

# SeCoST: SEQUENTIAL CO-SUPERVISION FOR LARGE SCALE WEAKLY LABELED AUDIO EVENT DETECTION

Anurag Kumar, Vamsi Krishna Ithapu

Facebook Reality Labs  
Redmond, WA USA

anuragkr@fb.com, vamsi.ithapu@oculus.com

## ABSTRACT

Weakly supervised learning algorithms are critical for scaling audio event detection to several hundreds of sound categories. Such learning models should not only disambiguate sound events efficiently with minimal class-specific annotation but also be robust to label noise, which is more apparent with weak labels instead of strong annotations. In this work, we propose a new framework for designing learning models with weak supervision by bridging ideas from sequential learning and knowledge distillation. We refer to the proposed methodology as SeCoST (*pronounced Sequest*) — Sequential Co-supervision for training generations of Students. SeCoST incrementally builds a cascade of student-teacher pairs via a novel knowledge transfer method. Our evaluations on Audioset (the largest weakly labeled dataset available) show that SeCoST achieves a mean average precision of **0.383** while outperforming prior state of the art by a considerable margin.

**Index Terms**— Audio Event Detection, Teacher Student Models, Weakly Labeled, Sequential Learning

## 1. INTRODUCTION

In the past decade, supervised learning has been extensively studied for audio event recognition and detection (AED) [1]. While several classical machine learning and deep learning methods have been developed for AED using strong labels [2, 1, 3, 4], much of the recent progress has focused on efficiently leveraging *weakly labeled* data [5]. In such a weak labeling paradigm, audio recordings are only tagged for the presence or absence of sound events; unlike the *strong labeling* alternative where explicit time stamps of sound events are required. Hence, the annotation efforts are substantially lower thereby giving the ability to scale AED to large datasets, *e.g.* Audioset [6]. It has also now become an important part of the annual DCASE challenge on sound events and scenes<sup>1</sup>.

Several authors have shown promising results on weakly labeled AED [7, 8, 9, 10]. A significant fraction of these works uses deep convolutional neural networks in one form or other. Some are driven by attention mechanisms in neural networks [11], so as to efficiently characterize the temporal occurrences of events in the audio recordings [8, 12, 13]. Other approaches have incorporated recurrent neural networks to model the temporal attributes of sound events [13, 14].

However, large scale AED using weakly labeled data remains an open problem. When the timestamps of event occurrences are not provided, one cannot use explicit example clips of sounds for training. This clearly makes it harder to learn the necessary features and

characteristics that disambiguate different sound events. Additionally, noise and the presence of other irrelevant sounds complicates the learning, in particular with long recordings. Lastly, one can observe that noisy labels are also of concern in weak label learning [15, 16].

In this work, we address some of the above issues by presenting a novel learning framework for sounds. The proposal derives ideas from two distinct, but partly related, learning paradigms — sequential learning [17] and knowledge sharing [18]. In the seminal work on *Sequence of Teaching Selves* [19], the authors hypothesize that the human learning goes through different stages of development, where each such stage is “guided” by previous stages. This is, in principle, similar to lifelong learning where new knowledge is accumulated while retaining previous (learned) experiences [18, 20]. The central theme of these works is sequential learning, or, learning over time. Alternatively, there is recent interest in knowledge distillation (KD) through *teacher - student* frameworks [21, 22, 23], where the main motivation is model compression *i.e.*, constructing a smaller, low-capacity, student model that emulates a high-capacity high-performance teacher. These student networks are optimized based on some carefully designed divergence measures [21, 24].

We tackle weakly labeled learning by constructing a sequence of reasonably well trained neural networks (on the weak labels), where each network in the series is designed to be better than the previous one. However, unlike the classical KD where “one-shot” distillation is done from a single teacher to a student with the goal of compressing the model, the student here aims to match the performance of the teacher *while* also correcting for teacher’s errors. This entails constructing a cascade of student-teacher pairs and allowing the student to learn from teacher’s mistakes over *multiple generations*.

