
A Sequential Self Teaching Approach for Improving Generalization in Sound Event Recognition

Anurag Kumar¹ Vamsi Krishna Ithapu¹

Abstract

An important problem in machine auditory perception is to recognize and detect sound events. In this paper, we propose a sequential self-teaching approach to learning sounds. Our main proposition is that it is harder to learn sounds in adverse situations such as from weakly labeled and/or noisy labeled data, and in these situations a single stage of learning is not sufficient. Our proposal is a sequential stage-wise learning process that improves generalization capabilities of a given modeling system. We justify this method via technical results and on Audioset, the largest sound events dataset, our sequential learning approach can lead to up to 9% improvement in performance. A comprehensive evaluation also shows that the method leads to improved transferability of knowledge from previously trained models, thereby leading to improved generalization capabilities on transfer learning tasks.

1. Introduction

Human interaction with the environment is driven by multi-sensory perception. Sounds and sound events, natural or otherwise, play a vital role in this first person interaction. To that end, it is imperative that we build acoustically intelligent devices and systems which can *recognize* and *understand* sounds. Although this aspect has been identified to an extent, and the field of Sound Event Recognition and detection (SER) is at least a couple of decades old (Xiong et al., 2003; Atrey et al., 2006), much of the progress has been in the last few years (Virtanen et al., 2018). Similar to related research domains in machine perception, like speech recognition, most of the early works in SER were fully supervised and driven by *strongly labeled data*. Here, audio recordings were carefully (and meticulously) annotated with time stamps of

sound events to produce exemplars. These exemplars then drive the training modules in supervised learning methods. Clearly, obtaining well annotated strongly labeled data is prohibitively expensive and cannot be scaled in practice. Hence, much of the recent progress on SER has focused on efficiently leveraging *weakly labeled* data (Kumar & Raj, 2016).

Weakly labeled audio recordings are only tagged with presence or absence of sounds (i.e., a binary label), and no temporal information about the event is provided. Although this has played a crucial role in scaling SER, large scale learning of sounds remains a challenging and open problem. This is mainly because, even in the presence of strong labels, large scale SER brings adverse learning conditions into the picture, either implicitly by design or explicitly because of the sheer number and variety of classes. This becomes more critical when we replace strong labels with weak labels. Tagging (*a.k.a.* weak labeling) very large number of sound categories in large number of recordings, often leads to considerable label noise in the training data. This is expected. Implicit noise via human annotation errors is clearly one of the primary factors contributing to this. *Audioset* (Gemmeke et al., 2017), currently the largest sound event dataset, suffers from this implicit label noise issue. Correcting for this implicit noise is naturally very expensive (one has to perform multiple re-labeling of the same dataset). Beyond this, there are more nuanced noise inducing attributes, which are outcomes of the number and variance of the classes themselves. For instance, real world sound events often overlap and as we increase the sound vocabulary and audio data size, the “mass” of overlapping audio in the training data can become large enough to start affecting the learning process. This is trickier to address in weakly labeled data where temporal locations of events are not available.

Lastly, when working with large real world datasets, one cannot avoid the noise in the inputs themselves. For SER these manifest either via signal corruption in the audio snippets themselves (i.e., acoustic noise), or signals from non-target sound events, both of which will interfere in the learning process. In weakly labeled setting, by definition, this noise level would be high, presenting harsher learning space for

¹Facebook Reality Labs, Redmond, USA. Correspondence to: Anurag Kumar <anuragkr@fb.com>.

networks. We need efficient SER methods that are sufficiently robust to the above three adverse learning conditions. In this work, we present an interesting take on large scale weakly supervised learning for sound events. Although we focus on SER in this work, we expect that the proposed framework is applicable for any supervised learning task.

The main idea behind our proposed framework is motivated by the attributes of human learning, and how humans adapt and learn when solving new tasks. An important characteristic of human’s ability to learn is that it is *not* a one-shot learning process, i.e., in general, we do not learn to solve a task in the first attempt. Our learning typically involves multiple stages of development where past experiences, and past failures or successes, “guide” the learning process at any given time. This idea of sequential learning in humans wherein each stage of learning is guided by previous stage(s) was referred to as *sequence of teaching selves* in (Minsky, 1994). Our proposal follows this meta principle of sequential learning, and at the core, it involves the concept of learning over time. Observe that this learning over time is rather different from, for instance, learning over iterations or epochs in stochastic gradients, and we make this distinction clear as we present our model. We also note that the notion of lifelong learning in humans, which has inspired *lifelong machine learning* (Silver et al., 2013; Parisi et al., 2019), is also, in principle, related to our framework.

Our proposed framework is called SeqUential Self TeAchING (SUSTAIN). We train a sequence of neural networks (designed for weakly labeled audio data) wherein the network at the current stage is guided by trained network(s) from the previous stage(s). The guidance from networks in previous stages comes in the form of “co-supervision”; i.e., the current stage network is trained using a *convex combination* of ground truth labels and the outputs from one or more networks from the previous stages. Clearly, this leads to a cascade of *teacher-student* networks. The student network trained in the current stage will become a teacher in the future stages. We note that this is also related to the recent work on knowledge distillation through teacher-student frameworks (Hinton et al., 2015; Ba & Caruana, 2014; Buciluă et al., 2006). However, unlike these, our aim is not to construct a smaller, compressed, model that emulates the performance of high capacity teacher. Instead, our SUSTAIN framework’s goal is to simply utilize the teacher’s knowledge better.

Specifically, the student network tries to correct the mistakes of the teachers, and this happens over multiple sequential stages of training and co-supervision with the aim of building better models as time progresses. We show that one can quantify the performance improvement, by explicitly controlling the transfer of knowledge from teacher to student over successive stages.

The **contributions** of this work include: **(a)** A sequential self-teaching framework based on co-supervision for improving learning over time, including few technical results characterizing the limits of this improved learnability; **(b)** A novel CNN for large scale weakly labeled SER, and **(c)** Extensive evaluations of the framework showing up to 9% performance improvement on Audioset, significantly outperforming existing procedures, and applicability to knowledge transfer.

The rest of the paper is organized as follows. We discuss some related work in Section 2. In Section 3, we introduce the sequential self-teaching framework and then discuss few technical results. In Section 4, we describe our novel CNN architecture for SER which learns from weakly labeled audio data. Sections 5 and 6 show our experimental results, and we conclude in Section 7.

2. Related Work

While earlier works on SER were primarily small scale (Couvreur et al., 1998), large scale SER has received considerable attention in the last few years. The possibility of learning from weakly labeled data (Kumar & Raj, 2016; Su et al., 2017) is the primary driver here, including availability of large scale weakly labeled datasets later on, like Audioset (Gemmeke et al., 2017). Several methods have been proposed for weakly labeled SER; (Kumar et al., 2018; Kong et al., 2019; Chou et al., 2018; McFee et al., 2018; Yu et al., 2018; Wang et al., 2018; Adavanne & Virtanen, 2017) to name a few. Most of these works employ deep convolutional neural networks (CNN). The inputs to CNNs are often time-frequency representations such as spectrograms, logmel spectrograms, constant-q spectrograms (Zhang et al., 2015; Kumar et al., 2018; Ye et al., 2015). Specifically, with respect to Audioset, some prior works, for example (Kong et al., 2019), have used features from a pre-trained network, trained on a massive amount of YouTube data (Hershey et al., 2017) for instance.

