



HAL
open science

AN EXPLAINABLE PROXY MODEL FOR MULTILABEL AUDIO SEGMENTATION

Théo Mariotte, Antonio Almudévar, Marie Tahon, Alfonso Ortega

► **To cite this version:**

Théo Mariotte, Antonio Almudévar, Marie Tahon, Alfonso Ortega. AN EXPLAINABLE PROXY MODEL FOR MULTILABEL AUDIO SEGMENTATION. International Conference on Acoustics Speech and Signal Processing, IEEE, Apr 2024, Seoul (Korea), France. hal-04393946

HAL Id: hal-04393946

<https://univ-lemans.hal.science/hal-04393946>

Submitted on 16 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

AN EXPLAINABLE PROXY MODEL FOR MULTILABEL AUDIO SEGMENTATION

Théo Mariotte¹, Antonio Almudévar², Marie Tahon¹, Alfonso Ortega²

¹LIUM, Le Mans Université

²ViVoLab, Aragón Institute for Engineering Research (I3A), University of Zaragoza
theo.mariotte@univ-lemans.fr

ABSTRACT

Audio signal segmentation is a key task for automatic audio indexing. It consists of detecting the boundaries of class-homogeneous segments in the signal. In many applications, explainable AI is a vital process for transparency of decision-making with machine learning. In this paper, we propose an explainable multilabel segmentation model that solves speech activity (SAD), music (MD), noise (ND), and overlapped speech detection (OSD) simultaneously. This proxy uses the non-negative matrix factorization (NMF) to map the embeddings used for the segmentation to the frequency domain. Experiments conducted on two datasets show similar performances as the pre-trained black box model while strong explainable features arise. Specifically, the frequency bins used for the decision can be easily identified at both the segment level (local explanations) and global level (class prototypes).

Index Terms— multilabel audio segmentation, explainability, non-negative matrix factorization, music detection, speech detection

1. INTRODUCTION

Audio segmentation is a key task for many speech technologies such as automatic speech recognition, speaker identification, and dialog monitoring in different multi-speaker scenarios, including TV/radio, meetings, and medical conversations. More precisely, these technologies must be aware of the presence of noisy environments (brouhaha, external noise), and how many speakers are active at each time. In many domains, such as health or human-machine interactions, the prediction of segment timestamps is not enough and it is necessary to include some explanations. Indeed, the current trend for explainable AI is a vital process for transparency of decision-making with machine learning: the user (a doctor, a judge, or a human scientist) has to justify the choice made based on the system output. Explainability for AI can be addressed at different stages of the process. Pre-hoc explainability intends to understand and describe data with explainable features and statistics. Another stage is to develop explainable-by-design models. The last stage consists of the extraction of explanations from a pre-trained model by the use of proxy models or perturbation mechanisms. Our work comes within the scope of such post-hoc explanations.

The proposed approach aims to train an explainable proxy model from a black-box segmentation system. While many works have

been conducted in explaining neural networks in vision [1] and NLP [2], the literature is limited in the audio domain. The first intents are focused on saliency map extraction to reveal what information is used in the output [1, 3]. Audio classifier decisions can be interpreted by post-hoc analysis with approaches such as SoundLime [4]. Recently, a few explainable models have been developed like APNet [5], which extends the training of prototypes to the audio domain, and post-hoc visualization of explanations obtained from Shapley values [6]. In the architecture proposed in [7], the authors explain a black box audio classifier with a proxy model which is optimized to classify audio scenes while reconstructing the audio with the non-negative matrix factorization framework [8].

This paper focuses on automatic audio segmentation in the context of large-scale multimedia data from archivists, i.e. the detection of homogeneous segments containing speech (SAD), music (MD), noise (ND), and overlapped speech (OSD) with a unified model. Early studies on speech [9], overlapped speech [10], and music [11] segmentation are focused on the statistical modeling of handcrafted acoustic features. Currently, the segmentation is mainly performed with neural networks and supervised learning. While each task has been generally solved independently as a binary frame-wise classification task (SAD [12], OSD [13, 14], MD [15, 16]), more recent approaches propose to solve multiple tasks simultaneously. The multiclass model predicts a single class and class intersection is empty. For example, in [17], authors propose to segment speech, music, and noise with a single multiclass model. A few works also report joint SAD and OSD [18–20]. In this paper, SAD, OSD, MD, and ND are simultaneously solved as a multilabel frame classification task. Thus, multiple classes can be predicted simultaneously, and the intersection between classes is not empty.