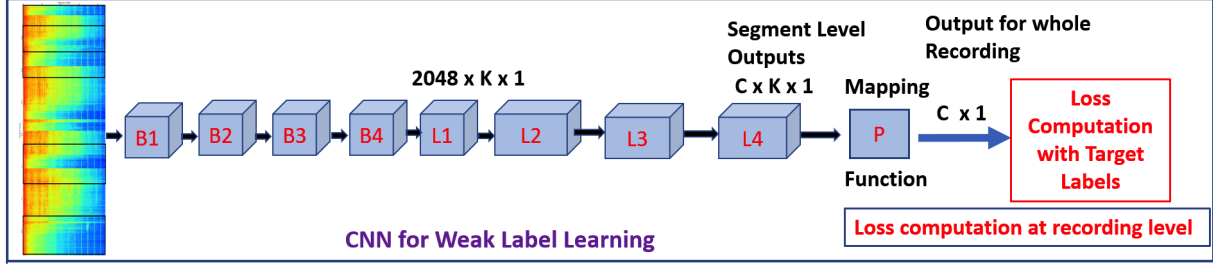
We do this by controlling the amount of transferable knowledge between consecutive generations. This helps in correcting the implicit noise associated with weak labels, while distilling the necessary knowledge needed for generalization. We refer to the proposed framework as **SeCoST** — *Sequential Co-Supervision* for training generations of Students from Teachers.

The rest of the presentation is as follows. In Section 2, we first describe the baseline deep convolutional network that drives the work-flow of the overall framework. We then describe our proposed **SeCoST** framework in Section 3, followed by experimental evaluations and discussion in Section 4. We conclude the paper in Section 5.

## 2. DEEP CNN FOR WEAKLY LABELED AED

**Notation:** Let  $S$  be the set of audio recordings.  $C$  denotes the set of labeled sound classes in these recordings. Each recording is rep-

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



**Fig. 1. WELS-Net:** Deep CNN for weakly labeled AED.  $|C|$ : number of classes.  $K$ : number of segments obtained for the given input.  $P$  maps the segment level output(s) to the recording level output.

represented by logmel features (denoted by  $X$ ). Let  $\mathbf{y} \in R^{|C|}$  be the (weak) label vector for the input  $X$ , and  $y_i = 1$  corresponds to  $i^{th}$  class being present (i.e., tagged).

We use a deep convolutional neural network as the prototypical learning model that drives our proposed framework [10, 25]. Given an input audio recording, the idea is to first produce segment level predictions. The segments are audio snippets of small duration (e.g., 1 second length). The resulting segment level outputs are then mapped to a recording level prediction. The appropriate loss is then calculated using this prediction and the recording level weak label. The mapping from segments to recording may be done via simple mean or max operation over segment level outputs or even by a neural network, if necessary.

Figure 1 and Table 1 summarize the network schematic and the specific architectural details.  $B1$  to  $B4$  blocks consist of two convolutional layers followed by a pooling layer. The size of segment level output at layer  $L4$  will be  $|C| \times K \times 1$ , where  $|C|$  is the number of classes and  $K$  denotes the number of possible segments for a given input. The recording level output for any given class is obtained by taking the average of segment level outputs (layer  $P$ ).

We will refer to the above network as WELS-Net – *WE*akly *L*abeled *S*ound *N*etwork. WELS-Net is flexible with respect to input size, thereby allowing processing of audio recordings of variable lengths. Moreover, it produces segment level outputs which can then be used for temporal localization of sound events in the audio recordings.

Note that, WELS-Net takes Logmel spectrograms as input. Specifically, we use 64 mel-filters and mel features are extracted using a 16 ms window moving by a stride of 10 ms. The sampling rate of all audio recordings is 16 kHz. Table 1 shows output sizes at each layer for a logmel input with 1024 frames (size  $1024 \times 64$ ). For this input, we get  $K = 30$  segments at layer  $L4$ . In the time domain, this network produces outputs for  $\sim 1$  second segments (with a stride of  $\sim 0.33$  seconds).

## 2.1. Network Training

Recall that the inputs are logmel features denoted by  $X$ . We assume a realistic scenario where  $X$  may be tagged with multiple labels, i.e., multiple sound events may be present in a single recording. The goal is to train a neural network  $\mathcal{N}(\theta)$  which can generalize well on unseen data. The network is trained by minimizing a loss function which measures divergence between the network outputs  $\mathcal{N}(\theta, X)$  and the target  $\mathbf{y}$ . Let  $\mathcal{L}(\mathcal{N}(\theta, X), \mathbf{y})$  denote the loss function. We

**Table 1.** CNN Architecture Used.  $F$ : Number of Filters.  $P$ : Padding Size.  $S$ : Stride.  $BN$ : Batch Normalization.  $|C|$ : Number of classes. Last column shows output sizes for an input with 1024 logmel frames, i.e  $1024 \times 64$  dimensional input.

