

# FRAMEWORK FOR EVALUATION OF SOUND EVENT DETECTION IN WEB VIDEOS

Rohan Badlani<sup>\*†</sup>   Ankit Shah<sup>†‡</sup>   Benjamin Elizalde<sup>†§</sup>   Anurag Kumar<sup>†</sup>   Bhiksha Raj<sup>†</sup>

<sup>†</sup> Language Technologies Institute, Carnegie Mellon University, Pittsburgh PA

<sup>\*</sup> Department of Computer Science, BITS Pilani, India

## ABSTRACT

The largest source of sound events is web videos. Most videos lack sound event labels at segment level, however, a significant number of them do respond to *text queries*, from a match found using metadata by search engines. In this paper we explore the extent to which a search query can be used as the true label for detection of sound events in videos. We present a framework for large-scale sound event recognition on web videos. The framework crawls videos using search queries corresponding to 78 sound event labels drawn from three datasets. The datasets are used to train three classifiers, and we obtain a prediction on 3.7 million web video segments. We evaluated performance using the search query as true label and compare it with human labeling. Both types of ground truth exhibited close performance, to within 10%, and similar performance trend with increasing number of evaluated segments. Hence, our experiments show potential for using search query as a preliminary true label for sound event recognition in web videos.

**Index Terms**— Sound Event Detection, Convolutional Neural Network, Large-Scale audio event detection, Video Content Analysis

## 1. INTRODUCTION

The Internet is being flooded with massive amount of multimedia data, mostly comprising of videos containing sound events which are often critical to understand the video content. Hence, it is necessary to automatically recognize sound events within the audio, *e.g. police siren, dishwasher or birds singing*. Sound event recognition has been applied to multiple forms, such as in conjunction with other modalities to retrieve and index consumer-generated videos based on content [1, 2, 3, 4], video surveillance (*e.g. detection of footsteps*) [5], human-robot interaction (*e.g. detection of choke*) [6, 7], wildlife monitoring (*e.g. detection of animals*) [8], and context-aware systems (*e.g. outdoors or home*) [9].

In recent years, the main sound recognition challenges: DCASE 2013 [10], 2016 [11] and 2017[12] have fostered research providing standard datasets, task guidelines, metrics

and benchmark performances. Although necessary, the literature tends to focus on DCASE-like datasets, which are audio-only recordings and smaller in scale. This leaves the primary source of sound events, the web, and its intrinsic problems less explored which makes it unclear how state of the art sound event recognition systems work on web videos. Although the YouTube based AudioSet [13] was recently released containing weak labels for sound events, our work explores mismatch conditions between existing audio-only research and datasets applied to YouTube videos.

Sound event recognition on large-scale web videos poses several challenges, mainly the lack of annotated audio recordings for sound events to train and evaluate systems. In order to exploit web recordings, unsupervised solutions have been explored, such as clustering [14] and sound diarization [15] or semi-supervised approaches [16, 17] or visual domain web video analysis [18, 19], which learn from a combination of labeled and unlabeled data sources. Another technique [20] relies on weak labels for learning where only the presence or absence of sounds in the recording is known. For web videos the primary idea is that associated metadata can be used to assign weak labels which can further be used for training classification models [21, 22]. However, the metadata, such as title, keywords and description, are noisy and often related to the visual information rather than the audio content, hence it remains to be seen how reliably it can be used as a true label or ground truth to train and evaluate sound recognition systems. This analysis forms the major contribution of the framework proposed in this paper.

In this paper, we first do an exploration to identify the extent to which search query, which relates to the textual metadata, can be used as a true label for sound events at segment level for YouTube videos. This study, to the best of our knowledge is unavailable in the literature. For our study, we developed a framework for large-scale sound event recognition on web videos consisting of three modules, *Crawl, Hear, Feedback*. In *Crawl*, YouTube videos were crawled using search queries corresponding to 78 sound event labels and the keyword *sound* (<sound event label> sound) drawn from three datasets. In *Hear*, the datasets are used to train three multi-class classifiers, which are used to obtain sound event label prediction on 3.7 million video segments. We evaluated performance using the search query as the true label