We propose to train an explainable proxy model from a pre-trained multilabel segmentation model (designated as the teacher). The architecture is inspired by [7]. The proxy is trained following a teacher-student approach, commonly used in knowledge distillation [21]. The teacher inputs Wavlm pre-trained features [22] and outputs the pseudo probability of each class. Two types of proxy models are investigated. The former inputs a spectrogram and the latter uses the teacher’s Wavlm outputs. Contrary to [7], we consider frame-level, i.e. time segmentation, instead of utterance-level classification. We demonstrate that proxy models provide similar or even better performance as the teacher while being smaller and explainable. These models provide local (segment-level) explanations by highlighting the salient frequencies for a given audio input and global explanations that can be seen as class prototypes. Contrary to state-of-the-art methods such as Shapley [6] or SoundLime [4], our contribution not only brings post-hoc explanations but also provides decisions. Finally, we show that the relevant information for classification can be mapped to the spectral domain. A classification score confirms the selection of relevant components. To the best of

This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 101007666. This work was performed using HPC resources from GENCI-IDRIS (Grant 2022-AD011012565). The research reported here was conducted at the 2023 Frederick Jelinek Memorial Summer Workshop on Speech and Language Technologies, hosted at Le Mans University (France) and sponsored by Johns Hopkins University.

our knowledge, this is the first intent to design an explainable neural multilabel segmentation model.

This paper is arranged as follows: sec. 2 presents the system along with the proxy model training strategy. Sec. 3 introduces the experimental protocol and sec. 4 the segmentation results. Sec. 5 demonstrates the explainability capacities of the proposed system.

2. NMF-BASED EXPLAINABLE MODEL

The architectures and training framework of the proxy models are described in this section.

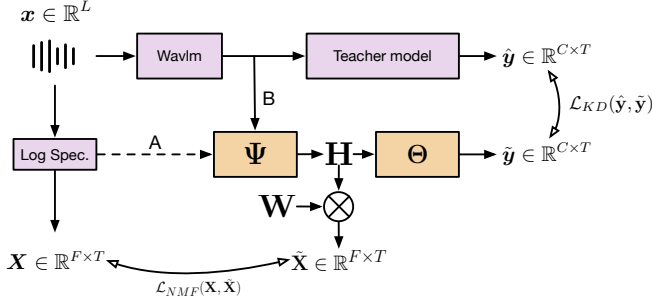


Fig. 1: Diagram of the proposed architecture for NMF-based multi-label segmentation explanation. The **A** and **B** branches represent the spectrogram and Wavlm-based proxy models respectively.

2.1. Multilabel audio segmentation formulation

Audio segmentation is solved as a multilabel framewise classification task. Let $\{\mathbf{S}, \mathbf{y}\}$ be a training set composed of sequences of acoustic features $\mathbf{S} \in \mathbb{R}^{D \times T}$, where D is the feature vector dimension and T the number of frames, and the aligned annotations $\mathbf{y} \in \mathbb{R}^{C \times T}$, with C being the number of classes. Let $f: \mathbb{R}^{D \times T} \rightarrow \mathbb{R}^{C \times T}$ be a parametric function that estimates the logits for each class from the sequence of features such as $\hat{\mathbf{y}} = f(\mathbf{S})$. The parameters of f are optimized to minimize a loss function $\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y})$ (here we used a cross-entropy), with an iterative algorithm. The sequence of feature \mathbf{S} is extracted from the audio signal by an additional function $\mathbf{S} = g(\mathbf{x})$.

