# LEARNING TO MATCH TRANSIENT SOUND EVENTS USING ATTENTIONAL SIMILARITY FOR FEW-SHOT SOUND RECOGNITION

Szu-Yu Chou<sup>1,2</sup>, Kai-Hsiang Cheng<sup>2</sup>, Jyh-Shing Roger Jang<sup>1</sup>, Yi-Hsuan Yang<sup>2</sup>

<sup>1</sup>Graduate Institute of Networking and Multimedia, National Taiwan University, Taipei, Taiwan <sup>2</sup>Research Center for Information Technology Innovation, Academia Sinica, Taipei, Taiwan

# ABSTRACT

In this paper, we introduce a novel attentional similarity module for the problem of few-shot sound recognition. Given a few examples of an unseen sound event, a classifier must be quickly adapted to recognize the new sound event without much fine-tuning. The proposed attentional similarity module can be plugged into any metric-based learning method for few-shot learning, allowing the resulting model to especially match related short sound events. Extensive experiments on two datasets show that the proposed module consistently improves the performance of five different metricbased learning methods for few-shot sound recognition. The relative improvement ranges from +4.1% to +7.7% for 5-shot 5-way accuracy for the ESC-50 dataset, and from +2.1% to +6.5% for noiseESC-50. Qualitative results demonstrate that our method contributes in particular to the recognition of transient sound events.

*Index Terms*— Few-shot learning, sound event detection, deep learning, transient sound event

# 1. INTRODUCTION

Understanding the surrounding environment through sounds has been considered as a major component in many daily applications, such as surveillance, smart city, and smart cars [1, 2]. Recent years have witnessed great progress in sound event detection and classification using deep learning techniques [3–7]. However, most prior arts rely on standard supervised learning algorithm and may not perform well for sound events with sparse training examples. While such a *few-shot learning* task has been increasingly studied in neighboring fields such as computer vision and natural language processing [8], little work if any has been done for few-shot sound event recognition, to our best knowledge.

Several metric-based learning methods for general fewshot learning have been proposed in the literature [8–11]. The *Matching Network* proposed by Vinyals *et al.* [8] uses the cosine similarity to measure the distance of a learned representation of the labeled set of examples with a giving unlabeled example for classification. Importantly, it proposes an *episode* procedure during training, which samples only a



**Fig. 1**: Illustration of using the proposed attentional similarity module (marked in red) for few-shot sound recognition. Compared to common similarity  $f_{sim}$ , the attentional similarity  $f_{att\_sim}$  performs better in matching transient (e.g., less than 1 second) sound events. Moreover, the attentional similarity module can be applied to any network for few-shot learning. The squares of different color denote different sound events. All the figures in this paper are best viewed in color.

few examples of each class as data points to simulate the fewshot learning scenario. Such a procedure allows the training phase to be close to the test phase in few-shot learning and accordingly improves the model's generalization ability. Snell *et al.* [11] follows the episodic procedure and proposes the *Prototypical Network.* They take the average of the learned representation of a few examples for each class as a classwise representation, and then classifies an input unlabeled datum by calculating the Euclidean distance between the input and the class-wise representations. Sung *et al.* [10] proposes the *Relation Network* to learn a non-linear distance metric for measuring the distance between the input unlabeled examples and a few examples from each class.

Most sound event recognition models are trained on datasets with clip-level labels, such as DCASE2016 [1], ESC-50 [12], and AudioSet [13]. Because the clip-level labels do not specify where the corresponding event actually takes place in an audio signal, this training strategy may make a

model overlook short or transient sound events [14]. To tackle this issue for few-shot sound recognition, in this paper, we propose a novel *attentional similarity* module to automatically guide the model to pay attention to specific segments of a long audio clip for recognizing relatively short or transient sound events. We show that our attentional similarity module can be learned relying on only clip-level annotation, and that it can be plugged into any existing methods to improve their performance for few-shot sound recognition.

# 2. APPROACH

## 2.1. Few-shot sound recognition

The goal of few-shot sound recognition is to learn a classifier which can quickly accommodate to unseen classes with only a few examples. In training, we are given a training set  $\mathcal{D} = \{(\chi, \mathbf{y}) \mid \chi \in \mathcal{S}, \mathbf{y} \in C_{train}\}$ , where  $\mathcal{S}$  is a support set, y denotes the class of the support set and  $C_{train}$  is the total number of classes in the training set. Generally, the support set  $\mathcal{S} = \{(s_i, ..., s_N) \mid s_i \in \{x_1, ..., x_{k \times c}, x_q\}\}$  consists of small support examples and a query example  $x_q$  (see Fig. 2), where x is the input feature and N is the number of support sets. The support examples are randomly sampled k examples from each of the c classes in the training set  $\mathcal{D}$ , and the query sample is randomly chosen from the remaining examples of c classes. The task is called c-way k-shot learning. And, k is often a small number from 1 to 5.