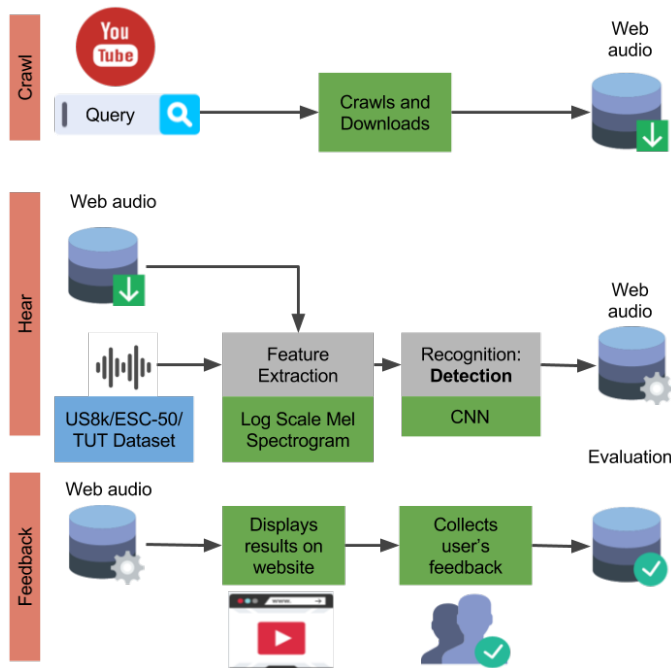
<sup>‡</sup> First two authors contributed equally

<sup>§</sup> Acknowledges CONACYT for his doctoral fellowship, No.343964

and compare it on a subset against human labeling which was collected in the *Feedback* module. Both types of ground truth exhibit similar performance trend. Hence, we show that search query provides a reasonable ground truth for large-scale sound event detection in web videos.

## 2. FRAMEWORK

The purpose of the framework is to use our sound event labels as search queries to crawl videos, which lack true labels at segment level; train classifiers using labeled audio to recognize sound events on the unlabeled crawled video segments; and evaluate the system performance using two types of ground truth, search query and human labeling collected through our website. The framework as described in following sections consists of three modules illustrated in Figure 1



**Fig. 1.** Our framework consists of three modules: Crawl, Hear and Feedback.

### 2.1. Crawl

The *Crawl* module employs search queries to scrape audio from YouTube videos using the Pafy API<sup>1</sup>. The queries are kept to use them later as true labels.

### 2.2. Hear

*Dataset Aggregator* organizes different annotated sound event datasets. The audio is then preprocessed and acoustic features are extracted in the *Feature Extractor* module. We run a *Sound Event Classifier*, the features are used to train classifiers, on unlabeled segments of the crawled videos. The per-

formance is evaluated using the previously used search query and with the human inspection carried on the next module.

### 2.3. Feedback

This module displays the classifier predictions along with corresponding audio segments on our website [nls.cs.cmu.edu](http://nls.cs.cmu.edu). Using our website, human feedback is collected [assumed as true label for an experiment] to evaluate classifier performance and compare the performance against search query as true label.

## 3. EXPERIMENTS AND EVALUATION

In this section, we explain how we use our framework to study the relation between the search query and the presence of sound events in video segments. To achieve our objective we trained sound event classifiers using labeled recordings sourced from three different audio-only datasets and run trained detectors on unlabeled crawled YouTube video segments. The performance was evaluated using two types of true labels - with the search query used to retrieve the videos and with the collected human inspection.

### 3.1. Crawl

In contrast to audio-only recordings, collecting audio from videos poses several challenges. YouTube contains massive amount of videos and a proper formulation of the search query is necessary to filter videos with higher chances of containing the desired sound event. Typing a query composed by a noun such as *air conditioner* will not necessarily fetch a video containing such sound event because the associated metadata often corresponds to the visual content; contrary to audio-only websites such as *freesounds.org*. Therefore, we modified the query to be a combination of keywords: “<sound event label> sound”, for example, “air conditioner sound”. Although the results empirically improved, the sound event was not always found to be occurring and even if it was present, sometimes it was present within a short duration. We discarded videos longer than ten minutes and shorter than three seconds because they were either likely to contain unrelated sounds or were too short to be processed.

