

A COMPARISON OF FIVE MULTIPLE INSTANCE LEARNING POOLING FUNCTIONS FOR SOUND EVENT DETECTION WITH WEAK LABELING

Yun Wang[†], Juncheng Li^{†,‡}, and Florian Metze[†]

[†]Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA, U.S.A.

[‡]Research and Technology Center, Robert Bosch LLC, Pittsburgh, PA, U.S.A.
maigoakisame@gmail.com, {junchenl, fmetze}@cs.cmu.edu

ABSTRACT

Sound event detection (SED) entails two subtasks: recognizing what types of sound events are present in an audio stream (audio tagging), and pinpointing their onset and offset times (localization). In the popular multiple instance learning (MIL) framework for SED with weak labeling, an important component is the pooling function. This paper compares five types of pooling functions both theoretically and experimentally, with special focus on their performance of localization. Although the attention pooling function is currently receiving the most attention, we find the linear softmax pooling function to perform the best among the five. Using this pooling function, we build a neural network called TALNet. It is the first system to reach state-of-the-art audio tagging performance on Audio Set, while exhibiting strong localization performance on the DCASE 2017 challenge at the same time.

Index Terms— Sound event detection (SED), weak labeling, multiple instance learning (MIL), pooling functions, attention

1. INTRODUCTION

Sound event detection (SED) is the task of detecting the type, onset time and offset time of sound events in audio stream. While some studies are satisfied with recognizing what types of sound events are present in a recording (*audio tagging*), this paper pays special attention to the *localization* of sound events.

Modern SED systems usually take the form of neural networks, with convolutional layers [2, 3], recurrent layers [4, 5, 6, 7], or both [8]. The networks predict the probability of each sound event type frame by frame; applying a threshold to these frame-level probabilities will then produce localized detections of sound events.

Traditionally, the training of SED models relied upon *strong labeling*, which specifies the type, onset time and offset time of each sound event occurrence. But such annotation is very tedious to obtain by hand. In order to scale SED up, researchers have turned to SED with *weak labeling*, which only specifies the types of sound events present in each training recording but does not provide any temporal information. In March 2017, Google released the weakly labeled Audio Set [9], which is by far the largest corpus available for SED. The DCASE challenge of 2017 [10] featured a task of SED with weak labeling, which used a subset of Audio Set.

A common framework for SED with weak labeling is *multiple instance learning* (MIL) [11], as shown in Fig. 1. In MIL, we do not know the ground truth label of every training instance; instead, the instances are grouped into bags, and we only know the label of

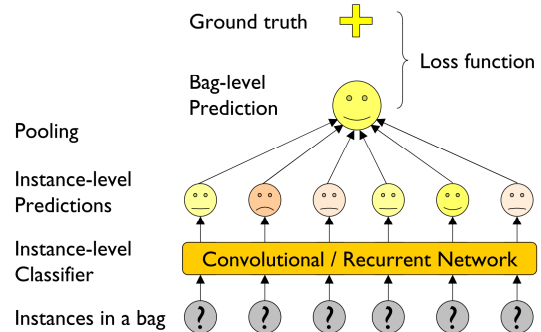


Fig. 1. Block diagram of a MIL system for SED with weak labeling.

bags. In the case of binary classification, the relationship between instance labels and bag labels often obey the *standard multiple instance (SMI) assumption*: the bag label is positive if and only if the bag contains at least one positive instance. In SED, each training recording is regarded as a bag, and its frames are regarded as instances. Each sound event type is considered independently, so SED becomes a binary classification problem for each sound event type. A neural network predicts the probability of each sound event type being active at each frame. Then, a *pooling function* aggregates the frame-level probabilities into a recording-level probability for each sound event type. The recording-level probabilities can be compared against the recording labels to compute a loss function, and the network can then be trained to minimize the loss.

The choice of the pooling function is an important decision. The default choice is the “max” pooling function [12, 13], which is faithful to the SMI assumption. A previous study of ours [14] has evaluated a “noisy-or” pooling function [15, 16, 17], and has shown it does not work for localization despite its nice probabilistic interpretation under the SMI assumption. Since the 2017 DCASE challenge, a number of other pooling functions have been reported to perform well even though they deviate from the SMI assumption. These include average pooling [18], two softmax pooling functions based on linear weighting [19] and exponential weighting [20], as well as an attention-based pooling function [21, 22]. The purpose of this paper is to compare these pooling functions against max pooling from two aspects: theoretically, we derive the gradient of the five pooling functions, and check if their signs lead the training down the right way; experimentally, we compare the five pooling functions on two SED corpora: the DCASE 2017 challenge [10] and Audio Set [9]. Although the attention pooling function appears to be the most favored by researchers, we demonstrate that it is the linear softmax pooling function that works best for localization.

