

AVASPEECH-SMAD: A STRONGLY LABELLED SPEECH AND MUSIC ACTIVITY DETECTION DATASET WITH LABEL CO-OCCURRENCE

Yun-Ning Hung¹ Karn N. Watcharasupat^{1,2} Chih-Wei Wu³ Iroro Orife³
Kelian Li¹ Pavan Seshadri¹ Junyoung Lee²

¹Center for Music Technology, Georgia Institute of Technology, USA

²School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore

³Netflix, Inc., USA

Corresponding author: amy.hung@gatech.edu (Y.-N. Hung)

ABSTRACT

We propose a dataset, AVASpeech-SMAD, to assist speech and music activity detection research. With frame-level music labels, the proposed dataset extends the existing AVASpeech dataset, which originally consists of 45 hours of audio and speech activity labels. To the best of our knowledge, the proposed AVASpeech-SMAD is the first open-source dataset that features strong *polyphonic* labels for both music and speech. The dataset was manually annotated and verified via an iterative cross-checking process. A simple automatic examination was also implemented to further improve the quality of the labels. Evaluation results from two state-of-the-art SMAD systems are also provided as a benchmark for future reference.

1. INTRODUCTION

Speech and music activity detection (SMAD) is a long-studied problem and has been included in the Music Information Retrieval Evaluation eXchange (MIREX) competition for several years. Like other Music Information Retrieval (MIR) tasks, the most recent improvements to SMAD rely on data-driven approaches [1–6]. However, due to copyright issues, many of these systems were trained on private datasets which impeded the reproducibility of the results (e.g., the radio datasets used in [2] and [1]). Although several publicly available datasets have been proposed for training or evaluation of SMAD systems, as shown in Table 2, they suffer from several drawbacks. GTZAN [7], MUSAN [8], SSMSC [9] and Muspeak [10] only have non-overlapping speech or music segments. OpenBMAT [11] and ORF TV [12] only support music labels, while the original AVASpeech dataset [13] only has speech labels. In the real-world where speech and

	Avg (%)	Min (%)	Max (%)
Speech	52.48	14.40	87.02
Music	43.40	0.29	100.00
Both speech and music	20.73	0.00	77.19

Table 1. Minimum, maximum, and average percentages of regions with speech only, with music only, and with both speech and music for each audio sample.

music co-occur regularly, these datasets might not be suitable training sources.

To solve the data limitation problem, we proposed a supplementary dataset, *AVASpeech-SMAD*. The dataset is the extension of the AVASpeech dataset proposed by Chaudhuri et al. [13]. The original AVASpeech dataset only contains speech activities and we extended it by manually labelling the music activities. We expect our proposed dataset to be used for training or evaluation of future SMAD systems. The dataset is open-sourced on the Github repository: https://github.com/biboamy/AVASpeech_Music_Labels

2. DATASET

The statistics of the dataset are shown in Table 2. Compared to other publicly available datasets, ours is the only one containing overlapping speech *and* music frame labels. The dataset includes a variety of content, languages, genres and production quality. The audio data, labels, and annotation process are discussed in the following sections.

2.1 Audio & Labels

The original dataset contains 160 excerpts of 15-minute clips taken from YouTube videos with a total duration of 45h. Each audio file is stereo and was sampled at 22 050 Hz with 16 bits per sample. Due to copyright issues, we could not distribute the audio files from this dataset. Instead, we include the scripts for downloading and preprocessing the audio files in our GitHub repository.

The speech labels were derived as-is from the original AVASpeech while the new music labels were manually annotated. The statistics of the labels are shown in Table 1. The dataset covers a variety of cases. Some samples barely have any music while some contain mostly music. Some



Dataset	Music Labels	Speech Labels	Overlap	# instances	Duration (hrs)
GTZAN Speech and Music [7]	✓	✓	No	128	1.1
MUSAN [8]	✓	✓	No	2016	109
Scheirer & Slaney Music Speech (SSMSC) [9]	✓	✓	No	245	1
Muspeak [10]	✓	✓	No	214	5
OpenBMAT [11]	✓	✗	Yes	1647	27.5
ORF TV [12]	✓	✗	Yes	13	9
AVASpeech-SMAD (Proposed)	✓	✓	Yes	160	45

Table 2. Metadata of the existing datasets compared to our proposed AVASpeech-SMAD dataset.




Figure 1. Overview of the annotation process.

of them do not have any region with overlapping music and speech, while some have over 70 % of the regions containing both music *and* speech. The detailed statistics of each clip can be found in our GitHub repository.

2.2 Annotation Process

The annotation process of the proposed dataset is shown in Fig. 1. The annotation process has four steps, namely music detection, first manual annotation check, cross-validate and second manual annotation check. Steps involving human annotators are highlighted in orange. Seven MIR students and researchers, all of whom are also musicians of varying experience levels, have volunteered as annotators. An internal algorithm is first used to pseudo-label the music regions. The annotators are then asked to manually annotate the regions with music activity, with the pseudo labels acting as a rough guide. Each annotators were assigned between 12 to 24 audio clips to annotate active music regions. Sonic Visualizer¹ was used to mark the active regions and all annotators used headphones during the annotation process. To ensure the consistency of the labels and avoid ambiguity, the guideline for determining music versus speech is listed in GitHub repository.