The weak label component of the learning process was earlier handled via *mean* or *max* global pooling (Su et al., 2017; Kumar et al., 2018). Recently, several authors proposed to use attention (Kong et al., 2019; Wang et al., 2018; Chen et al., 2018), recurrent neural networks (Adavanne & Virtanen, 2017), adaptive pooling (McFee et al., 2018). Some works have tried to understand adverse learning conditions in weakly supervised learning of sounds (Shah et al., 2018; Kumar et al., 2019), although it still is an open problem. Recently, problems related to learning from noisy labels have been included in the annual DCASE challenge on sound event classification (Fonseca et al., 2019b)¹.

Sequential learning, and more generally, learning over time,

¹<http://dcase.community/challenge2019/>

is being actively studied recently (Parisi et al., 2019), starting from the seminal work (Minsky, 1994). Building cascades of models has also been tied to lifelong learning (Silver et al., 2013; Ruvolo & Eaton, 2013). Further, several authors have looked at the teacher-student paradigm in a variety of contexts including knowledge distillation (Hinton et al., 2015; Furlanello et al., 2018; Chen et al., 2017; Mirzadeh et al., 2019), compression (Polino et al., 2018) and transfer learning (Yim et al., 2017; Weinshall et al., 2018). (Furlanello et al., 2018) in particular show that it is possible to sequentially distill knowledge from neural networks and improve performance. Our work builds on top of (Kumar & Ithapu, 2020), and proposes to learn a sequence of self-teacher(s) to improve generalizability in adverse learning conditions. This is done by co-supervising the network training along with available labels and controlling the knowledge transfer from the teacher to the student.

3. Sequential Self-Teaching (SUSTAIN)

3.1. SUSTAIN Framework

Notation : Let $\mathcal{D} := \{\mathbf{x}^s, \mathbf{y}^s\}$ ($s = 1, \dots, S$) denote the dataset we want to learn with S training pairs. \mathbf{x}^s are the inputs to the learning algorithms and $\mathbf{y}^s \in \{0, 1\}^C$ are the desired outputs. C is the number of classes. $y_c^s = 1$ indicates the presence of c^{th} class in the input \mathbf{x}^s . Note that $\mathbf{y}_c^s \forall c$ are the observed labels and may have noise.

For the rest of the paper, we restrict ourselves to the binary cross-entropy loss function. However, in general, the method is applicable to other loss functions as well, such as mean squared error loss.

If $\mathbf{p}^s = [p_1^s, \dots, p_C^s]$ is the predicted output, then the loss is

$$\mathcal{L}(\mathbf{p}^s, \mathbf{y}^s) = \frac{1}{C} \sum_{c=1}^C \ell(p_c^s, y_c^s) \quad \text{where} \quad (1)$$

$$\ell(p_c^s, y_c^s) = -y_c^s \log(p_c^s) - (1 - y_c^s) \log(1 - p_c^s) \quad (2)$$

With this notation, we will now formalize the ideas motivated in Section 1. The learning process entails T stages indexed by $t = 0, \dots, T$. The goal is to train a cascade of learning models denoted by $\mathcal{N}^0, \dots, \mathcal{N}^T$ at each stage. The final model of interest is \mathcal{N}^T . Zeroth stage serves as an *initialization* for this cascade. It is the default teacher that learns from the available labels \mathbf{y}^s . Once \mathcal{N}^0 is trained, we can get the predictions $\hat{\mathbf{p}}_0^s \forall s$ (note the $\hat{\cdot}$ here).

The learning in each of the later stages is co-supervised by the already trained network(s) from previous stages, i.e., at t^{th} stage, $\mathcal{N}^0, \dots, \mathcal{N}^{t-1}$ guide the training of \mathcal{N}^t . This guidance is done via replacing the original labels (\mathbf{y}^s) with a convex combination of the predictions from the teacher network(s) and \mathbf{y}^s , which will be the new targets for training

Algorithm 1 SUSTAIN: Single Teacher Per Stage

Input : \mathcal{D} , #stages T , $\{\alpha_t, t = 0, \dots, T - 1\}$
Output : Trained Network \mathcal{N}^T after T stages
 1: Train default teacher \mathcal{N}^0 using $\mathcal{D} := \{\mathbf{x}^s, \mathbf{y}^s\} \forall s$
 2: **for** $t = 1, \dots, T$ **do**
 3: Compute new target $\bar{\mathbf{y}}_t^s (\forall s)$ using Eq. 4
 4: Train \mathcal{N}^t using new target $\mathcal{D} := \{\mathbf{x}^s, \bar{\mathbf{y}}_t^s\} \forall s$
 5: **end for**
 6: Return \mathcal{N}^T

\mathcal{N}^t . In the most general case, if all networks from previous stages are used for teaching, the new target at t^{th} stage is,

$$\bar{\mathbf{y}}_t^s = \alpha_0 \mathbf{y}^s + \sum_{\tilde{t}=1}^t \alpha_{\tilde{t}} \hat{\mathbf{p}}_{\tilde{t}-1}^s \quad s.t. \quad \sum_{\tilde{t}=0}^t \alpha_{\tilde{t}} = 1 \quad (3)$$

More practically, the network from only last stage will be used, in which case,

$$\bar{\mathbf{y}}_t^s = \alpha_0 \mathbf{y}^s + (1 - \alpha_0) \hat{\mathbf{p}}_{t-1}^s \quad (4)$$

or the students from previous m stages will co-supervise the learning at stage t , which will lead to $\bar{\mathbf{y}}_t^s = \alpha_0 \mathbf{y}^s + \sum_{\tilde{t}=1}^m \alpha_{\tilde{t}} \hat{\mathbf{p}}_{t-\tilde{t}}^s$, $s.t. \sum_{\tilde{t}=0}^m \alpha_{\tilde{t}} = 1$.

Algorithm 1 summarizes this self teaching approach driven by co-supervision with single teacher per stage. It is easy to extend it to m teachers per stage, driven by appropriately chosen α^s 's.

3.2. Analyzing SUSTAIN w.r.t to label noise

In this section, we provide some insights into our SUSTAIN method with respect to label noise, a common problem in large scale learning of sound events. $\mathbf{y}_c^s \forall c$ denote our noisy observed labels. Let y_c^{*s} be the corresponding true label parameterized as follows,

$$y_c^s = \begin{cases} y_c^{*s} & w.p. \delta_c \\ 1 - y_c^{*s} & else \end{cases} \quad (5)$$

Within the context of learning sounds, in the simplest case, δ_c characterizes the per-class noise in labeling process. Nevertheless, depending on the nature of the labels themselves, it may represent something more general like sensor noise, overlapping speakers and sounds etc.