Our experiments also result in a convolutional and recurrent neural network (CRNN) which is the first system within our knowledge

The first two authors were supported by a graduate research fellowship award from Robert Bosch LLC; the first author was also supported by a faculty research award from Google. This work used the “comet” and “bridges” clusters of the XSEDE environment [1], supported by NSF grant number ACI-1548562.

that exhibits strong performance on audio tagging and localization at the same time. We name this network ‘‘TALNet’’, where ‘‘TAL’’ stands for ‘‘tagging and localization’’. This network closely matches the current state-of-the-art audio tagging performance on Audio Set, while achieving competitive localization performance on the DCASE 2017 challenge without any finetuning.

2. THEORETICAL COMPARISON OF THE FIVE POOLING FUNCTIONS

2.1. Definition of the Pooling Functions

Let $y_i \in [0, 1]$ be the predicted probability of a certain event type at the i -th frame, and $y \in [0, 1]$ be the aggregated recording-level probability of the same event. We list the definitions of the five pooling functions to be compared in Table 1.

The max pooling function simply takes the largest y_i to be y . If the same threshold is applied to the recording-level and frame-level probabilities, then the frame-level predictions and recording-level prediction are guaranteed to be consistent with the SMI assumption. However, the max pooling function has a defect that only one frame in a recording can receive an error signal. As a consequence, if an event occurs multiple times in a recording, the occurrences that do not cover this frame may be easily missed. All the other four pooling functions try to alleviate this problem by assigning some weight to smaller y_i ’s when aggregating them to produce y .

The average pooling function [18] assigns an equal weight to all frames. The equation appears to defy the SMI assumption, but it is reported to perform better than the max pooling function in [18].

The two softmax pooling functions compute y as a weighted average of the y_i ’s, where larger y_i ’s receive larger weights. In this way, the recording-level probability is still mainly determined by the larger frame-level probabilities, but frames with smaller probabilities get a chance to receive an error signal. The linear softmax function [19] assigns weights equal to the frame-level probabilities y_i themselves, while the exponential softmax function [20] assigns a weight of $\exp(y_i)$ to the frame-level probability y_i .

Finally, in the attention pooling function [21, 22], the weights for each frame w_i are learned with a dedicated layer in the network. The recording-level probability y is then computed using the general weighted average formula. The attention pooling function appears to be most favored by researchers because of its flexibility, and variants have emerged such as the multi-level attention in [23].

2.2. Gradient of the Pooling Functions

In this section, we analyze the gradient of the loss function w.r.t. the frame-level probabilities y_i (and, in the case of attention, also the weights w_i). Let $t \in \{0, 1\}$ be the recording-level ground truth. The loss function is usually the cross entropy:

$$L = -t \log y - (1 - t) \log(1 - y) \quad (1)$$

We decompose its gradient with respect to the frame-level probabilities y_i (and the frame-level weights w_i) using the chain rule:

$$\frac{\partial L}{\partial y_i} = \frac{\partial L}{\partial y} \frac{\partial y}{\partial y_i}, \quad \frac{\partial L}{\partial w_i} = \frac{\partial L}{\partial y} \frac{\partial y}{\partial w_i} \quad (2)$$

The first term,

$$\frac{\partial L}{\partial y} = -\frac{t}{y} + \frac{1-t}{1-y} \quad (3)$$

does not depend on the choice of the pooling function. It is negative when the recording label is positive ($t = 1$), and positive when the