In our work, we use the simple yet powerful ConvNet architecture as our feature learning model  $f_{cnn}(\cdot)$ . The model can be learnt by minimizing the following objective function:

$$\min_{\theta} \sum_{(\chi, \mathbf{y}) \in \mathcal{D}} \mathcal{L}(f_{cnn}(\chi), \mathbf{y}) + \mathcal{R}(\theta), \qquad (1)$$

where  $\mathcal{L}$  is a loss function,  $\theta$  is the parameters of network and  $\mathcal{R}(\theta)$  is a regularization term for avoiding overfitting.

Similar to the sate-of-the-art algorithms [8,10,11] for fewshot learning, our loss function is based on the cross entropy:

$$\mathcal{L} = \frac{\exp(\sum_{\mathbf{X}_j \in S_y} f_{sim}(\mathbf{X}_q, \mathbf{X}_j))}{\sum_{j=1}^c \exp(\sum_{\mathbf{X}_j \in S_j} f_{sim}(\mathbf{X}_q, \mathbf{X}_j))}, \qquad (2)$$

where **X** is the output (a.k.a. a *feature map*) from the last convolutional layer of  $f_{cnn}$ ,  $S_y$  denotes the set of inputs labeled with class y and  $f_{sim}$  is the similarity function for measuring the distance between two inputs.

### 2.2. Attentional similarity

To deal with variable-length inputs, most approaches [8, 11, 15–17] use pooling functions to aggregate the feature maps  $\mathbf{X} \in \mathbb{R}^{M \times T}$  to yield a fixed-length vector  $\mathbb{R}^{M \times 1}$ , where M is the number of channels and T is the number of the temporal dimension. The similarity function can be written as:

$$f_{sim}(\mathbf{X}_q, \mathbf{X}_j) = dist(pool(\mathbf{X}_q), pool(\mathbf{X}_j)), \qquad (3)$$



**Fig. 2**: The training and test sets for few-shot sound recognition. The classes of test set are not seen during training.

where  $pool(\cdot)$  is the pooling function and  $dist(\cdot)$  any distance function between two vectors, such as the inner product.

Second-order similarity: A recent work on second-order similarity estimation [15] computes the segment-by-segment (second-order) similarity  $\mathbf{X}_q^T \mathbf{X}_j \in \mathbb{R}^{T_q \times T_j}$  between two inputs using the feature maps of the last layer of the ConvNet before pooling. Compared with the fixed-length vector (clip-level feature), this method allows the model to use segment-level feature to learn the temporal correlation between two inputs. The second-order similarity can be written as follows:

$$f_{order2\_sim}(\mathbf{X}_q, \mathbf{X}_j) = \mathbf{X}_q^T \mathbf{X}_j.$$
(4)

Being inspired by this method, we propose to learn a weight  $W_{qj} \in \mathbb{R}^{T_q \times T_j}$  to generate the attentional second-order similarity to capture the importance of segment-by-segment similarity. We can rewrite Eq. (4) as:

$$f_{order2\_wsim}(\mathbf{X}_q, \mathbf{X}_j) = \mathbf{X}_q^T \mathbf{X}_j W_{qj} \,.$$
 (5)

Following [18], we can compute attentional similarity by approximating weight  $W_{qj}$  as a rank-1 approximation,  $W = \mathbf{A}_j \mathbf{A}_q^T$  where  $\mathbf{A}_j, \mathbf{A}_q \in \mathbb{R}^{T_j \times 1}, \mathbb{R}^{T_q \times 1}$ . Then, we can derive the following *attentional similarity* function:

$$f_{att\_sim}(\mathbf{X}_q, \mathbf{X}_j) = Tr(\mathbf{X}_q^T \mathbf{X}_j \mathbf{A}_j \mathbf{A}_q^T), \qquad (6)$$

$$=Tr(\mathbf{A}_{q}^{T}\mathbf{X}_{q}^{T}\mathbf{X}_{j}\mathbf{A}_{j}), \qquad (7)$$