To analyze our approach and to derive some technical guarantees on performance, we assume a trained default teacher \mathcal{N}^0 and a new student to be learned (i.e., $T = 1$). The new training targets in this case are given by

$$\bar{\mathbf{y}}_1^s = \alpha_0 \mathbf{y}^s + (1 - \alpha_0) \hat{\mathbf{p}}_0^s \quad (6)$$

Recall from Eq. 5 that δ_c parameterizes the error in \mathbf{y}^s vs. the unknown truth \mathbf{y}^{*s} . Similarly, we define $\bar{\delta}_c$ to parameter-

ize the error in $\hat{\mathbf{p}}^s$ vs. \mathbf{y}^{*s} i.e., noise in teacher's predictions w.r.t the true unobserved labels.

$$\hat{p}_{0,c}^s = \begin{cases} y_c^{*s} & w.p. \bar{\delta}_c \\ 1 - y_c^{*s} & else \end{cases} \quad (7)$$

The interplay between δ_c and $\bar{\delta}_c$ in tandem with the performance accuracy of \mathcal{N}^0 will help us evaluate the gain in performance for \mathcal{N}^1 versus \mathcal{N}^0 . To theoretically assess this performance gain, we consider the case of uniform noise $\forall c$ followed by a commentary on class-dependent noise. Further, we explicitly focus the technical results on high noise setting and revisit the low-to-medium noise setup in evaluations in Section 5.

3.2.1. UNIFORM NOISE: $\delta_c = \delta \forall c$

This is the simpler setting where the apriori noise in classes is uniform across all categories with $\delta_c = \delta \forall c$. We have the following result.

Proposition 1. *Let \mathcal{N}^1 be trained using $\{\mathbf{x}^s, \bar{\mathbf{y}}^s\} \forall s$ using binary cross-entropy loss, and let ϵ_c denote the average accuracy of \mathcal{N}^0 for class c . Then, we have*

$$\bar{\delta}_c = \epsilon_c \delta + (1 - \epsilon_c)(1 - \delta) \forall c \quad (8)$$

and whenever $\delta < \frac{1}{2}$, \mathcal{N}^1 improves performance over \mathcal{N}^0 . The per class performance gain is $(1 - \epsilon_c)(1 - 2\delta)$

Proof. Recall the entropy loss from Eq. 1, for a given s and c . Using the definition of the new label from Eq. 6, we get the following

$$\ell(p_c^s, \bar{y}_c^s) = \alpha_0 \ell(p_c^s, y_c^s) + (1 - \alpha_0) \ell(p_c^s, \hat{p}_c^s) \quad (9)$$

Now, Eq. 5 says that w.p. δ (recall $\delta_c = \delta \forall c$ here), $\ell(p_c^s, y_c^s) = \ell(p_c^s, y_c^{*s})$, else $\ell(p_c^s, y_c^s) = \ell(p_c^s, 1 - y_c^{*s})$. Hence, using Eq. 5 and Eq. 7, and using the resulting equations in Eq. 9 we have the following

$$\begin{aligned} \mathbb{E}_s \ell(p_c^s, y_c^s) &= \delta \sum_{s=1}^S \ell(p_c^s, y_c^{*s}) + (1 - \delta) \sum_{s=1}^S \ell(p_c^s, 1 - y_c^{*s}) \\ \mathbb{E}_s \ell(p_c^s, \hat{p}_c^s) &= \bar{\delta}_c \sum_{s=1}^S \ell(p_c^s, y_c^{*s}) + (1 - \bar{\delta}_c) \sum_{s=1}^S \ell(p_c^s, 1 - y_c^{*s}) \\ \mathbb{E}_s \ell(p_c^s, \bar{y}_c^s) &= (\alpha_0 \delta + (1 - \alpha_0) \bar{\delta}_c) \sum_{s=1}^S \ell(p_c^s, y_c^{*s}) \\ &\quad + (\alpha_0(1 - \delta) + (1 - \alpha_0)(1 - \bar{\delta}_c)) \sum_{s=1}^S \ell(p_c^s, 1 - y_c^{*s}) \end{aligned}$$

If $(\alpha_0 \delta + (1 - \alpha_0) \bar{\delta}_c) > \delta$ then we can ensure that using \bar{y}_c^s as targets is better than using y_c^s . Now given the accuracy of \mathcal{N}^0 denoted by $\epsilon_c \forall c$, combining Eq. 5 and Eq. 7, we

can see that $\bar{\delta}_c = \epsilon_c \delta + (1 - \epsilon_c)(1 - \delta)$. Using this, for \mathcal{N}^1 to be better than \mathcal{N}^0 , we need

$$\alpha_0 \delta + (1 - \alpha_0)(\epsilon_c \delta + (1 - \epsilon_c)(1 - \delta)) > \delta \quad (10)$$

which requires $\delta < \frac{1}{2}$. And the gain is simply $\alpha_0 \delta + (1 - \alpha_0) \bar{\delta}_c - \delta$ which reduces to $(1 - \epsilon_c)(1 - 2\delta)$. \square

3.2.2. REMARKS

The above proposition is fairly intuitive and summarizes a core aspect of the proposed framework. Observe that, Proposition 1 is rather conservative in the sense that we are claiming \mathcal{N}^1 is better than \mathcal{N}^0 only if Eq. 10 holds for all classes, i.e., performance improves for all classes. This may be relaxed, and we may care more about some specific classes. We discuss this below, for the high and low noise scenarios separately.

High noise $\delta < \frac{1}{2}$: The given labels y_c^s are wrong more than half of the time, and with such high noise, we expect \mathcal{N}^0 to have high error i.e., $\hat{p}_{0,c}^s$ and y_c^s do not match. Putting these together, as Proposition 1 suggests, the probability that $\hat{p}_{0,c}^s$ matches the truth y_c^{*s} is implicitly large, leading to $\bar{\delta}_c > \delta$. Note that we cannot just flip **all** predictions i.e., $p_{0,c}^s = 1 - y_c^s$ would be infeasible, and there is some trade-off between \mathcal{N}^0 's predictions and given labels. Thereby, the choice of α_0 then becomes critical (which we discuss further in Section 3.2.3). Beyond this interpretation, we show extensive results in Section 5 supporting this.

Low-to-medium noise $\delta > \frac{1}{2}$: When $\delta \gg \frac{1}{2}$, \mathcal{N}^0 is expected to perform well, and $\hat{p}_{0,c}^s$ matches y_c^s , which in turn matches y_c^{*s} since the noise is low. Hence, \mathcal{N}^1 's role of combining \mathcal{N}^0 's output with y_c^s becomes rather moot, because on average, for most cases, they are same. For medium noise settings with $1 \gg \delta > \frac{1}{2}$, proposition 1 does not infer anything specific. Nevertheless, via extensive set of experiments, we show in section 5 that \mathcal{N}^1 still improves over \mathcal{N}^0 in some cases.