The defined search query was used to crawl videos corresponding to 78 sound event labels described in the following Section 3.2.1. Around 260 hours of video was processed, equally distributed per audio event, which corresponds to over 3.7 million video segments (90% overlap) of 2.3 seconds each. The segments were converted to 16-bit encoding, mono-channel, and 44.1 kHz sampling rate WAV files. Note that, the acoustic content from these videos is unstructured and the target sound is often overlapping by other audio, such as noise, speech or music.

### 3.2. Hear

In this subsection, we explain how we used three labeled datasets to train our three sound event classifiers and run them on the unlabeled crawled YouTube video segments.

<sup>1</sup><https://pypi.python.org/pypi/pafy>

### 3.2.1. Dataset Aggregator

The 78 sound events come from 3 publicly available annotated datasets - ESC50, US8k and TUT. We partitioned each dataset into 60% training, 20% validation and 20% testing sets to avoid dealing with the costly process of cross-fold validation during testing of 3.7 million segments.

*ESC-50 or Environmental Sound Classification* [23] has 50 classes from five categories: animals, natural soundscapes and water sounds, human non-speech sounds, interior/domestic sounds and exterior sounds. ESC-50 consists of 2,000 audio segments with an average duration of 5 seconds.

*The US8K or UrbanSounds8K* [24] has 10 classes: *air conditioner, car horn, children playing, dog bark, street music, gun shot, drilling, engine idling, siren, jackhammer*. UrbanSounds8k consists of 8,732 audio segments with an average duration of 3.5 seconds.

*TUT 2016* [25] has 18 classes like *car passing by, bird singing, door banging* from two major sound contexts namely home context and residential area. TUT dataset consists of 954 audio segments with an average duration of 5 seconds.

### 3.2.2. Feature Extractor

We extracted features for all audio recordings in the datasets based on the work in [23] because they provided near to state-of-the-art performance at the time of developing our experimental results. Our pipeline is agnostic of the classifier used. The audio recordings were re-sampled into 16-bit encoding, mono channel at 44.1 kHz sampling rate as a standard format for all experiments. We feed in two channels to our learning model. The first channel comprises of log-scaled mel-spectrograms with 60 mel-bands with a window size of 1024 (23 ms) and hop size is 512 and the second channel comprises of delta coefficients for mel-spectrograms.

### 3.2.3. Sound Event Classifiers

We used multi-class classifiers using Convolutional Neural Networks (CNNs) for each of the datasets based on the work in [23]. Thus, we trained 3 CNN models for classification of 50, 18 and 10 sound events from ESC-50, TUT and US8k respectively. We used different models for each dataset because using a single model for 78 audio events presented many challenges like dealing with unbalanced classes, inconsistency in feature normalization and doing so resulted in low performance (15% lower accuracy).

The CNN architecture consisted of the following layer parameters and optimizations done using the validation set. The input to the CNN is  $60 \times 101 \times 2$ . We used 60 mel-filters, 101 number of frames (approximately 2.3 seconds of data) and 2 channels - mel-spectra and delta features for mel-spectrograms. The input window length of 101 frames is moved by 10 frames (90% overlap). Hence, we trained and predicted on audio segments of approximately 2.3 secs. The first convolutional ReLU layer consisted of 80 filters of rectangular shape (57x6 size, 1x1 stride) allowing for slight frequency invariance. Max pooling was applied with a pool

shape of 4x3 and a stride of 1x3. A second convolutional ReLU layer consisted of 80 filters (1x3 size, 1x1 stride) with max pooling (1x3 pool size, 1x3 pool slide). Further processing was applied through two fully connected hidden layers of 5000 neurons with ReLU non-linearity. The final output layer is a softmax layer. Training was performed using Keras implementation of mini batch stochastic gradient descent even with shuffled sequential batches (batch size 1000) and a nestrov momentum of 0.9. We used L2 weight decay of 0.001 for each layer and dropout probability of 0.5 for all layers.

### 3.3. Evaluation of Classifiers Performance

YouTube videos at segment level lack of true labels for sound events. Hence, we evaluated the classification performance with two types of references or ground truth - the search query used to retrieve the videos and the human inspection collected with the website described in the *Feedback* module.