2.2. Non-negative matrix factorization (NMF)

NMF has been extensively used for audio signal processing [23]. This approach factorizes a given non-negative matrix $\mathbf{X} \in \mathbb{R}_+^{F \times T}$ into two non-negative matrices: $\mathbf{W} \in \mathbb{R}_+^{F \times K}$, usually denoted as *dictionary*, and $\mathbf{H} \in \mathbb{R}_+^{K \times T}$, denoted as *activations*. K represents the rank of the factorization. Both matrices are jointly learned by solving

$$\mathbf{W}, \mathbf{H} = \arg \min_{\mathbf{W}, \mathbf{H}} \|\mathbf{X} - \mathbf{WH}\|_2^2, \quad (1)$$

with a two-step optimization process [8]. In this work, we consider the sparse NMF implementation [24] as suggested in [7]. In our segmentation proxy model, the dictionary \mathbf{W} is pre-learned while the activation \mathbf{H} is extracted by a neural model Ψ and referred to as an embedding. The proxy model architecture and the way NMF is integrated is described in the following subsection.

2.3. Proxy model framework

In this work, the f model is pre-trained with frozen weights and serves as a teacher for the proxy model. We use a similar approach

as [7] where the proxy model is composed of two functions. Let Ψ be a function that maps a sequence of D -dimension feature vectors $\mathbf{S} \in \mathbb{R}^{D \times T}$ to the embedding $\mathbf{H} \in \mathbb{R}_+^{K \times T}$. The proxy model logits $\hat{\mathbf{y}} \in \mathbb{R}^{C \times T}$ are obtained with an additional Θ function such as $\tilde{\mathbf{y}} = \Theta(\mathbf{H})$. The log-spectrogram of the input audio signal $\tilde{\mathbf{X}}$ is also reconstructed from \mathbf{H} with the \mathbf{W} dictionary: $\tilde{\mathbf{X}} = \mathbf{WH}$. Non-negativity is ensured by a ReLU activation function σ , e.g. $\sigma(\mathbf{H})$. Note that \mathbf{W} can also be trained simultaneously with the model.

The training objective of such a model is composed of 3 loss terms. The first trains the proxy model to mimic the teacher's decisions. Considering multilabel segmentation, and following the common approaches in knowledge distillation [21], we use the binary Kullback-Leiber (KL) divergence between the teacher and the proxy model output distributions, denoted as $\mathcal{L}_{KD}(\hat{\mathbf{y}}, \tilde{\mathbf{y}})$.

The second loss term constrains the \mathbf{H} embedding to minimize the NMF-based spectrogram reconstruction and is implemented as the squared L2-norm between the target spectrogram \mathbf{X} and the reconstruction $\tilde{\mathbf{X}}$:

$$\mathcal{L}_{NMF}(\mathbf{X}, \tilde{\mathbf{X}}) = \|\mathbf{X} - \mathbf{WH}\|_2^2. \quad (2)$$

The last loss term minimizes the L1-norm of the \mathbf{H} embedding to enforce its sparsity. Having a sparse embedding reduces the number of active components, and makes the explanation step easier (Sec. 5). Finally, the global training objective is the weighted sum of the 3 terms given in eq. 3 where (α, β, γ) is a triplet of hyperparameters to weight each term of the loss.

$$\mathcal{L} = \alpha \mathcal{L}_{KD}(\hat{\mathbf{y}}, \tilde{\mathbf{y}}) + \beta \mathcal{L}_{NMF}(\mathbf{X}, \tilde{\mathbf{X}}) + \gamma \|\mathbf{H}\|_1. \quad (3)$$

3. EXPERIMENTAL PROTOCOL

3.1. Datasets

In the current literature, there is no audio data annotated to perform SAD, OSD, MD, and ND simultaneously. Therefore, models are trained on 4 datasets listed in the table 1 to perform multilabel audio segmentation so that all the classes are represented. The teacher is optimized on the 4 training subsets described in the papers associated with each dataset. This represents about 300 hours of annotated audio for training. Due to the large amount of data, the teacher pre-training requires a lot of resources and a lengthy training process to converge. We propose to train the proxy model on a subset of the training set. Only Albayzin (AragonRadio train and 324TV) and DiHard train data are used for the knowledge distillation step, which represents 87 hours of audio.

To evaluate the segmentation performance, both teacher and proxy models are evaluated following the same protocol. The models are evaluated on Aragon Radio and DiHard III test sets. The test partitions follow the split proposed for the original challenges [25, 26].

