of the time (e.g., $\gamma_i = 0.9$). This can be achieved by weighting the loss terms corresponding to frames where a source is active by the inverse of the source's prior probability of being active, and similarly for the frames where the source is inactive. We define the loss weight for the i-th source as

$$W_{i,\tau} = \begin{cases} \gamma_i^{-1} & i \in \mathcal{A}_\tau, \\ (1 - \gamma_i)^{-1} & i \notin \mathcal{A}_\tau, \end{cases} \quad (16)$$

where $\mathcal{A}_\tau$ is the set of active source indices in time frame τ. When using such weights in a loss term, the expected number

TABLE I
SUMMARY OF FULLY-SUPERVISED AND WEAKLY-SUPERVISED LOSS FUNCTIONS.

Label type	Eqn.	Loss function
Strong	(3)	$\mathcal{L}_{\text{mi}} = \sum_{i,\omega,\tau} \lvert X_{\omega,\tau} \hat{M}_{i,\omega,\tau} - S_{i,\omega,\tau} \rvert$
Frame	(5)	$\mathcal{L}_{\text{f-mix}} = \sum_{\omega,\tau} \lvert X_{\omega,\tau} - \sum_{i \in \mathcal{A}_\tau} \hat{S}_{i,\omega,\tau} \rvert + \sum_{\omega,\tau} \sum_{i \notin \mathcal{A}_\tau} \lvert \hat{S}_{i,\omega,\tau} \rvert$
	(10)	$\mathcal{L}_{\text{f-class}}^{\text{total}}(\tau) = \sum_i H(l_{i,\tau}, p_{i,\tau}(\boldsymbol{X})) + \sum_i (H(l_{i,\tau}, p_{i,\tau}(\hat{\boldsymbol{S}}_i)) + \sum_{j \neq i} H(0, p_{j,\tau}(\hat{\boldsymbol{S}}_i)))$
	(11)	$\mathcal{L}_{\text{f}} = \sum_\tau \mathcal{L}_{\text{f-class}}^{\text{total}}(\tau) + \alpha \mathcal{L}_{\text{f-mix}}$
Clip	(6)	$\mathcal{L}_{\text{c-mix}} = \sum_{\omega,\tau} \lvert X_{\omega,\tau} - \sum_{i \in \mathcal{A}} \hat{S}_{i,\omega,\tau} \rvert + \sum_{\omega,\tau} \sum_{i \notin \mathcal{A}} \lvert \hat{S}_{i,\omega,\tau} \rvert$
	(12)	$\mathcal{L}_{\text{c-class}}^{\text{total}} = \sum_i H(l_i, p_i(\boldsymbol{X})) + \sum_i (H(l_i, p_i(\hat{\boldsymbol{S}}_i)) + \sum_{j \neq i} H(0, p_j(\hat{\boldsymbol{S}}_i)))$
	(15)	$\mathcal{L}_{\text{c}} = \mathcal{L}_{\text{c-class}}^{\text{total}} + \alpha \mathcal{L}_{\text{c-mix}}$

of frames contributing to that loss is not only the same for active and inactive regions of a given source, but also the same across all sources.

We can incorporate these weights in the fully supervised mask inference loss (3) as

$$\mathcal{L}_{\text{mi},W} = \sum_{i,\omega,\tau} W_{i,\tau} \left| X_{\omega,\tau} \hat{M}_{i,\omega,\tau} - S_{i,\omega,\tau} \right|. \qquad (17)$$

We can also incorporate these weights in the case of frame-level weak labels, reformulating the classification loss functions from (8) and (9) as follows:

$$\mathcal{L}_{\text{f-class},W}(\boldsymbol{X}, \tau) = \sum_{i=1}^{n} W_{i,\tau} H(l_{i,\tau}, p_{i,\tau}(\boldsymbol{X})), \qquad (18)$$

$$\mathcal{L}_{\text{f-class},W}(\hat{\boldsymbol{S}}_i, \tau) = W_{i,\tau} H(l_{i,\tau}, p_{i,\tau}(\hat{\boldsymbol{S}}_i))$$
$$+ \sum_{j \neq i} W_{j,\tau} H(0, p_{j,\tau}(\hat{\boldsymbol{S}}_i)). \qquad (19)$$

The total classification loss $\mathcal{L}_{\text{f-class}}^{\text{total}}(\tau)$ in (10) and the overall loss function $\mathcal{L}_{\text{f}}$ in (11) are modified accordingly, leading to $\mathcal{L}_{\text{f-class},W}^{\text{total}}(\tau)$ and $\mathcal{L}_{\text{f},W}$.

In the clip-level scenario, we no longer have access to the prior knowledge regarding sound class activities at a frame-level granularity, but we can similarly use clip-level weights if the sound classes are not uniformly distributed at the clip level. We did not consider this in our experiments as we assumed they were uniformly distributed at the clip level (equal probability of being active within a clip).

F. Network Architecture

The architecture of the separator block used in our experiments is depicted in Fig. 2(a). It is composed of a 3-layer bidirectional long short-term memory (BLSTM) network, with each layer including 600 nodes in each direction. A fully connected layer maps the output of the BLSTM network to n masks with the same size as the input mixture. Activation functions of all BLSTM units are *tanh*, while the dense layer outputs go through *sigmoid* functions, so that the mask values are always in [0,1].

To design a frame-level classifier, we explored a number of architectures, ranging from very simple, such as a small stack of fully connected layers, to increasingly more sophisticated ones, such as convolutional recurrent neural networks

(CRNNs) [50], [51]. The clip-level classifier in this work is a simple extension of the frame-level classifier. It is built by adding a max-pooling operator to the output of the frame-level classifier for each sound class, in order to perform frame-level to clip-level mapping of sound presence probabilities. More specifically, denoting by $p_{i,\tau}(\boldsymbol{Y})$ the prediction of our frame-level classifier at frame τ for sound class i with an input spectrogram $\boldsymbol{Y}$, our clip-level classifier prediction $p_i(\boldsymbol{Y})$ for class i and input $\boldsymbol{Y}$ is obtained as:

$$p_i(\boldsymbol{Y}) = \max_\tau p_{i,\tau}(\boldsymbol{Y}). \qquad (20)$$

We leave the investigation of separation performance for some of the more advanced temporal pooling operations explored in [32] and [33] to future work.