Layers	Layer Parameters	Output size
Input	-	$1 \times 1024 \times 64$
Block B1	$3 \times 3$ conv, F: 64, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $3 \times 3$ conv, F: 64, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $4 \times 4$ pool, S:4	$64 \times 1024 \times 64$ $64 \times 1024 \times 64$ $64 \times 256 \times 16$
Block B2	$3 \times 3$ conv, F: 128, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $3 \times 3$ conv, F: 128, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $2 \times 2$ pool, S:2	$128 \times 256 \times 16$ $128 \times 256 \times 16$ $128 \times 128 \times 8$
Block B3	$3 \times 3$ conv, F: 256, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $3 \times 3$ conv, F: 256, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $2 \times 2$ pool, S:2	$256 \times 128 \times 8$ $256 \times 128 \times 8$ $256 \times 64 \times 4$
Block B4	$3 \times 3$ conv, F: 512, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $3 \times 3$ conv, F: 512, P:1, S:1 $\rightarrow$ BN $\rightarrow$ ReLU $2 \times 2$ pool, S:2	$512 \times 64 \times 4$ $512 \times 64 \times 4$ $512 \times 32 \times 2$
Layer L1	$3 \times 2$ conv, F: 2048, P:0, S:1 $\rightarrow$ BN $\rightarrow$ ReLU	$2048 \times 30 \times 1$
Layer L2	$1 \times 1$ conv, F: 1024, P:0, S:1 $\rightarrow$ BN $\rightarrow$ ReLU	$1024 \times 30 \times 1$
Layer L3	$1 \times 1$ conv, F: 1024, P:0, S:1 $\rightarrow$ BN $\rightarrow$ ReLU	$1024 \times 30 \times 1$
Layer L4	$1 \times 1$ conv, F: C, P:0, S:1 $\rightarrow$ BN $\rightarrow$ ReLU	$ C  \times 30 \times 1$
P	Global Average Pooling	$ C  \times 1$

use binary-cross entropy loss. Thereby, we have

$$l(p_i, y_i) = -y_i \log(p_i) - (1 - y_i) \log(1 - p_i)$$

$$\mathcal{L}(\mathcal{N}(\theta, X), \mathbf{y}) = \frac{1}{|C|} \sum_{i=1}^{|C|} l(y_i, p_i) \quad (1)$$

$p_i$  is the output of the network for the  $i^{th}$  class.  $l(y_i, p_i)$  is the loss for this  $i^{th}$  class, and  $\mathcal{L}(\mathcal{N}(\theta, X), \mathbf{y})$  computes the overall loss for the input  $X$  and the corresponding target  $\mathbf{y}$ .

## 3. SECoST FRAMEWORK

Using the network architecture presented in Section 2, we will now describe the proposed sequential co-supervision learning. As mentioned in Section 1, SeCoST follows the principle of sequence of teaching selves. A sequence of learners (neural networks here) are trained. At each stage, the learning of a new network is supervised by the already trained network(s) from previous stages. First recall that training a network with the loss function in Eq 1 corresponds to learning from available ground truth labels  $\mathbf{y}$ . This is our initialization a.k.a. the first teacher. We denote this *base model* by  $\mathcal{N}^{T_0}$ . Once it is learned, we propose that it can co-supervise training a new network (identical in architecture to the teacher) from scratch. In other

---

**Algorithm 1 SeCoST:**


---

**Input:** Training data  $\mathcal{D} = \{X^i, \mathbf{y}^i\}$ , Number of stages  $S$ ,  $\{\alpha^s$  for each stage  $s = 1$  to  $S\}$

**Output:** Trained Network after  $S$  stages

- 1: Train base WELS-Net or Teacher-0 ( $\mathcal{N}^{T_0}$ ) using  $\mathcal{D} = \{X^i, \mathbf{y}^i\}$
  - 2: **for**  $s = 1, 2, \dots, S$  **do**
  - 3:    Compute new target  $\bar{\mathbf{y}}^i$  for all training points ( $X^i, \mathbf{y}^i$ ) from  $\mathbf{y}^i, \alpha^s$  and prediction of  $\mathcal{N}^{T_{s-1}}$  on  $X^i$  using Eq 2
  - 4:    Train new WELS-Net ( $\mathcal{N}^{S_s}$ ) for current stage using  $\mathcal{D} = \{X^i, \bar{\mathbf{y}}^i\}$
  - 5:     $\mathcal{N}^{T_s} = \mathcal{N}^{S_s}$  // Student becomes teacher for next stage
  - 6: **end for**
  - 7: return  $\mathcal{N}^{T_S}$
- 

words, the outputs of the teacher networks drive the supervision in future generations. We now formalize this process.

