



HAL
open science

Tackling Interpretability in Audio Classification Networks with Non-negative Matrix Factorization

Jayneel Parekh, Sanjeel Parekh, Pavlo Mozharovskyi, Gael Richard, Florence d'Alché-Buc

► **To cite this version:**

Jayneel Parekh, Sanjeel Parekh, Pavlo Mozharovskyi, Gael Richard, Florence d'Alché-Buc. Tackling Interpretability in Audio Classification Networks with Non-negative Matrix Factorization. IEEE/ACM Transactions on Audio, Speech and Language Processing, 2024, 32, pp.1392–1405. 10.1109/TASLP.2024.3358049 . hal-04539879

HAL Id: hal-04539879

<https://hal.science/hal-04539879>

Submitted on 9 Apr 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.

Tackling Interpretability in Audio Classification Networks with Non-negative Matrix Factorization

Jayneel Parekh, Sanjeel Parekh, Pavlo Mozharovskyi, Gaël Richard, Florence d’Alché-Buc

Abstract—This paper tackles two major problem settings for interpretability of audio processing networks, *post-hoc* and *by-design* interpretation. For *post-hoc* interpretation, we aim to interpret decisions of a network in terms of high-level audio objects that are also listenable for the end-user. This is extended to present an inherently interpretable model with high performance. To this end, we propose a novel interpreter design that incorporates non-negative matrix factorization (NMF). In particular, an interpreter is trained to generate a regularized intermediate embedding from hidden layers of a target network, learnt as time-activations of a pre-learnt NMF dictionary. Our methodology allows us to generate intuitive audio-based interpretations that explicitly enhance parts of the input signal most relevant for a network’s decision. We demonstrate our method’s applicability on a variety of classification tasks, including multi-label data for real-world audio and music.

Index Terms—Audio interpretability, explainability, by-design interpretable models, audio convolutional networks, non-negative matrix factorization

I. INTRODUCTION

Deep learning models, while state-of-the-art for several tasks in domains such as computer vision, natural language processing and audio, are typically not interpretable. Their increasing use, in decision-critical domains, in every-day applications, raises the problem of interpreting their decisions. This issue has been grappled with in two different ways by the research community. The first one relies on the availability of a predictive model trained and optimized for performance but not for interpretability and aims to devise an additional approach to interpret the given model. This problem setting is usually referred to *post-hoc interpretation*. While highly popular in the research community, *post-hoc* interpretation methods, have been criticised recently for being potentially misleading or unreliable in a variety of contexts [1]–[4]. Consequently, it has led to increasing importance of another problem setting with the goal to build predictive models that are interpretable. This has specially been argued to be crucial for critical decision-making applications [5], [6]. The challenge then is to demonstrate high classification performance while maintaining interpretability in the same model. This setting is often referred to as *by-design interpretation* problem. Real-world scenarios of utilizing interpretability of networks can occur under variety of constraints and demands regarding

deployment, level of interpretability and performance. Thus, from a practical standpoint, both problem settings hold independent value.

In this paper, our aim is to address both problems for audio classification networks while proposing a system more suited for understanding interpretations for audio modality than other common methods in literature.

An ideal interpreter is supposed to offer insights about a model’s decision in an understandable fashion to humans. In the case of audio classification, there are certain desirable traits for an interpreter that effectively help to fulfil this purpose. Firstly, we advocate that the interpretations should be generated in terms of high-level audio objects. Even more importantly, the interpretation should be listenable for an end-user. The rationale behind posing these traits as desirable is as follows: Audio scenes are often composed of multiple high-level audio objects [7]. Moreover, understanding events/scenes through the notion of audio objects also aligns with cognitive development in human and animals [8], [9]. Listenability is essential since it is significantly more intuitive and easier to listen to an interpretation rather than visualizing it in its time-frequency representation (eg. spectrogram). Its utility to users, even for multimodal applications has previously been argued in [10]. Usefulness for both the traits can be reinforced through an example. Imagine an audio-based surveillance system for a house raising an alarm for break-in. An interpreter can be expected to be able to localize the event among a host of concurrent events that triggered the alarm. If for example ‘glass-breaking’ is the triggering event that the interpreter recognizes in the input, a human would find it easier to understand, if they can hear the interpretation rather than visualize it on a spectrogram.

To this end, we propose an interpreter that relies on processing selected hidden representations of the classifier by a neural network to extract an intermediate embedding. This intermediate encoding is regularized in multiple ways, the two essential ones being: (i) Mimicking the classifier output to be able to interpret its decisions, and (ii) Reconstruct the input through the help of a dictionary of spectral patterns. The latter loss and its design is strongly inspired by the structure in Non-negative Matrix Factorization (NMF, [11]), known to provide part-based decompositions. The loss is crucial in imposing a highly understandable meaning of “time activations” on the intermediate embedding. This decomposition structure also allows the interpreter to benefit from filtering information from the input. It’s worth emphasizing that audio interpretability is not the same as classical tasks of separation or denoising. These tasks involve recovering complete object of interest in

Jayneel Parekh, Sanjeel Parekh, Pavlo Mozharovskyi, Gaël Richard, Florence d’Alché-Buc are with the Laboratoire de Traitement et Communication de l’Information (LTCI), Télécom Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France. (e-mail: {jayneel.parekh,pavlo.mozharovskyi,florence.dalche,gael.richard}@telecom-paris.fr, sanjeelparekh@gmail.com)

the output audio. On the other hand, a classifier network might focus more on salient regions. When interpreting its decision and making it listenable we expect to uncover such regions and not necessarily the complete object of interest.

This paper builds upon our work on post-hoc interpretability [12]. It differs in two significant ways:

- In addition to post-hoc, we present a novel by-design interpretable neural architecture that is one of the first such systems apart from prototypical networks [13]. Our proposed network achieves state-of-the-art performance among interpretable models for audio. Moreover, it marks an important step to bridge the performance gap between state-of-the-art audio classification models and interpretable audio classifiers on highly complex tasks. To do so, we modify the training procedure to allow the hidden layers of the classifier to be fine-tuned, and apply a classification loss at its output.
- Our prior application focus was limited to object/event detection tasks to outdoor audio. We extend the applicability of our design and demonstrate the utility of the proposed methods on music data. In particular, we present results on the largest non-synthetic publicly available polyphonic (non-synthetic) music instrument recognition dataset, OpenMIC-2018 [14].

In summary, our aforementioned contributions help us in building a holistic approach that generates listenable concept-based interpretations to tackle post-hoc and by-design interpretability for audio classification networks. Our system’s learning relies on an original formulation that constrains the interpreter encoding through two loss functions, one for input reconstruction through NMF dictionary and the other for fidelity to the network’s decision. From a learning perspective, it is a novel way to link NMF with deep neural networks, especially for generating interpretations. We extensively evaluate on a diverse set of audio event analysis benchmarks, tackling both multi-class and multi-label classification tasks. Our method’s design allows us to simulate feature removal and perform *faithfulness* evaluation.