Here, we present the two architectures that performed best in our experiments:

i) *RNN*: A 2-layer BLSTM network, with each layer including 100 nodes in each direction, followed by a fully connected layer that maps the BLSTM output for every time frame to n class probabilities. Activation functions of all BLSTM units are again *tanh*. Since the classifier is expected to detect the presence of multiple overlapping sound classes independently from one another, its output for each class is mapped to probability values through a *sigmoid* function. Figure 2(b) illustrates this architecture.

ii) *2D-CRNN*: A CRNN architecture composed of a 3-layer 2D convolutional network including max-pooling after each layer, followed by a BLSTM layer and a fully connected layer, which maps the BLSTM output to class probabilities. Activation functions of convolutional, BLSTM, and fully connected layers are *relu*, *tanh*, and *sigmoid*, respectively. The output of each convolutional layer is batch normalized prior to the application of the activation function. Figure 2(c) illustrates this architecture in detail. This network is a slightly modified version of the SED model proposed in [33]. Note that the second and third pooling operations in the convolutional network are applied across both frequency and time axes, which results in a downsampled version of frame-level predicted probabilities. To match this coarser time resolution while computing the frame-level loss values, we also downsample the true weak labels via max-pooling.
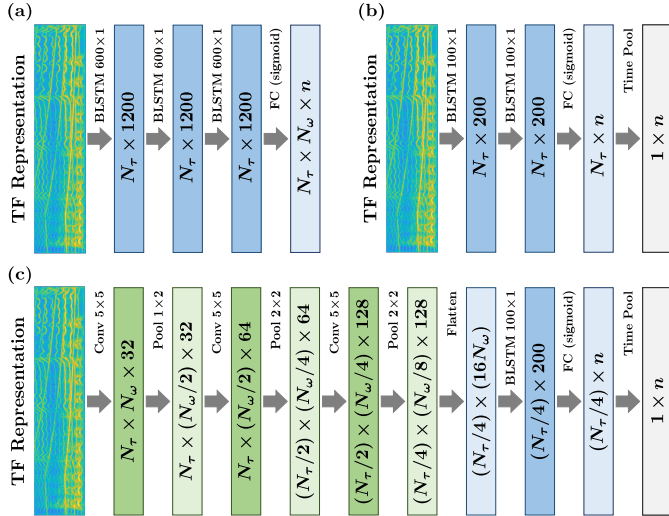
Fig. 2. Architectures of (a) the separator, (b) the RNN classifier, and (c) the 2D-CRNN classifier. N_τ and N_ω denote the number of time frames and frequency bins in the input representation, respectively. n is the total number of sound classes.

III. EXPERIMENTS

In this section, we present the results of our experiments, and compare our proposed weakly supervised method to the fully supervised approach using strong labels. We also discuss our observations regarding the importance of different model components and parameter setups.

A. Sound Event Dataset

UrbanSound8K[1] [27] is a dataset of 8732 sound excerpts of length ≤ 4 s, taken from field recordings. The dataset contains 10 sound classes: air conditioner, car horn, children playing, dog bark, drilling, engine idling, gun shot, jackhammer, siren, and street music. The audio excerpts are labeled based on the sound classes to which they belong as well as their salience in the auditory scene (foreground or background).

In our experiments, five classes are included in mixture generation: car horn, dog bark, gun shot, jackhammer, and siren. The class selection was made based on two criteria: i) audio examples in one class should contain mostly the sound corresponding to the class label, with a reasonable salience level, and ii) audio examples from different classes should be acoustically distinct enough so that they are at least recognizable as different sounds by human listeners. The air conditioner, children playing, and street music classes do not meet the first criterion, as their examples contain target sounds that are either in the background and barely audible, or accompanied by sounds from other classes. The drilling, jackhammer, and engine idling classes include many examples that sound very similar, thus only one of them was selected.

Audio mixtures in our dataset are 4 seconds long and sampled at 16 kHz. Each mixture is composed of at least one *sound event* (i.e., a single occurrence of a sound class) from one of the five selected classes. The total number of sound events per mixture is sampled from a zero-truncated Poisson

[1] https://urbansounddataset.weebly.com/urbansound8k.html

distribution with an expected value of λ. It is important to note that this number can include multiple sound events from one class, which are grouped together and regarded as one source while generating the weak labels. Thus, the value of λ determines how crowded the auditory scene is. For instance, $\lambda = 10$ means there are on average 10 sound events (from any class) per mixture. For each event, we first select one of the five classes uniformly at random, and then sample the actual sound event from all sounds of that class uniformly as well. Sound events are of arbitrary lengths, ranging from 0.5 s to 4 s, with a start time sampled uniformly at random under the constraint that the event fits entirely in the 4 s clip. The level of each sound event is randomly sampled from a uniform distribution of -30 to -25 loudness units full-scale (LUFS) [52].

UrbanSound8K is distributed with the data split into 10 folds. We use folds 1-6 for creating the training set, folds 7-8 for the validation set, and folds 9-10 for the test set. Our training, validation, and test sets include 20K, 5K, and 5K mixtures, respectively. The frame-level prior probabilities of activity γ_i (see Section II-E) for the five sound classes and λ values of 5 and 10 are presented in Table II. Since all classes were sampled uniformly during training, the clip-level prior probabilities of activity are uniformly distributed and thus not reported. To gain an idea of the amount of overlap between sources, we have also computed the distribution of frames and clips containing different numbers of sources in the entire training set. This information is provided in Table III.

TABLE II
FRAME-LEVEL PRIOR PROBABILITIES OF ACTIVITY γ_i FOR THE FIVE
SELECTED SOUND CLASSES. THE PROBABILITIES ARE COMPUTED FOR
TRAINING DATASETS WITH DIFFERENT λ VALUES.

	Sound class				
λ	Car horn	Dog bark	Gun shot	Jackhammer	Siren
5	0.26	0.36	0.27	0.40	0.40
10	0.44	0.57	0.45	0.62	0.63

TABLE III
DISTRIBUTION OF FRAMES AND CLIPS CONTAINING DIFFERENT NUMBERS
OF SOURCES IN TRAINING DATASETS WITH DIFFERENT λ VALUES.

	Number of sources per frame					
λ	0	1	2	3	4	5
5	0.17	0.28	0.30	0.18	0.06	0.01
10	0.07	0.13	0.21	0.28	0.23	0.08

	Number of sources per clip					
λ	0	1	2	3	4	5
5	0.00	0.06	0.20	0.34	0.30	0.10
10	0.00	0.00	0.02	0.12	0.38	0.48

B. Training Setup

In all training sessions, we used the ADAM optimizer, with a learning rate of 10^{-4}, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. The batch size was set to 10 in all experiments, except in the experiment investigating the effect of a shorter window size (8 ms) where the batch size was set to 8 (see Table X). We train all networks until the loss on the validation set stops

TABLE IV
FRAME-LEVEL SOUND SOURCE CLASSIFICATION PERFORMANCE IN
TERMS OF F-MEASURE. THE CLASSIFIERS ARE TRAINED AND TESTED ON
DATASETS WITH $\lambda = 5$.

