

META LEARNING BASED AUDIO TAGGING

Kele Xu¹, Boqing Zhu¹, Dezhi Wang², Yuxing Peng¹, Huaimin Wang¹, Lilun Zhang², Bo Li³

¹ National University of Defense Technology, Computer Dept., Changsha, China,
kelele.xu@gmail.com, zhuboqing09@nudt.edu.cn, pengyuxing@aliyun.com, whm.w@163.com

² National University of Defense Technology, College of Meteorology and Oceanography,
Changsha, China,

wang_dezhi@hotmail.com, zll0434@163.com

³ Beijing University of Posts and Telecommunications, Automation Dept., Beijing, 100876, China,
deepblue.lb@gmail.com

ABSTRACT

In this paper, we describe our solution for the general-purpose audio tagging task, which belongs to one of the subtasks in the DCASE 2018 challenge. For the solution, we employed both deep learning methods and statistic features-based shallow architecture learners. For single model, different deep convolutional neural network architectures are tested with different kinds of input, which ranges from the raw-signal, log-scaled Mel-spectrograms (log Mel) to Mel Frequency Cepstral Coefficients (MFCC). For log Mel and MFCC, the delta and delta-delta information are also used to formulate three-channel features, while mixup is used for the data augmentation.

Using ResNeXt, our best single convolutional neural network architecture provides a mAP@3 of 0.967 on the public Kaggle leaderboard, 0.939 on the private leaderboard. Moreover, to improve the accuracy further, we also propose a meta learning-based ensemble method. By employing the diversities between different architectures, the meta learning-based model can provide higher prediction accuracy and robustness with comparison to the single model. Our solution achieves a mAP@3 of 0.977 on the public leaderboard and 0.951 as our best on the private leaderboard, while the baseline gives a mAP@3 of 0.704.

Index Terms— Audio tagging, convolutional neural networks, meta-learning, mixup

1. INTRODUCTION

With the increase of smart mobile devices in recent years, huge amounts of user generated sound recordings are uploaded to the web every day [1]. Thus, the demand for analyzing these audio signals is increasing dramatically, for example, audio scene classification [2], automatic audio tagging [3]. Indeed, audio tagging task, the problem of predicting the presence or absence of certain acoustic events in the acoustic scenes, has drawn lots of attention during the last several years, due to its widely applications.

Historically, audio tagging has been addressed with different handcrafted features and shallow-architecture classifiers. The classifiers include: GMMs, HMMs, NMF and SVMs [4, 5, 6, 7]. Since the developments are rapid in signal processing and machine learning domains, there is an increasing interest in applying deep learning approaches for the audio tagging task [8, 9, 10, 11]. However, it remains challenging and falls short of accuracy and efficiency, and

no reliable automatic general-purpose audio tagging systems exists. We argue that several factors lead to the phenomenon:

(1). Due to the lack of large-scale labeled data, the progress of audio tagging is far behind the analogous problem in the computer vision field. The audio-based approaches have been under-explored, and the state-of-the-art audio-based techniques are not able to achieve the comparable performance to its image/video counterpart. In fact, audios can sometimes be more descriptive than videos/images, especially when it comes to the description of an event.

(2). For the sound data, the number of sound events is huge, and the data quality is also of great diversity. Thus, it is important to handle the noisy training data, as the reliability of annotations is varying.

(3). Both shallow-architecture classifier using handcrafted features and deep learning approach should be employed together, which is under-explored. As demonstrated in many previous studies, efficient fusion between different models can boost the performance dramatically.

In this paper, we aim to build a general-purpose audio tagging system that can categorize an audio clip [12] as belonging to one of a set of 41 categories drawn from the AudioSet Ontology [13] (e.g., applause, bark, bus, animals, etc.). In more detail, our system has two levels: single deep model in the first level and the meta-learning in the second level. For single model in the first level, only convolutional neural networks are investigated, and different network architectures are tested with different kinds of input, which ranges from the raw-signal, log-scaled Mel-spectrograms (log Mel) to Mel Frequency Cepstral Coefficients (MFCC). For log Mel and MFCC, the delta and delta-delta information are also used to formulate three-channel features. Inception, ResNet, ResNeXt, Dual Path Networks (DPNs) are selected as the neural network architectures, while mixup is used for the data augmentation.