$$= \mathbf{A}_{q}^{T} (\mathbf{X}_{q}^{T} \mathbf{X}_{j}) \mathbf{A}_{j}, \qquad (8)$$

where  $Tr(\cdot)$  is the trace operator and  $\mathbf{A}_q$  is the attention vector computed by using another stack of convolutional layers  $f_{att}(\cdot)$  by feeding  $\mathbf{X}_q$  to find the important of segments. Eq. (8) uses the attention vector  $\mathbf{A}_q$  and  $\mathbf{A}_j$  to compute a weighted average of segment-by-segment similarity. More importantly, Eq. (8) can be rewritten as follows:

$$f_{att\_sim}(\mathbf{X}_q, \mathbf{X}_j) = \mathbf{A}_q^T(\mathbf{X}_q^T \mathbf{X}_j) \mathbf{A}_j, \qquad (9)$$

$$= (\mathbf{X}_q \mathbf{A}_q)^T (\mathbf{X}_j \mathbf{A}_j) \,. \tag{10}$$

Model	Attentional	Danth	th Param.	5-way Acc		10-way Acc	
	similarity	Depui		1-shot	5-shot	1-shot	5-shot
SIAMESE NETWORK [9]		3	1.18M	43.5%	50.9%	26.1%	31.6%
MATCHING NETWORK [8]		3	1.18M	53.7%	67.0%	34.5%	47.9%
Relation Network [10]		7	4.78M	60.0%	70.3%	41.7%	52.0%
SIMILARITY EMBEDDING NETWORK [15]		8	1.61M	61.0%	78.1%	45.2%	65.7%
PROTOTYPICAL NETWORK [11]		3	1.18M	67.9%	83.0%	46.2%	74.2%
SIAMESE NETWORK [9]	$\checkmark$	3+1	2.50M	49.3%	58.6%	29.0%	39.0%
MATCHING NETWORK [8]	$\checkmark$	3+1	2.50M	59.0%	74.0%	38.8%	55.3%
RELATION NETWORK [10]	$\checkmark$	7+1	6.11M	64.0%	74.4%	46.0%	57.0%
SIMILARITY EMBEDDING NETWORK [15]	$\checkmark$	8+1	3.40M	71.2%	82.0%	<b>56.9</b> %	71.0%
PROTOTYPICAL NETWORK [11]	$\checkmark$	3+1	2.50M	74.0%	87.7%	55.0%	76.5%

**Table 1**: The result of few-shot sound recognition (in %) on ESC-50. All the baselines reported here are based on our implementation. We indicate whether a method uses attention similarity, the network depth, and the number of model parameters.

Model	Attentional	Denth	Depth Param.	5-way Acc		10-way Acc	
	similarity	Depui		1-shot	5-shot	1-shot	5-shot
SIAMESE NETWORK [9]		3	1.18M	38.2%	43.5%	25.0%	28.0%
MATCHING NETWORK [8]		3	1.18M	51.0%	61.5%	31.7%	43.0%
RELATION NETWORK [10]		7	4.78M	56.2%	74.5%	39.2%	52.5%
SIMILARITY EMBEDDING NETWORK [15]		8	1.61M	63.2%	78.5%	44.2%	62.0%
PROTOTYPICAL NETWORK [11]		3	1.18M	66.2%	83.0%	46.5%	72.2%
SIAMESE NETWORK [9]	$\checkmark$	3+1	2.50M	46.0%	50.0%	29.0%	29.7%
MATCHING NETWORK [8]	$\checkmark$	3+1	2.50M	52.7%	66.5%	36.2%	48.2%
Relation Network [10]	$\checkmark$	7+1	6.11M	61.0%	76.2%	40.0%	59.2%
SIMILARITY EMBEDDING NETWORK [15]	$\checkmark$	8+1	3.40M	70.2%	83.2%	49.2%	67.2%
PROTOTYPICAL NETWORK [11]	$\checkmark$	3+1	2.50M	69.7%	85.7%	51.5%	73.5%

Table 2: The result of few-shot sound recognition performance (in %) on noiseESC-50. All our implementation.

The final equation can be interpreted as we compute the similarity score by using the inner product between two attentional vector  $\mathbf{X}_q \mathbf{A}_q$  and  $\mathbf{X}_j \mathbf{A}_j$ . This allows us to replace the inner product with common distance functions (e.g., cosine similarity or Euclidean distance) to measure the distance between two attentional vectors. So, in general the attentional similarity can also be computed as  $dist(\mathbf{X}_q \mathbf{A}_q, \mathbf{X}_j \mathbf{A}_j)$ .