	Sound class				
Classifier	Car horn	Dog bark	Gun shot	Jackhammer	Siren
RNN (mel-40)	0.850	0.813	0.809	0.915	0.811
2D-CRNN (STFT)	0.948	0.870	0.856	0.940	0.876

TABLE V
CLIP-LEVEL SOUND SOURCE CLASSIFICATION PERFORMANCE IN TERMS
OF F-MEASURE. THE CLASSIFIERS ARE TRAINED AND TESTED ON
DATASETS WITH $\lambda = 5$.

	Sound class				
Classifier	Car horn	Dog bark	Gun shot	Jackhammer	Siren
RNN (mel-40)	0.914	0.915	0.915	0.934	0.864
2D-CRNN (STFT)	0.958	0.924	0.949	0.922	0.914

improving for five consecutive epochs, with a maximum of 50 epochs. The separator takes the log-magnitude STFT of a mixture as input using the square root of a Hann window of size 32 ms and a hop size of 8 ms. To provide an upper bound for the separation performance, we trained a separator network as described in Section II-F on strong labels (i.e., target sources) with the weighted version of the fully supervised mask inference loss function $\mathcal{L}_{\mathrm{mi},W}$ in (17). In the weak label cases, we considered three training strategies: i) training the separator and classifier jointly from scratch using (11), ii) pre-training the classifier until convergence using (8), then training the separator through the pre-trained classifier while the classifier is being fine-tuned using (11), and iii) pre-training the classifier until convergence using (8), then training the separator through the pre-trained and fixed classifier using (11), where only the term involving the estimated sources contributes to the gradient. Our most effective setup used the 2D-CRNN classifier shown in Fig. 2(c), with linear magnitude STFT features as classifier input, where we first pre-trained the classifier on the mixtures, then trained the separator through the fixed pre-trained classifier using a mixture loss weight $\alpha = 100$ in (11) and (15). We use this as our default setup, and explore the importance of these choices in Section III-D.

C. Sound Event Separation Results

We evaluate the performance of the classifier in terms of F-measure $\mathcal{F} = \frac{2\mathcal{P}\mathcal{R}}{\mathcal{P}+\mathcal{R}}$, the harmonic mean of precision $\mathcal{P} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}}$ and recall $\mathcal{R} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}}$, where TP, FP, and FN respectively denote the number of true positives, false positives, and false negatives in the classification results. To measure the quality of the separated sources, we use the scale-invariant source-to-distortion ratio (SI-SDR) [12], [53], which has been shown to be more appropriate for single-channel instantaneous separation evaluation than the original SDR [54]. When computing SI-SDR over the test set, we ignore silent sources as well as any mixtures that contain isolated sources, which can happen occasionally for $\lambda = 5$ (see Table III).

Tables IV and V present the average F-measure for frame-level and clip-level sound classification, respectively. The input to the RNN classifier is a magnitude mel spectrogram with 40 filters and the 2D-CRNN input is a magnitude STFT with linear frequency. It can be observed that, at the frame level, the 2D-CRNN classifier outperforms the RNN classifier by a large margin in identifying all sound sources. The two classifiers perform more similarly at the clip level, with 2D-CRNN working slightly worse than RNN for the jackhammer class, but still better than RNN for all other classes.

Source separation results for strong labels (fully supervised upper bound) and weak labels are shown in Table VI and Fig. 3, where the weak label results are obtained with our default setup described above, using a pre-trained 2D-CRNN classifier. In Table VI, we present both means and medians of SI-SDR values, since the former measure is commonly used for reporting separation results in the literature, and the latter is better suited to the non-Gaussian distributions with a large number of outliers, which is particularly the case for the strong label results in Fig. 3. Since the overall trends between mean and median results in Table VI are similar, for clarity we report only mean results for the ablation studies in Section III-D. From these summary statistics, we see SI-SDR improvements with respect to the mixture for all classes with both frame- and clip-level weak labels. The smallest and largest SI-SDR improvements in Table VI are for the siren and gun shot classes, respectively. The siren class in our dataset contains a more diverse set of sounds compared to other classes (e.g., police siren versus ambulance siren), which is likely the reason why it is the most difficult sound type to separate even when strong labels are used.

The scatter plots of separation results, shown in Fig. 3, allow a more detailed comparison between the performance of separators trained through different types of labels. We note that all test mixtures included in these plots contain at least two sound sources. Each panel shows the amount of SI-SDR improvement versus input SI-SDR for all test set examples of one sound class. The input SI-SDR refers to the SI-SDR obtained when considering the input mixture as the estimate for the target source. One common trend observed in all cases is the downward tilted shape of the data distribution, which is also typically observed in speech separation [12], [55]. This pattern indicates that the highest SI-SDR improvement is achieved for low-SI-SDR inputs and the amount of improvement shrinks when using inputs with higher SI-SDR values.

When going from strong to weak labels in all sound classes, an obvious trend is a decrease in the number of points in the higher end of SI-SDR improvement values. For example, in the plot corresponding to the results for the car horn class and strong labels (top row, leftmost panel), there are several points with SI-SDR improvements above 50 dB. When using frame-level labels, the highest SI-SDR improvement drops to around 40 dB, and it decreases even further down to 30 dB when using clip-level labels. Interestingly, however, the high-density regions of the distributions in each class seem to remain largely similar, contrary to what one may have expected given the