For the second level, to improve the classification further, we explore the use of meta-learning based method for the component classifier ensemble. Moreover, we propose to add the hand-crafted statistic features into the second level. In our experiments, this kind of ensemble method can provide superior accuracy and robustness. The paper is organized as follows. Section 2 gives the data augmentation method, while the brief introduction of the employed single models is presented in section 3. Section 4 describes the proposed meta-learning method and the brief experimental results.

2. DATA AUGMENTATION

The disadvantage of small dataset is that the model tends to overfitting. Currently, most publicly available audio tagging datasets have limited sizes [14]. To overcome this problem, we randomly extract a snippet of the original audio signal with equal length, 1.5 seconds. In this paper, we explore the use of mixup data augmentation [15]. In more detail, virtual training examples can be constructed by using the following formula:

$$x = \alpha \times x_i + (1 - \alpha) \times x_j \quad (1)$$

$$y = \alpha \times y_i + (1 - \alpha) \times y_j \quad (2)$$

where (x_i, y_i) and (x_j, y_j) are two examples randomly selected from the training batch. α is the mixed ratio. In our experiments, $\alpha \in \text{Beta}(3, 3)$. It is worthwhile to notice that the training samples can be either the raw wave signal segments or the time-frequency representations of the signal segment.

3. SINGLE MODEL

We used three kinds of inputs to train the network: raw wave signal, log-mel of the audio segment, and MFCC of the audio segment. In the papers [16, 17], we observed the complementarity of different features, so different features are used to improve performance. We select a 1.5s section randomly from the audio and input it into the network. The selected section is different in each epoch. When we take raw wave as input. We directly input $1.5 \times 44100 = 66150$ samples. When we take log-mel or MFCC as input, we extract a 64-dimensional log-mel and MFCC feature with a frame width of 80ms and a frame shift of 10ms, then we calculate the delta and delta-delta features of log-mel and MFCC with a window size of 9. Then we concatenate log-mel or MFCC with delta and delta-delta features to form a $3 \times 64 \times 150$ dimension input [18]. Two different ways are used to train the model: using ImageNet-based pre-trained model to initialize the weights, and training the weight from scratch. For the neural network architectures, 6 different model architectures are used.

3.1. Xception

Xception [19] is a deep convolutional neural network architecture inspired by Inception, where Inception modules have been replaced with depth-wise separable convolutions. In our experiments, Inception-V3 is employed with the log mel as input.

3.2. ResNet

ResNet makes the network deeper through a residual learning [20]. Instead of expecting each few stacked layers directly fit a desired underlying mapping, ResNet explicitly let these layers fit a residual mapping. Formally, denoting the desired underlying mapping as $H(x)$, the stacked nonlinear layers fit another mapping of $F(x) := H(x) - x$. The original mapping is recast into $F(x) + x$.

3.3. ResNeXt

ResNeXt [21] is a successful improvement based on ResNet. In more detail, ResNeXt is constructed by repeating a building block that aggregates a set of transformations with the same topology. Experiments on image classification demonstrate that increasing cardinality is a more effective way of gaining accuracy than going deeper

or wider, especially when depth and width starts to give diminishing returns for existing models. The cardinality and the width of bottleneck are chosen as 32 and 4 respectively.

3.4. SE-ResNeXt

By introducing a new architectural unit, which we term the Squeeze-and-Excitation (SE) block [22], networks could improve the representational power by explicitly modeling the interdependencies between the channels of its convolutional features. The SE block takes into account another relationship besides spatial relations: the channel relationship. It allows the network to perform feature recalibration, through which it can learn to use global information to selectively emphasize informative features and suppress less useful ones. We apply the SE block on the ResNeXt, which is denoted as SE-ResNeXt in Table 1.