## 3. EXPERIMENTS

## 3.1. Experimental Settings

**Dataset**: We conduct experiments on few-shot sound recognition using two datasets: ESC-50 and *noise*ESC-50. The ESC-50 dataset [12] contains 2,000 5-seconds audio clips labeled with 50 classes, each having 40 examples. The sound categories cover sounds of Animals, Natural, Human and Ambient noises. To evaluate the models under background noise conditions, following [2], we create the second dataset, coined *noise*ESC-50, by augmenting ESC-50 audio clips with additive background noise randomly selected from audio recordings of 15 different acoustic scenes from the DCASE2016 dataset [1]. Such a synthetic strategy allows to generate artificially noisy audio examples that reflect the sound recordings in everyday environment. Therefore, evaluation on *noise*ESC-50 may better measure how the models perform in real-world applications.

We note that, although the size and vocabulary of ESC-50 is small and limited, it is appropriate to use it as a public benchmark dataset for few-shot sound recognition. Larger benchmark datasets such as AudioSet [13] may suffer from the openness issue of audio data [19] and class imbalance problems [20].

**Data preparation**: In order to directly compare our model against strong baselines for few-shot learning, our experiment uses the similar splits proposed by [8]. The 50 sound event classes from ESC-50 dataset are divided into 35 classes for training and 10 classes for test. We train our model and baselines on 35 classes and use the remaining 5 validation classes for selecting the final model.

Feature extraction: To speed up model training, all audio clips from the ESC-50 and DCASE2016 datasets are down-sampled from 44.1 kHz to 16 kHz. We extract the 128-bin log mel-spectrogram from raw audio as the input feature to the neural networks. The librosa library [21] is used for



**Fig. 3**: Qualitative examples of few-shot sound recognition for 1-shot 5-way task on *noise*ESC-50. Note that the sound events in the test set are never seen during training. The examples marked with red and cyan colors are matched by using the prototypical network with and without attentional similarity, respectively. We manually add the bars in green and gray colors underneath the spectrograms to indicate parts of the audio clips that comprise sounds of interest and the background noise, respectively.

feature extraction. Before training a model, the input features are z-score normalized using the mean and standard deviation coefficients computed from the training set.

**Network design**: The backbone network is based on the simple yet powerful CNN structure, which has been widely used in audio tasks [16, 17]. The input feature of network is a mel-spectrograms, with 128 frequency bins and 160 frames. Our backbone network consists of a stack of blocks, each of which has a  $3 \times 3$  convolutional layer followed by batch normalization [22], a ReLU activation layer and a  $4 \times 4$  maxpooling layer. For optimization, we use stochastic gradient descent (SGD) and initial learning rate of 0.01. The learning rate is divided by 10 every 20 epochs for annealing, and we set the maximal number of epochs to 60. Moreover, we set the weight decay to 1e-4 to avoid overfitting.

Please note that all the baselines and our model are used the same backbone network. In our pilot studies, we have also explored using other advanced CNN structure such as ResNet [23] or VGG [24] as our backbone network but seen no much improvement. This may be due to the moderate size of the ESC-50 dataset. Code are available at https://github. com/kevinco27/attentional-similarity.

## 3.2. Experimental Results

Table 1 compares the performance of our own implementation of five metric-based learning methods for few-shot sound recognition with and without the proposed attentional similarity module on ESC-50. We can see that using attentional similarity clearly improves all the existing methods, giving rise to +4% to +7.1% relative improvement in 5-way 5-shot learning, a large performance gain. According to our experiment, the prototypical network turns out to be more effective and efficient than the relation network and the similarity embedding network for few-shot sound recognition. This seems to support similar findings in computer vision tasks [10, 11].

Table 2 shows the experimental result on *noise*ESC-50. Again, we see that the attentional similarity module consistently improves the result of existing methods, and the relative performance gain ranges from +2.1% to +6.5% in 5-way 5-shot learning. Compared with ESC-50, the performance gain decreases slightly, possibly because sound event recognition over *noise*ESC-50 is more challenging. Finally, similar to the case in ESC-50, the attentional similarity-empowered prototypical network achieves the best result among the evaluated methods by a great margin in 5-shot learning.