II. RELATED WORKS

We organize the discussion about literature in three parts. We first cover interpretability works in a broad sense, followed by audio specific methods. We conclude the section by discussion on applications of NMF to audio signal processing.

A. General Interpretability literature

Feature attribution: The vast majority of interpretability literature is covered under feature attribution methods. They are a class of methods which offer interpretations through input feature importance or selection. In case of post-hoc interpretation, this includes perturbation based approaches [15]–[17], as well as saliency map based approaches [18]–[22]. Perturbation based approaches rely on observing model output on many locally perturbed versions of the input to determine importance of each individual feature for the decision. Part of the research challenge for these methods is to define “meaningful” perturbations [23]. Saliency map based

approaches typically generate interpretations through modified gradient backpropagation [18], [21], but some also utilize upsampled versions of intermediate activation maps [20], [22]. This form of interpretation is also common for by-design interpretable models. Common ways of training such models involve modifying the architecture, loss function (or both) to incorporate interpretability in the model.

Interpretation beyond attribution: Feature attribution methods have been under scrutiny for their robustness and faithfulness [1], [2]. This has overlapped with increasing amount of research for both post-hoc and by-design interpretation models to develop systems that offer interpretations through different means. Prior research has now proposed systems that provide interpretations via logical rules [24], counterfactuals [25] and even natural language [26]. Each of them can potentially be a more suited choice for certain problem domains and use-cases. Our approach uses high-level audio objects for interpretation. This aspect renders it closer to prototype and concept-based approaches. Prototype-based approaches [27], [28] tackle by-design interpretability by learning a predefined number of embeddings which are encouraged to represent singular training datapoints. The final classification decision is made based on similarity of given test sample to all the prototypes. Concept-based approaches aim to represent high-level concepts explicitly and subsequently offer interpretation in terms of them. An interesting approach is that of FLINT [29] with whom we share the idea of utilizing the hidden layers and loss functions to encourage interpretability. However we crucially differ from FLINT and other related approaches in concept representation and their applicability for audio interpretations. FLINT represents concepts by a dictionary of attribute functions over input space. The learnt concepts are not obviously comprehensible to a user, requiring a separate visualization pipeline to get insights. Approaches based on TCAV [30], such as ACE [31], ConceptSHAP [32], define concept using a set of images and learn a representation for it in terms of hidden layers of the network, termed as concept activation vector (CAV). These designs for concepts are not related to our NMF-inspired dictionary representation. Importantly, none of the above mentioned approaches can generate listenable interpretations which is key for understandability of audio processing networks.

B. Audio Interpretability

Compared to literature in other domains like text or images, work on interpretability with audio signals is considerably sparse. Usefulness of saliency-map based methods for audio interpretability has been illustrated by prior works. Becker et al. [33] demonstrate the use of popular Layerwise Relevance Propagation (LRP) [21] algorithm for audio digit recognition task. Won et al. [34] highlight the use of attention for music tagging task. Muckenhirn et al. [35] use GuidedBackprop [18] to analyze CNNs operating on 1D waveform. However, these methods do not address the issue of listenability of interpretations and moreover are limited to post-hoc interpretation. A few works based on the LIME algorithm [15] have attempted to address this issue. SLIME [36], [37] proposed to segment

the input along time or frequency. The input is perturbed by switching "on/off" the individual segments. AudioLIME [38], [39] proposed to separate the input using predefined sources to create the simplified representation. AudioLIME arguably generates more meaningful interpretations than SLIME as it relies on audio objects readily listenable for end-users. However, it suffers from limited applicability, requiring existence of known and meaningful predefined sources that compose the input audio. Schuller et al. [10] discussed in detail the various research themes surrounding sonification of explanations. More recently, [40] extended the idea of TCAV to represent concepts in music data. The supervised approach requires the overhead of human annotation of concepts, whereas the unsupervised approached based on non-negative tensor decomposition faces the challenge of meaningful learning of concepts. APNet [13] proposes an interpretable system by-design that extends prototypical networks [27], [28] for audio input by defining a more suitable distance measure for audio prototypes. An advantage of our approach over all the previous methods is its ability to address both post-hoc and by-design problem settings. Another unique feature of the approach is its means of generating interpretations, relying on a dictionary of NMF components for the same. From the point of view of interpretability, this is a novel strategy to gain insights about a model.

C. NMF applications for audio

NMF is a data decomposition technique popularized by Lee and Seung [11] as a method to learn "parts of an object". It has since been used widely within the audio community to tackle source separation [41], denoising [42], inpainting [43] and transcription [44], [45]. Typically, a nonnegative time-frequency audio representation is decomposed into two nonnegative factors namely, spectral patterns or dictionary matrix and their time activations. Its traditional usage as a supervised dictionary or feature learning method involves learning class-wise dictionaries over training data [46]. Time activations, the so-called features, can then be generated for any input by projecting it onto the learnt dictionaries. These features can subsequently be used for downstream tasks such as classification. In this section we focus on prior works that combine NMF with deep architectures.

Bisot et al. [47] couple NMF-based features with neural networks to boost performance of acoustic scene classification. NMF has also been successfully employed with audio-visual deep learning models for separation [48] and classification [49].

Iterations of NMF optimization algorithms can be unfolded as novel deep neural networks. This observation has led to development of "Deep NMF" methods. In particular, Le Roux et al. [50] unfold the multiplicative updates of NMF parameters into a deep network for speech separation. Wisdom et al. [51] apply this strategy to iterative soft thresholding algorithm to propose deep recurrent NMF.

While these works share with us the high-level idea of combining neural networks and NMF, there is no overlap between our goals and methodologies. Unlike aforementioned studies, we wish to investigate a classifier's decision using NMF as a

regularizer. Furthermore, to our best knowledge, attempting to regress temporal activations of a fixed NMF dictionary by accessing intermediate layers of an audio classification network is novel even within the NMF literature.

III. SYSTEM DESIGN

We organize this section as follows: We start with a brief note on notation used throughout the paper. We describe the setup of our framework to address post-hoc interpretation in section III-B. This is extended to address by-design interpretation in section III-C. We expound on the specific architectural details common to both problem settings in section III-D and conclude the section by detailing how we generate interpretations with our design in section III-E.

A. Data Notation

We denote a training dataset by $\mathcal{S} := (x, y)_{i=1}^N$, where $x \in \mathcal{X}$ is the time domain audio signal and $y \in \mathcal{Y}$, a label vector. The label vector could be a one-hot or binary encoding depending upon a multi-class or multi-label dataset, respectively. For listenable interpretations through NMF, we favor a representation of x that can be easily inverted back to the time-domain and use a log-magnitude spectrogram $\mathbf{X} \in \mathbb{R}^{F \times T}$ that is computed by applying an element-wise transformation $x_0 \rightarrow \log(1 + x_0)$ on the magnitude spectrogram with F frequency bins and T time frames. This is preferred over using magnitude spectrograms as it corresponds more closely to human perception of sound intensity [52]. A deep neural network classifier for post-hoc interpretations is denoted as $f : \mathcal{X} \rightarrow \mathcal{Y}$.

