

TRANSFERRING VOICE KNOWLEDGE FOR ACOUSTIC EVENT DETECTION: AN EMPIRICAL STUDY

Dawei Liang^{**}, Yangyang Shi[†], Yun Wang[†], Nayan Singhal[†], Alex Xiao[†],
Jonathan Shaw[†], Edison Thomaz^{*}, Ozlem Kalinli[†], Mike Seltzer[†]

^{*} University of Texas at Austin, USA

[†] Facebook AI, USA

ABSTRACT

Detection of common events and scenes from audio is useful for extracting and understanding human contexts in daily life. Prior studies have shown that leveraging knowledge from a relevant domain is beneficial for a target acoustic event detection (AED) process. Inspired by the observation that many human-centered acoustic events in daily life involve voice elements, this paper investigates the potential of transferring high-level voice representations extracted from a public speaker dataset to enrich an AED pipeline. Towards this end, we develop a dual-branch neural network architecture for the joint learning of voice and acoustic features during an AED process and conduct thorough empirical studies to examine the performance on the public AudioSet [1] with different types of inputs. Our main observations are that: 1) Joint learning of audio and voice inputs improves the AED performance (mean average precision) for both a CNN baseline (0.292 vs 0.134 mAP) and a TALNet [2] baseline (0.361 vs 0.351 mAP); 2) Augmenting the extra voice features is critical to maximize the model performance with dual inputs.

Index Terms— Acoustic event detection, transfer learning, feature fusion, data augmentation, speaker recognition

1. INTRODUCTION

With the development of modern smart devices with listening capabilities such as voice assistants, smartphones, and wearable devices, audio has been increasingly used as a modality for the inference of human activities, and contexts [3, 4]. Acoustic event detection (AED) is the process of detecting the type and temporal onset/offset of acoustic events within an audio stream. While existing AED models have been advanced for modeling the target audio, studies have shown that knowledge transfer from a relevant domain is beneficial to boost the learning process further and the model capabilities [5, 6]. For example, knowledge transfer is especially useful when data in the target domain is not sufficient for model generalization [7], which is a typical case when dealing with real-world audio.

Many acoustic event types captured in daily life are related to human voice or contain voice elements, such as conversations, TV/radio sounds, music, or sounds from a crowd. To the best of our knowledge, however, very few prior attempts have studied the opportunities of transferring and incorporating voice knowledge into a general AED process. This paper investigates this opportunity by jointly training audio and high-level voice representations for an AED model based on dual-branch neural network architecture. We observe the benefits of adding extra voice inputs to a convolutional neural network (CNN) and a TALNet [2] baseline by using AudioSet [1] as the test dataset. Specifically, our study demonstrates a few strategies that bridge the learning gap between the audio and the voice features.

2. RELATED WORK

The general process of AED is to build a classification model where the existence of an acoustic class is determined by the output class probability from the model. Conventional approaches include the usage of statistical models based on hand-crafted features [8, 9]. Recent work has increasingly focused on using neural networks for the modeling of audio [10, 11, 12, 13] given their success in computer vision.

Due to the increasing scale of audio data exposure, modern audio datasets are typically weak-labeled by annotators. The dataset only gives the label for the whole recording without detailed frame-level annotation. However, in practice, frame-level labeling is required. The following work [2, 14, 15, 16] has addressed the issue. The TALNet [2] is one of the state-of-the-art efforts for AED with weakly labeled audio inputs, which has demonstrated strong performance for acoustic event tagging and localization at the same time.

Transferring knowledge from a source domain to a target task can be a useful way to enrich the learning of the target dataset [7]. It is particularly meaningful in real-world audio analysis where the target audio accessibility can be limited due to challenges such as scalability [17] and privacy constraints [18]. For acoustic classification, transfer learning has been successfully applied both across tasks [19, 20, 21, 22, 23] and across modalities [5, 24]. Specifically, the extrac-

^{*}Work performed at Facebook.