#### 3.3.1. Evaluation assuming Search Query as Ground Truth

In this evaluation process, all the segments of a retrieved video using a given search query, such as *dog barking sound* are labeled to contain *dog barking*, even if this might not necessarily be true. Motivation is that search query is a reflection of the accumulated metadata tags such as title, description and keywords and hence, we wanted to see to what degree the query relates to the acoustic content of the segments.

#### 3.3.2. Evaluation using Human Feedback as Ground Truth

Human inspection is needed to provide the most reliable ground truth (true label). Hence, the 3.7 million predicted segments were sorted based on classifier confidence (probability) and were evaluated by a group of experts on tasks related to sound recognition. The top 40 segments for each of the 78 classes were distributed randomly among 5 human evaluators and at least 3 people evaluated each segment to reduce human bias and decide based on majority vote. The segments were displayed using a similar web interface as in the main page of nels.cs.cmu.edu, but the difference is that only the audio was displayed in lieu of video in order to avoid revealing other cues, such as images or title. The evaluators had to choose between two options, *Correct* or *Incorrect*, whether the evaluator claims that the system's predicted class was present within the segment or not.

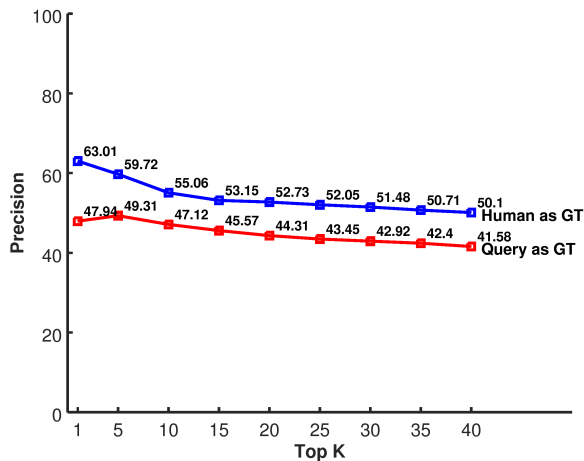
## 4. RESULTS AND DISCUSSION

### 4.1. Results on Crawled YouTube Videos

This main takeaway of our study is the exhibited correlation between the presence of sound events in video segments and their corresponding search query (including the keyword sound), illustrated in Figures 3 and 2. Note that human inspection is the most reliable ground truth while the search query is an assumption of true class because it is based on metadata, which may be based on visual content. Thus, the precision with human feedback was expected to be higher than the one with the query, but it was uncertain how big that

gap could be. The query-based performance was better than we expected considering the uncertainty of the audio content in web videos. Moreover, the performance follows a similar trend to the one from human feedback with a relatively close precision of less than an absolute 10%, which shows potential for the search query (including the keyword sound) to be used as the class label in lieu of human annotations.

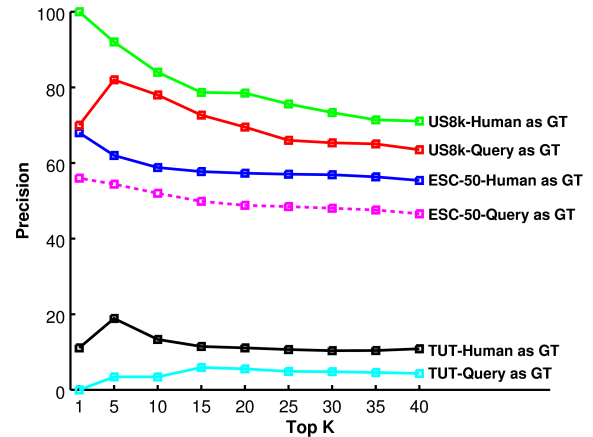
The performance of the three classifiers on the video segments evaluated for both types of ground truth, search query and human feedback, are shown in both Figures 2 (combined weighted-average) and 3 (individual performance). The y-axis, has the performance in terms of Precision@K (a common retrieval metric), which is the precision of  $k$  high-confidence (probability) ranked segments. The x-axis has the Top  $K$  high-confidence segments yielded by our systems. In both figures, the results for  $K = 1-5$  is unstable as the number of audio segments is small and could vary depending on the selected audio segments, however performance stabilizes as  $K$  grows. We stopped at Top 40 results because a YouTube user, for example, tend to focus on the home page of results which translates to  $K$  equals to 10-20. Further, we evaluated all the 3.7 million segments using search query as ground truth and obtained precision scores of 15.43% for ESC-50, 33.58% for US8k and 7.43% for TUT datasets. Future work involves using crowd-sourcing to collect more human feedback to determine whether performance based on human inspection would remain within 10% precision.