B. Post-hoc Interpretation

When addressing the problem of post-hoc interpretation, the classifier f will be pre-trained and then fixed throughout. We describe now the components of the interpreter and what are its inputs and outputs.

Overview The system design is illustrated in figure 1. The interpreter is designed to have access to hidden representations of the classifier and is tasked to produce an intermediate embedding through the function Ψ . This embedding is placed under certain constraints via function Θ and a pre-learned dictionary of NMF components \mathbf{W} . These constraints impose a highly meaningful structure on the embedding and help in interpreting the decision $f(x)$. We discuss the constraints, which form the core of our approach, in this subsection. The precise architectures of Ψ and Θ and optimization problem used to pre-learn \mathbf{W} are covered later in Sec. III-D.

Specifically, hidden layer outputs of the classifier f , taken as input by the interpreter, are denoted as $f_{\mathcal{I}}(x) \in \mathcal{Z}$. They are processed through the function $\Psi : \mathcal{Z} \rightarrow \mathbb{R}_+^{K \times T}$, modelled as a neural network. This produces an intermediate encoding. For simplicity, we will denote this encoding generated from hidden layers as $\mathbf{H}_{\mathcal{I}}(x) = \Psi \circ f_{\mathcal{I}}(x)$, a function over input x . The constraints on this encoding, implemented as loss functions are as follows:

Loss 1 (Fidelity loss): To be able to identify the relevant signal for interpretation, we constrain $\mathbf{H}_{\mathcal{I}}(x)$ to approximate

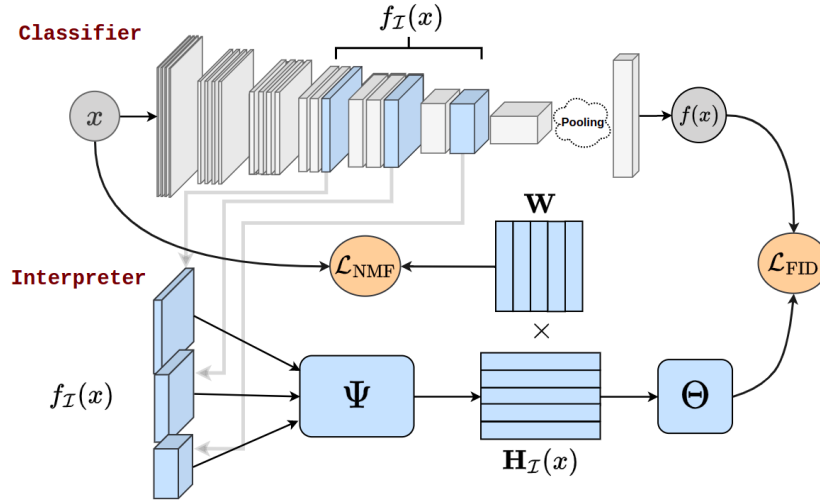


Fig. 1: **System overview:** The core design common to both post-hoc and by-design interpretation. The interpreter (indicated in blue) accesses hidden layer outputs $f_I(x)$ of the classifier. These are used to predict an intermediate encoding $\mathbf{H}_I(x)$ via a network Ψ . Through regularization terms \mathcal{L}_{NMF} and \mathcal{L}_{FID} , we encourage this encoding to both mimic the classifier’s output (through the network Θ) and also serve as the time activations of a pre-learned NMF dictionary \mathbf{W} . In post-hoc interpretation, the classifier is pre-trained and fixed, and only the interpreter is trained. For by-design interpretation we train both jointly and make final predictions using output of interpreter.

classifiers output probabilities $f(x)$ through the function $\Theta : \mathbb{R}_+^{K \times T} \rightarrow \mathcal{Y}$. The term $\Theta(\mathbf{H}_I(x))$ is also referred to as interpreter’s output. We implement this constraint as a loss function by minimizing the generalized cross-entropy loss between $\Theta(\mathbf{H}_I(x))$ and $f(x)$. We refer to it as the fidelity loss \mathcal{L}_{FID} . Denoting the parameters of Ψ, Θ as V_Ψ, V_Θ , for multi-class classification the loss can be written as,

$$\mathcal{L}_{\text{FID}}(x, V_\Psi, V_\Theta) = -f(x)^\top \log(\Theta(\mathbf{H}_I(x))) \quad (1)$$

On the other hand, for multi-label classification this loss reads,

$$\mathcal{L}_{\text{FID}}(x, V_\Psi, V_\Theta) = -\sum f(x) \odot \log(\Theta(\mathbf{H}_I(x))) + (1 - f(x)) \odot \log(1 - \Theta(\mathbf{H}_I(x))). \quad (2)$$

Here \odot denotes element-wise multiplication.

Loss 2 (Reconstruction loss): We additionally constrain $\mathbf{H}_I(x)$ to be able to reconstruct the input audio using pre-learned dictionary $\mathbf{W} \in \mathbb{R}_+^{F \times K}$. This constraint asks to decompose input log-magnitude spectrogram as $\mathbf{X} \approx \mathbf{W}\mathbf{H}_I(x)$, that is, a product of two non-negative matrices. This loss is based on popular non-negative matrix factorization. Crucially, this allows us to consider $\mathbf{H}_I(x)$ as a time activation matrix for \mathbf{W} . We refer to this as the reconstruction loss, denoted as \mathcal{L}_{NMF} .

$$\mathcal{L}_{\text{NMF}}(x, V_\Psi) = \|\mathbf{X} - \mathbf{W}\mathbf{H}_I(x)\|_2^2. \quad (3)$$

Loss 3 (Sparsity loss): In addition to \mathcal{L}_{FID} and \mathcal{L}_{NMF} , we impose ℓ_1 regularization on $\mathbf{H}_I(x)$ to encourage well-behavedness, especially for large dictionary sizes [53].