Class-Specific Noise: $\delta_c \neq \delta \forall c$ It is reasonable to assume that in practice there are specific classes of interest that we desire to be more accurately predictable than others, including the fact that annotation is more carefully done for such classes. One can generalize Proposition 1 for this class-dependent δ_c s, by putting some reasonable lower bound on loss of accuracy for undesired classes ϵ_c s. We leave such technical details to a follow-up work, and now address the issue of choosing α s for learning.

3.2.3. INTERPLAY OF α_t AND T

Recall that the main hyperparameters of SUSTAIN are the weights $\alpha_0, \dots, \alpha_T$ and T , and the main unknowns are the

noise levels in the dataset (δ_c). We now suggest that Algorithm 1 is implicitly robust to these unknowns and provides an empirical strategy to choose the hyperparameters as well. We have the following result focusing on a given class c . \bar{T}_c and \bar{T} denote the optimal number of stages per class c and across all classes respectively. The proof is in supplement.

Corollary 1. *Let ϵ_c^t denote the accuracy of \mathcal{N}^t for class c . Given some δ , there exists an optimal \bar{T}^c such that $\epsilon_c^{\bar{T}^c} \geq \epsilon_c^t$.*

Remarks. The main observation here is that \bar{T}_c might be very different for each c , and it may be possible that $\bar{T}_c = 0$ in certain cases, i.e., the teacher is already better than any student. In principle, there may exist an optimal \bar{T} that is class independent for the given dataset, but it is rather hard to comment about its behaviour in general without explicitly accounting for the individual class-specific accuracies ϵ_c s. This is simply because correcting for noisy labels in one class may have the outcome of corrupting another class. Lastly, it should be apparent that the gain in performance per class c has diminishing returns as T increases.

Choosing \bar{T} and α_t s: Corollary 1 is an existence result and does not give us a procedure to compute \bar{T}_c s (or \bar{T}) and the corresponding α s. In practice, there is a simple strategy one can follow. At stage 0, we train \mathcal{N}^0 and we record its average across-class performance $\epsilon^0 = \frac{1}{C} \sum_{c=1}^C \epsilon_c^0$. At stage 1, we empirically select the best α_0 that results in maximal ϵ^1 . If $\epsilon^1 \leq \epsilon^0$, then we stop and declare $\bar{T} = 0$ i.e., no student needed. On the other hand, if $\epsilon^1 \geq \epsilon^0$, then we continue to stage $t = 2$. And repeat this process until the accuracy ϵ^t saturates or starts to decrease. This averages out the per-class influence on \bar{T} .

4. CNN for Weakly Labeled SER

We now evaluate Algorithm 1 for weakly labeled SER, as motivated in Section 1. We first propose a novel architecture for the problem and then study SUSTAIN using this network. Observe that most of the existing approaches to SER are variants of Multiple Instance Learning (MIL) (Dietterich et al., 1997), the first proposed framework being (Kumar & Raj, 2016).

Our key novelty is to include a class-specific ‘‘attention’’ learning mechanism within the MIL framework. We introduce some brief notation followed by presenting the model. In MIL, the training data \mathcal{D} is made available via *Bags* \mathcal{B}_i , with each bag corresponding to a collection of m_i training instances $\{\mathbf{x}_i^1, \dots, \mathbf{x}_i^{m_i}\}$. Each \mathcal{B}_i has one label vector \mathbf{z}_i . $\mathbf{z}_{i,c} = 1$ for class c if at least one of the m_i instances is positive, otherwise $\mathbf{z}_{i,c} = 0$.

The key idea in MIL is that the learner first predicts on instances, and then maps (accumulates) these instance-level predictions to a bag-level prediction. For instance, a widely

Stage	Layers	Output Size
Input	Unless specified – (S)tride = 1, (P)adding = 1	$1 \times 1024 \times 64$
Block B1	Conv: 64, 3×3 Conv: 64, 3×3 Pool: 4×4 (S:4)	$64 \times 1024 \times 64$ $64 \times 1024 \times 64$ $64 \times 256 \times 16$
Block B2	Conv: 128, 3×3 Conv: 128, 3×3 Pool: 2×2 (S:2)	$128 \times 256 \times 16$ $128 \times 256 \times 16$ $128 \times 128 \times 8$
Block B3	Conv: 256, 3×3 Conv: 256, 3×3 Pool: 2×2 (S:2)	$256 \times 128 \times 8$ $256 \times 128 \times 8$ $256 \times 64 \times 4$
Block B4	Conv: 512, 3×3 Conv: 512, 3×3 Pool: 2×2 (S:2)	$512 \times 64 \times 4$ $512 \times 64 \times 4$ $512 \times 32 \times 2$
Block B5	Conv: 2048, 3×2 (P:0)	$2048 \times 30 \times 1$
Block B6	Conv: 1024, 1×1	$1024 \times 30 \times 1$
Block B7	Conv: 1024, 1×1	$1024 \times 30 \times 1$
Block B8	Conv: C, 1×1	$C \times 30 \times 1$
$g()$	$\mathbf{W}_\Phi (C \times C)$	$C \times 1$

Table 1. WEANET: All convolutional layers (except B8) are followed by batch norm and ReLU; Sigmoid activation follows B8.

used SVM based MIL (Andrews et al., 2002) uses this principle, using *max* operator as the mapping function. Based on similar principle, we formulate the learning process as follows:

$$\sum_{i=1}^N \ell(g_\Phi(f_\Theta(\mathbf{x}_i^1), \dots, f_\Theta(\mathbf{x}_i^{m_i})), \mathbf{z}_i) \quad (11)$$

$f(\cdot)$, parameterized by Θ , is the learner and does the instance level prediction of outputs, and $g(\cdot)$ maps these $f_\Theta(\mathbf{x}_i^s)$ to bag level predictions. For weakly labeled SER, \mathcal{B}_i s are full audio recordings and instances are short duration segments of the recordings. We design a CNN which takes in Log-scaled Melfilter-bank feature representations of the entire audio recording, produces instance (i.e., segment) level predictions which are then mapped to recording level predictions.

The inputs are computed as follows: 64 Mel-filter-bank is obtained for each 16ms window of audio, and the window moves by 10ms, leading to 100 Logmel frames per second of audio (with a sampling rate of 16KHz for the audio recordings).

The proposed architecture, referred to as WEakly labeled Attention NETwork (WEANET), is shown in Table 1. Example output sizes for an input with 1024 Logmel frames (approx. 10 seconds long audio) is shown in Table 1. The first few convolutional layers (B1 to B5) produce 2048 dimensional bag representations for the input at Block B5. B6-B8 are 1×1 convolutional layers that produce instance level predictions of size $C \times K$, C is number of classes and K is number of segments obtained for a given input. The network is designed such that the receptive field of each segment (i.e., instance) is ~ 1 second (96 frames), and the

segments themselves move by 0.33 seconds (32 frames). The instance level predictions are then used to produce bag (i.e., recording) level predictions using $g(\cdot)$. An easy parameter free way of doing this is to use *mean* (or *max*) functions which will simply take average (or maximum) over segment level predictions from B8.