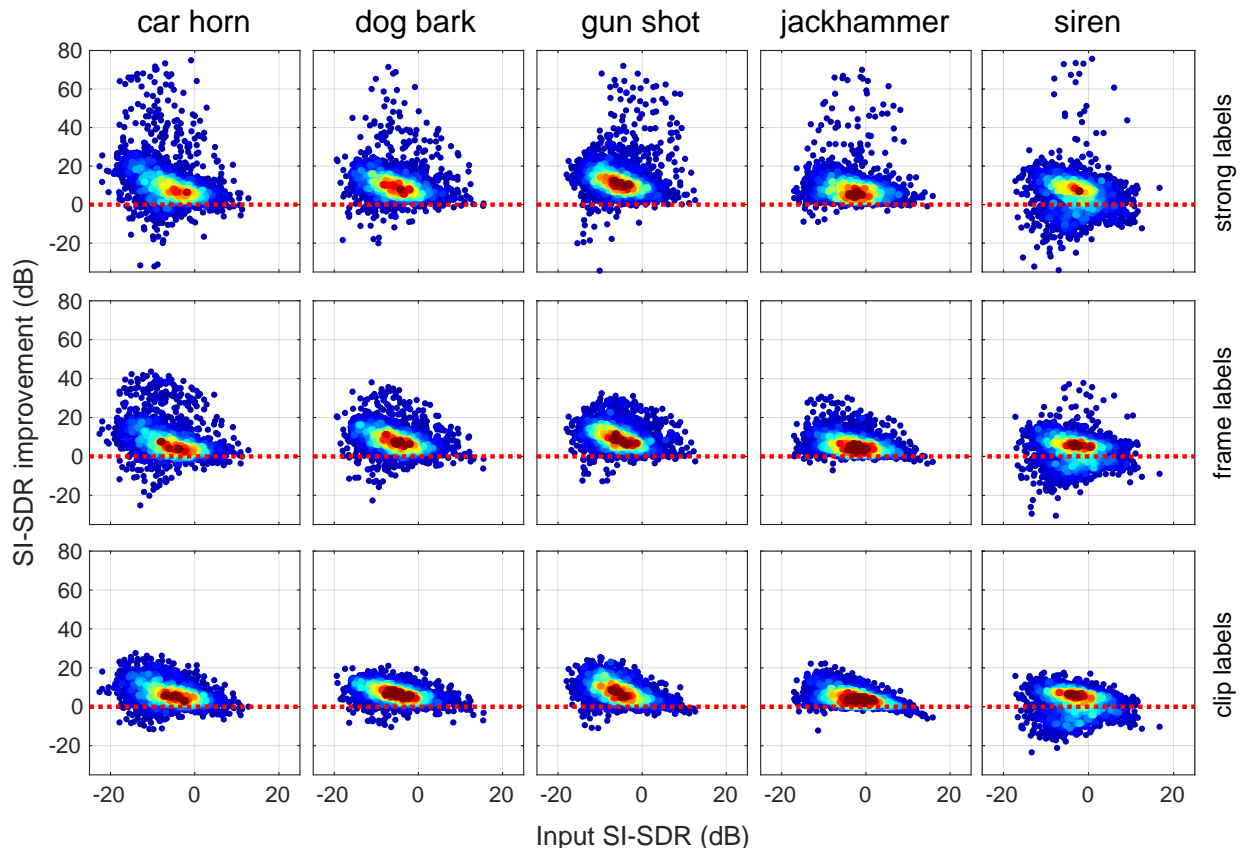|	Sound class										
	Car horn		Dog bark		Gun shot		Jackhammer		Siren		Overall
Input SI-SDR	-5.8 /	-5.7	-5.4 /	-5.4	-5.5 /	-5.8	-2.9 /	-2.8	-3.0 /	-3.0	-4.5 / -4.5
ΔSI-SDR-strong	9.9 /	7.9	10.0 /	9.2	12.5 /	11.2	7.8 /	6.8	4.9 /	6.2	9.0 / 8.3
ΔSI-SDR-frame	7.0 /	5.4	8.3 /	7.7	9.7 /	9.1	5.7 /	4.9	3.1 /	4.3	6.8 / 6.2
ΔSI-SDR-clip	6.5 /	5.7	6.4 /	6.1	8.8 /	8.3	4.6 /	4.0	1.8 /	3.5	5.6 / 5.5

Fig. 3. Separation results for all sound classes when the separator is trained on strong labels (top row), frame-level weak labels (middle row), and clip-level weak labels (bottom row). All panels show SI-SDR improvement versus input SI-SDR values. The 2D-CRNN classifier and the magnitude STFT input are used in experiments with both frame-level and clip-level labels. Warmer colors mean higher densities of data points.

difference in the strength of labels used for training. Although frame-level labels yield better results than clip-level labels in general, the distribution of output SI-SDRs for these two label types seems to be very similar in all cases. Furthermore, both weak-label distributions seem to have large amounts of overlap with strong-label distributions and to provide significant SI-SDR improvement over the input SI-SDRs.

D. Ablation studies

1) Training strategies: The effect of different training strategies on separation performance can be observed in Table VII. The joint separation-classification model was trained on frame-level weak labels under the three training strategies listed in Section II-C. Regardless of the classifier architecture, the best separation results are achieved when the classifier is pre-trained on mixtures and its parameters are then fixed when training the separator. Training the separator and classifier

jointly from scratch, or fine-tuning the classifier when the separator is being trained always resulted in a worse separation performance in our experiments. We hypothesize that this behavior is due to the co-adaptation of the two networks, where the classifier can adapt its weights to correctly classify errors made by the separator, rather than forcing the separator to output estimated sources that match the previously learned representation for each sound class. In other words, this co-adaptation weakens the ability of the classifier to objectively assess the performance of the separator.

2) Mixture loss: To investigate the effect of the mixture loss term, we trained the separator network using different α values in the overall frame-level loss of (11). The SI-SDR improvement results, presented in Table VIII, clearly show the importance of this loss term for the separation task. A similar trend is observed for both classifiers. When $\alpha = 0$, only the classification loss is used to train the separator, which leads

TABLE VII
MEAN SI-SDR IMPROVEMENT (dB) FOR DIFFERENT TRAINING
STRATEGIES, OVER ALL CLASSES. THE MODELS ARE TRAINED ON
FRAME-LEVEL LABELS. IN ALL CASES, $\alpha = 100$, $\lambda = 5$, AND THE
AVERAGE INPUT SI-SDR IS -4.5 dB.

Classifier	Training strategy		
	Joint	Fine-tune classifier	Fix classifier
RNN (mel-40)	−4.4	5.5	**6.2**
2D-CRNN (STFT)	−0.2	1.3	**6.8**

to poor separation performance as the separator network only needs to isolate the TF features necessary for classification, not signal reconstruction. Conversely, a comparatively very low contribution of the classification loss term (e.g., $\alpha = 10^4$) results in degraded performance as the separator only needs to reconstruct the mixture without isolating the individual sound sources. A good balance between the two loss terms (e.g., $\alpha = 10^2$), is essential to obtain high SI-SDR gains.

TABLE VIII
MEAN SI-SDR IMPROVEMENT (dB) USING DIFFERENT MIXTURE LOSS
WEIGHTS, OVER ALL CLASSES. THE MODELS ARE TRAINED ON
FRAME-LEVEL LABELS. IN ALL CASES, $\lambda = 5$ AND THE AVERAGE INPUT
SI-SDR IS -4.5 dB.

Classifier	Mixture loss weight (α)				
	0	10	10^2	10^3	10^4
RNN (mel-40)	1.9	4.6	**6.2**	5.3	1.9
2D-CRNN (STFT)	0.9	3.9	**6.8**	5.1	1.1

3) TF representation: The properties of the audio representation input to the classifier, such as frequency scaling and resolution, proved to have a considerable impact on the separation results in our experiments. The performance of the separator is essentially correlated with the efficacy of the classifier in capturing the spectro-temporal patterns that distinguish each sound class from the others. For instance, a classifier that depends only on a few frequency bins to identify a sound will output accurate class probabilities as long as the separator assigns correct amounts of energy to those bins. Using such a classifier, the model could correctly identify an impulsive, broadband sound (e.g., gun shot) in a mixture, even if the separated source estimate includes only a small portion of the actual spectral content.