3.2. Model architectures

The teacher model is similar to [14] and is composed of two main parts: feature extraction and sequence modeling. The former is performed using pre-trained Wavlm Large [22] that outputs a sequence of 1024-dimension vectors. The latter transforms the sequence of features to predict the segmentation. It is composed of a 64-channel bottleneck layer followed by 3 TCN blocks [30] composed of 5 1-D convolutional layers with exponentially increasing dilatation. A 1-D convolution layer processes the output of the TCN blocks to project

Table 1: Datasets used to train both teacher (T) and proxy (P) models with the available labels in each of them. Missing annotations for a given class are removed from the classification loss.

Dataset	Model		Available label			
	T	P	SAD	MD	ND	OSD
Albayzin [25, 27]	✓	✓	✓	✓	✓	
OpenBMAT [28]	✓			✓		
ALLIES [29]	✓		✓			✓
DiHard III [26]	✓	✓	✓			✓

the hidden representations to the logits space. A sigmoid activation function is then applied to get normalized scores for each class.

As described in section 2, two types of proxy models are investigated. In the first approach, the proxy model (denoted as Spec. in Table 2) inputs a log-spectrogram extracted on 64ms sliding windows with 20ms step. In that case, $\mathbf{S} = \mathbf{X}$. The Ψ function is implemented with the same TCN-based architecture as the teacher but with 4 TCN blocks of 4 convolutional layers. The bottleneck and hidden layers are composed of 128 and 256 channels respectively. The second approach uses Wavlm features (denoted as Wavlm) from the teacher as input. In this case, the Ψ model is smaller to reduce the number of trainable parameters w.r.t. the teacher. The architecture is the same as the teacher’s but with only 2 TCN blocks of 3 convolutional layers. The Θ function is implemented as a linear layer with no bias of K input dimensions and C outputs, followed by a sigmoid.

3.3. NMF pre-training and hyperparameters

As previously introduced, the \mathbf{W} dictionary is pre-trained to map the embedding to the frequency domain. The pre-training is performed on a subset of Aragon Radio using Sparse NMF [24]. We select segments between 1s and 4s representing each type of class. In practice, we found that building a subset with 1200 segments containing 16% of speech and 42% of music and noise offers the best reconstruction quality for each type of signal. This approach however showed limited reconstruction quality with the Wavlm-based model. To tackle this, we explored training the \mathbf{W} dictionary along with the proxy model (denoted as Wavlm (\mathbf{W} trained)). In this case, the dictionary is implemented as a linear layer with no bias followed by a ReLU activation. For each approach, the NMF rank is fixed at $K = 256$ since it offers the best performance.

The proxy model is optimized with the ADAM optimizer with a learning rate of 10^{-3} and batches of 64 segments. We set $(\alpha, \beta, \gamma) = (10, 5, 0.1)$ to scale the classification and reconstruction losses while lowering the impact of the sparsity term.

4. SEGMENTATION PERFORMANCE

This section compares the performances in terms of the F1-score of each proxy model to the teacher model which is considered as our baseline. The SAD, ND, MD, and OSD performances are presented in the table 2. The proxy models are expected to reach similar, eventually slightly lower, performances as the teacher. The score obtained on OSD by the teacher model compares with the state of the art which is 63.4% on DiHard III [14].

The spectrogram-based system (spec.) shows 96.3% and 95.5% F1-score on the SAD task on DiHard and AragonRadio respectively. This represents a little degradation concerning the teacher. The ND (73.0%) and MD (87.6%) are also degraded by an absolute -5,6%

Table 2: F1-score (%) obtained for each segmentation task on both AragonRadio and DiHard III evaluation sets. Bold values indicate the significantly best scores.

Model	Aragon Radio			DiHard III	
	SAD	ND	MD	SAD	OSD
Teacher	96.8	78.6	93.2	96.9	60.7
Spec.	96.3	73.0	87.6	95.5	40.1
Wavlm (\mathbf{W} trained)	96.7	78.5	93.0	96.8	60.8
Wavlm	96.8	79.5	93.1	96.9	61.4

compared to the teacher. On OSD, the performance is strongly degraded with a 40.1% F1-score. The spectrogram-based model degrades the segmentation performance. However, this result is expected since the teacher is trained on Wavlm pre-trained features which perform better than the spectrogram. Considering the cost of Wavlm pre-training and the inference time, the spectrogram-based model offers a convincing performance.