**Fig. 2.** Performance for the combination (weighted average) of the three classifiers. (GT = ground truth)

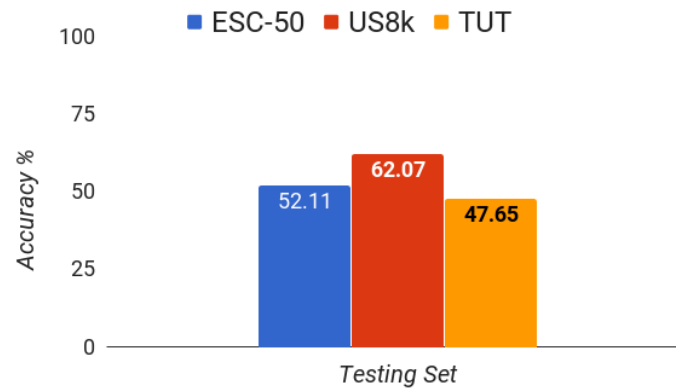
#### 4.2. Results on Datasets

The classification accuracy of our three systems on their corresponding testing sets is shown in order to establish their reliable performance in match conditions. Although the three datasets are well explored in the field, we split them in a different manner to avoid cross-fold experiments with 3.7 million segments. The classification accuracy for ESC-50 52.11%, US8k 62.07% and TUT 47.65% was considerably



**Fig. 3.** For the three classifiers, the search query-based performance follows a similar and close trend to the one based on human feedback. (GT = ground truth)

better than their corresponding random performance: 2%, 10%, 5.5% as shown in Figure 4.



**Fig. 4.** Classification accuracy for each of the three classifiers trained on the each of the three datasets showed reliable results on match conditions.

## 5. CONCLUSIONS

Web videos have no established true sound event labels at segment level. Thus, we studied the relation between search queries, based on sound event labels, and the presence of the corresponding sound event. We developed a framework to crawl videos using search queries, trained classifiers with audio-only datasets on video segments and evaluate performance with two types of ground truth - the search query and the collected human inspection. They showed a correlation between the search query (including the keyword sound) and the presence of sound events in video segments. Our results encourage further exploration of the search query as a preliminary label of the true class to evaluate sound event classification at a large-scale.