	Pooling Function	Gradient
Max pooling	$y = \max_i y_i$	$\frac{\partial y}{\partial y_i} = \begin{cases} 1, & \text{if } y_i = y \\ 0, & \text{otherwise} \end{cases}$
Average pooling	$y = \frac{1}{n} \sum_i y_i$	$\frac{\partial y}{\partial y_i} = \frac{1}{n}$
Linear softmax	$y = \frac{\sum_i y_i^2}{\sum_i y_i}$	$\frac{\partial y}{\partial y_i} = \frac{2y_i - y}{\sum_j y_j}$
Exp. softmax	$y = \frac{\sum_i y_i \exp(y_i)}{\sum_i \exp(y_i)}$	$\frac{\partial y}{\partial y_i} = (1 - y + y_i) \cdot \frac{\exp(y_i)}{\sum_j \exp(y_j)}$
Attention	$y = \frac{\sum_i y_i w_i}{\sum_i w_i}$	$\frac{\partial y}{\partial y_i} = \frac{w_i}{\sum_j w_j}, \quad \frac{\partial y}{\partial w_i} = \frac{y_i - y}{\sum_j w_j}$

Table 1. The five pooling functions and their gradients. n is the number of frames in a recording.

recording label is negative ($t = 0$). The second term, $\partial y / \partial y_i$ (and $\partial y / \partial w_i$), is calculated for each pooling function in Table 1.

With the max pooling function, $\partial y / \partial y_i$ equals 1 for the frame with the largest probability and 0 elsewhere. The fact that only one frame receives a non-zero gradient may cause many frame-level false negatives. The gradient for this single frame, though, does have the correct sign: when $t = 1$, the gradient $\partial L / \partial y_i$ is negative, so the frame-level probability y_i will be boosted in order to reduce the loss; when $t = 0$, the gradient is positive, so y_i will be suppressed.

With the average pooling function, $\partial y / \partial y_i$ equals $1/n$ regardless of the value of y_i . This means the gradient is distributed evenly across all frames. For negative recordings, this will suppress the probability y_i of all frames, and this is correct behavior. For positive recordings, however, not all frames should be boosted, and the average pooling function can produce a lot of false positive frames.

With the linear softmax pooling function, $\partial y / \partial y_i$ is positive where $y_i > y/2$, which gives rise to complicated and interesting behavior. For positive recordings ($t = 1$), the gradient is negative where $y_i > y/2$, and positive where $y_i < y/2$. As a result, larger y_i ’s will be boosted, while smaller y_i ’s will be suppressed. This is exactly the desired behavior under the SMI assumption: the frame-level probabilities are driven to the extremes 0 and 1, resulting in well-localized detections of sound events. For negative recordings ($t = 0$), the gradient is positive where $y_i > y/2$, and negative where $y_i < y/2$. This means all frame-level probabilities will be pushed toward $y/2$. Considering that y is a weighted average of the y_i ’s, given enough iterations, all the y_i ’s will converge to zero as desired.

With the exponential pooling function, $\partial y / \partial y_i$ is always positive, just like with the average pooling function. As a result, the exponential pooling function also has the concern of producing too many false positive frames. Nevertheless, the problem will be less serious, because smaller y_i ’s receive smaller gradients.

With the attention pooling function, the term $\partial y / \partial y_i$ is always positive. Therefore, the frame-level probabilities will be boosted or suppressed according to the recording label, with strengths proportional to the learned weights. This is correct behavior if frames with larger probabilities y_i also get larger weights w_i . However, because the weights w_i are also learned, we should also consider $\partial y / \partial w_i$, the gradient of the loss function w.r.t. the weights: this term is positive where $y_i > y$. When the recording is positive, this will cause the weight w_i to rise where the frame-level probability y_i is large and to shrink where the y_i is small, agreeing with the motivation that frames with larger probabilities y_i should get larger weights w_i . When the recording is negative, however, the opposite phenomenon will happen: larger weights will concentrate upon frames with smaller probabilities. This has a serious consequence: while the recording-level probability y will indeed be small, there




Fig. 2. Structures of the networks used in Secs. 3.1 (left) and 3.2 (right). The shape is specified as “frames * frequency bins * feature maps” for 3-D tensors (shaded), and “frames * feature maps” for 2-D tensors. “conv $n*m$ ” stands for a convolutional layer with the specified kernel size and ReLU activation; “(*2)” means the layer is repeated twice. “pool $n*m$ ” stands for a max pooling layer with the specified stride. “FC” is short for “fully connected”. At the output end, the “attention weights” block is only used with the attention pooling function.

will be frames with large probabilities y_i and small weights w_i . This means the recording-level prediction and frame-level predictions will be inconsistent with the SMI assumption, and the frames with large probabilities will end up being false positives for localization.