Using Wavlm features as input of the proxy model drastically improves the segmentation performance. When the NMF dictionary is trained simultaneously as the segmentation task, the proxy model reaches 96.1% and 96.7% F1-score on SAD. The ND performance remains lower than the teachers with a 72.6% F1-score, which represents an absolute -6% degradation. The MD and OSD performances are similar to the teacher with 93.0% and 60.5% F1-score respectively. While showing similar performance as the teacher, training the NMF dictionary degrades noise detection. When the NMF dictionary is pre-learned, the proxy model delivers similar or even better segmentation performance as the teacher. The proxy model reaches a 79.5% F1-score on ND (+0.9% absolute improvement w.r.t. the teacher), and 61.4% F1-score on OSD. This system shows the best segmentation performance among all the models.

The Wavlm-based proxy models deliver the best performance when the dictionary is pre-learned. This model is kept for explanation extraction in the next section.

5. DECISION EXPLANATION

This section describes the explanation extraction process to identify the relevant frequency bins for the segmentation.

5.1. Explanation extraction

The NMF framework allows to map the $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_t, \dots, \mathbf{h}_T]$ where $\mathbf{h}_t \in \mathbb{R}^K$ embedding to the frequency domain. Furthermore, the segmentation prediction is obtained from this embedding with the Θ linear transformation. The first step in the explanation process is to identify the $k \in [1, \dots, K]$ NMF components that are the most relevant for each classification task. We first apply a pooling operation to the embedding by averaging it over the time dimension: $z_k = \frac{1}{T} \sum_{t=1}^T h_{k,t}$. To identify the relevant components to detect the class c , we define a relevance vector $\mathbf{r}_c = [r_{1,c}, \dots, r_{k,c}, \dots, r_{K,c}]$ in which each element is computed following (4), where $\theta_{k,c}$ is the k -th weight of the linear layer associated to class c . The highest values in \mathbf{r} correspond to the most relevant component. Therefore, the most relevant components are selected by applying a threshold τ . A filtered relevance vector $\mathbf{R}_{c,\tau}$ is obtained, in which the k -th element is defined as:

$$R_{k,c}(\tau) = \begin{cases} r_{k,c} = z_k \times \theta_{k,c} & \text{if } r_{k,c} > \tau \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$




Fig. 2: (Top) segment-level (local) relevant frequency bins for speech and music segmentation according to the relevance threshold τ . (Bottom) Classification scores for (—) speech, (---) music, and (· · ·) overlapped speech with each selected components. Two audio samples from AragonRadio eval set with speech only (left) and music only (right)

5.2. Segment-level (local) explanation

The filtered relevance vector $\mathbf{R}_c(\tau)$ belongs to the same space as the \mathbf{H} embedding. Hence, it can be projected to the frequency domain by the following NMF linear transformation:

$$\mathbf{X}_c(\tau) = \mathbf{W}\mathbf{R}_c(\tau), \quad (5)$$

where $\mathbf{X}_c(\tau) \in \mathbb{R}^K$ is the projection of the relevant components in the frequency domain given class c and threshold τ . This representation highlights the relevant frequency bins to detect class c .

After identifying the relevant components, it is necessary to measure how confident the model is regarding this selection. Model confidence is evaluated with the classification scores. The score is obtained by only keeping the relevant components in the \mathbf{H} embedding and forwarding this filtered embedding $\mathbf{H}_c(\tau)$ through the Θ layer such that $\tilde{\mathbf{y}}_c(\tau) = \Theta(\mathbf{H}_c(\tau))$. Thus, the frequency-domain explanation $\mathbf{X}_c(\tau)$ can be compared to the model output for a given threshold τ .

Figure 2 shows the relevant components projected in the frequency domain with (5) for two types of segments: speech and music. The relevant components are presented as a function of the relevance threshold: the higher τ the fewer components are selected. The classification score obtained from the selected components is also presented. The figure shows that the relevant components for speech are located between 100Hz and 1kHz. While $\tau > 0.40$, the most relevant components for speech classification remain, *e.g.*, around 200Hz. As we are in a multilabel framework, the presence of speech can also imply the presence of overlap, therefore overlap prediction is also affected by these components. This can be explained by the presence of common relevant frequency bins between SAD and OSD (Fig. 3) as described in the next section.