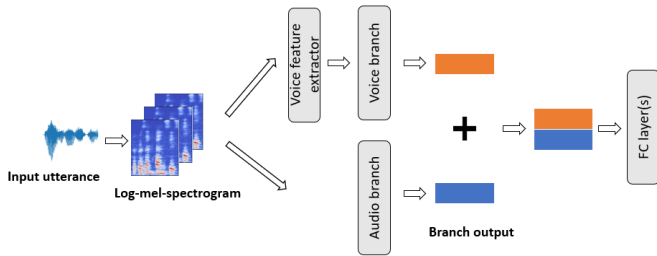


Fig. 1. Overall pipeline of our study. The acoustic classifier is a neural network with two input branches, taking as input the log-mel features and extracted voice embeddings of an utterance respectively.

tion and leveraging of pre-trained neural network embeddings is a common way of audio knowledge transfer [2, 19, 25]. Compared to conventional hand-crafted voice features such as i-vectors or the mel features [26, 27, 28, 29], voice embeddings are directly obtained from a neural network trained for speaker voice classification. The voice embedding represents the knowledge that the network has learned to identify the speaker patterns [30, 31, 32, 33]. Our study aims to leverage voice embeddings extracted from an existing speaker dataset to enrich the AED process. As far as we know, this is the first effort to incorporate knowledge from voice inputs for AED on the AudioSet corpus.

3. ARCHITECTURE

3.1. Overall pipeline

Fig. 1 shows the overall pipeline of our study. The pipeline consists of two steps - feature extraction and acoustic event detection. The first step extracts the log-mel features of an audio input utterance. In addition to leveraging the log-mel features as the AED inputs, we apply an extra pre-trained model to extract voice embedding representations from the audio. Applying voice embedding transfers pre-trained voice knowledge of the feature extractor to the target audio, and it applies to both vocal and non-vocal input. In the second step, the acoustic classifier is a neural network architecture with two input branches: an audio branch with log-mel inputs and a voice branch with the generated voice feature inputs. The outputs of both branches are concatenated along the feature dimension and fed to the final fully connected layer(s). We did not apply early fusion of the features, because feature fusion at an intermediate layer left us more flexibility to optimize the two input branches separately [34]. In our study, the voice feature extractor was trained with an existing speaker dataset. Once pre-trained, the parameters of the feature extractor were fixed, and only the audio and voice branches were trained for the target AED task.




Fig. 2. Architecture of the voice branch. It consists of 1D convolutional layers on the time (t) dimension followed by a uni-directional GRU layer.

3.2. Audio branch architecture

The audio branch was developed for the log-mel inputs. In our study, we started with a shallow CNN baseline and then the TALNet. The CNN architecture is as follows:

Input \rightarrow **Conv1[64]** \rightarrow **Conv2[128]** \rightarrow **Conv3[256]** \rightarrow **Conv4[256]** \rightarrow **FC[2048]** \rightarrow **FC[1024]** \rightarrow **FC[527]**

where ConvX[K] denotes a 2D convolutional layer with the ReLU [35] activation and K channels. The kernel size, padding, and stride were (3×3) , (1×1) , and (1×1) . Besides, a max-pooling of size (2×2) , (2×2) , and (1×2) was added for Conv1, Conv2 and Conv4, respectively. FC[K] denotes a fully connected layer of size K with the ReLU activation. We adopted the model architectures for both baselines, excluding their fully connected layer(s) as our audio branch. We then added back the fully connected layer(s) after feature fusion.

3.3. Voice branch architecture

Unlike the audio branch, we did not apply convolution on the feature dimension of the voice inputs since adjacent elements of an embedding may not have spatial correlation as the log-mel vectors do. Hence, our voice branch consists of 1D convolutional layers along the temporal dimension of the voice embeddings. The feature dimension of the embeddings is mapped to the channel dimension of the convolutional layers. Fig. 2 shows such a process. In such a design, reducing the number of network channels of each convolutional layer essentially reduces the size of the feature dimension. We added an extra uni-directional GRU layer following the convolutional layers to improve the learning performance.