3. EXPERIMENTAL COMPARISON OF THE FIVE POOLING FUNCTIONS

3.1. The DCASE 2017 Challenge

We first compared the five pooling functions on Task 4 of the DCASE 2017 challenge [10]. The task involves 17 types of vehicle and warning sounds, and evaluates both tagging and localization.

The data used in the task is a subset of Audio Set [9]. It consists of a training set (51,172 recordings), a public test set (488 recordings), and a private evaluation set (1,103 recordings). All the recordings are 10-second excerpts from YouTube videos. The test and evaluation sets are strongly labeled so they can be used to evaluate both audio tagging and localization, but the training set only comes with weak labeling. Because we did not have access to the ground truth of the evaluation set, we report the performance of our systems on the test set. The test and evaluation sets have balanced numbers of the events, but the training set is unbalanced. We set aside 1,142 recordings from the training set to make a balanced validation set, and used the remaining 50,030 recordings for training.

We implemented a convolutional and recurrent neural network (CRNN), whose structure is shown in Fig. 2 (left). The input is a matrix of filterbank features; it has 400 frames and 64 frequency bins. The convolutional and pooling layers reduce the frame rate from 40 Hz to 10 Hz. At the output end, a fully connected layer with sigmoid activation produces frame-level predictions, which are then aggregated across time into recording-level predictions using any of the five pooling functions. If the attention pooling function is used, a separate fully connected layer with exponential activation is used to generate the weights. The network was trained using the PyTorch toolkit [24]. We applied data balancing so each minibatch contained roughly equal numbers of recordings of each event type.

The performance of audio tagging was evaluated with the micro-average F_1 on the recording level; localization was evaluated with the micro-average error rate (ER) and F_1 on 1-second segments. The F_1 is the larger the better while the error rate is the smaller the better; refer to [25] for detailed definitions of these evaluation metrics. To make binary predictions on the recording level, we thresholded the recording-level probabilities with class-specific thresholds, which were tuned to optimize the audio tagging F_1 on the validation set. To make binary predictions on 1-second segments, we first computed segment-level probabilities by aggregating the 10 frame-level probabilities within each segment, then thresholded the segment-level

	Max Pool.	Ave. Pool.	Lin. Soft.	Exp. Soft.	Attention
Audio Tagging					
TP	284	297	317	298	301
FN	322	309	289	308	305
FP	364	285	359	324	317
Precision	43.8	51.0	46.9	47.9	48.7
Recall	46.9	49.0	52.3	49.2	49.7
F_1	45.3	50.0	49.5	48.5	49.2
Localization					
TP	1,206	2,114	1,832	2,121	1,926
FN	3,154	2,246	2,528	2,239	2,434
FP	1,253	3,758	2,187	3,437	3,309
Precision	49.0	36.0	45.6	38.2	36.8
Recall	27.7	48.5	42.0	48.6	44.2
F_1	35.4	41.3	43.7	42.8	40.1
Localization					
Sub.	712	1,385	1,040	1,292	1,275
Del.	2,442	861	1,488	947	1,159
Ins.	541	2,373	1,147	2,145	2,034
Error Rate	84.7	105.9	84.3	100.6	102.5

Table 2. Detailed performance of the five systems on Task 4 of the DCASE 2017 challenge. Error rates and F_1 ’s are in percentages.

probabilities using the same class-specific thresholds.

Table 2 compares the performance of the five pooling functions. All the four new pooling functions outperform max pooling in terms of F_1 for both audio tagging and localization. In terms of localization error rate, however, only the linear softmax system slightly outperforms max pooling; the average pooling, exponential softmax and attention systems yield error rates over 100%.

Table 2 also includes a breakdown of the error types. All the five pooling functions achieve a reasonable balance between false negatives and false positives for audio tagging. However, the breakdown of errors for localization reveals that only the linear softmax pooling function maintains a good balance. As analyzed in Sec. 2.2, the max pooling system makes too many false negatives, which result in a low recall and a low F_1 ; the average, exponential softmax and attention pooling functions make too many false positives, which result in a high insertion rate and consequently a high error rate.