In the case of MD, the relevant frequency bins are located between 50Hz and 200Hz. The components located in the band [60,80]Hz are highly relevant for music detection. When those are removed ($\tau=0.9$), the music prediction drops. Speech and overlap scores remain similar for each τ since the frequency bins are different between music and these classes.




Fig. 3: Global relevant components for speech (sp), music (mu), and overlapped speech (ov).

The quality of the reconstruction in the frequency domain limits the explanation extraction. In the case of the Wavlm-based proxy model, the reconstruction is of low quality in high frequencies. Thus, no explanation can be extracted in this part of the spectrum.

5.3. Global explanation

The previous explanation is obtained at the segment level, meaning that the relevant components are identified locally. However, there is no guarantee that the components used to detect a class c in a segment are the same across an entire dataset. Thus, we propose a global explanation, $\bar{\mathbf{r}}_c$ vector, which highlights the global relevant components to detect the class c . It averages the relevance vectors \mathbf{r}_c for a set of N segments $\mathcal{D}_c = \{\mathbf{x}_{1,c}, \dots, \mathbf{x}_{n,c}, \dots, \mathbf{x}_{N,c}\}$ containing the class of interest c : $\bar{\mathbf{r}}_c = \frac{1}{N} \sum_{\mathbf{x}_c \in \mathcal{D}_c} \mathbf{r}_c$.

To compute the average relevance $\bar{\mathbf{r}}_c$, $N = 100$ segments containing exclusively each class of interest are selected to build each subset \mathcal{D}_c . To better balance the data towards the overlap class, two-speaker artificial speech mixtures are generated by summing two randomly selected speech segments. The noise class is not represented since too few segments contain exclusively noise.

Figure 3 presents the average relevance $\bar{\mathbf{r}}_c$ for each k component, computed for SAD, OSD, and MD. It shows that some components are typical for SAD and OSD. For example, the component of index 50 is activated for both classes. The figure also demonstrates that the components related to music are different from speech and overlap. This confirms the behavior observed at the segment level in Fig. 2. Finally, Fig. 3 shows negative values that are due to negative weights in the Θ layer. This may highlight some components that are very discriminant between classes, *e.g.* around $k = 70$ where the relevance is negative for speech and positive for overlap.

6. CONCLUSIONS

This paper proposes a set of explainable proxy models for multilabel audio segmentation. The proxy is trained to fit the logit distribution of a pre-trained black-box teacher model. Non-negative Matrix Factorization (NMF) is the core of this approach. It maps the embedding used for the decision-making process to the frequency domain. Hence, one can easily identify the relevant frequency bins used to predict the segmentation. Experiments conducted on AragonRadio and DiHard III datasets show that the proxy model offers the same performance as the teacher while providing strong explainability capacities. In the future, we plan to work on the evaluation of the extracted explanations