If the vector  $\hat{\mathbf{y}}$  denotes the output of the teacher with input  $X$ , then the new target denoted by  $\bar{\mathbf{y}}$ , for the same input  $X$ , is given by

$$\bar{\mathbf{y}} = \alpha \mathbf{y} + (1 - \alpha) \hat{\mathbf{y}} \quad (2)$$

where  $\alpha$  is a hyper-parameter that controls the contribution of the teacher network’s supervision.

With a single teacher, the class-wise target for the student model is  $\bar{y}_i = \alpha y_i + (1 - \alpha) \hat{y}_i$ . The corresponding loss for the  $i^{\text{th}}$  class now becomes

$$\begin{aligned} l(p_i, \bar{y}_i) &= -\bar{y}_i \log(p_i) - (1 - \bar{y}_i) \log(1 - p_i) \\ &= -\alpha y_i p_i - (1 - \alpha y_i) \log(1 - p_i) + (1 - \alpha) \hat{y}_i \log \frac{1 - p_i}{p_i} \\ l(p_i, \bar{y}_i) &= l(p_i, \alpha y_i) + (1 - \alpha) \hat{y}_i \log \frac{1 - p_i}{p_i} \end{aligned} \quad (3)$$

Using this new class-wise loss, the overall new loss function for the student network, denoted by  $\mathcal{L}(\cdot, \bar{\mathbf{y}})$ , is

$$\begin{aligned} \mathcal{L}(\cdot, \bar{\mathbf{y}}) &= \frac{1}{|\mathcal{C}|} \sum_{i=1}^{|\mathcal{C}|} [l(\alpha y_i, p_i) + (1 - \alpha) \hat{y}_i \log \frac{1 - p_i}{p_i}] \\ &= \mathcal{L}(\cdot, \alpha \mathbf{y}) + (1 - \alpha) \frac{1}{|\mathcal{C}|} \sum_{i=1}^{|\mathcal{C}|} \hat{y}_i \log \frac{1 - p_i}{p_i} \end{aligned} \quad (4)$$

which is a combination of loss w.r.t. ground truth  $\mathbf{y}$ , although, weighted by a factor of  $\alpha < 1$ ; and a term representing supervision from the teacher. We can rewrite this as

$$\mathcal{L}(\cdot, \bar{\mathbf{y}}) = \mathcal{L}(\cdot, \mathbf{y}) + (1 - \alpha) \frac{1}{|\mathcal{C}|} \sum_{i=1}^{|\mathcal{C}|} (y_i - \hat{y}_i) \log \frac{1 - p_i}{p_i} \quad (5)$$

Here, the first term is same as computing the loss w.r.t.  $\mathbf{y}$ . The second term now involves both the ground truth labels and the teacher’s predictions. Hence, the additional supervision for student is entirely determined by how much the teacher’s predictions ( $\hat{y}_i$ ) differ from the true labels  $y_i$ , i.e., the student is trying to learn from the mistakes made by the teacher network. This additional information is hypothesized to improve the generalization capabilities.

This overall procedure can now be emulated to multiple stages  $S$ . Algorithm 1 summarizes this sequential procedure over  $S$  stages. The output of this procedure is a network trained over  $S$  generations with one or more teachers in each stage. Note that we expect the improvement in generalization to have diminishing returns after some stages, and we discuss more about this behaviour in Section 4.