Instead, we propose an attention mechanism here which aims to appropriately weigh each segment’s contribution in the final recording level prediction. Moreover, this is done in a class-specific manner as different sounds might be located at different places in the recording. More formally, $g(\cdot)$ is parameterized as follows:

$$\mathbf{A} = \tilde{\sigma}(\mathbf{W}_\Phi \mathbf{S}) \quad (12)$$

$$\mathbf{o} = \sum_{k=1}^K \tilde{\mathbf{O}}_k \quad \text{s.t.} \quad \tilde{\mathbf{O}} = \mathbf{A} \odot \mathbf{S} \quad (13)$$

where $\mathbf{W}_\Phi \in \mathbb{R}^{C \times C}$. $\mathbf{S} \in \mathbb{R}^{C \times K}$ denotes the segment level predictions. $\tilde{\sigma}$ is the softmax function applied across segments, and $\mathbf{A} = \tilde{\sigma}(\mathbf{W}_\Phi \mathbf{S})$ gives us the attention weights for each segment and class. \odot is element wise multiplication and $\tilde{\mathbf{O}}_k$ is k^{th} column of $\tilde{\mathbf{O}}$, which represents the *weighted* predictions for each class in k^{th} segment. All these are then pooled into \mathbf{o} , which represents the recording level prediction for the input. \mathbf{W}_Φ is learned along with rest of the parameters of the WEANET. Note that, the size of attention parameter \mathbf{W}_Φ is independent of the number of segments obtained for an input or in other words the duration of the input. It depends on the number of classes in the dataset.

5. Experiments and Results

5.1. Datasets and Experimental Setup

Audioset: (Gemmeke et al., 2017) is very challenging dataset in terms of adverse learning conditions outlined in Section 1. It is the largest dataset for sound events with weakly labeled YouTube clips for 527 sound classes. Each recording is ~ 10 seconds long and on an average, there are 2.7 labels per recording. The training and evaluation sets consist of ~ 2 million and $\sim 20,000$ recordings respectively. The dataset is highly unbalanced with the number of training examples varying from close to 1 million for classes such as *Music* and *Speech* to < 100 for classes such as *Screech* and *Toothbrush*. The evaluation set has at least 59 examples for each class. A sample of $\sim 25,000$ videos from the training set are sampled out for validation.

An analysis of label noise was done by the authors by sampling 10 examples for each class and sending them for expert label reviewing. This puts label noise at broad range of 0 to 80-90% across classes. Note however that this is an extremely rough estimate for a dataset of this size.

FSDKaggle: (Fonseca et al., 2019b) is a dataset of 80

sound events. It has 2 training sets: a *Curated* set with 4970 recordings and a *Noisy* set with 19,815 audio recordings. The *Curated* set is a clean training set which has been carefully annotated by humans to ensure minimal to no label noise. The *Noisy* training set is obtained from Flickr videos and not labeled by humans. They contain considerable amount of label noise. The evaluation set has 3361 recordings. We use the *Public* test set with 1120 recordings for validation.

(Fonseca et al., 2019b) does a more thorough examination of label noise. The estimated per-class label noise roughly ranges from 20% to 80% and overall around 60% of the labels show some type of label noise. While the *Curated*, validation and test sets are sourced from freesounds.org (and then labeled by humans), the *Noisy* training set recordings are sourced from Flickr. This heavy mismatch in domain adds on to the already difficult learning conditions for the *Noisy* training set and leads to considerable impact on performance.

ESC-50: (Piczak, 2015) This dataset consists of 2000 recordings from 50 sound classes. Each sound class has 40 audio recordings and all recordings are 5 seconds long. We use this dataset primarily in our transfer learning experiments in Section 6. It comes with 5 pre-defined sets and we follow the same setup in our experiments as in prior works such as (Kumar et al., 2018).

Experimental Setup: All of our experiments uses Pytorch (Paszke et al., 2017) for neural network implementations. Adam optimizer is used, and networks are trained for 20 epochs. Minibatch size is set to 144. Hyperparameters such as learning rates and the best model during training is selected using the validation set. The attention weight parameter \mathbf{W}_Φ is initialized with 0’s such that the initial attention weights come out to be equal for all segments for all classes, $\mathbf{A}_{ck} = 1/K, \forall c, k$. The updates for attention weight parameter is turned on from fifth epoch. For Audioset, given its highly unbalanced nature, we use a weighted loss for each class. This weight for class c is given by $w_c = 1 + \log_2(\gamma_c)$, where γ_c is the inverse of the class prior in the training set. The training set up is consistent across all stages of SUSTAIN and only the teacher(s) and the parameter α changes.

Similar to prior works on Audioset, Average Precision (AP) and Area under ROC curves (AUC) are used to measure performance. Mean AP (mAP) and mean AUC over all classes are used as overall metrics for performance assessment. For FSDkaggle dataset, the metric used is a label-weighted label-ranking average precision (lwlrap) (Fonseca et al., 2019b). Given the smaller size of this dataset, we use a lighter version of WEANET. The details of this lighter WEANET are provided in the supplementary material. For ESC-50 dataset accuracy is used as the performance metric.

Method	mAP	mAUC
(Kong et al., 2019)-1	0.361	0.969
(Wang et al., 2019)-1	0.354	0.963
(Kong et al., 2019)-2	0.369	0.969
(Wang et al., 2019)-2	0.362	0.965
WEANET ($g(\cdot) = avg(\cdot)$)	0.352	0.970
WEANET ($g(\cdot; \mathbf{W}_\Phi)$)	0.366	0.958

Table 2. Comparison of WEANET with other attention architectures on Audioset dataset.

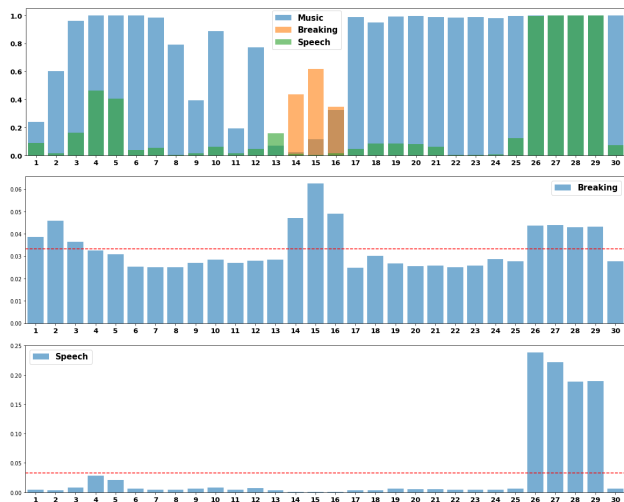


Figure 1. An example of WEANET outputs on a recording from test set. **Top:** Segment level probability outputs for three classes present in the recording. **Mid:** Segment wise attention weights (\mathbf{A}) for the Breaking Sound. **Bottom:** Segment wise attention weights (\mathbf{A}) for the Speech Sound. Red line denotes if all segments were given equal weights (1/30).