Training optimization. The complete loss function over our training dataset \mathcal{S} can thus be given as:

$$\mathcal{L}(V_\Psi, V_\Theta) = \sum_{x \in \mathcal{S}} \mathcal{L}_{\text{FID}}(x, V_\Psi, V_\Theta) + \alpha \mathcal{L}_{\text{NMF}}(x, V_\Psi) + \beta \|\mathbf{H}_I(x)\|_1 \quad (4)$$

where $\alpha, \beta \geq 0$ are loss hyperparameters. All the parameters of the system are constituted in the functions Ψ, Θ and dictionary \mathbf{W} . Since \mathbf{W} is pre-learned and fixed, the training loss \mathcal{L} is optimized only w.r.t V_Ψ, V_Θ . As a reminder, when training the interpreter for post-hoc analysis, the classifier network is kept fixed. The final optimization problem addressed for post-hoc interpretation writes as follows:

$$\hat{V}_\Psi, \hat{V}_\Theta = \arg \min_{V_\Psi, V_\Theta} \mathcal{L}(V_\Psi, V_\Theta) \quad (5)$$

C. By-design Interpretation

Interestingly, the same framework can also be utilized to train an inherently interpretable model. As a first step, we propose the following function to be used for making final predictions

$$g : \mathcal{X} \rightarrow \mathcal{Y}, g(x) = \Theta \circ \Psi \circ f_I(x)$$

which is a mixture of interpreter and classifier layers. One might be tempted to employ the same training mechanism for by-design problem as done for post-hoc interpretation. However, there is a difference in the problem setting that poses a crucial challenge. Namely, the classifier layers are not trained for prediction as before. This implies that we cannot simply aim to generate meaningful representations from it as is.

To remedy the difficulty, we modify the training in two different ways: (i) Layers of f are now modified by back-propagating all interpreter losses to them, and thus are now jointly trained with the interpreter. (ii) We modify the loss function to include an additional prediction loss on the output $f(x)$ to train all the layers in f . The training loss function and optimization problem are written as follows:

$$\begin{aligned}
\mathcal{L}_f(x, V_f) &= -y^\top \log(f(x)) \\
\mathcal{L}_{\text{NMF}}(x, V_\Psi, V_f) &= \|\mathbf{X} - \mathbf{W}\mathbf{H}_\mathcal{I}(x)\|_2^2 \\
\mathcal{L}_{\text{FID}}(x, V_\Psi, V_\Theta, V_f) &= -f(x)^\top \log(\Theta(\mathbf{H}_\mathcal{I}(x))) \\
\mathcal{L}(V_\Psi, V_\Theta, V_f) &= \sum_{x \in \mathcal{S}} \mathcal{L}_f(x, V_f) + \gamma \mathcal{L}_{\text{FID}}(x, V_\Psi, V_\Theta, V_f) \\
&\quad + \alpha \mathcal{L}_{\text{NMF}}(x, V_\Psi, V_f) + \beta \|\mathbf{H}_\mathcal{I}(x)\|_1 \\
\hat{\Psi}, \hat{\Theta}, \hat{f} &= \arg \min_{\Psi, \Theta, f} \mathcal{L}(V_\Psi, V_\Theta, V_f)
\end{aligned}$$

A reader might question the need for applying a prediction loss at output of f when the function g described above is proposed to make final predictions. This is indeed a reasonable variant of our current choice and we resolve this issue by comparing the performance of both systems in experiments in section V-D

D. Filling the gaps

It should be noted that the network architectures and other implementation details remain the same in both problem settings. We now cover the remaining architectural details of Ψ , Θ and the algorithm for pre-learning \mathbf{W} .

Design of Ψ . The network Ψ is tasked with producing the encoding $\mathbf{H}_\mathcal{I}(x) \in \mathbb{R}_+^{K \times T}$ from the set of convolutional feature maps of the classifier, given by $f_\mathcal{I}(x)$. These feature maps potentially originate from different layers and thus can be of different resolutions. To perform joint processing on them, each one is first appropriately transformed to ensure same width and height dimensions. The subsequent layers process these maps through some convolutional (with ReLU activation) and resampling layers. However, this composition is based on certain important aspects. Firstly, audio feature maps of CNNs with spectrogram-like inputs contain the notion of time and frequency along the width and height dimensions. Secondly, our goal with this network is to process a 3D representation of feature patterns across time and frequency, and convert it to a 2D intermediate encoding that can serve as time activation matrix of size $K \times T$. To achieve this, the subsequent convolutional layers continuously decrease resolution on the frequency axis and increase resolution the time axis to T frames. Furthermore, the input axis for number of feature maps corresponds to the axis of number of components K in output of Ψ , equal to the number of components in dictionary \mathbf{W} .

Design of Θ . The goal of this network is to mimic the output $f(x)$ by processing $\mathbf{H}_\mathcal{I}(x)$. This directly helps in shaping $\mathbf{H}_\mathcal{I}(x)$ to interpret $f(x)$. An important consideration for designing Θ was to keep its operations on $\mathbf{H}_\mathcal{I}(x)$ interpretable. This helps during the interpretation phase in easily quantifying how different parts of $\mathbf{H}_\mathcal{I}(x)$ influence the interpreters output. It is thus composed of two parts. The first part pools activations $\mathbf{H}_\mathcal{I}(x)$ across time. This pooling can be implemented in multiple ways, for eg. max or average pooling. However, we opt for an intermediate style of attention-based pooling [54], i.e., $\mathbf{z} = \sum_{t=1}^T \mathbf{H}_\mathcal{I}(x) \mathbf{a}$, where $\mathbf{a} \in \mathbb{R}^T$ are the attention weights and $\mathbf{z} \in \mathbb{R}^K$ is the pooled vector. In the scenario of a multi-label task with many classes and potentially high overlap, we

allow the network to attend to activations $\mathbf{H}_\mathcal{I}(x)$ separately for prediction of each class by generating a class-specific attention vector. The pooled representation vector is passed through a linear layer. This is followed by an appropriate activation function to convert its output to probabilities, that is, softmax for multi-class classification and sigmoid for multi-label classification.

Pre-learning \mathbf{W} . The non-negative matrix \mathbf{W} forms an integral part of the interpreter design. It is pre-learned from the input data, and essential in formulating the reconstruction loss \mathcal{L}_{NMF} . We employ Sparse-NMF [53] for the pre-learning. The following optimization problem is solved through multiplicative updates to pre-learn \mathbf{W} :

$$\begin{aligned}
\min \quad & D(\mathbf{X}_{\text{train}} | \mathbf{W}\mathbf{H}) + \mu \|\mathbf{H}\|_1 \\
\text{subject to} \quad & \mathbf{W} \geq 0, \mathbf{H} \geq 0, \\
& \|\mathbf{w}_k\| = 1, \forall k.
\end{aligned} \tag{6}$$

where $\mathbf{X}_{\text{train}}$ is a subset of the training data \mathcal{S} . Note that its construction is dataset dependent and will be covered in experiments. Here $D(\cdot)$ is a divergence cost function. In practice, euclidean distance is used. Training audio files are converted into log-magnitude spectrogram space for factorization.