**Table 2.** Comparing SeCoST with state of the art methods on Audioset.

Method	mAP	mAUC	Method	mAP	mAUC
[7] - Pooling	0.343	0.966	[7] - Attention	0.361	0.969
[7] - Attention-Large	0.369	0.969	TALNet [13] - Pooling	0.361	0.966
TALNet [13] - Attention	0.362	0.965	WELS-Net (Our’s Base)	0.352	0.962
<b>SeCoST (fixed <math>\alpha = 0.3</math>)</b>	<b>0.379</b>	<b>0.970</b>	<b>SeCoST (Variable <math>\alpha</math>)</b>	<b>0.383</b>	<b>0.971</b>

**Multiple Teachers per Stage:** Observe that Eqs 2 and 5 are parameterizing learning from a single teacher at a given stage. The above procedure can also be extended to incorporate multiple teachers (denoted by  $T$ ) per stage. The new target is given by a convex combination of all available supervision, from all the  $T$  teachers, as shown below.

$$\bar{\mathbf{y}} = \sum_{k=0}^T \alpha_k \hat{\mathbf{y}}^k \quad \text{and} \quad \sum_{k=0}^T \alpha_k = 1 \quad (6)$$

where  $\hat{\mathbf{y}}^0 = \mathbf{y}$  represents the ground truth labels, and  $\alpha_k, k = 0$  to  $T$ , parameterize the contribution of the ground truth and the  $T$  teacher networks respectively. The class-wise loss  $l(p_i, \bar{y}_i)$  becomes


$$l(p_i, \bar{y}_i) = l(p_i, \alpha_0 y_i) + \sum_{k=1}^T \alpha_k \hat{y}_i^k \log \frac{1 - p_i}{p_i} \quad (7)$$

## 4. EXPERIMENTS AND RESULTS


We evaluate SeCoST using Audioset (the largest available dataset for sound events) [6]. It has weakly labeled examples for 527 sound events, with approximately 2 million *training* examples and 20k *evaluation* recordings. We use 25k samples from the training set for validation. Each recording is 10 seconds long (and we resample them at 16kHz). Audioset is multi-label in nature with  $\sim 2.7$  labels per recording. PyTorch is used for implementing the networks [26]. Training utilizes Adam [27] where hyperparameters (like learning rates) are tuned using the validation set. Similar to existing AED works [6, 10, 7], Average precision (AP) and area under ROC (AUC) are used to measure the performance [28, 29]. Further, mean average precision (mAP) and mean AUC (mAUC) over all classes summarizes the overall performances. In this work, we use a single teacher at each stage (Alg. 1), leaving evaluation of multiple teacher per stage for future work.

**Performance Comparison:** Table 2 compares **SeCoST** with existing state-of-the-art methods on Audioset. Note that the authors of [7] use embeddings from a network trained on a very large database of audio recordings (YouTube-70M) [30], thereby the resulting feature representations already lead to enhanced performance. However, we work *directly* with audio recordings and use their logmel representations. Our base WELS-Net model ( $\mathcal{N}^{T_0}$ ) trained on the ground truth labels gives an mAP of 0.352 over all the 527 events. SeCoST gives an mAP of **0.383**, improving WELS-Net by **8.8%**. Notably, it is also **3.8%** better than the best reported performance in literature. Our best performance of 0.383 mAP is obtained by increasing the contribution of teacher ( $1 - \alpha$ ) as the sequence progresses. Using a fixed  $\alpha$  of 0.3, we improve the base WELS-Net by 7.7% (from 0.352 to 0.379).

**Effect of  $\alpha$ :** Figure 2 (Left) analyzes the influence of teacher’s contribution (parameterized by  $1 - \alpha$ ) in SeCoST (see Alg 1). Here a single stage of SeCoST is done with base WELS-Net as the teacher.  $1 - \alpha = 0$  represents no contribution from the teacher, i.e training only on ground truth labels. We can see that the performance improves as teacher’s contribution increases, but only up to a certain point. This occurs at  $1 - \alpha = 0.3$ , with the corresponding mAP of



**Fig. 2. Left:** Single stage SeCoST with varying contribution ( $1 - \alpha$ ) from teacher. Base WELS-Net ( $\mathcal{N}^{T_0}$ ) as teacher. **Right:** 4 stages of SeCoST showing performance after each stage. Teacher’s contribution fixed at 0.7.




**Fig. 3. Left:** 5 stage SeCoST with variable teacher’s contribution.  $x$ -axis shows teacher’s contribution at each stage. **Right:** SeCoST with only teacher networks ( $\alpha = 0$ )

0.374. Thereby, improving the base WELS-Net by 6.3%. From this, we can argue that teacher’s supervision should be well proportioned along with ground truth supervision.