One way to address this problem is to force the classifier to produce predictions based on broader frequency ranges by decreasing the frequency resolution of mixtures and estimated sources prior to feeding them to the classifier. To lower the frequency resolution, we consider applying a mel-frequency filterbank to the magnitude STFTs to be used as classifier inputs. The mel-frequency filterbank also has the advantage of changing the frequency resolution logarithmically, with a grid that is finer across lower frequencies, maintaining most of the information necessary to distinguish harmonic sources, and grows coarser as the frequency increases. We investigate the effect of frequency scaling and resolution on the quality of spectral patterns learned by the classifier, which in turn impacts the separation quality, by using two different repre-

sentations as the classifier input: a linear magnitude STFT and a linear magnitude mel spectrogram with a varying number of mel filters. The STFT parameters (window size and hop length) are the same for the separator and classifier inputs.

The amount of SI-SDR improvement for different mel-frequency filterbanks (featuring different numbers of filters and different center frequencies) are provided in Table IX. The results for the linear magnitude STFT, corresponding to the linear frequency case (no filterbank), are reported as "Linear Freq." in the table. As can be seen, changing the frequency scale and resolution of the classifier input can make a difference of up to 2 dB in the average SI-SDR improvement. The performance of a model using the RNN classifier can be improved up to 0.4 dB by using a mel spectrogram as input. The best number of mel filters, however, seems to be difficult to choose without running a grid search. The 2D-CRNN classifier, on the other hand, provides the highest improvement when the original magnitude STFT is used as input. We hypothesize that since convolutional networks inherently downsample frequency, using an input with low frequency resolution (e.g., mel) is more harmful than beneficial, while the RNN architecture, which performs no implicit downsampling, benefits from using mel spectrogram inputs.

We also note that the 2D-CRNN classifier consistently outperforms the RNN classifier in terms of separation performance in Tables VII-IX. Since the 2D-CRNN classifier also provided the best classification performance for most classes in Tables IV and V, these results imply that better classification performance is correlated with better separation performance when training a weakly labeled separation system. Further refinements of the classifier may thus lead to improved separation quality.

TABLE IX
MEAN SI-SDR IMPROVEMENT (dB) USING DIFFERENT FREQUENCY
SCALES AND RESOLUTIONS, OVER ALL CLASSES. THE MODELS ARE
TRAINED ON FRAME-LEVEL LABELS. IN ALL CASES, $\alpha = 100$, $\lambda = 5$ AND
THE AVERAGE INPUT SI-SDR IS -4.5 dB.

Classifier	Number of mel filters					
	10	20	30	40	56	Linear freq.
RNN	5.5	**6.4**	6.0	6.2	5.0	6.1
2D-CRNN	5.3	4.8	5.9	6.0	6.2	**6.8**

We further investigate the effect of frequency and time resolutions by varying the window size of the magnitude STFTs input to both separator and classifier blocks. We run this experiment using the best performing model thus far, which includes a 2D-CRNN classifier with magnitude STFT (linear frequency scale) as input. The results, provided in Table X, demonstrate that decreasing the time resolution by using windows longer than 32 ms degrades the results, while decreasing the window size provides performance improvement. The limit on the improvement, however, seems to be reached when the window size is 16 ms, as using a window size of 8 ms results in worse performance, likely due to the poor frequency resolution.

4) Source density: Next, we compare how the density of sound sources in the scene impacts network performance.

TABLE X
MEAN SI-SDR IMPROVEMENT (dB) FOR DIFFERENT STFT WINDOW
SIZES, OVER ALL CLASSES. IN ALL CASES, THE 2D-CRNN CLASSIFIER IS
USED WITH MAGNITUDE STFT INPUT FEATURES. THE MODELS ARE
TRAINED ON FRAME-LEVEL LABELS. IN ALL CASES, $\alpha = 100$, $\lambda = 5$, THE
OVERLAP BETWEEN WINDOWS IS 75%, AND THE AVERAGE INPUT SI-SDR
IS -4.5 dB.

	Sound class				
Win. size	Car horn	Dog bark	Gun shot	Jackhammer	Siren
8 ms	5.8	6.8	10.0	4.9	3.0
16 ms	**7.8**	**8.7**	**10.4**	**6.1**	**4.1**
32 ms	7.0	8.3	9.7	5.7	3.2
64 ms	5.1	5.8	7.0	4.1	2.8

As mentioned in Section III-A, the parameter λ used when creating our mixtures determines the expected number of events in each four-second scene. The left four columns in Table XI compare separation performance between training using strong labels (fully supervised upper bound) and both frame- and clip-level weak labels for different λ values, where we only report results with the 2D-CRNN for brevity. In the fully supervised case, training with more difficult mixtures ($\lambda = 10$) leads to improved separation performance compared to training with easier mixtures ($\lambda = 5$). However, for frame-level weak labels, training with $\lambda = 10$ leads to slightly worse performance than training with $\lambda = 5$, and for clip-level weak labels training with $\lambda = 10$ causes a larger performance drop. Revisiting Table III, we see that when $\lambda = 10$ the training set contains no clips with only a single active source, compared to 6% of the clips for $\lambda = 5$, while there are some single source frames for both λ values. Therefore, we hypothesize that the drop in clip-level weak-label performance for the $\lambda = 10$ training set in Table XI is due to the lack of any training data containing labeled regions with isolated sources. While the higher SI-SDR numbers in Table XI for both frame-level and clip-level weak labels using the $\lambda = 5$ training set indicate that the network does likely make use of labeled regions containing isolated sources, the method can still work in their absence, as shown by the SI-SDR improvements of 4.9 dB and 4.4 dB for the most difficult case of clip-level weak-labels and the $\lambda = 10$ training set.

5) Class weights: Finally, we investigate the importance of the loss weights, computed based on the prior probability of sound class activities, in separation training (see Section II-E). In our training setup, the loss weights are used in the case of strong labels and frame-level labels, in the formulation of the mask inference loss and the classification loss, respectively. The results presented in Table XI (right four columns) are obtained using the same setup as in the previous section, with the only difference that the loss terms are not weighted. Removing the weights when a less dense training dataset (e.g., $\lambda = 5$) is used results in a larger performance drop in both strong and frame-level cases. Intuitively, this behavior is expected as in such scenarios, the prior probabilities are much smaller than 0.5 for sparser classes, such as gun shot (see Table II), and hence these classes are assigned larger weights

than less sparse classes, such as jackhammer. However, as the prior probabilities get closer to 0.5 in a dataset with $\lambda = 10$, the difference between class weights becomes smaller and their effect on separation results less noticeable. Further, it can be observed that removing the weights has a more dramatic impact on the strong label than frame-level label results. This can be explained by considering the fact that in the frame-level label case, the weights are only incorporated in the classification loss. Therefore, as long as the classification performance does not degrade significantly (which was the case in our experiments), the separation performance is anticipated not to vary dramatically.