E. Generating Interpretations

Having described the goals and details of all components of our framework, we finally discuss how the interpretations are generated. To generate audio that interprets the classifier’s decision for a sample x and a predicted class c , we follow a two-step procedure: The first step consists of identifying the components which are considered “important” for the prediction. This is determined by estimating their relevance using the pooled time activations in Θ and the weights for linear layer. Precisely, given a sample x , the pooled activations are computed as $\mathbf{z} = \mathbf{H}_\mathcal{I}(x) \mathbf{a}$. Denoting the weights for class c in the linear layer as θ_c^w , the relevance of component k is estimated as $r_{k,c,x} = \frac{(\mathbf{z}_k \theta_{c,k}^w)}{\max_l |\mathbf{z}_l \theta_{c,l}^w|}$. This is essentially the normalized contribution of component k in the output logit for class c . To select the “important” components, we simply threshold the relevance via a parameter $\tau \in (0, 1)$ as, $L_{c,x} = \{k : r_{k,c,x} > \tau\}$.

The second step consists of estimating a time domain signal for each relevant component $k \in L_{c,x}$ and also for set $L_{c,x}$ as a whole. In this paper, we refer to the latter as the generated interpretation audio, x_{int} . For certain classes, it may also be meaningful to listen to each individual component, x_k . As discussed earlier under NMF basics, estimating time domain signals from spectral patterns and their activations typically involves a soft-masking and inverse STFT procedure. We detail this step with appropriate equations in Algorithm 1.

IV. EXPERIMENTAL DESIGN

Most of the experimental settings remain the same for post-hoc, and by-design interpretations since the underlying architecture and the loss functions directly affecting interpreter are identical. Thus, the datasets, audio representation used by the network and the learnt dictionaries remain unchanged.

Algorithm 1 Audio interpretation generation

```

1: Input: log-magnitude spectrogram  $\mathbf{X}$ , input phase  $\mathbf{P}_x$ 
   components  $\mathbf{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_K\}$ , time activations
    $\mathbf{H}_T(x) = [\mathbf{h}_1^T(x), \dots, \mathbf{h}_K^T(x)]^\top$ , set of selected compo-
   nents  $L_{c,x} = \{k_1, \dots, k_B\}$ .
2: for all  $k \in L_{c,x}$  do
3:    $\mathbf{X}_k \leftarrow \frac{\mathbf{w}_k \mathbf{h}_k^T(x)^\top}{\sum_{l=1}^K \mathbf{w}_l \mathbf{h}_l^T(x)^\top} \odot \mathbf{X}$            { // Soft masking }
4:    $x_k = \text{INV}(\mathbf{X}_k, \mathbf{P}_x)$                                        { // Inverse STFT }
5: end for
6:  $\mathbf{X}_{\text{int}} \leftarrow \sum_{k \in L_{c,x}} \mathbf{X}_k$ 
7:  $x_{\text{int}} = \text{INV}(\mathbf{X}_{\text{int}}, \mathbf{P}_x)$ 
8: Output:  $\{x_{k_1}, \dots, x_{k_B}\}, x_{\text{int}}$ 

```

However, there are some differences in training and evaluation that will be discussed explicitly. We start by covering the above details in section IV-A-IV-B. We discuss the interpretation evaluation strategies relevant for both problems in section IV-C, including all the systems evaluated.

A. Datasets

We experiment with three datasets covering different types of learning tasks, source data etc. We discuss each of them in greater detail below.

1) *ESC50*: ESC-50 [55] is a popular benchmark for environmental sound classification task. It is a multi-class dataset that contains 2000 audio recordings of 50 different environmental sounds. The classes are broadly arranged in five categories namely, animals, natural soundscapes/water sounds, human/non-speech sounds, interior/domestic sounds, exterior/urban noises. Each clip is five-seconds long and extracted from publicly available recordings on the `freesound.org` project. The dataset is prearranged into 5 folds. The authors also conducted a subjective evaluation to estimate human performance for this task. Quality control was maintained via participant pre-screening and ongoing monitoring through randomly inserted test questions with an expected answer. Around 4000 judgements were collated after filtering for potential outliers. The reported human accuracy was 81.3%.

2) *SONYC-UST*: The DCASE task used a very challenging real-world dataset called Sounds of New York City-Urban Sound Tagging (SONYC-UST) [56]. It contains audio collected from multiple sensors placed in the New York City to monitor noise pollution. It consists of eight coarse-level and 20 fine-level labels. We opt for the coarse-level labeling task that involves multi-label classification into: ‘engine’, ‘machinery-impact’, ‘non-machinery-impact’, ‘powered-saw’, ‘alert-signals’, ‘music’, ‘human-voice’, ‘dog’. This task is highly challenging for several reasons: (i) since it is real-world audio, the samples contain a very high level of background noise, (ii) the audio sources corresponding to the classes are often weak in intensity, as they are not necessarily close to the sensors, (iii) some classes may also be highly localized in time and more challenging to detect, (iv) lastly, noisy audio

also makes it difficult to annotate, leading to labeling noise. This is especially true for training data labeled by volunteers.

3) *OpenMIC-2018*: The OpenMIC-2018 dataset [14] is composed of 20000 polyphonic audio recordings annotated with weak labels from among 20 instrument classes. The dataset was created by querying the content available on Free Music Archive under the Creatives Commons license with AudioSet concept ontology and using a multi-instrument estimator model trained on AudioSet data to suggest candidates for annotation. Each recording/clip is 10 seconds long. A single sample generally consists of weak labels of only a small subset of classes. Each instrument class has at least 500 positive and 1500 total annotated samples. Compared to SONYC-UST, the number of positive samples intra class and inter class are considerably more balanced. It is currently the only large publicly available dataset with multi-label annotation for polyphonic audio.

B. Implementation details

1) *Classification network*: We interpret a VGG-style convolutional neural network proposed by Kumar et al. [57]. This network was chosen due to its popularity and applicability for various audio scene and event classification tasks. It can process variable length audio and has been pretrained on AudioSet [58], a large-scale weakly labeled dataset for sound events. It takes as input a log-mel spectrogram. The architecture broadly consists of six convolutional blocks (B1–B6) followed by a convolutional layer with pooling for final prediction. Most convolutional blocks consist of two sets of conv2D + batch norm + ReLU layers followed by a max pooling layer. We fine-tune this network on each dataset separately before training our system for any post-hoc interpretations. For ESC-50, following the original paper [57], we modify only fully-connected layers after the convolutional blocks while for SONYC-UST and OpenMIC-2018, we modify all the layers during fine-tuning.