5.2. WEANET Model

We first provide some results on WEANET. Table 2 shows performance of WEANET framework and compares it with respect to some other attention frameworks for weakly labeled SER. Note that, (Kong et al., 2019) uses embeddings for audio recordings from a network trained on a very large database (YouTube-70M) (Hershey et al., 2017). These pre-trained representations lead to enhanced performance on Audioset. We (and also (Wang et al., 2018)) work with the actual audio recordings and use logmel feature representations. In summary, WEANET performs better than other attention frameworks. (Kong et al., 2019)-2 performs slightly better but uses pre-trained embeddings as just mentioned. WEANET with class-specific attention is 4% better than WEANET with $g(\cdot)$ as simple average pooling.

The major advantage of having class-specific attention learning is for localization of events. Figure 1 shows segment level outputs for 3 sounds present in a specific recording

Method	mAP	mAUC
(Kong et al., 2019) - Small	0.361	0.969
(Kong et al., 2019) - Large	0.369	0.969
(Wang et al., 2019) - TALNet (exp. pooling)	0.362	0.965
(Wang et al., 2019) - TALNet (Attention)	0.354	0.963
(Ford et al., 2019) - ResNet-34 (Attention)	0.360	0.966
(Ford et al., 2019) - ResNet-101 (Attention)	0.380	0.970
WEANET	0.366	0.958
SUSTAIN - Single Teacher	0.394	0.972
SUSTAIN - 2 Teachers	0.398	0.972

Table 3. Comparison with state-of-the-art methods on Audioset

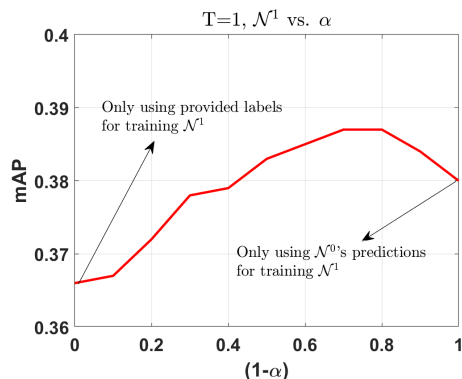


Figure 2. Single teacher: \mathcal{N}^1 vs. α_0

from the test set. Note that we matched with the location of the events in the actual recording. The lower two figures show attention weights for the two events (*Breaking* sound and *Speech* sound) that are highly localized in the recording. Observe that the weights are much higher than average for segments where the event is actually located. For speech in particular, segments 1 – 6 show high probability of presence even though actually speech is not present. However, the class-specific attention framework is capable of flagging this false positive and assigns very low weights to them.

5.3. SUSTAIN Framework

Comparison with state-of-the-art: Table 3 compares performance of our SUSTAIN framework with state-of-the-art methods on Audioset. “SUSTAIN-Single Teacher” uses 1 teacher at each stage, specifically the network trained in the previous stage (as in Eq. 4). “SUSTAIN-2 Teachers” uses 2 networks learned in the previous 2 consecutive stages as teachers. SUSTAIN learning leads to superior performance over all prior methods. The ResNet based architectures in Table 3 are much larger compared to our WEANET and have several times more parameters. Our method outperforms the previous best method (ResNet-50) by 4.7%. Note that, (Ford et al., 2019) also reports a performance of 0.392 but that is obtained through ensemble of models, by averaging outputs of multiple models.

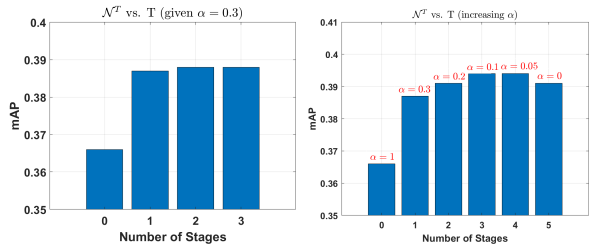


Figure 3. Performance of students as T increases: (a) $\alpha_0 = 0.3$, (b) decreasing α_0 as T increases.

Stage(T)	Teach.	α_1, α_2	α_0	Stud.	Stud. Perf.
0	-	-	1.0	\mathcal{N}^0	0.366
1	$\mathcal{N}^0, -$	0.7, -	0.3	\mathcal{N}^1	0.387
2	$\mathcal{N}^0, \mathcal{N}^1$	0.3, 0.5	0.2	\mathcal{N}^2	0.393
3	$\mathcal{N}^1, \mathcal{N}^2$	0.4, 0.5	0.1	\mathcal{N}^3	0.396
4	$\mathcal{N}^2, \mathcal{N}^3$	0.45, 0.5	0.05	\mathcal{N}^4	0.398
5	$\mathcal{N}^3, \mathcal{N}^4$	0.45, 0.53	0.03	\mathcal{N}^5	0.398

Table 4. 2 teachers at each stage (weights: α_1 and α_2).

Focusing primarily on the SUSTAIN learning, we notice that it can lead up to 8 – 9% improvement in results for the WEANET model. Thus, the same architecture WEANET, generalizes much better after a few stages of SUSTAIN learning as opposed to just training it on the available labels.

Single Stage ($T = 1$) vs. varying α_0 : Figure 2 shows that α_0 influences mAP, as suggested by in Section 3.2.2. The two extremes of $\alpha_0 = 0$ (only using \mathcal{N}^0 's predicted labels) and $\alpha_0 = 1$ (only using provided labels \mathbf{y}^s for learning) perform worse than learning using a combination of the two. This asserts our primary claim in Proposition 1. Depending on the weight $(1 - \alpha_0)$ given to the teacher, even a single stage of SUSTAIN can lead to up to 5.7% improvement in performance.

Multiple Stages ($T > 1$) vs. α_0 : For a fixed α_0 , Figure 3(a) shows that as T increases mAP starts to increase and then quickly saturates, showing evidence for Corollary 1. Note that this setup corresponds to using the last trained network as teacher (i.e., \mathcal{N}^3 uses \mathcal{N}^2 as the teacher). It is reasonable to expect that \mathcal{N}^t is better than \mathcal{N}^{t-1} , and so, one can put more confidence in the predicted labels of latest teachers than teachers from earlier stages. This is validated in Figure 3(b) where \mathcal{N}^0 uses only predicted labels ($\alpha_0 = 1$) and as T increases, we reduce α_0 , putting more confidence on teacher's predictions and achieve better mAP. Unlike the fixed alpha case, considerable improvement is obtained from stage 1 to 2 and then stage 2 to 3 by increasing the weight given to teacher's predictions.

Multiple Stages and Multiple Teachers: Table 4 shows the performance as T increases with 2 teachers. Each row

Method	lwrap	Remarks
(Fonseca et al., 2019b)	0.312	-
WEANET ^L - (T = 0)	0.436	-
SUSTAIN - (T = 1)	0.454	$\alpha = 0.5$
SUSTAIN - (T = 2)	0.456	$\alpha = 0.3$
SUSTAIN - (T = 3)	0.462	$\alpha = 0.2$
SUSTAIN - (T = 4)	0.470	$\alpha = 0.15$
SUSTAIN - (T = 5)	0.472	$\alpha = 0.05$
SUSTAIN - (T = 6)	0.472	$\alpha = 0.05$

Table 5. Performance when trained on FSDKaggle-Noisy set. First row is baseline. Last column shows α for each stage. Single Teacher (\mathcal{N}^{T-1}) at each stage of training.

corresponds to one stage and \mathcal{N}^0 is as usual the default teacher. As expected, the mAP improves as T increases. In particular, observe that after 5 stages we reach 0.398 mAP here, compared to the single teacher setup where we get 0.392 after 5 stages (refer to Figure 3(b)), a 1.5% relative improvement. We also see the saturation of performance as T increases, further supporting our results from Section 3.2.