E. Unsuccessful attempts

In addition to using the RNN and 2D-CRNN sound event classifiers shown in Figs. 2(b) and 2(c) for frequency pooling, we explored simple frequency pooling (e.g., average pooling over frequency), a learned linear transform, or a feedforward deep network (without memory). In all cases, these frequency pooling approaches failed to learn to separate. Furthermore, we experimented with an "idempotent" loss, where an additional loss function term enforced the separation network to pass estimated separated sources unchanged. We found that this loss term only hurt separation performance. Our hypothesis is that in most cases this constraint was redundant with information provided by the weak labels. Finally, we explored using log magnitude STFT features (as opposed to linear magnitude) as input to the classifier. Although log features gave slightly better classification performance, our separator networks did not train reliably, even when regularizing the log (i.e., adding a small positive value to the log input).

F. Application to Music Source Separation

1) Dataset: MUSDB18 [28] currently serves as the standard dataset for benchmarking music source separation algorithms [20]. It consists of 150 songs split into a 100 song training set and 50 song test set. Each song consists of a mixture, and four stems that compose the mixture, containing isolated audio tracks for *bass, drums, vocals* and *other* source classes. These isolated sources serve as training targets (i.e., strong labels) for fully-supervised separation. The lack of precision in the definition of the *other* class makes training a classifier to reliably predict the presence or absence of this class difficult, so in the present study we remove this source and consider three-source mixtures of *bass, drums*, and *vocals*.

We further divide the training set into a subset of 86 songs used for training and another with the remaining 14 songs used as a validation set. Each song in the training and validation sets is further split into four second non-overlapping chunks (matching the clip-length of our sound event dataset described in Section III-A), which are used as input to our system. The stems for each four-second chunk are used as strong labels, and the presence or absence of energy in that stem is used as a clip-level weak label. We only consider clip-level weak labels for MUSDB18 since several of the stems in MUSDB18 contain bleed from other instruments or high levels

TABLE XI
MEAN SI-SDR IMPROVEMENT (dB) FOR SEPARATORS TRAINED USING STRONG LABELS, FRAME-LEVEL WEAK LABELS, CLIP-LEVEL WEAK LABELS, DATASETS WITH DIFFERENT λ VALUES, AND WEIGHTED (LEFT FOUR COLUMNS) OR NOT WEIGHTED (RIGHT FOUR COLUMNS) LOSS FUNCTIONS. ALL SEPARATORS ARE TRAINED USING THE 2D-CRNN CLASSIFIER.

	With class weights				Without class weights			
Training λ	5		10		5		10	
Testing λ	5	10	5	10	5	10	5	10
Input SI-SDR	-4.5	-6.2	-4.5	-6.2	-4.5	-6.2	-4.5	-6.2
ΔSI-SDR-strong	9.0	7.0	9.4	7.3	3.2	1.8	9.3	7.4
ΔSI-SDR-frame	6.8	5.2	6.4	5.4	6.3	4.9	6.1	5.3
ΔSI-SDR-clip	5.6	4.7	4.9	4.4	5.6	4.7	4.9	4.4

TABLE XII
MEAN/MEDIAN SI-SDR VALUES (dB) FOR MUSDB18 USING STRONG AND WEAK LABELS. ΔSI-SDR INDICATES THE SI-SDR IMPROVEMENT. THE LAST COLUMN SHOWS THE RESULTS OVER ALL SAMPLES AND ALL CLASSES.

	Sound class			
	Bass	Drums	Vocals	Overall
Input SI-SDR	-4.3 / -3.9	-1.8 / -2.1	-4.9 / -4.0	-3.7 / -3.6
ΔSI-SDR-strong	6.2 / 6.6	6.0 / 6.2	13.6 / 13.9	8.6 / 8.1
ΔSI-SDR-clip	1.5 / 2.1	2.7 / 2.7	8.9 / 8.8	4.4 / 3.6

of background noise that make obtaining accurate frame-level labels ambiguous.

All audio is downsampled to 16 kHz, and our training setup matches exactly the best performing setup used for the sound event dataset described in Section III-B, i.e., 2D-CRNN classifier with STFT features and mixture loss weight $\alpha = 100$. We made two modifications to the network architecture described in Section II-F for MUSDB: i) as mentioned in [32] because music sources tend to be consistently active/inactive over periods of several seconds, we replaced the max-pooling operation at the output of our clip-level classifier with an average pooling operation, and ii) we add a batch normalization layer at the input of the separation network, which while not beneficial for urban sound event separation improved performance for MUSDB18.

2) Results: We evaluate performance by computing SI-SDR on each of the 50 songs in the test set. We run each entire song through the separator network, i.e., we do not divide the songs in the test set into chunks. Table XII compares the overall performance in terms of mean and median SDR for strong labels and clip-level weak labels, while Fig. 4 shows corresponding scatter plots. From Table XII, we observe improvements in SI-SDR for all source classes and note that the 4.2 dB difference between the overall mean SDR improvements for the strong and clip-level weak labels (8.6 dB vs. 4.4 dB) for MUSDB18 is comparable to the 3.4 dB difference observed for sound events in Table VI. We also note similarities in the overall shape of the strong and weak-label scatter plots of Fig. 4, especially for vocals.

Compared to vocals, both bass and drums exhibit a larger drop in performance between the strong and weak label cases. This is most likely due to these sources almost always playing at the same time. Thus, without any prior information it is extremely difficult for a classifier trained on mixtures to learn the difference between these co-occurring classes. We also mention that when trying to include an overly broad class in

our experiments, such as *other* in MUSDB18, our separation networks tended to put most of the energy into the output corresponding to this class. Tackling this issue and improving performance for consistently co-occurring sources such as bass and drums are important extensions of the current work. However, given the difficult copyright issues necessary to obtain isolated sources for training music separation networks, we believe the results presented here represent a potentially important path in scaling up music source separation.
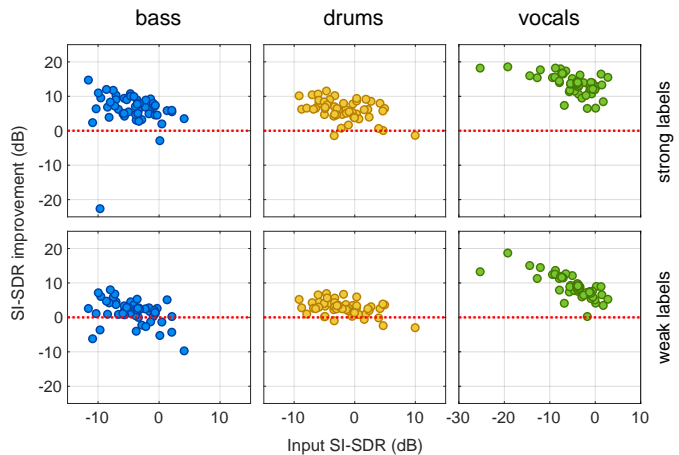