Fig. 3 illustrates the false positives made by the average pooling, exponential softmax and attention systems. The recording contains speech and bicycle noise, but no bus noise. On the recording level, all five systems correctly detect the bicycle sound and deny the existence of bus sounds. On the frame level, the max and linear softmax systems predict low probabilities that safely stay below the threshold for bus throughout the recording. In the average pooling and exponential softmax systems, however, some frame-level probabilities exceed the threshold, even though other frames with lower probabilities keep the recording-level probability under control. In the attention system, we see exactly what we have anticipated in Sec. 2.2: the attention (light blue line) mostly focuses on regions where the frame-level probabilities are low (8.2s). This

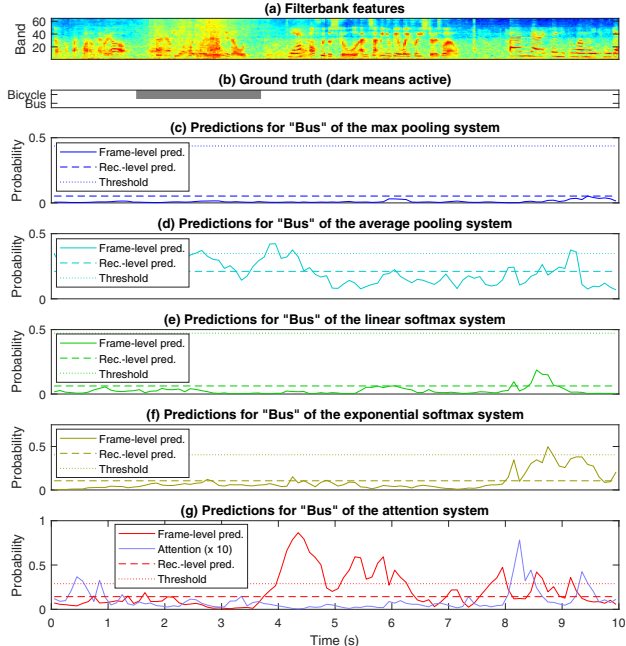


Fig. 3. The frame-level predictions of the five systems for the bus event on the test recording “-nqm_RJ2xj8” (unfortunately, this recording is no longer available on YouTube). Best viewed in color.

correctly produces a negative recording-level prediction, but lets many frame-level false positives (4~6s) get away unconstrained. The false positives shown in Fig. 3 are common throughout the data.

Considering the balance between false negatives and false positives for localization, as well as the agreement between recording-level and frame-level predictions, we recommend the linear softmax pooling function among all the pooling functions we have studied.

3.2. TALNet: Joint Tagging and Localization on Audio Set

We also compared the five pooling functions on the entire Audio Set [9]. This corpus provides a training set of over 2 million recordings, and an evaluation set of 20,371 recordings. The recordings in both sets are 10-second YouTube video excerpts, labeled with the presence or absence of 527 types of sound events.

We trained a CRNN with the structure shown in Fig. 2 (right). We name the network “TALNet”, where “TAL” stands for “tagging and localization”. The network has 10 convolutional layers, 5 pooling layers, and 1 recurrent layer. We applied data balancing during training; we also found it essential to apply batch normalization [26] before the ReLU activation of each convolutional layer.

Audio Set only provides evaluation metrics for audio tagging. These include the mean average precision (MAP), mean area under the curve (MAUC), and d-prime (d'); all these metrics are the larger the better. To evaluate the performance of localization, we applied TALNet to the DCASE 2017 challenge directly, and measured the same metrics as in Sec. 3.1. The results are listed in the top row of Table 3. Although the linear softmax system is not the best in terms of all the evaluation metrics, it is the only system that achieves a low error rate and a high F_1 for localization. The max pooling system falls behind on the F_1 , while the other three systems exhibit excessively high error rates.

In Table 3 we also list the results on Audio Set reported in all the literature that we can find. Our system closely matches the system of Yu *et al.* [23], and outperforms all other systems by a large margin.

System	# Train Recs.	Audio Set			DCASE 2017		
		MAP	MAUC	d'	Tag.F1	Loc.ER	Loc.F1
Max pooling	2M	0.351	0.961	2.497	52.6	81.5	42.2
Average pooling		0.361	0.966	2.574	53.8	101.8	46.8
Linear softmax		0.359	0.966	2.575	52.3	78.9	45.4
Exp. softmax		0.362	0.965	2.554	52.3	89.2	46.2
Attention		0.354	0.963	2.531	51.4	92.0	45.5
Hershey [27, 9]	1M	0.314	0.959	2.452			
Kumar [28]	22k	0.213	0.927				
Shah [18]	22k	0.229	0.927				
Wu [29]	22k		0.927				
Kong [22]	2M	0.327	0.965	2.558			
Yu [23]	2M	0.360	0.970	2.660			
Chen [30]	600k	0.316					
Chou [31]	1M	0.327	0.951				

Table 3. The performance of TALNet on both Audio Set and the DCASE 2017 challenge, compared with various systems in the literature (not all used the full training set). Bold font indicates the best performance in each group.