Figure 3 gives a qualitative comparison of the result of the prototypical network with and without attentional similarity for 3 query examples from the *noise*ESC-50 test set. Those picked by the model without attentional similarity (i.e., marked in cyan) do not share the same class as the queries; they are picked possibly because both the query and the picked one have a long silence. In contrast, the model with attentional similarity finds correct matches (marked in red).

## 4. CONCLUSION

We have introduced a simple module of attentional similarity for few-shot sound recognition to generate an attentional representation of inputs. It allows the model to ignore the unrelated background noise while matching relative short sound events. Extensive experiments show that attentional similarity consistently improves the performance of various existing methods on datasets of either noise-free or noisy clips. In the future, we plan to extend the model to adopt a multi-label learning setting for few-shot sound recognition.

### 5. REFERENCES

- Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen, "TUT database for acoustic scene classification and sound event detection," in *Proc. EUSIPCO*, 2016, pp. 1128–1132.
- [2] Annamaria Mesaros, Toni Heittola, Aleksandr Diment, Benjamin Elizalde, Ankit Shah, Emmanuel Vincent, Bhiksha Raj, and Tuomas Virtanen, "DCASE 2017 Challenge setup: Tasks, datasets and baseline system," in *Proc. DCASE*, 2017.
- [3] Jen-Yu Liu and Yi-Hsuan Yang, "Event localization in music auto-tagging," in *Proc. ACM MM*, 2016.
- [4] Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron Weiss, and Kevin Wilson, "CNN architectures for large-scale audio classification," in *Proc. ICASSP*. 2017.
- [5] Toan Vu, An Dang, and Jia-Ching Wang, "Deep learning for DCASE2017 Challenge," in *Proc. DCASE*, 2017.
- [6] Donmoon Lee, Subin Lee, Yoonchang Han, and Kyogu Lee, "Ensemble of convolutional neural networks for weakly-supervised sound event detection using multiple scale input," in *Proc. DCASE*, 2017.
- [7] Ting-Wei Su, Jen-Yu Liu, and Yi-Hsuan Yang, "Weaklysupervised audio event detection using event-specific gaussian filters and fully convolutional networks," in *Proc. ICASSP*, 2017.
- [8] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra, "Matching networks for one shot learning," in *Proc. NIPS*, pp. 3630– 3638. 2016.
- [9] Richard Zemel Gregory Koch and Ruslan Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *Proc. ICML Deep Learning workshop*, 2015.
- [10] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales, "Learning to compare: Relation network for few-shot learning," in *Proc. CVPR*, 2018.
- [11] Jack Snell, Kevin Swersky, and Richard Zemel, "Prototypical networks for few-shot learning," in *Proc. NIPS*, 2017.
- [12] Karol J. Piczak, "ESC: Dataset for Environmental Sound Classification," in *Proc. ACM MM*, 2015, pp. 1015–1018.

- [13] Jort F. Gemmeke, Daniel P. W. 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 *Proc. ICASSP*, 2017.
- [14] Szu-Yu Chou, Jyh-Shing Roger Jang, and Yi-Hsuan Yang, "Learning to recognize transient sound events using attentional supervision," in *Proc. IJCAI*, 2018, pp. 3336–3342.
- [15] Yu-Siang Huang, Szu-Yu Chou, and Yi-Hsuan Yang, "Generating music medleys via playing music puzzle games," in *Proc. AAAI*, 2018.
- [16] Aaron van den Oord, Sander Dieleman, and Benjamin Schrauwen, "Deep content-based music recommendation," in *Proc. NIPS*, 2013.
- [17] Anurag Kumar, Maksim Khadkevich, and Christian Fügen, "Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes," in *Proc. ICASSP*, 2018.
- [18] Rohit Girdhar and Deva Ramanan, "Attentional pooling for action recognition," in *Proc. NIPS*, 2017.
- [19] Eduardo Fonseca, Jordi Pons, Xavier Favory, Frederic Font, Dmitry Bogdanov, Andrés Ferraro, Sergio Oramas, Alastair Porter, and Xavier Serra, "Freesound datasets: A platform for the creation of open audio datasets," in *Proc. ISMIR*, 2017, pp. 486–493.
- [20] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang, "Learning deep representation for imbalanced classification," in *Proc. CVPR*, 2016.
- [21] Brian McFee, Colin Raffel, Dawen Liang, Daniel P. W. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nietok, "librosa: Audio and music signal analysis in Python," in *Proc. scipy*, 2015.
- [22] Sergey Ioffe and Christian Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML*, 2015.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016.
- [24] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.