3.5. Wave-ResNeXt

To process the raw wave-form, we use a one-dimensional convolution to simulate a band-pass filter to extract features. Moreover, in order to obtain more complementary features, we use multi-scale convolution to the original signal. Just like the multi-scale feature extraction process we designed in [17], the backend network is replaced by the ResNeXt.

3.6. DPN

Residual Network (ResNet) enables feature re-usage while Densely Convolutional Network (DenseNet) enables new features exploration which are both important for learning good representations. To enjoy the benefits from both path topologies, Dual Path Network (DPN) [23] shares common features while maintaining the flexibility to explore new features through dual path architectures.

The detail performances of different single models are given in Table 1. The number following the network name represents the number of layers in the network, for example ResNet50 means ResNet with the configuration of 50 layers. As can be seen from Table 1, we found ImageNet-based pre-trained model can improve the performances, but the generalization error between the train and test data is increased. Further, the Log-Mel feature can achieve better performance than waveform or MFCC generally. Moreover, it is worthwhile to notice that mixup can boost the performance without any exception.

4. META LEARNING-BASED ENSEMBLE AND EXPERIMENTAL RESULTS

It is widely known that ensemble diverse classifiers can improve the accuracy and robustness for the classification task. However, the ensemble learning has been under-explored for the audio tagging task. Previous efforts employ linear regression for the ensemble learning. Here, unlike previous attempts, we explore the use of stacked generalization in multiple levels to improve accuracy and robustness in this multi-class classification problem. The framework is computational, scalable and it have been tested on multiple machine learning tasks. Fig. 1 shows the proposed stacking architecture used in our task, which is composed of two levels. We randomly split the data into 5 folds in our experiments. For each CNN, we run 5 individual CNN models for each fold, and one model to predict the probabilities for each sample in the validating set by using the whole training

Table 1: Performance comparison between different single models

Network architecture	Pretraied	Input	Data augmentation	Public mAP@3	Private mAP@3
Wave-ResNext	Yes	wave	-	0.938	0.918
Wave-ResNext	No	wave	-	0.910	0.902
Xception	Yes	log mel	mixup	0.917	0.906
ResNet50	Yes	MFCC	mixup	0.932	0.914
ResNet50	Yes	log mel	mixup	0.950	0.932
ResNeXt101	Yes	log mel	-	0.935	0.924
ResNeXt101	Yes	log mel	mixup	0.967	0.939
ResNeXt101	No	log mel	mixup	0.921	0.887
SE-ResNeXt101	Yes	log mel	-	0.939	0.920
SE-ResNeXt101	Yes	log mel	mixup	0.950	0.927
DPN68	Yes	log mel	-	0.939	0.925
DPN68	Yes	log mel	mixup	0.950	0.922
DPN96	Yes	log mel	-	0.937	0.926
DPN96	Yes	log mel	mixup	0.964	0.936
DPN96	No	log mel	mixup	0.917	0.886
DPN107	Yes	log mel	mixup	0.957	0.938

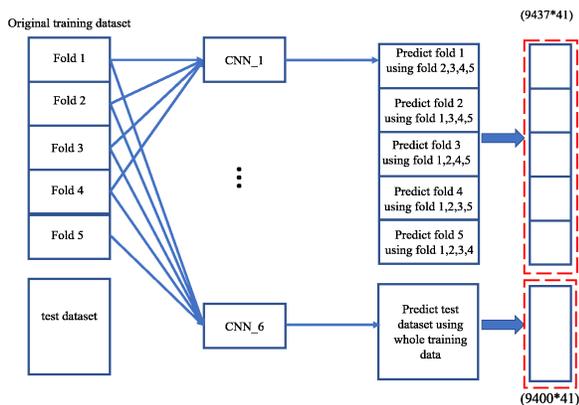


Figure 1: Meta-features construction using CNNs.

dataset. The predicted probabilities of different classes will be concatenated to generate meta-features. For each classifier, the probabilities for 41 classes will be used as the meta-features, which will be concatenated to generate the new training dataset (as can be seen in Fig.1), and the meta features will be used as the input for level 2.