4. VOICE REPRESENTATIONS

4.1. Pre-training of voice models

To develop the voice feature extractor, we constructed a speaker recognition task where a network was trained to classify given speaker voice classes. The task was built on the public VoxCeleb1 [33] speaker dataset. The dataset consists of audio utterances of over 1K celebrities from public YouTube videos of varying lengths. Specifically, we leveraged 1,211 speakers in the dataset for our model training and

Arch1	Arch2
conv2D (96, 3×3, 1, 1)	conv2D (32, 3×3, 1, 1)
mpool (2×2)	mpool (2×2)
conv2D (256, 3×3, 1, 1)	conv2D (32, 3×3, 1, 1)
mpool (2×2)	conv2D (64, 3×3, 1, 1)
conv2D (384, 3×3, 1, 1)	mpool (2×2)
conv2D×2 (256, 3×3, 1, 1)	conv2D (64, 3×3, 1, 1)
mpool (1×2)	flatten → (batch×100×1024)
conv2D (1024, 1×8, 1, 0)	biGRU (512×2)
fc×2 (1024 / 1211)	fc (1211)

Table 1. Architecture of our voice models. conv2D: 2D convolutional layers (channel, kernel size, stride, padding); mpool: max pooling; fc: fully-connected layers.

validation. For each speaker, ten utterances were randomly selected for model validation, and the rest were used for training, resulting in an average of 109 utterances per speaker in our training set.

The audio utterances were sampled at 16kHz and truncated or padded to 10 seconds. We then extracted 64D log-mel features using a frame length of 64ms and a frame shift of 25ms. The resulting log-mel features of a minibatch of input to our voice models had a shape of (batch×1×400×64). The features were normalized per dimension by the mean and standard deviation calculated over the entire training set. Two voice models were then applied, inspired from [33] (*Arch1*) and a simplified version of TALNet (*Arch2*). We slightly modified the model parameters to fit our requirements, as shown in Table 1. The models were deployed using PyTorch [36]. ReLU activation and batch normalization were applied except for the last fc layer. Besides, the “ceil” mode was enabled for the max-pooling layers. For Arch1, the output of the last conv2D layer was re-shaped from (batch×1024×100×1) to (batch×100×1024). The temporal dimension (100) was then aggregated by average pooling. For Arch2, the temporal dimension was aggregated after the fc layer.

To set up training, we used an initial learning rate of 2×10^{-4} and a batch size of 25. We applied the same class balancing strategy as in [2] to account for class imbalance. The learning rate was shrunk by a factor of 0.9 when the validation accuracy plateaued, and training was stopped when the learning rate reached 1×10^{-6} . We used the Adam [37] optimizer and the cross-entropy loss.

4.2. Voice feature extraction

The best validation accuracy we obtained was 95.8% and 93.4% for Arch1 and Arch2, respectively. The high accuracy values were as expected because of the large utterance size we used. The voice embeddings were then extracted before average pooling for Arch1 and following the GRU layer for Arch2, which maintains the temporal information of the input utterance at a resolution of 10Hz. The resulting embeddings

were then of shape (batch×100×1024) for a batch. In the following sections, we will refer to the embeddings from the two models as *emb_1* and *emb_2* for convenience.

5. EXPERIMENTS

5.1. Training setup

We leveraged AudioSet for our AED study. The dataset consists of over 2 million 10-second audio utterances of 527 annotated acoustic classes, including vocal and non-vocal sounds extracted from public YouTube videos. We used the same evaluation set containing 24,832 utterances.

We followed our speaker recognition steps to derive the same type of log-mel features as inputs to the AED part. We then tested the two baselines of the audio branch independently. In deployment, we removed Conv4 of the CNN baseline for our dual-branch tests to maintain a similar model size with and without voice inputs. Besides, we used a hidden size of 768 for the GRU layer of TALNet. For both baselines, we enabled the “ceil” mode of the max-pooling layers in PyTorch. The output of the audio branch was consistently in a shape of (batch×100×768), where the temporal size was 100 (0.1s resolution).

For the voice branch, we applied two 1D convolutional layers so that the parameter size of the voice branch could be less considerable (<1M) compared to the baseline models. The kernel size, padding, and stride were 3, 1, and 1, respectively, with ReLU activation. Batch normalization was also added before the activation. In the convolutional layers, the number of channels was 256 and 64, respectively. For the GRU layer, the hidden size was 64. Hence, the outputs of the voice branch were of shape (batch×100×64).