7. REFERENCES

- [1] M. T. Ribeiro, S. Singh, and C. Guestrin, ““Why Should I Trust You?”: Explaining the Predictions of Any Classifier,” in *Proceedings of the 22nd ACM SIGKDD*, New York, NY, USA, Aug. 2016, KDD '16, pp. 1135–1144.
- [2] B. Murphy, P. Talukdar, and T. Mitchell, “Learning Effective and Interpretable Semantic Models using Non-Negative Sparse Embedding,” in *Proceedings of COLING 2012*, Mumbai, India, 2012, pp. 1933–1950.
- [3] Z. Qi, S. Khorrarn, and F. Li, “Visualizing deep networks by optimizing with integrated gradients,” in *CVPR Workshops*, 2019, vol. 2, pp. 1–4.
- [4] S. Mishra, B. L. Sturm, and S. Dixon, “Local interpretable model-agnostic explanations for music content analysis,” in *ISMIR*, 2017, vol. 53, pp. 537–543.
- [5] P. Zinemanas, M. Rocamora, M. Miron, et al., “An interpretable deep learning model for automatic sound classification,” *Electronics*, vol. 10, no. 7, pp. 850, 2021.
- [6] W. Ge, J. Patino, M. Todisco, and N. Evans, “Explaining deep learning models for spoofing and deepfake detection with shapley additive explanations,” in *ICASSP 2022*. IEEE, 2022, pp. 6387–6391.
- [7] J. Parekh, S. Parekh, P. Mozharovskiy, et al., “Tackling interpretability in audio classification networks with non-negative matrix factorization,” *arXiv preprint arXiv:2305.07132*, 2023.
- [8] D. Lee and H. S. Seung, “Algorithms for non-negative matrix factorization,” *Advances in neural information processing systems*, vol. 13, 2000.
- [9] J. Sohn, N. S. Kim, and W. Sung, “A statistical model-based voice activity detection,” *IEEE signal processing letters*, vol. 6, no. 1, pp. 1–3, 1999.
- [10] D. Charlet, C. Barras, and J-S. Lienard, “Impact of overlapping speech detection on speaker diarization for broadcast news and debates,” in *ICASSP*, 2013, pp. 7707–7711.
- [11] Y. Lavner and D. Ruinskiy, “A decision-tree-based algorithm for speech/music classification and segmentation,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2009, pp. 1–14, 2009.
- [12] M. Lavechin, M-P. Gill, R. Bousbib, H. Bredin, and L. P. Garcia-Perera, “End-to-End Domain-Adversarial Voice Activity Detection,” in *Interspeech*, 2020, pp. 3685–3689.
- [13] L. Bullock, H. Bredin, and P. Garcia-Perera, “Overlap-Aware Diarization: Resegmentation Using Neural End-to-End Overlapped Speech Detection,” in *ICASSP*, 2020, pp. 7114–7118.
- [14] M. Lebourdais, M. Tahon, et al., “Overlapped speech and gender detection with WavLM pre-trained features,” in *Interspeech*, 2022, pp. 5010–5014.
- [15] 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,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2019, no. 1, pp. 1–12, 2019.
- [16] D. de Benito-Gorron, A. Lozano-Diez, et al., “Exploring convolutional, recurrent, and hybrid deep neural networks for speech and music detection in a large audio dataset,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2019, no. 1, pp. 1–18, 2019.
- [17] P. Gimeno, I. Viñals, A. Ortega, A. Miguel, and E. Lleida, “Multiclass audio segmentation based on recurrent neural networks for broadcast domain data,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2020, pp. 1–19, 2020.
- [18] J-W. Jung, H-S. Heo, Y. Kwon, J. Son Chung, and B-J. Lee, “Three-Class Overlapped Speech Detection Using a Convolutional Recurrent Neural Network,” in *Interspeech*, 2021, pp. 3086–3090.
- [19] H. Bredin and A. Laurent, “End-To-End Speaker Segmentation for Overlap-Aware Resegmentation,” in *Interspeech*, 2021, pp. 3111–3115.
- [20] M. Lebourdais, T. Mariotte, M. Tahon, A. Larcher, et al., “Joint speech and overlap detection: a benchmark over multiple audio setup and speech domains,” *arXiv preprint arXiv:2307.13012*, 2023.
- [21] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [22] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, et al., “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [23] G. Richard, P. Smaragdis, S. Gannot, P. A. Naylor, et al., “Audio signal processing in the 21st century,” *IEEE Signal Processing Magazine*, 2023.
- [24] J. Le Roux, F. J. Wengler, and J. R. Hershey, “Sparse nmf-half-baked or well done?,” *Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA, Tech. Rep., no. TR2015-023*, vol. 11, pp. 13–15, 2015.
- [25] A. Ortega, D. Castan, A. Miguel, and E. Lleida, “The Albayzín 2012 audio segmentation evaluation,” in *iberspeech*, 2012, pp. 283–289.
- [26] N. Ryant, P. Singh, V. Krishnamohan, R. Varma, et al., “The third dihard diarization challenge,” in *Interspeech*, Brno, Czechia, 2021, pp. 3570–3574.
- [27] T. Butko and C. Nadeu, “Audio segmentation of broadcast news in the Albayzín-2010 evaluation: overview, results, and discussion,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2011, no. 1, pp. 1–10, 2011.
- [28] 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.
- [29] A. Larcher, A. Mehrish, M. Tahon, S. Meignier, et al., “Speaker Embedding For Diarization Of Broadcast Data In The ALLIES Challenge,” in *ICASSP*, Toronto, Canada, 2021, pp. 5799–5803.
- [30] S. Bai, J. Z. Kolter, and V. Koltun, “An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling,” *arXiv:1803.01271 [cs]*, 2018.