In our experiments, the first layers are composed of 5 different CNN architectures: ResNeXt using the log mel with mixup, ResNeXt using the raw wave without mixup; ResNet using the log mel without mixup, ResNet using the wave with mixup, DPN using the log mel with mixup (as can be seen from Fig.2).

Except for the deep learning-based meta features, we also employ the traditional handcrafted features. In more detail, we calculate the max value, min value, variance value, skewness for the MFCC of the audio signal segment. And the statistical features are also used as the meta-features.

For level 2, we employ the widely used method gradient tree boosting machine for the multi-class classification task. The eXtreme Gradient Boosting method (XGBoost) [24, 25] library, a tree boosting machine based classification implementation, is selected as the benchmark. It is because that compared to other approaches (such as, linear regression, Support Vector Machine, Random For-

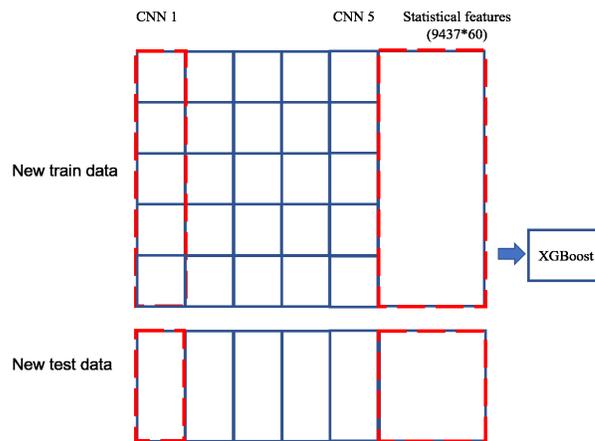


Figure 2: Meta-learning for audio tagging.

est [26]), XGBoost provides better classification performance in our experiment, and its power has furthermore been validated on several public machine learning challenges. We use the default hyper parameters for the XGBoost, and the maximum depth is set to 3 to prevent overfitting.

Using the proposed meta-learning method, our solution achieves a mAP@3 of 0.977 on the public leaderboard and 0.951 as our best on the private leaderboard, while the baseline gives a mAP@3 of 0.70.

5. CONCLUSION

In this article, we proposed an effective meta-learning system employing both deep learning architectures and statistic features-based learners to achieve a successful solution for the general-purpose audio tagging task in DCASE 2018. A comparative study of the performance of several well-developed convolutional neural network architectures with different types of input were conducted to obtain excellent single models for the subsequent meta-learning. Mixup

technique was also implemented in the training process which constantly improves the model performance as expected. The XGBoost approach was applied on a hybrid combination of meta-features including deep-learning features and statistical features, which have a superb classification performance. The final results put us in the first place on the task public leaderboard with a mAP@3 of 0.977 and the fourth place on the private leaderboard. In future, we would like to further evaluate the performance of our method on the Google AudioSet.