The final shape of features was (batch×100×832) after fusion. The last fully connected layer was of size 527 with the sigmoid activation. The per-class binary predictions were aggregated for an utterance by linear softmax pooling on the frame-level probability outputs. We used the same learning rate and optimization setup for training as in the speaker recognition task, but the validation metric was switched to the mean average precision (mAP) score which AudioSet also used. Besides, we used the binary cross-entropy loss.

5.2. Experiments and result discussions

We first examined the maximum performance of the models with augmented voice inputs. Inspired by common augmentation strategies for acoustic features, we processed the voice embeddings with three strategies – time masking [38], mixup [39], and adding dropout [40] to the voice branch. Specifically, we randomly masked 40 voice embedding frames out of the total 100 for each input utterance. The mixup was applied by mixing up labels and the corresponding input features at a batch level. We applied this for both branches at the same time with an alpha value of 1. Besides, we dropped out the voice features with a probability of 0.5.

Combination	mAP	mAUC	d -prime
CNN	0.134	0.903	1.840
CNN+emb1	0.292	0.951	2.338
CNN+emb2	0.256	0.950	2.325
TALNet	0.351	0.966	2.584
TALNet+emb1	0.360	0.962	2.506
TALNet+emb2	0.361	0.962	2.517

Table 2. Overall AED results with different input types and base architectures of the audio branch.

Table 2 shows the overall AED results based on combinations of the input types and the audio branch architectures. In addition to the best mAP values, we also reported the best mean area under the curve (mAUC) and d -prime metrics to show the performance better. The results with both baselines show that adding the extra voice branch improves the AED performance regarding the mAP metric, and the observation was consistent for both types of voice inputs. As expected, the performance jump was much bigger for the CNN baseline than for TALNet, since TALNet is a more sophisticated model for the AudioSet AED task, and the supplementary information carried by the voice inputs can be relatively marginal. Interestingly, the mAUC and d -prime metrics degraded for TALNet with the voice inputs, possibly because the training was only optimized for the validation mAP. We also tried doubling the number of convolutional layers for the voice branch (with 1/2 of the channel size per layer to maintain the same output shape). However, no improvement was observed for either type of embeddings (0.357 mAP for both emb1 and emb2), indicating that 2 layers of CNN were sufficient for learning from the high-level voice features.

We then studied the AED effects with and without augmentation applied to the voice inputs. Table 3 shows the best validation epochs, the corresponding loss values and the mAP for the test cases. From the comparisons, it can be clearly seen that adding augmentation is critical to improve joint training of the two branches and maximize the model performance. Table 4 further demonstrates the effects of individual types of augmenting strategies with the TALNet-based audio branch. While all techniques improved the joint training process, we can see that mixup augmentation yielded the biggest improvement. For a better comparison, we also ran a test for the original TALNet baseline with the same setup of mixup augmentation on the log-mel inputs. However, no improvement was observed in our test (0.342 mAP), indicating that the augmentation process contributed more significantly with the extra voice inputs than with the original baseline model.

By checking the class-wise predictions of the TALNet-based models, we found that the predicting performance led most (≥ 0.06 mAP) with the voice embeddings compared to the baseline model for the 10 classes: *Hoot*, *Dental-drill*, *Yodeling*, *Air-horn*, *Gobble*, *Roar*, *Ringtone*, *Tele-*

Combination	Epoch	Train / Val loss	mAP
CNN+emb1, no aug	11	9.7 / 13.3	0.229
CNN+emb1, with aug	39	12.7 / 12.2	0.264
CNN+emb2, no aug	13	8.9 / 12.8	0.256
CNN+emb2, with aug	45	12.6 / 11.7	0.292
TALNet+emb1, no aug	25	6.7 / 11.8	0.328
TALNet+emb1, with aug	56	11.4 / 10.6	0.360
TALNet+emb2, no aug	19	7.6 / 11.5	0.331
TALNet+emb2, with aug	71	11.4 / 10.6	0.361

Table 3. Results with and without augmentation on the voice features. Feature augmentation is a critical factor for better joint training of the two input branches.