5.4. Noisy label vs Clean Label Conditions:

We now try to specifically look into clean and noisy label learning conditions using FSDKaggle2019 dataset. The *Noisy* and *Curated* training sets of this dataset (refer to their descriptions from Section 5.1) are used as noisy (i.e., hard) and clean (i.e., easy) learning conditions. The test set remains same for the two cases. We use a lighter version of WEANET model (WEANET^L) for these experiments, the details of which are available in the supplementary material. For all these experiments, only one teacher is used per stage; the network trained in the previous stage. Most prior works on FSDKaggle have relied heavily on different forms of data augmentation on the *Curated* set for improved performance. We do not do any data augmentation and instead focus on easy and hard conditions. We use the performance reported by the dataset paper, (Fonseca et al., 2019a), on each training set as the baseline.

Table 5 summarizes the results for *Noisy* training set. We observed that for the *Noisy* training set, the trends of results (for different parameters such as α s, stages etc.) are similar to those observed for Audioset. Overall SUSTAIN leads to around 8% improvement. Even just 1 stage of SUSTAIN, leads to almost 4.1% improvement in performance over base WEANET^L (T = 0). The performance saturates after 5 stages of SUSTAIN.

The *Curated* set presents a different picture though. While one stage of SUSTAIN still leads to small improvement, any further co-supervision leads to deterioration in performance. This is the expected behavior our technical results claimed in Section 3.2 i.e., SUSTAIN learning primarily helps in

Method	lwrap	Remarks
Baseline - (Fonseca et al., 2019b)	0.542	-
WEANET ^L - (T = 0)	0.619	-
SUSTAIN - (T = 1)	0.619	$\alpha = 0.9$
SUSTAIN - (T = 1)	0.625	$\alpha = 0.7$
SUSTAIN - (T = 1)	0.632	$\alpha = 0.5$
SUSTAIN - (T = 1)	0.622	$\alpha = 0.3$
SUSTAIN - (T = 1)	0.622	$\alpha = 0.1$
SUSTAIN - (T = 2)	0.624	$\alpha = \{0.5, 0.9\}$
SUSTAIN - (T = 2)	0.623	$\alpha = \{0.5, 0.7\}$
SUSTAIN - (T = 2)	0.627	$\alpha = \{0.5, 0.5\}$
SUSTAIN - (T = 2)	0.624	$\alpha = \{0.5, 0.3\}$
SUSTAIN - (T = 2)	0.625	$\alpha = \{0.5, 0.1\}$

Table 6. Performance when trained on FSDKaggle-Curated set. Last row last column shows α for each stage.

adverse learning conditions.

5.5. Class-specific Performance Gains

On the Audioset dataset, we observed that for almost 85% of all the classes (527 total), the performance improved with SUSTAIN learned model \mathcal{N}^4 (from Table 4), compared to base WEANET model (\mathcal{N}^0 from Table 4). Most classes have under 25% relative improvement, and 69 of the classes get $> 25\%$ improvement, and this reaches up to 100% for classes like *Squeal* and *Rattle*. Maximum drop in performance (down by 30%) is observed for *Gurgling* class. We also see that low performing classes ($AP < 0.1$) have more improvements in relative sense. On average, AP of these classes (44 of them) improve by 23%, while classes with high AP (> 0.5 , 146 in number), we see 6% gain in performance. Class-specific performance plots are shown in supplementary material. Overall, *Bagpipes* sounds are easiest to recognize and we achieve an AP of 0.931 for it. *Squish* on the other hand is hardest to recognize with an AP of 0.02.

6. Knowledge Transfer using SUSTAIN

In the preceding sections, we showed that the generalizability of a model can be improved through the proposed SUSTAIN learning. We now ask, are the models obtained from SUSTAIN learning more suitable for transfer learning? We study whether WEANET obtained after T stages of training (\mathcal{N}^T) is more suited for transfer learning compared to the one just trained on the available labels. Since SUSTAIN is not explicitly designed to handle this, the transfer learning question reveals the learnability power of the proposed framework.

We pick \mathcal{N}^4 and \mathcal{N}^0 WEANET models from Table 4 for this analysis, with \mathcal{N}^0 being the base model trained only on available labels and \mathcal{N}^4 being a SUSTAIN trained model. These WEANET models trained on Audioset are used to

ESC-50		FSDKaggle	
Method	Acc. (%)	Method	lwrap
(Sailor et al., 2017)	86.5	Noisy, WEANET (\mathcal{N}^0)	0.486
(Guzhov et al., 2020)	91.5	Noisy, WEANET (\mathcal{N}^4)	0.503
WEANET (\mathcal{N}^0)	92.6	Curated, WEANET (\mathcal{N}^0)	0.712
WEANET (\mathcal{N}^4)	94.1	Curated, WEANET (\mathcal{N}^4)	0.728

Table 7. Transfer Learning from SUSTAIN Models trained on Audioset. Results on ESC-50 and FSDKaggle dataset.

obtain representations for the audio recordings in the given target tasks. Outputs after Block B5 (refer to WEANET model from Table 2) are used as feature representations for the audio recordings. Recall that, Block B5 produces representations for 1 second long audio every 0.33 sec. These segment level representations are simply max-pooled across all segments to get a fixed 2048-dimensional vector for all audio recordings.

We study these transfer learning tasks on FSDKaggle and ESC-50 datasets. A simple linear classifier is trained on the feature representations obtained for the audio recordings.

Table 7 shows the results for these transfer learning tasks. We see that the representations from SUSTAIN framework leads to significantly improved feature learning for all datasets. For the clean conditions (ESC-50 and FSDKaggle-Curated), we see 1.5-2.2% improvement whereas for the noisy learning conditions we see up to 3.5% improvement in performance. For the ESC-50 dataset, this transfer learning also outperforms previous state-of-the-art results by a considerable margin (2.8% relative).

7. Conclusions

Designing robust learning models for weakly labelled datasets while also scaling them to large scale and ensuring good generalization is a hard problem, and is an open question. We addressed this problem in this paper. We proposed a sequential self-teaching framework that utilizes co-supervision across trained models to improve generalization. We specifically show promising results on sound event recognition and detection, in particular in large scale weakly labelled settings. We also proposed a novel architecture for learning sounds which incorporates class-specific attention learning. A better theoretical understanding of the role α and T play in different adverse learning conditions can lead to an enhanced understanding of SUSTAIN. We will explore these directions in future works.