After the annotations are completed, all the labels were cross-validated with the speech labels provided by original AVASpeech dataset. Since the original labels contain “Speech with music” and “Clean speech” classes, we can expect that the former contains music while the latter do not contains music. Any discrepancies between our labels and the original labels were algorithmically detected: the region that was labeled as music in the original labels but not in our labels, and the region that was labeled no music in the original labels but labeled as music in our labels. If the discrepancies are less than 3 s, we automatically modify the labels based on the original labels. However, if the discrepancies are greater than three seconds, we conducted

¹ <https://www.sonicvisualiser.org/>, last accessed 04/30/2021

Algo.	Music			Speech		
	F1	Precision	Recall	F1	Precision	Recall
[1]	80.2	69.8	94.3	77.6	82.7	73.1
[5]	62.2	85.7	48.8	79.1	83.1	75.4

Table 3. The segment-level evaluation (%) on our AVASpeech-SMAD dataset by using two SOTA systems.

a manual review of the regions. Each region is randomly assigned to additional two annotators and majority vote is considered to determine the final labels of the regions.

3. BENCHMARK

We choose two existing SOTA systems to evaluate on our proposed dataset. First is the CRNN-based detector proposed by [1]. The method is trained on the combination of synthetic data and radio broadcast. The second system is the InaSpeechSegmentor proposed by [5]. The segmentor has a CNN architecture and can split audio into homogeneous music, speech and noise region. The `sed_eval` toolbox is used to perform segment-level evaluation as used in the MIREX 2018 competition². The result is shown in Table 3 for future reference.

4. CONCLUSION

In this work, we proposed a supplementary dataset for SMAD. The dataset not only contains overlapping speech and music frame labels, but also includes a variety of content. Based on our benchmark experiments, one of the SOTA systems achieved F1-scores (80% and 77 % for music and speech, respectively) that are slightly lower than the reported 85% for both music and speech on MIREX test sets in [1]. This result suggests that our proposed dataset is able to present new challenges to the model and potentially complement the existing datasets (e.g., MIREX test sets) with diverse content and frame-level labels. We expect this dataset to serve as a new resource and reference for future SMAD research.

5. ACKNOWLEDGEMENT

K. N. Watcharasupat and J. Lee acknowledge the support from the CN Yang Scholars Programme, Nanyang Tech-

² https://www.music-ir.org/mirex/wiki/2018:Music_and/or_Speech_Detection, last accessed 04/30/2021

nological University, Singapore. We also gratefully acknowledge Professor Alexander Lerch and Music Informatics Group supported this research by providing a Titan X GPU for computing the experiment.

6. REFERENCES

- [1] S. Venkatesh, D. Moffat, and E. R. Miranda, “Investigating the effects of training set synthesis for audio segmentation of radio broadcast,” *Electronics*, vol. 10, no. 7, p. 827, 2021.
- [2] Q. Lemaire and A. Holzapfel, “Temporal convolutional networks for speech and music detection in radio broadcast,” in *Proceedings of the 20th International Society for Music Information Retrieval Conference*, 2019, pp. 229–236.
- [3] D. de Benito-Gorron, A. Lozano-Diez, D. T. Toledano, and J. Gonzalez-Rodriguez, “Exploring convolutional, recurrent, and hybrid deep neural networks for speech and music detection in a large audio dataset,” *Journal on Audio, Speech, and Music Processing*, vol. 2019, no. 1, pp. 1–18, 2019.
- [4] B.-Y. Jang, W.-H. Heo, J.-H. Kim, and O.-W. Kwon, “Music detection from broadcast contents using convolutional neural networks with a mel-scale kernel,” *Journal on Audio, Speech, and Music Processing*, vol. 2019, no. 1, pp. 1–12, 2019.
- [5] D. Doukhan, E. Lechapt, M. Evrard, and J. Carrive, “INA’s MIREX 2018 music and speech detection system,” in *14th Music Information Retrieval Evaluation eXchange*, 2018.
- [6] M. Papakostas and T. Giannakopoulos, “Speech-music discrimination using deep visual feature extractors,” *Expert Systems with Applications*, vol. 114, pp. 334–344, 2018.
- [7] G. Tzanetakis and P. Cook, “Marsyas: A framework for audio analysis,” *Organised Sound*, vol. 4, no. 3, pp. 169–175, 2000.
- [8] D. Snyder, G. Chen, and D. Povey, “MUSAN: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [9] E. Scheirer and M. Slaney, “Construction and evaluation of a robust multifeature speech/music discriminator,” vol. 2, 1997, pp. 1331–1334.
- [10] D. Wolff, T. Weyde, E. Benetos, and D. Tidhar, “MIREX muspeak sample dataset,” <http://mirg.city.ac.uk/datasets/muspeak/>, 2015, accessed: 2020-09-30.
- [11] B. Meléndez-Catalán, E. Molina, and E. Gómez, “Open broadcast media audio from tv: A dataset of tv broadcast audio with relative music loudness annotations,” *Transactions of the International Society for Music Information Retrieval*, vol. 2, no. 1, 2019.
- [12] K. Seyerlehner, T. Pohle, M. Schedl, and G. Widmer, “Automatic music detection in television productions,” in *Proceedings of the 10th International Conference on Digital Audio Effects (DAFx)*, 2007.
- [13] S. Chaudhuri, J. Roth, D. P. Ellis, A. Gallagher, L. Kaver, R. Marvin, C. Pantofaru, N. Reale, L. G. Reid, K. Wilson *et al.*, “AVA-speech: A densely labeled dataset of speech activity in movies,” in *Proceedings of the 19th Annual Conference of the International Speech Communication Association*, 2018.