Strategy	Epoch	Train / Val loss	mAP
emb1+dp	21	7.1 / 11.4	0.334
emb1+tmask	25	6.8 / 11.7	0.331
emb1+mixup	62	10.6 / 10.8	0.350
emb2+dp	20	7.3 / 11.4	0.334
emb2+tmask	19	7.8 / 11.5	0.334
emb2+mixup	68	11.6 / 10.7	0.356

Table 4. Results with a single type of augmenting strategies. *dp*: dropout; *tmask*: time-mask augmentation; *mixup*: mixup augmentation. Mixup augmentation is the most effective approach to improve the training.

vision, *Snoring*, *Video_game_music*. An interesting finding was that adding the extra voice inputs did not guarantee a better prediction for the voice-related AudioSet classes. For example, the mAP for *Speech* was 0.776 / 0.778 / 0.782 with the original TALNet / TALNet+emb1 / TALNet+emb2, but the values were 0.320 / 0.314 / 0.336 for *Singing* and 0.209 / 0.186 / 0.206 for *Conversation*. A possible explanation is that the training of the two branches was an integrated process and the models were optimized for the global performance, sometimes at a compromise of such individual classes.

6. CONCLUSIONS

This paper explored a novel approach for acoustic event detection by incorporating pre-trained voice embeddings into an AED pipeline. Towards this end, we developed a dual-branch neural network architecture for joint training of the inputs. We then reported the overall and class-wise performance with a CNN baseline and a strong TALNet baseline developed on AudioSet. Our results showed the benefits of adding extra voice inputs to the tested models (0.292 vs 0.134 mAP for the CNN baseline and 0.361 vs 0.351 mAP for TALNet baseline). Furthermore, we showed that adding augmentation and dropout on the voice inputs is critical to maximize the model performance with dual inputs.