We would like to point out that the systems in the literature either do not perform localization well, or do not perform localization at all. For example, the system of Kong *et al.* [22] uses the attention pooling function. As we have demonstrated, this pooling function can cause many false positives on the frame level, and suffer from a high error rate. The system of Yu *et al.* [23] uses multi-level attention: attention layers are built upon multiple hidden layers, whose outputs are concatenated and further processed by a fully connected layer to yield a recording-level prediction. No frame-level predictions at all are made in this process. In contrast, our TALNet is the first system we know that achieves good performance for both audio tagging and localization at the same time.

4. CONCLUSION AND DISCUSSION

In this paper we have compared five pooling functions, and shown linear softmax to be the best among the five. The linear softmax pooling function has the following advantages: (1) it allows the gradient to flow unobstructedly; (2) it achieves a balance between false negatives and false positives for localization; (3) its predictions on the recording level and the frame level are relatively consistent. Using the linear softmax pooling function, we have built TALNet, which is the first network to achieve a strong performance for both audio tagging and localization at the same time. Our findings may not be limited to SED, but can apply generally to any MIL problem.

Nevertheless, linear softmax is by no means the ultimate optimal pooling function. An adaptive pooling function has been proposed in [32]; it gives a weight of $\exp(\alpha y_i)$ to the frame-probability y_i , and may be considered a generalization of the exponential softmax pooling function. Along this line, we may also consider a weighting scheme of $y_i^\beta \exp(\alpha y_i)$, which would subsume both the linear softmax and the exponential softmax pooling functions.

At the same time, the flexibility of the attention pooling function to learn the weights on the fly is still attractive, despite its excessive false positives on the frame level. We have found that the false positives are caused by the attention focusing on frames with low probabilities. We believe the attention pooling function is still promising if we could somehow impose a constraint of monotonicity that frames with larger probabilities must receive larger weights.

For more details about the experiments, such as the data balancing algorithm, how the class-specific thresholds were tuned, and the hyperparameters for training, please refer to Chapter 3 of the first author’s PhD thesis [33]. The code and acoustic features for the experiments are available at <https://github.com/MaigoAkisame/cmu-thesis>.