Fig. 4. Separation results for all sound classes in music mixtures when the separator is trained on strong labels (top row) and clip-level labels (bottom row). All panels show SI-SDR improvement versus input SI-SDR values. Each panel contains 50 data points, each point corresponding to an entire test song. The 2D-CRNN classifier and the magnitude STFT input are used in experiments with clip-level labels.

IV. CONCLUSION AND FUTURE WORKS

We have presented an algorithm for training a source separation system with weak labels, where isolated sources are not required for the training process. In our proposed model, an SED classifier is employed as the principal metric for loss calculation while training the separator. The model is trained

to minimize an objective function that requires the separator to produce source estimates that are identifiable by the classifier. The objective function also enforces the estimated sources to sum up to the mixture. Our experiments yielded promising results and showed significant SI-SDR improvement even when using weak labels on a very coarse-resolution time grid.

In the present work, we only explored a source separation algorithm based on magnitude masking in the spectral domain. Moving forward, we could extend our weak label separation objectives to systems based on phase sensitive masking [9], complex masking [56], phase estimation [49], time domain separation [15], and/or deep TF embeddings [11]. Combining the discriminative approach presented here with the generative approaches of [41], [43], [45] while still minimizing the amount of required supervision is also a potential avenue for future exploration. Finally, this work only considered mixtures of labelled sounds from a given set of classes, whereas real-world sound mixtures are likely to also contain unlabelled sounds from other classes. Moreover, we considered all instances of the same class as a single source, whereas one may in general be interested in further separating each instance. Dealing with such limitations is an important topic for future work.

ACKNOWLEDGMENT

REFERENCES

[1] F. Ehlers and H. G. Schuster, "Blind separation of convolutive mixtures and an application in automatic speech recognition in a noisy environment," *IEEE Transactions on Signal Processing*, vol. 45, no. 10, pp. 2608–2612, 1997.

[2] M. Cooke, J. R. Hershey, and S. J. Rennie, "Monaural speech separation and recognition challenge," *Computer speech & language*, vol. 24, no. 1, pp. 1–15, 2010.

[3] J. F. Woodruff, B. Pardo, and R. B. Dannenberg, "Remixing stereo music with score-informed source separation." in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, Oct. 2006, pp. 314–319.

[4] R. Lyon, "A computational model of binaural localization and separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 1983, pp. 1148–1151.

[5] M. Weintraub, "A theory and computational model of monaural auditory sound separation," Ph.D. dissertation, Ph. D. dissertation, Stanford Univ, 1985.

[6] D. Wang and G. J. Brown, *Computational Auditory Scene Analysis:Principles, Algorithms, and Applications.* Wiley-IEEE press, 2006.

[7] F. Weninger, F. Eyben, and B. Schuller, "Single-channel speech separation with memory-enhanced recurrent neural networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 3709–3713.

[8] F. Weninger, J. Le Roux, J. R. Hershey, and B. Schuller, "Discriminatively trained recurrent neural networks for single-channel speech separation," in *Proc. IEEE GlobalSIP Symposium on Machine Learning Applications in Speech Processing*, Dec. 2014.

[9] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2015, pp. 708–712.

[10] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.

[11] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2016, pp. 31–35.

[12] Y. Isik, J. Le Roux, Z. Chen, S. Watanabe, and J. R. Hershey, "Single-channel multi-speaker separation using deep clustering," in *Proc. ISCA Interspeech*, Sep. 2016, pp. 545–549.

[13] M. Kolbæk, D. Yu, Z.-H. Tan, and J. Jensen, "Multi-talker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.

[14] Z.-Q. Wang, J. Le Roux, and J. R. Hershey, "Alternative objective functions for deep clustering," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018.

[15] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 27, no. 8, pp. 1256–1266, 2019.

[16] Y. Luo, Z. Chen, J. R. Hershey, J. Le Roux, and N. Mesgarani, "Deep clustering and conventional networks for music separation: Stronger together," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2017, pp. 61–65.

[17] N. Takahashi, N. Goswami, and Y. Mitsufuji, "MMDenseLSTM: An efficient combination of convolutional and recurrent neural networks for audio source separation," in *Proc. IEEE International Workshop on Acoustic Signal Enhancement (IWAENC)*, Sep. 2018.

[18] P. Seetharaman, G. Wichern, S. Venkataramani, and J. Le Roux, "Class-conditional embeddings for music source separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 301–305.

[19] 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. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Oct. 2019.

[20] 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.

[21] R. Kumar, Y. Luo, and N. Mesgarani, "Music source activity detection and separation using deep attractor network." in *Proc. ISCA Interspeech*, Sep. 2018, pp. 347–351.

[22] I. Kavalerov, S. Wisdom, H. Erdogan, B. Patton, K. Wilson, J. Le Roux, and J. R. Hershey, "Universal sound separation," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Oct. 2019.

[23] E. Humphrey, S. Durand, and B. McFee, "OpenMIC-2018: An open data-set for multiple instrument recognition." in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, Sep. 2018, pp. 438–444.

[24] M. Cartwright, G. Dove, A. E. Méndez Méndez, J. P. Bello, and O. Nov, "Crowdsourcing multi-label audio annotation tasks with citizen scientists," in *Proc. ACM CHI*, 2019, p. 292.

[25] A. Mesaros, A. Diment, B. Elizalde, T. Heittola, E. Vincent, B. Raj, and T. Virtanen, "Sound event detection in the DCASE 2017 challenge," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 6, pp. 992–1006, 2019.

[26] F. Pishdadian, G. Wichern, and J. Le Roux, "Learning to separate sounds from weakly labeled scenes," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020, pp. 91–95.

[27] J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," in *Proc. ACM Multimedia*, Nov. 2014, pp. 1041–1044.

[28] Z. Rafii, A. Liutkus, F.-R. Stöter, S. I. Mimilakis, and R. Bittner, "The MUSDB18 corpus for music separation," Dec. 2017. [Online]. Available: https://doi.org/10.5281/zenodo.1117372

[29] A. Kumar and B. Raj, "Audio event detection using weakly labeled data," in *Proc. ACM Multimedia*, 2016, pp. 1038–1047.