Classifier performance. We benchmark the classifier performance with other baselines in Tab. I to establish that the model we interpret in the post-hoc phase is a complex system with strong performance. On ESC-50, the classifier is evaluated using 5-fold cross-validation. It achieves an accuracy of $82.5 \pm 1.9\%$ over the 5 folds, higher than the average human accuracy of 81.3%. SONYC-UST is an unbalanced multi-label dataset. The evaluation is done using AUPRC based metrics. Our fine-tuned classifier achieves a macro-AUPRC (official metric for DCASE 2020 challenge) of 0.601. This is higher than the DCASE baseline performance of 0.510 and comparable to the best performing system macro-AUPRC of 0.649 [59]. Note that it is obtained without use of data augmentation or additional strategies to improve performance. OpenMIC-2018 is a relatively balanced multi-label dataset. To evaluate our trained classifier, we use the weighted average F1-score metric, proposed in the original paper. The metric computes for each class a weighted average of F1-scores over the positive and negative samples. The final score is the average over 20 classes. Our classifier achieves final score of 0.83, better than the VGGish based baseline score of 0.78 and competitive with other recent models. As noted earlier, the pre-training is only executed for post-hoc interpretations.

System	ESC-50 (in %)	SONYC-UST	OpenMIC-2018
	top-1	macro-AUPRC	avg-weighted-F1
Human accuracy [55]	81.3	×	×
ESC50-CNN baseline [55]	64.5	×	×
Arnault et al. [59]	×	0.649	×
Koutini et al. [60]	×	×	0.822
VGGish [14], [61]	×	0.510	0.785
Current- f	82.5	0.601	0.831

TABLE I: Benchmarking performance of pre-trained classifier f for post-hoc interpretation.

2) *Audio time-frequency representation*: For both the tasks, we perform the same audio pre-processing steps. All audio files are sampled at 44.1kHz. STFT is computed with a 1024-pt FFT and 512 sample hop size, which corresponds to about 23ms window size and 11.5ms hopsize. The log-mel spectrogram is extracted using 128 mel-bands.

3) *Dictionary learning*: The matrix on which we apply sparse-NMF to learn \mathbf{W} , $\mathbf{X}_{\text{train}}$, is constructed differently for each dataset due to their specific properties. For ESC-50, $\mathbf{X}_{\text{train}}$ is constructed by concatenating the log-magnitude spectrograms corresponding to each sample in the training data of the cross-validation fold (1600 samples for each fold). SONYC-UST however, is an imbalanced multilabel dataset with very strong presence of background noise. A procedure to learn components, as for ESC-50, yields many components capturing significant background noise, affecting understandability of interpretations. Hence, we process this dataset differently. We first learn $\mathbf{W}_{\text{noise}}$, that is, a set of 10 components to model noise using training samples with no positive label. Then, for each class, we randomly select 700 positively-labeled samples from all training data and learn 10 new components (per class) with $\mathbf{W}_{\text{noise}}$ held fixed for noise modeling. All $10 \times 8 = 80$ components are stacked column-wise to build our dictionary \mathbf{W} . While this strategy helps us reduce the number of noise-like components in the final dictionary, it does not completely avoid it. OpenMIC is instead a balanced multilabel dataset for rare noise presence. We simply select random 500 positively labeled samples for each of the 20 classes and learn 15 components. All of them are stacked together to create $\mathbf{X}_{\text{train}}$.

4) *Hyperparameters*: The hidden layers input to the interpreter module are selected from the convolutional block outputs. As is often the case with CNNs, the latter layers are expected to capture higher-order features. We thus select the last three convolutional block outputs as input to the network Ψ . The loss weights and number of components used for post-hoc interpretation are summarized in table II. Ablation studies about all the hyperparameters and justification of their choices will be presented in the next section. The hyperparameters for by-design interpretation are guided by choices in post-hoc interpretation and are tabulated in table III.

5) *Optimization*: All the networks are optimized using Adam [62] with learning rate 2×10^{-4} .

Dataset	α	β	K	# of epochs
ESC-50	10.0	0.8	100	35
SONYC-UST	10.0	0.8	80	21
OpenMIC-2018	5.0	0.2	300	21

TABLE II: Hyperparameters for all datasets for post-hoc interpretation

Dataset	γ	α	β	K	# of epochs
ESC-50	1.0	3.0	0.2	100	51
SONYC-UST	1.0	4.0	0.2	80	21
OpenMIC-2018	1.0	3.0	0.2	300	21

TABLE III: Hyperparameters for all datasets for by-design interpretation

C. Evaluating Interpretations

Quantifying different aspects of interpretability has been a challenging research question recently. This challenge stems from the inherent subjectivity involved in its definition. Our unique style of “concept-like” basis for interpretation and global approximation of the base model results in a testing situation to conduct its evaluation, wherein no other method can be directly compared to it. We resolve this hurdle by evaluating different aspects of the interpretation separately. We first discuss quantitative metrics for post-hoc and by-design interpretation along with their goals, followed by discussion on subjective evaluation of interpretations.

Metrics and baselines (Post-hoc). The simplest aspect to evaluate is how well does the interpreter agree with the classifier’s output. We refer to this metric as the *fidelity* metric. To do so for any given task, we utilize the same metric used to evaluate the classifier performance but instead treat classifiers output as ground truth and evaluate the interpreter’s approximation $\Theta(\mathbf{H}_{\mathcal{I}}(x))$ w.r.t to it. Thus, for multi-class classification, this is done by computing fraction of samples where the class predicted by f is among the top- k classes predicted by the interpreter, referred to as *top- k fidelity*. For multi-label classification tasks with unbalanced number of positive samples of classes, we compute Area Under Precision-Recall Curve (AUPRC) based metrics. In case of balanced classes, we compute F1-score based metrics. We denote our proposed Listen to Interpret (L2I) system, with attention based pooling in Θ by L2I w/ Θ_{ATT} . The most suitable baselines to benchmark its fidelity are *post-hoc* methods that approximate the classifier over input space with a single surrogate model. We select two state-of-the-art systems, FLINT [29] and VIBI [63]. A variant of our own proposed method, L2I w/ Θ_{MAX} , is also evaluated. Herein, attention is replaced with 1D max-pooling operation.

We also conduct a *faithfulness* evaluation for our interpretations. In general for any interpretability method, *faithfulness* tries to assess if the features identified to be of high relevance are *truly* important in classifier’s prediction [64]. Since a “ground-truth” importance measure for features is rarely available, attribution based methods evaluate faithfulness by performing feature removal (generally by setting feature value to 0) and observing the change in classifier’s output [64].

However, it is hard to conduct such evaluation for non-attribution or concept based interpretation methods on data modalities like image/audio, as simulating feature removal from input is not evident in these cases.

Interestingly, our interpretation module design allows us to simulate removal of a set of components from the input. Given any sample x with predicted class c , we remove the set of relevant components $L_{c,x} = \{k : r_{k,c,x} > \tau\}$ by creating a new time domain signal $x_2 = \text{INV}(\mathbf{X}_2, \mathbf{P}_x)$, where $\mathbf{X}_2 = \mathbf{X} - \sum_{l \in L_{c,x}} \mathbf{X}_l$. We define faithfulness of the interpretation to classifier f for sample x with:

$$\text{FF}_x = f(x)_c - f(x_2)_c \quad (7)$$

where $f(x)_c, f(x_2)_c$ denote the output probabilities for class c . It should be noted that this strategy to simulate removal may introduce artifacts in the input that can affect the classifier’s output unpredictably. Also, interpretations on samples with poor fidelity can lead to negative FF_x . Both of these observations point to the potential instability and outlying values for this metric. Thus, we report the final faithfulness of the system as median of FF_x over test set, denoted by $\text{FF}_{\text{median}}$. A positive $\text{FF}_{\text{median}}$ would signify that interpretations generally tend to be faithful to the classifier.