References

Adavanne, S. and Virtanen, T. Sound event detection using weakly labeled dataset with stacked convolutional and recurrent neural network. *arXiv preprint arXiv:1710.02998*, 2017.

- Andrews, S., Tsochantaridis, I., and Hofmann, T. Support vector machines for multiple-instance learning. In *Advances in neural information processing systems*, pp. 561–568, 2002.
- Atrey, P. K., Maddage, N. C., and Kankanhalli, M. S. Audio based event detection for multimedia surveillance. In *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, volume 5. IEEE, 2006.
- Ba, J. and Caruana, R. Do deep nets really need to be deep? In *Advances in neural information processing systems*, pp. 2654–2662, 2014.
- Buciluă, C., Caruana, R., and Niculescu-Mizil, A. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 535–541. ACM, 2006.
- Chen, G., Choi, W., Yu, X., Han, T., and Chandraker, M. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems*, pp. 742–751, 2017.
- Chen, S., Chen, J., Jin, Q., and Hauptmann, A. Class-aware self-attention for audio event recognition. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, pp. 28–36, 2018.
- Chou, S.-Y., Jang, J.-S. R., and Yang, Y.-H. Learning to recognize transient sound events using attentional supervision. In *IJCAI*, pp. 3336–3342, 2018.
- Couvreur, C., Fontaine, V., Gaunard, P., and Mubikangiey, C. G. Automatic classification of environmental noise events by hidden markov models. *Applied Acoustics*, 54(3):187–206, 1998.
- Dietterich, T. G., Lathrop, R. H., and Lozano-Pérez, T. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71, 1997.
- Fonseca, E., Plakal, M., Ellis, D. P., Font, F., Favory, X., and Serra, X. Learning sound event classifiers from web audio with noisy labels. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 21–25. IEEE, 2019a.
- Fonseca, E., Plakal, M., Font, F., Ellis, D. P., and Serra, X. Audio tagging with noisy labels and minimal supervision. *arXiv preprint arXiv:1906.02975*, 2019b.
- Ford, L., Tang, H., Grondin, F., and Glass, J. A deep residual network for large-scale acoustic scene analysis. In *Interspeech*, 2019.
- Furlanello, T., Lipton, Z. C., Tschannen, M., Itti, L., and Anandkumar, A. Born again neural networks. *arXiv preprint arXiv:1805.04770*, 2018.
- Gemmeke, J. F., Ellis, D. P., Freedman, D., Jansen, A., Lawrence, W., Moore, R. C., Plakal, M., and Ritter, M. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 776–780. IEEE, 2017.
- Guzhov, A., Raue, F., Hees, J., and Dengel, A. Esresnet: Environmental sound classification based on visual domain models. *arXiv preprint arXiv:2004.07301*, 2020.
- Hershey, S., Chaudhuri, S., Ellis, D. P., Gemmeke, J. F., Jansen, A., Moore, R. C., Plakal, M., Platt, D., Saurous, R. A., Seybold, B., et al. Cnn architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 131–135. IEEE, 2017.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Kong, Q., Yu, C., Xu, Y., Iqbal, T., Wang, W., and Plumbley, M. D. Weakly labelled audioset tagging with attention neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(11):1791–1802, 2019.
- Kumar, A. and Ithapu, V. K. Secost:: Sequential co-supervision for large scale weakly labeled audio event detection. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 666–670. IEEE, 2020.
- Kumar, A. and Raj, B. Audio event detection using weakly labeled data. In *24th ACM International Conference on Multimedia*. ACM Multimedia, 2016.
- Kumar, A., Khadkevich, M., and Fugen, C. Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes. In *Acoustics, Speech and Signal Processing (ICASSP), 2018 IEEE International Conference on*. IEEE, 2018.
- Kumar, A., Shah, A., Hauptmann, A. G., and Raj, B. Learning sound events from webly labeled data. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pp. 2772–2778, 2019.
- McFee, B., Salamon, J., and Bello, J. P. Adaptive pooling operators for weakly labeled sound event detection. *arXiv preprint arXiv:1804.10070*, 2018.

- Minsky, M. Society of mind: a response to four reviews. 1994.
- Mirzadeh, S.-I., Farajtabar, M., Li, A., and Ghasemzadeh, H. Improved knowledge distillation via teacher assistant: Bridging the gap between student and teacher. *arXiv preprint arXiv:1902.03393*, 2019.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- Piczak, K. J. Esc: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, pp. 1015–1018. ACM, 2015.
- Polino, A., Pascanu, R., and Alistarh, D. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018.
- Ruvolo, P. and Eaton, E. Ella: An efficient lifelong learning algorithm. In *International Conference on Machine Learning*, pp. 507–515, 2013.
- Sailor, H. B., Agrawal, D. M., and Patil, H. A. Unsupervised filterbank learning using convolutional restricted boltzmann machine for environmental sound classification. In *INTERSPEECH*, pp. 3107–3111, 2017.
- Shah, A., Kumar, A., Hauptmann, A. G., and Raj, B. A closer look at weak label learning for audio events. *arXiv preprint arXiv:1804.09288*, 2018.
- Silver, D. L., Yang, Q., and Li, L. Lifelong machine learning systems: Beyond learning algorithms. In *2013 AAAI spring symposium series*, 2013.
- Su, T.-W., Liu, J.-Y., and Yang, Y.-H. Weakly-supervised audio event detection using event-specific gaussian filters and fully convolutional networks. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 791–795, 2017.
- Virtanen, T., Plumbley, M. D., and Ellis, D. *Computational analysis of sound scenes and events*. Springer, 2018.
- Wang, Y., Li, J., and Metze, F. Comparing the max and noisy-or pooling functions in multiple instance learning for weakly supervised sequence learning tasks. *arXiv preprint arXiv:1804.01146*, 2018.
- Wang, Y., Li, J., and Metze, F. A comparison of five multiple instance learning pooling functions for sound event detection with weak labeling. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 31–35. IEEE, 2019.
- Weinshall, D., Cohen, G., and Amir, D. Curriculum learning by transfer learning: Theory and experiments with deep networks. *arXiv preprint arXiv:1802.03796*, 2018.
- Xiong, Z., Radhakrishnan, R., Divakaran, A., and Huang, T. S. Audio events detection based highlights extraction from baseball, golf and soccer games in a unified framework. In *Multimedia and Expo, 2003. ICME'03. Proceedings. 2003 International Conference on*, volume 3, pp. III–401. IEEE, 2003.
- Ye, J., Kobayashi, T., Murakawa, M., and Higuchi, T. Acoustic scene classification based on sound textures and events. In *Proceedings of the 23rd ACM international conference on Multimedia*, pp. 1291–1294, 2015.
- Yim, J., Joo, D., Bae, J., and Kim, J. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Yu, C., Barsim, K. S., Kong, Q., and Yang, B. Multi-level attention model for weakly supervised audio classification. *arXiv preprint arXiv:1803.02353*, 2018.
- Zhang, H., McLoughlin, I., and Song, Y. Robust sound event recognition using convolutional neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 559–563. IEEE, 2015.