## 6. REFERENCES

- [1] Shouo-I Yu, Lu Jiang, Zexi Mao, Xiaojun Chang, Xingzhong Du, Chuang Gan, Zhenzhong Lan, Zhongwen Xu, Xuanchong Li, Yang Cai, et al., “Informedia@ trecvid 2014 med and mer,” in *NIST TRECVID Video Retrieval Evaluation Workshop*, 2014, vol. 24.
- [2] Yu-Gang Jiang, Xiaohong Zeng, Guangnan Ye, Dan Ellis, Shih-Fu Chang, Subhabrata Bhattacharya, and Mubarak Shah, “Columbia-ucf trecvid2010 multimedia event detection: Combining multiple modalities, contextual concepts, and temporal matching,” in *TRECVID*, 2010, vol. 2, pp. 3–2.
- [3] Zhen-zhong Lan, Lei Bao, Shouo-I Yu, Wei Liu, and Alexander Hauptmann, “Double fusion for multimedia event detection,” *Advances in Multimedia Modeling*, pp. 173–185, 2012.
- [4] Hui Cheng, Jingen Liu, Saad Ali, Omar Javed, Qian Yu, Amir Tamrakar, Ajay Divakaran, Harpreet S Sawhney, R Manmatha, James Allan, et al., “Sri-sarnoff aurora system at trecvid 2012: Multimedia event detection and recounting,” in *Proceedings of TRECVID*, 2012.
- [5] Pradeep K Atrey, Namunu C Maddage, and Mohan S Kankanhalli, “Audio based event detection for multimedia surveillance,” in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*. IEEE, 2006, vol. 5.
- [6] Janvier Maxime, Xavier Alameda-Pineda, Laurent Girin, and Radu Horaud, “Sound representation and classification benchmark for domestic robots,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6285–6292.
- [7] Maxime Janvier, Xavier Alameda-Pineda, Laurent Girin, and Radu Horaud, “Sound-event recognition with a companion humanoid,” in *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)*. IEEE, 2012, pp. 104–111.
- [8] Jose F Ruiz-Muñoz, Mauricio Orozco Alzate, and Germán Castellanos-Domínguez, “Multiple instance learning-based birdsong classification using unsupervised recording segmentation,” in *24th International Joint Conference on Artificial Intelligence*, 2015.
- [9] Antti J Eronen, Vesa T Peltonen, Juha T Tuomi, Anssi P Klauri, Seppo Fagerlund, Timo Sorsa, Gaëtan Lorho, and Jyri Huopaniemi, “Audio-based context recognition,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 1, pp. 321–329, 2006.
- [10] Tuomas Virtanen, Annamaria Mesaros, Toni Heittola, Mark D. Plumbley, Peter Foster, Emmanouil Benetos, and Mathieu Lagrange, *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2016 Workshop (DCASE2016)*, Tampere University of Technology. Department of Signal Processing, 2016.
- [11] Dimitrios Giannoulis, Emmanouil Benetos, Dan Stowell, Mathias Rossignol, Mathieu Lagrange, and Mark D Plumbley, “Detection and classification of acoustic scenes and events: an IEEE AASP challenge,” in *2013 IEEE WASPAA*. IEEE, 2013, pp. 1–4.
- [12] 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)*, November 2017, submitted.
- [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. IEEE ICASSP 2017*, New Orleans, LA, 2017.
- [14] Justin Salamon and Juan Pablo Bello, “Unsupervised feature learning for urban sound classification,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015.
- [15] Benjamin Elizalde, Gerald Friedland, Howard Lei, and Ajay Divakaran, “There is no data like less data: Percepts for video concept detection on consumer-produced media,” in *Proceedings of the 2012 ACM international workshop on Audio and multimedia methods for large-scale video analysis*. ACM, 2012, pp. 27–32.
- [16] Wenjing Han, Eduardo Coutinho, Huabin Ruan, Haifeng Li, Björn Schuller, Xiaojie Yu, and Xuan Zhu, “Semi-supervised active learning for sound classification in hybrid learning environments,” *PloS one*, vol. 11, no. 9, pp. e0162075, 2016.
- [17] Ankit Shah, Rohan Badlani, Anurag Kumar, Benjamin Elizalde, and Bhiksha Raj, “An approach for self-training audio event detectors using web data,” *arXiv preprint arXiv:1609.06026*, 2016.
- [18] C. Gan, T. Yao, K. Yang, Y. Yang, and T. Mei, “You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 923–932.
- [19] Chuang Gan, Chen Sun, Lixin Duan, and Boqing Gong, “Webly-supervised video recognition by mutually voting for relevant web images and web video frames,” in *European Conference on Computer Vision (ECCV)*, 2016.
- [20] Anurag Kumar and Bhiksha Raj, “Audio event detection using weakly labeled data,” in *24th ACM International Conference on Multimedia*. ACM Multimedia, 2016.
- [21] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al., “Cnn architectures for large-scale audio classification,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 131–135.
- [22] Anurag Kumar and Bhiksha Raj, “Deep cnn framework for audio event recognition using weakly labeled web data,” *arXiv preprint arXiv:1707.02530v2*, 2016.
- [23] Karol J Piczak, “Environmental sound classification with convolutional neural networks,” in *2015 IEEE 25th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2015, pp. 1–6.
- [24] J. Salamon, C. Jacoby, and J. P. Bello, “A dataset and taxonomy for urban sound research,” in *22st ACM International Conference on Multimedia (ACM-MM’14)*, Orlando, FL, USA, Nov. 2014.
- [25] Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen, “TUT database for acoustic scene classification and sound event detection,” in *24th European Signal Processing Conference 2016 (EUSIPCO 2016)*, Budapest, Hungary, 2016.