As already discussed, it is not possible to measure faithfulness for concept-based *post-hoc* interpretability approaches. While measurement for input attribution based approaches is possible, the interpretations themselves and the feature removal strategies are different, making comparisons with our system significantly less meaningful. We thus compare our faithfulness against a *Random Baseline*, wherein the less-important components, those not present in $L_{c,x}$, are randomly removed. To compare fairly, we remove the same number of components that are present in $L_{c,x}$ on average. This would validate that, if the interpreter selects *truly* important components for the classifier’s decision, then randomly removing the less important ones should not cause a drop in the predicted class probability.

We also emphasize at this point that works related to audio interpretability (see Sec. II-B), are not suitable for comparison on these metrics. Particularly, APNet [13] is not designed for *post-hoc* interpretations. AudioLIME [38] is not applicable on our tasks as it requires known predefined audio sources. Moreover, SLIME [37] and AudioLIME still rely on LIME [15] for interpretations. It is a feature-attribution method that approximates a classifier for *each* sample separately. As discussed before, these characteristics are not suitable for comparison on our metrics.

Separate from the quantitative metrics, we conducted a subjective evaluation to evaluate quality and understandability of interpretations. Our design for the same was based on qualitative understanding of saliency maps for images. Attribution maps in images are qualitatively judged by observing the visual overlap in input with the given class being interpreted. In similar spirit, our design was based on providing the user with input and class being interpreted and asking them to rate auditory overlap of the interpretation and part of input audio corresponding to the class. Further details and results

are covered in the next section. Apart from evaluating understandability, we also extensively analyze our interpretations qualitatively.

Metrics and baselines (By-design). For by-design interpretation, the faithfulness metric is much less significant. This is because the final classification output is generated by the interpreter itself and thus faithfulness is ensured by-design. The classification performance of the interpreter is the primary metric, similar in spirit to fidelity evaluation for post-hoc interpretations. We compare this with several baselines to (i) benchmark performance of our by-design interpretable network, and (ii) to evaluate the two key modifications introduced in the learning problem while extending from post-hoc to by-design interpretation (section III-C). Specifically, the hidden layers of f are not pre-trained on the given dataset in by-design problem and updated jointly with interpreter layers. And secondly, applying an additional classification loss on $f(x)$ to affect the hidden layers. The various baselines and the reasons to include them are the following:

- Audio prototypical networks (APNet) [13] act as a primary baseline from literature. It is an audio processing by-design interpretable network. While it generates interpretation differently from us, it is the only system in the literature addressing by-design interpretation for audio modality. Note that the dedicated post-hoc interpretation systems VIBI and SLIME are not relevant for this problem. For fair comparison, we use the same number of prototypes in their network as our number of components.
- In order to ascertain that using a CNN based representation for NMF offer advantage over typical NMF based representations in terms of prediction performance, we also evaluate performance of two NMF variants: Unsupervised NMF based classification and the Task-driven Dictionary Learning (TDL)-NMF model [47]. The unsupervised NMF model simply learns a dictionary on training data, computes average time activations on test samples and makes predictions using a linear model. The TDL-NMF model instead updates the initial learnt dictionary with classification loss from the linear model and thus learns them jointly. For both the systems, we experiment with use of two data types to learn NMF-dictionaries. The first is log-magnitude spectrograms and second is power mel-spectrogram (with a square root transformation). We vary dictionary sizes from 64 to 512 components and report results for best performance.
- Given the framework level similarities between FLINT and L2I, we also evaluate the performance of the FLINT interpreter when trained for by-design interpretation. As before, we again emphasize that FLINT is not suitable for audio interpretations, but provides a interpretable network design for comparison of performance.
- Variants of L2I: We denote our proposed version of L2I for by-design interpretation as L2I_{BD} w/ Θ_{ATT} . We further evaluate two variants of our proposed classification network $g(x)$. The first variant “ $\text{L2I}_{\text{BD}}\text{-NoPred}$ ” does not include a classification loss applied to $f(x)$ and instead applies it directly to $g(x)$. The second variant

System	ESC-50 (in %)		SONYC-UST		OpenMIC-2018
	top-1	top-5	macro-AUPRC	micro-AUPRC	avg-weighted-F1
L2I w/ Θ_{ATT}	65.7 \pm 2.8	88.2 \pm 1.7	0.909 \pm 0.011	0.917 \pm 0.008	0.920 \pm 0.004
L2I w/ Θ_{MAX}	73.3 \pm 2.3	92.7 \pm 1.2	0.866 \pm 0.014	0.913 \pm 0.012	0.906 \pm 0.004
FLINT [29]	73.5 \pm 2.3	93.4 \pm 0.9	0.816 \pm 0.013	0.907 \pm 0.011	0.907 \pm 0.004
VIBI [63]	27.7 \pm 2.3	53.0 \pm 1.8	0.608 \pm 0.027	0.575 \pm 0.019	0.581 \pm 0.037

TABLE IV: Fidelity results for the interpreter w.r.t classifier’s output on all datasets. We report top-1 and top-5 fidelity (in %) for ESC-50 (all five folds), AUPRC-based metrics for SONYC-UST and weighted F1-score averaged over all classes for OpenMIC-2018. All results contain mean and variance over three runs. Values in bold indicate maximum of the metric among all the evaluated systems (incl. baselines).

“L2I-PostHoc” is simply the interpreter trained for post-hoc interpretation. We compare with these variants to gain perspective on effect of differences between our formulations of post-hoc and by-design problems.

The implementation details of all the baselines (post-hoc and by-design) can be found on our companion website.¹

V. RESULTS AND DISCUSSION

A. Post-hoc Interpretation

1) *Fidelity*: As discussed previously, to quantify fidelity, we use the same respective metrics as done to benchmark classifier performance but evaluate them for interpreter output w.r.t classifier output. For ESC-50, mean and standard deviation of top- k fidelity is calculated over the 5 folds. We show these results for $k = 1, 5$. For SONYC-UST, we report the macro-AUPRC, micro-AUPRC and max-F1 for the interpreter output w.r.t classifier. For fairness, we ignore the class ‘non-machinery impact’ from all class-wise evaluations involved in fidelity (*i.e.* macro-AUPRC) or faithfulness. This is because the classifier predicts only one sample in test set with positive label for this class, causing AUPRC scores to vary widely for different interpreters. For OpenMIC-2018, we report the Fidelity weighted F1-score for each system. All the above results are available in Tab. IV.