6. REFERENCES

- [1] E. Marchi, D. Tonelli, X. Xu, F. Ringeval, J. Deng, S. Squartini, and B. Schuller, "Pairwise decomposition with deep neural networks and multiscale kernel subspace learning for acoustic scene classification," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2016 Workshop (DCASE2016)*, 2016, pp. 65–69.
- [2] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," 2018, submitted to DCASE2018 Workshop. [Online]. Available: <https://arxiv.org/abs/1807.09840>
- [3] E. Fonseca, M. Plakal, F. Font, D. P. W. Ellis, X. Favory, J. Pons, and X. Serra, "General-purpose tagging of freesound audio with audioset labels: Task description, dataset, and baseline," 2018, submitted to DCASE2018 Workshop. [Online]. Available: <https://arxiv.org/abs/1807.09902>
- [4] P. Dhanalakshmi, S. Palanivel, and V. Ramalingam, "Classification of audio signals using aann and gmm," *Applied soft computing*, vol. 11, no. 1, pp. 716–723, 2011.
- [5] I. Karpov and D. Subramanian, "Hidden markov classification for musical genres," *Course Project*, 2002.
- [6] P. Giannoulis, G. Potamianos, and P. Maragos, "On the joint use of nmf and classification for overlapping acoustic event detection," in *Multidisciplinary Digital Publishing Institute Proceedings*, vol. 2, no. 2, 2018, p. 90.
- [7] K. 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.
- [8] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, *et al.*, "Cnn architectures for large-scale audio classification," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 131–135.
- [9] E. Fonseca, R. Gong, and X. Serra, "A simple fusion of deep and shallow learning for acoustic scene classification," *arXiv preprint arXiv:1806.07506*, 2018.
- [10] Y. Xu, Q. Huang, W. Wang, P. Foster, S. Sigtia, P. J. Jackson, and M. D. Plumbley, "Unsupervised feature learning based on deep models for environmental audio tagging," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1230–1241, 2017.
- [11] T. Hayashi, S. Watanabe, T. Toda, T. Hori, J. Le Roux, K. Takeda, T. Hayashi, S. Watanabe, T. Toda, T. Hori, *et al.*, "Duration-controlled lstm for polyphonic sound event detection," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 25, no. 11, pp. 2059–2070, 2017.
- [12] F. Font, G. Roma, and X. Serra, "Freesound technical demo," in *Proceedings of the 21st ACM international conference on Multimedia*. ACM, 2013, pp. 411–412.
- [13] J. F. Gemmeke, D. P. 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 *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 776–780.
- [14] J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," in *ACM International Conference on Multimedia*, 2014, pp. 1041–1044.
- [15] K. Xu, D. Feng, H. Mi, B. Zhu, D. Wang, L. Zhang, H. Cai, and S. Liu, "Mixup-based acoustic scene classification using multi-channel convolutional neural network," *arXiv preprint arXiv:1805.07319*, 2018.
- [16] Y. Tokozume and T. Harada, "Learning environmental sounds with end-to-end convolutional neural network," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 2721–2725.
- [17] B. Zhu, C. Wang, F. Liu, J. Lei, Z. Lu, and Y. Peng, "Learning environmental sounds with multi-scale convolutional neural network," *arXiv preprint arXiv:1803.10219*, 2018.
- [18] D. Feng, K. Xu, H. Mi, F. Liao, and Y. Zhou, "Sample dropout for audio scene classification using multi-scale dense connected convolutional neural network," *arXiv preprint arXiv:1806.04422*, 2018.
- [19] J. Carreira, H. Madeira, and J. G. Silva, "Xception: A technique for the experimental evaluation of dependability in modern computers," *IEEE Transactions on Software Engineering*, vol. 24, no. 2, pp. 125–136, 1998.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [21] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*. IEEE, 2017, pp. 5987–5995.
- [22] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," *arXiv preprint arXiv:1709.01507*, vol. 7, 2017.
- [23] Y. Chen, J. Li, H. Xiao, X. Jin, S. Yan, and J. Feng, "Dual path networks," in *Advances in Neural Information Processing Systems*, 2017, pp. 4467–4475.
- [24] T. Chen, T. He, M. Benesty, *et al.*, "Xgboost: extreme gradient boosting," *R package version 0.4-2*, pp. 1–4, 2015.
- [25] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. ACM, 2016, pp. 785–794.
- [26] A. Liaw, M. Wiener, *et al.*, "Classification and regression by randomforest," *R news*, vol. 2, no. 3, pp. 18–22, 2002.