7. REFERENCES

- [1] Jort F Gemmeke et al., “Audio set: An ontology and human-labeled dataset for audio events,” in *ICASSP*, 2017.
- [2] 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.
- [3] Nicholas D Lane, Petko Georgiev, and Lorena Qendro, “Deep-ear: robust smartphone audio sensing in unconstrained acoustic environments using deep learning,” in *UbiComp*, 2015.
- [4] Dawei Liang and Edison Thomaz, “Audio-based activities of daily living (adl) recognition with large-scale acoustic embeddings from online videos,” *IMWUT*, 2019.
- [5] Yusuf Aytar, Carl Vondrick, et al., “Soundnet: Learning sound representations from unlabeled video,” *NIPS*, 2016.
- [6] Qiuqiang Kong, Yin Cao, et al., “Panns: Large-scale pre-trained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [7] Chuanqi Tan et al., “A survey on deep transfer learning,” in *ICANN*, 2018.
- [8] Lode Vuegen, BVD Broeck, et al., “An mfcc-gmm approach for event detection and classification,” in *WASPAA*, 2013.
- [9] Burak Uzkent, Buket D Barkana, et al., “Non-speech environmental sound classification using svms with a new set of features,” *International Journal of Innovative Computing, Information and Control*, vol. 8, no. 5, pp. 3511–3524, 2012.
- [10] Arseniy Gorin, Nurtas Makhzhanov, and Nickolay Shmyrev, “Dcase 2016 sound event detection system based on convolutional neural network,” *IEEE AASP Challenge: Detection and Classification of Acoustic Scenes and Events*, 2016.
- [11] Miquel Espi et al., “Exploiting spectro-temporal locality in deep learning based acoustic event detection,” *EURASIP Journal on Audio, Speech, and Music Processing*, 2015.
- [12] Yun Wang, Leonardo Neves, and Florian Metze, “Audio-based multimedia event detection using deep recurrent neural networks,” in *ICASSP*. IEEE, 2016, pp. 2742–2746.
- [13] Giambattista Parascandolo, Heikki Huttunen, and Tuomas Virtanen, “Recurrent neural networks for polyphonic sound event detection in real life recordings,” in *ICASSP*, 2016.
- [14] Changsong Yu et al., “Multi-level attention model for weakly supervised audio classification,” *arXiv:1803.02353*, 2018.
- [15] Qiuqiang Kong et al., “Audio set classification with attention model: A probabilistic perspective,” in *ICASSP*, 2018.
- [16] Szu-Yu Chou, Jyh-Shing Roger Jang, and Yi-Hsuan Yang, “Learning to recognize transient sound events using attentional supervision,” in *IJCAI*, 2018.
- [17] Yohan Chon et al., “Understanding the coverage and scalability of place-centric crowdsensing,” in *UbiComp*, 2013.
- [18] Dawei Liang, Wenting Song, and Edison Thomaz, “Characterizing the effect of audio degradation on privacy perception and inference performance in audio-based human activity recognition,” in *MobileHCI*, 2020.
- [19] Qiuqiang Kong, Yin Cao, et al., “Cross-task learning for audio tagging, sound event detection and spatial localization: Dcase 2019 baseline systems,” *arXiv:1904.03476*, 2019.
- [20] Vipula Dissanayake et al., “Speech emotion recognition in the wild using an autoencoder,” in *INTERSPEECH*, 2020.
- [21] Niko Moritz, Gordon Wichern, et al., “All-in-one transformer: Unifying speech recognition, audio tagging, and event detection,” *Interspeech*, 2020.
- [22] Noriyuki Tonami et al., “Joint analysis of acoustic events and scenes based on multitask learning,” in *IEEE WASPAA*, 2019.
- [23] Ruixiong Zhang, Wei Zou, and Xiangang Li, “Cross-task pre-training for on-device acoustic scene classification,” *arXiv:1910.09935*, 2019.
- [24] Heysem Kaya, Furkan Gürpınar, and Albert Ali Salah, “Video-based emotion recognition in the wild using deep transfer learning and score fusion,” *Image and Vision Computing*, 2017.
- [25] Shawn Hershey et al., “Cnn architectures for large-scale audio classification,” in *icassp*, 2017.
- [26] Xiaojia Zhao and DeLiang Wang, “Analyzing noise robustness of mfcc and gfcc features in speaker identification,” in *ICASSP*, 2013, pp. 7204–7208.
- [27] Michael Schmidt and Herbert Gish, “Speaker identification via support vector classifiers,” in *ICASSP*, 1996.
- [28] Ludwig Schmidt, Matthew Sharifi, and Ignacio Lopez Moreno, “Large-scale speaker identification,” in *ICASSP*, 2014.
- [29] Musab TS Al-Kaltakchi, Wai L Woo, et al., “Comparison of i-vector and gmm-ubm approaches to speaker identification with timit and nist 2008 databases in challenging environments,” in *EUSIPCO*, 2017.
- [30] David Snyder et al., “X-vectors: Robust dnn embeddings for speaker recognition,” in *ICASSP*, 2018.
- [31] Yanick X Lukic et al., “Learning embeddings for speaker clustering based on voice equality,” in *IEEE MLSP*, 2017.
- [32] Li Wan et al., “Generalized end-to-end loss for speaker verification,” in *ICASSP*. IEEE, 2018.
- [33] Arsha Nagrani, Joon Son Chung, and Andrew Senior, “Voxceleb: a large-scale speaker identification dataset,” *arXiv:1706.08612*, 2017.
- [34] Bahador Khaleghi et al., “Multisensor data fusion: A review of the state-of-the-art,” *Information fusion*, 2013.
- [35] Abien Fred Agarap, “Deep learning using rectified linear units (relu),” *arXiv:1803.08375*, 2018.
- [36] Adam Paszke et al., “Pytorch: An imperative style, high-performance deep learning library,” in *NIPS*. 2019.
- [37] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv:1412.6980*, 2014.
- [38] Daniel S Park et al., “SpecAugment: A simple data augmentation method for automatic speech recognition,” *arXiv:1904.08779*, 2019.
- [39] Hongyi Zhang et al., “mixup: Beyond empirical risk minimization,” *arXiv:1710.09412*, 2017.
- [40] Nitish Srivastava et al., “Dropout: a simple way to prevent neural networks from overfitting,” *The journal of machine learning research*, 2014.