[30] T.-W. Su, J.-Y. Liu, and Y.-H. Yang, "Weakly-supervised audio event detection using event-specific gaussian filters and fully convolutional networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2017, pp. 791–795.

[31] Y. Xu, Q. Kong, W. Wang, and M. D. Plumbley, "Large-scale weakly supervised audio classification using gated convolutional neural network," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 121–125.

[32] B. McFee, J. Salamon, and J. P. Bello, "Adaptive pooling operators for weakly labeled sound event detection," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 26, no. 11, pp. 2180–2193, 2018.

[33] Y. Wang, J. Li, and F. Metze, "A comparison of five multiple instance learning pooling functions for sound event detection with weak labeling," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 31–35.

[34] B. Kim and B. Pardo, "Sound event detection using point-labeled data," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Oct. 2019.

[35] D. Little and B. Pardo, "Learning musical instruments from mixtures of audio with weak labels." in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, Sep. 2008, pp. 127–132.

[36] M. I. Mandel and D. P. Ellis, "Multiple-instance learning for music information retrieval," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, Sep. 2008, pp. 577–582.

[37] J. Schlüter, "Learning to pinpoint singing voice from weakly labeled examples." in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, Aug. 2016, pp. 44–50.

[38] S. Uhlich, M. Porcu, F. Giron, M. Enenkl, T. Kemp, N. Takahashi, and Y. Mitsufuji, "Improving music source separation based on deep neural networks through data augmentation and network blending," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2017, pp. 261–265.

[39] O. Slizovskaia, L. Kim, G. Haro, and E. Gomez, "End-to-end sound source separation conditioned on instrument labels," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 306–310.

[40] N. Zhang, J. Yan, and Y. Zhou, "Weakly supervised audio source separation via spectrum energy preserved Wasserstein learning," *arXiv preprint arXiv:1711.04121*, 2017.

[41] D. Stoller, S. Ewert, and S. Dixon, "Adversarial semi-supervised audio source separation applied to singing voice extraction," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2018, pp. 2391–2395.

[42] D. Stowell and R. E. Turner, "Denoising without access to clean data using a partitioned autoencoder," *arXiv preprint arXiv:1509.05982*, 2015.

[43] M. Michelashvili, S. Benaim, and L. Wolf, "Semi-supervised monaural singing voice separation with a masking network trained on synthetic mixtures," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019.

[44] S. Ewert and M. B. Sandler, "Structured dropout for weak label and multi-instance learning and its application to score-informed source separation," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Mar. 2017, pp. 2277–2281.

[45] E. Karamatli, A. T. Cemgil, and S. Kirbiz, "Audio source separation using variational autoencoders and weak class supervision," *IEEE Signal Processing Letters*, vol. 26, no. 9, pp. 1349–1353, 2019.

[46] R. Gao and K. Grauman, "Co-separating sounds of visual objects," *arXiv preprint arXiv:1904.07750*, 2019.

[47] T. Heittola, A. Mesaros, T. Virtanen, and A. Eronen, "Sound event detection in multisource environments using source separation," in *Proc. Workshop on Machine Listening in Multisource Environments*, 2011, pp. 36–40.

[48] Q. Kong, Y. Xu, I. Sobieraj, W. Wang, and M. D. Plumbley, "Sound event detection and time–frequency segmentation from weakly labelled data," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 27, no. 4, pp. 777–787, 2019.

[49] J. Le Roux, G. Wichern, S. Watanabe, A. Sarroff, and J. R. Hershey, "Phasebook and friends: Leveraging discrete representations for source separation," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 370–382, 2019.

[50] E. Cakır, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, "Convolutional recurrent neural networks for polyphonic sound event detection," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1291–1303, 2017.

[51] S. Adavanne and T. Virtanen, "Sound event detection using weakly labeled dataset with stacked convolutional and recurrent neural network," in *Proc. Detection and Classification of Acoustic Scenes and Events Workshop (DCASE)*, Nov. 2017.

[52] E. Grimm, R. Van Everdingen, and M. 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.

[53] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR–half-baked or well done?" in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 626–630.

[54] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–1469, 2006.

[55] G. Wichern, E. McQuinn, J. Antognini, M. Flynn, R. Zhu, D. Crow, E. Manilow, and J. Le Roux, "WHAM!: Extending speech separation to noisy environments," in *Proc. ISCA Interspeech*, Sep. 2019.

[56] D. S. Williamson, Y. Wang, and D. Wang, "Complex ratio masking for monaural speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 3, pp. 483–492, 2016.
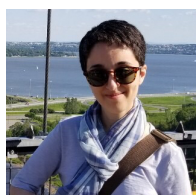
Fatemeh Pishdadian is a Ph.D. candidate in Computer Science at Northwestern University. She received her B.Sc. in Electrical Engineering from Ferdowsi University, Mashhad, Iran, and her M.Sc. in Electrical and Computer Engineering from George Mason University. Her research interest lies in the application of signal processing and machine learning methods to the analysis of audio/music. More specifically, the focus of her doctoral research has been on audio/music source separation.

Gordon Wichern is a Principal Research Scientist at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, Massachusetts. He received his B.Sc. and M.Sc. degrees from Colorado State University in electrical engineering and his Ph.D. from Arizona State University in electrical engineering with a concentration in arts, media and engineering, where he was supported by a National Science Foundation (NSF) Integrative Graduate Education and Research Traineeship (IGERT) for his work on environmental sound recognition. He was previously a member of the research team at iZotope, inc. where he focused on applying novel signal processing and machine learning techniques to music and post production software, and a member of the Technical Staff at MIT Lincoln Laboratory where he worked on radar signal processing. His research interests include audio, music, and speech signal processing, machine learning, and psychoacoustics.
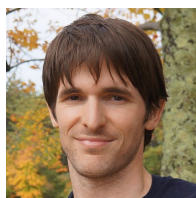
Jonathan Le Roux is a Senior Principal Research Scientist and the Speech and Audio Senior Team Leader at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, Massachusetts. He completed his B.Sc. and M.Sc. degrees in Mathematics at the Ecole Normale Supérieure (Paris, France), his Ph.D. degree at the University of Tokyo (Japan) and the Université Pierre et Marie Curie (Paris, France), and worked as a postdoctoral researcher at NTTs Communication Science Laboratories from 2009 to 2011. His research interests are in signal processing and machine learning applied to speech and audio. He has contributed to more than 100 peer-reviewed papers and 20 granted patents in these fields. He is a founder and chair of the Speech and Audio in the Northeast (SANE) series of workshops, and a Senior Member of the IEEE.