**SeCoST Stage-wise performance:** Fig 2 (Right) evaluates the sequential training aspect of SeCoST. We do 4 stages of training in SeCoST ( $S = 4$  in Alg 1). The contribution of teacher in the supervision remains same for all 4 stages with  $1 - \alpha_s = 0.7$ , for  $s = 1$  to 4. We can see that sequential co-supervision leads to improvements, although with diminishing returns after each stage. After the first stage of co-supervision ( $\mathcal{N}^{T_0}$  as teacher), the mAP improves from 35.2 to 37.4, a 6.3% improvement over the base WELS-Net. Using this improved network as teacher in the second stage of SeCoST, we see a further improvement of 1.3% (overall a 7.6% improvement from WELS-Net). The performance then saturates and we do not see any additional improvements in future stages.

We saw above that, SeCoST in general works better with lower  $\alpha$  i.e., larger weight to teacher’s contributions. This suggests that as newer generations of networks become teachers, it might be helpful to increase their contribution in co-supervision. To evaluate this, we run SeCoST for 5 stages with  $\{\alpha_1 = 0.3, \alpha_2 = 0.3, \alpha_3 = 0.2, \alpha_4 = 0.1, \alpha_5 = 0.05\}$ . The performance at different stages is shown in Figure 3. As teacher’s contribution is increased from 0.7 to 0.8 at stage 3, we see an improvement in performance, unlike the case where it was fixed at 0.7 and the performance remained same at stage 3 (refer back to Figure 2 (left)). At stage 4, with 90% supervision from teacher, we get an improved mAP of 0.383 (overall 8.8% improvement on base WELS-Net).

**Only Teacher Networks ( $\alpha = 0$ ) in SeCoST:** In Figure 2 (left), we showed that training the student using only the teacher’s output ( $1 - \alpha = 1$ ) as the target already gives a better model. We get 0.366 mAP versus 0.352 when training only on ground truth labels. A similar observations has been made in [21]. This may be because the soft probability outputs from the teacher network provides richer information compared to the ‘hard’ ground truth labels. However,



**Fig. 4. Left:** Relative Improvement in mean AP within different ranges. For instance, 0 – 0.1 represents sound classes for which AP using base WELS-Net is  $< 0.1$ .  $y$ -axis shows relative improvement in mean AP. Number at top of each bar shows number of classes with AP in that range.

as shown in Figure 3 (right), this teacher-alone strategy does not work well for future generations. Here we run SeCoST for 3 stages using only teacher’s supervision. We see that the first stage leads to an improvement over the base network, and then there is saturation. This shows that co-supervision is necessary. Hence, the teacher’s knowledge needs to be coupled with the ground truth labels while training the sequence of students.

**Some Class Specific Analysis:** We observe that SeCoST improves performance for  $> 85\%$  percent of the sound classes in Audioset (448 out of 527). Of these, for 110 classes we get  $> 20\%$  improvement in AP. Specifically, for *Crushing, Harmonic and Mouse* sounds in Audioset vocabulary, we observe more than 100% improvement using SeCoST. On the other hand, there are only 12 classes with more than 10% drop in performance. Sound class *Squish* has maximum drop in performance, around 19%. Note that, these summaries are based on the best SeCoST model with mAP of 0.383 and base WELS-Net with 0.352 mAP. To further analyze class specific performance, we try to see whether the improvements are coming for classes where the base model already does well or if the classes with low APs are actually improving. Figure 4 shows this relative mAP improvement for classes with APs within a specified range. For classes with the APs  $< 0.1$ , the mean AP improves by  $\sim 28\%$ . For classes with APs  $\in [0.1, 0.2)$ , the gain is 21%. This show that SeCoST leads to considerable improvement in classes which are harder to learn for WELS-Net. For relatively easy classes (WELS-Net AP  $> 0.5$ ), we see  $\sim 4.5\%$  improvement.

## 5. CONCLUSION

We proposed a sequential co-supervision learning framework for audio event detection. Our proposal, SeCoST, builds a generation of networks by designing student models that learn to predict a convex combination of teachers’ predictions instead of the given ground truth. We showed that SeCoST gives a considerably better performance on Audioset compared to baseline model and state of the art performance. We note that the proposed framework is generally applicable to learning from noisy, weak labels, and we intend to future investigate the theoretical merits of the model in the future.