5. REFERENCES

- [1] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, V. Hazlewood, S. Lathrop, D. Lifka, G. D. Peterson, R. Roskies, J. R. Scott, and N. Wilkins-Diehr, "XSEDE: Accelerating scientific discovery," *Computing in Science & Engineering*, vol. 16, no. 5, pp. 62–74, 2014.
- [2] A. Gorin, N. Makhazhanov, and N. Shmyrev, "DCASE 2016 sound event detection system based on convolutional neural network," DCASE2016 Challenge, Tech. Rep., 2016.
- [3] M. Espi, M. Fujimoto, K. Kinoshita, and T. Nakatani, "Exploiting spectro-temporal locality in deep learning based acoustic event detection," *EURASIP Journal on Audio, Speech, and Music Processing*, 2015.
- [4] Y. Wang, L. Neves, and F. Metze, "Audio-based multimedia event detection using deep recurrent neural networks," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2016, pp. 2742–2746.
- [5] G. Parascandolo, H. Huttunen, and T. Virtanen, "Recurrent neural networks for polyphonic sound event detection in real life recordings," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2016, pp. 6440–6444.
- [6] S. Adavanne, G. Parascandolo, P. Pertilä, T. Heittola, and T. Virtanen, "Sound event detection in multichannel audio using spatial and harmonic features," in *Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, IEEE, 2016, pp. 6–10.
- [7] T. Hayashi, S. Watanabe, T. Toda, T. Hori, J. Le Roux, and K. Takeda, "Bidirectional LSTM-HMM hybrid system for polyphonic sound event detection," in *Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, IEEE, 2016, pp. 35–39.
- [8] E. Çakır, 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.
- [9] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio Set: An ontology and human-labeled dataset for audio events," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2017, pp. 776–780.
- [10] A. Mesaros, T. Heittola, A. Diment, B. Elizalde, A. Shah, E. Vincent, B. Raj, and T. Virtanen, "DCASE 2017 challenge setup: Tasks, datasets and baseline system," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE2017)*, 2017.
- [11] J. Amores, "Multiple instance classification: Review, taxonomy and comparative study," *Artificial Intelligence*, vol. 201, pp. 81–105, 2013.
- [12] 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 *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2017, pp. 791–795.
- [13] A. Kumar and B. Raj, "Audio event detection using weakly labeled data," in *Multimedia Conference*, ACM, 2016, pp. 1038–1047.
- [14] Y. Wang, J. Li, and F. Metze, "Comparing the max and noisy-or pooling functions in multiple instance learning for weakly supervised sequence learning tasks," in *Proceedings of Interspeech*, ISCA, 2018, pp. 1339–1343.
- [15] O. Maron and T. Lozano-Pérez, "A framework for multiple-instance learning," in *Advances in Neural Information Processing Systems (NIPS)*, 1998, pp. 570–576.
- [16] C. Zhang, J. C. Platt, and P. A. Viola, "Multiple instance boosting for object detection," in *Advances in Neural Information Processing Systems (NIPS)*, 2006, pp. 1417–1424.
- [17] B. Babenko, P. Dollár, Z. Tu, and S. Belongie, "Simultaneous learning and alignment: Multi-instance and multi-pose learning," in *Workshop on Faces in Real-Life Images: Detection, Alignment, and Recognition*, 2008.
- [18] A. Shah, A. Kumar, A. G. Hauptmann, and B. Raj, "A closer look at weak label learning for audio events," *arXiv e-prints*, 2018. [Online]. Available: <http://arxiv.org/abs/1804.09288>.
- [19] A. Dang, T. H. Vu, and J.-C. Wang, "Deep learning for DCASE2017 challenge," DCASE2017 Challenge, Tech. Rep., 2017.
- [20] J. Salamon, B. McFee, and P. Li, "DCASE 2017 submission: Multiple instance learning for sound event detection," DCASE2017 Challenge, Tech. Rep., 2017.
- [21] Y. Xu, Q. Kong, W. Wang, and M. D. Plumbley, "Large-scale weakly supervised audio classification using gated convolutional neural network," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2018, pp. 121–125.
- [22] Q. Kong, Y. Xu, W. Wang, and M. D. Plumbley, "Audio set classification with attention model: A probabilistic perspective," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2018, pp. 316–320.
- [23] C. Yu, K. S. Barsim, Q. Kong, and B. Yang, "Multi-level attention model for weakly supervised audio classification," *arXiv e-prints*, 2018. [Online]. Available: <http://arxiv.org/abs/1803.02353>.
- [24] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in PyTorch," in *NIPS Workshop*, 2017.
- [25] A. Mesaros, T. Heittola, and T. Virtanen, "Metrics for polyphonic sound event detection," *Applied Sciences*, vol. 6, no. 6, pp. 162–178, 2016.
- [26] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning (ICML)*, ACM, 2015, pp. 448–456.
- [27] S. Hershey *et al.*, "CNN architectures for large-scale audio classification," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2017, pp. 131–135.
- [28] A. Kumar, M. Khadkevich, and C. Fügen, "Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2018, pp. 326–330.
- [29] Y. Wu and T. Lee, "Reducing model complexity for DNN based large-scale audio classification," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2018, pp. 331–335.
- [30] S. Chen, J. Chen, Q. Jin, and A. Hauptmann, "Class-aware self-attention for audio event recognition," in *International Conference on Multimedia Retrieval (ICMR)*, ACM, 2018, pp. 28–36.
- [31] S.-Y. Chou, J.-S. R. Jang, and Y.-H. Yang, "Learning to recognize transient sound events using attentional supervision," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018, pp. 3336–3342.
- [32] B. McFee, J. Salamon, and J. P. Bello, "Adaptive pooling operators for weakly labeled sound event detection," *arXiv e-prints*, 2018. [Online]. Available: <http://arxiv.org/abs/1804.10070>.
- [33] Y. Wang, "Polyphonic sound event detection with weak labeling," PhD thesis, Carnegie Mellon University, 2018.