## 6. REFERENCES

[1] Tuomas Virtanen, Mark D Plumbley, and Dan Ellis, *Computational analysis of sound scenes and events*, Springer, 2018.

- [2] D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M.D. Plumbley, "Detection and classification of acoustic scenes and events," *Multimedia, IEEE Transactions on*, vol. PP, 2015.
- [3] Karol J Piczak, "Environmental sound classification with convolutional neural networks," in *Machine Learning for Signal Processing (MLSP), 2015 IEEE 25th International Workshop on*. IEEE, 2015, pp. 1–6.
- [4] Ian McLoughlin, Haomin Zhang, Zhipeng Xie, Yan Song, and Wei Xiao, "Robust sound event classification using deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 3, pp. 540–552, 2015.
- [5] Anurag Kumar and Bhiksha Raj, "Audio event detection using weakly labeled data," in *24th ACM International Conference on Multimedia*. ACM Multimedia, 2016.
- [6] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [7] Qiuqiang Kong, Changsong Yu, Yong Xu, Turab Iqbal, Wenwu Wang, and Mark D Plumbley, "Weakly labelled audioset tagging with attention neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1791–1802, 2019.
- [8] Szu-Yu Chou, Jyh-Shing Roger Jang, and Yi-Hsuan Yang, "Learning to recognize transient sound events using attentional supervision.," in *IJCAI*, 2018, pp. 3336–3342.
- [9] Brian McFee, Justin Salamon, and Juan Pablo Bello, "Adaptive pooling operators for weakly labeled sound event detection," *arXiv preprint arXiv:1804.10070*, 2018.
- [10] Anurag Kumar, M. Khadkevich, and C. Fugen, "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.
- [11] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [12] Changsong Yu, Karim Said Barsim, Qiuqiang Kong, and Bin Yang, "Multi-level attention model for weakly supervised audio classification," *arXiv preprint arXiv:1803.02353*, 2018.
- [13] Yun Wang, Juncheng Li, and Florian Metze, "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)*. IEEE, 2019, pp. 31–35.
- [14] Sharath Adavanne and Tuomas Virtanen, "Sound event detection using weakly labeled dataset with stacked convolutional and recurrent neural network," *arXiv preprint arXiv:1710.02998*, 2017.
- [15] Anurag Kumar, Ankit Shah, Alexander G. Hauptmann, and Bhiksha Raj, "Learning sound events from weakly labeled data," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 2019, pp. 2772–2778.
- [16] Eduardo Fonseca, Manoj Plakal, Daniel PW Ellis, Frederic Font, Xavier Favory, and Xavier Serra, "Learning sound event classifiers from web audio with noisy labels," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 21–25.
- [17] Guang-Bin Huang and Lei Chen, "Convex incremental extreme learning machine," *Neurocomputing*, vol. 70, no. 16-18, pp. 3056–3062, 2007.
- [18] Daniel L Silver, Qiang Yang, and Lianghao Li, "Lifelong machine learning systems: Beyond learning algorithms," in *2013 AAAI spring symposium series*, 2013.
- [19] Marvin Minsky, "Society of mind: a response to four reviews," 1994.
- [20] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter, "Continual lifelong learning with neural networks: A review," *Neural Networks*, 2019.
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [22] Jimmy Ba and Rich Caruana, "Do deep nets really need to be deep?," in *Advances in neural information processing systems*, 2014, pp. 2654–2662.
- [23] Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2006, pp. 535–541.
- [24] Tommaso Furlanello, Zachary C Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar, "Born again neural networks," *arXiv preprint arXiv:1805.04770*, 2018.
- [25] Ankit Shah, Anurag Kumar, Alexander G Hauptmann, and Bhiksha Raj, "A closer look at weak label learning for audio events," *arXiv preprint arXiv:1804.09288*, 2018.
- [26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer, "Automatic differentiation in PyTorch," in *NIPS Autodiff Workshop*, 2017.
- [27] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [28] Tom Fawcett, "Roc graphs: Notes and practical considerations for researchers," *Machine learning*, vol. 31, no. 1, pp. 1–38, 2004.
- [29] Chris Buckley and Ellen M Voorhees, "Retrieval evaluation with incomplete information," in *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2004, pp. 25–32.
- [30] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al., "Cnn architectures for large-scale audio classification," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 131–135.