Among the four systems, VIBI performs the worst in terms of fidelity. This is very likely because it treats the classifier as a black-box, while the other three systems access its hidden representations. This strongly indicates that accessing hidden layers can be beneficial for fidelity of interpreters. While on ESC50, FLINT achieves the best fidelity, L2I w/ Θ_{ATT} outperforms all systems on the other datasets. It should be noted that our system variants distinctly hold the advantage of generating listenable interpretations over FLINT and VIBI. Nevertheless, these systems form strong baselines for fidelity and the results demonstrate that our interpreter can generate high-fidelity *post-hoc* interpretations. Moreover, its design is flexible w.r.t different pooling functions. Results for a significance test between each pair of systems are available in supplementary discussion on our website²

System	Threshold τ	FF _{median}
L2I w/ Θ_{ATT}	$\tau = 0.9$	0.002
	$\tau = 0.7$	0.004
	$\tau = 0.5$	0.012
	$\tau = 0.3$	0.040
	$\tau = 0.1$	0.113
Random Baseline	$\tau = 0.1$	$< 10^{-4}$

TABLE V: Faithfulness results on ESC-50 for different thresholds, τ . We report FF_{median} for proposed L2I w/ Θ_{ATT} and the Random Baseline.

2) *Faithfulness*: In Table V, we report median faithfulness FF_{median} on ESC-50 for our primary system L2I w/ Θ_{ATT} at different thresholds τ averaged over the five folds. Smaller τ corresponds to higher $|L_{c,x}|$, which denotes the number of components being used for generating interpretations. Thus, for Random Baseline, we report FF_{median} at the lowest threshold $\tau = 0.1$, to ensure removal of maximal number of components. To recall the definition of Random Baseline, please refer to Sec. IV-C. FF_{median} for L2I w/ Θ_{ATT} is positive for all thresholds. It is also significantly higher than the Random Baseline, indicating faithfulness of interpretations.

The results for class-wise faithfulness on SONYC-UST and OpenMIC are illustrated in Fig. 2 and 3 respectively. We show FF_{median} (absolute drop in probability) for our system and the Random Baseline. For most classes, interpretations can be considered faithful, with a significantly positive median compared to random baseline results, which are very close to 0.

3) *Subjective evaluation*: The test was conducted with 15 participants. Each participant was provided with 10 input samples, a predicted class by the classifier for each sample and the corresponding interpretation audios from SLIME and L2I. They were asked to rate the interpretations on a scale of 0-100 for the following question: “How well does the interpretation correspond to the part of input audio associated with the given class?”. The 10 samples were randomly selected from a set of 36 (5-6 random test examples per class). For each sample, we ensured that the predicted class was both, present in the ground-truth and audible in input. Class-wise preference

¹<https://jayneelparekh.github.io/listen2interpretV2/>

²<https://jayneelparekh.github.io/listen2interpretV2/appendix.pdf>

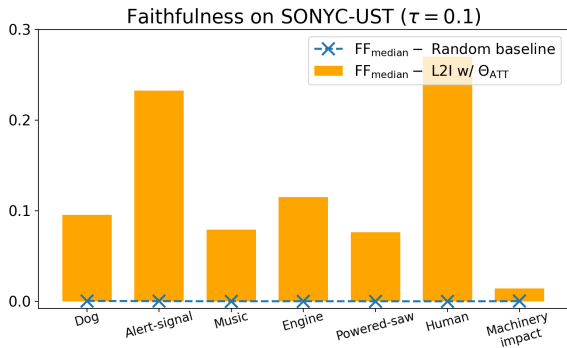


Fig. 2: Faithfulness (absolute drop in probability value) results for SONYC-UST arranged class-wise for threshold, $\tau = 0.1$

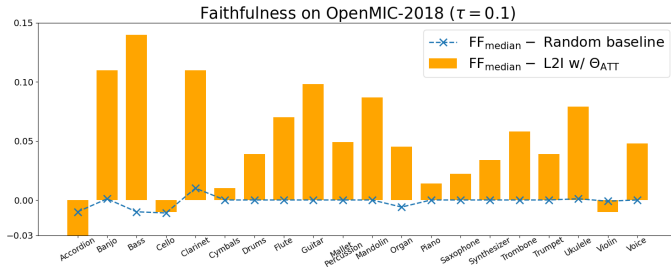


Fig. 3: Faithfulness (absolute drop in probability value) results for OpenMIC-2018 arranged class-wise for threshold, $\tau = 0.1$

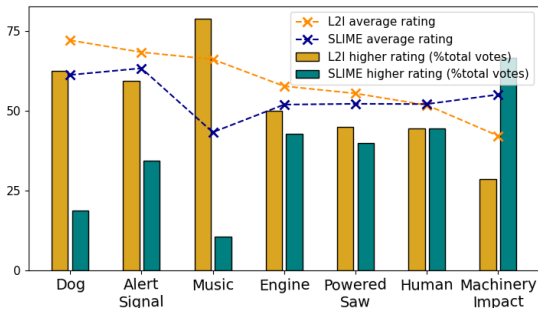


Fig. 4: Subjective evaluation results. Average scores for L2I and SLIME and fraction of votes in favour of each system

results and average ratings are shown in Fig. 4. L2I is preferred for ‘music’, ‘dog’ & ‘alert-signal’, SLIME is preferred for ‘machinery-impact’, no clear preference for others.

B. Qualitative analysis of interpretations

Qualitatively we observe that our interpretations are capable of emphasizing the object of interest and are insightful for an end-user to understand the classifier’s prediction. We share multiple examples on our companion website.¹ Samples in case of SONYC-UST and OpenMIC are often already challenging with the presence of other sources of audio. In case of ESC50, to create more interesting and challenging scenarios we devise an experiment described below

Audio corruption experiment: interpretability illustration. For ESC50, we generate interpretations after corrupting the testing data for fold-1 in two different ways (i) either

with white noise at 0dB SNR (signal-to-noise ratio), (ii) or mixing it with a sample of a different class. It should be noted that in both these cases the system is exactly the same as before and **not** trained with corrupted samples. Some examples, covering both types of corruptions are shared on our companion website.¹

For SONYC-UST, we observe good interpretations for classes ‘alert-signal’, ‘dog’ and ‘music’. For them, the background noise is significantly suppressed and the interpretations mainly focus on the object of interest. Interpretations for class ‘human’ are also able to suppress noise to a certain extent and focus on parts of human voices. However, for this class, we found presence of some signal from other audio sources too. For the remaining classes, namely ‘Engine’, ‘Powered-saw’ and ‘Machinery-impact’ the quality of the interpretation is more sample dependent. This is due to their acoustic similarity with the background noise. We provide example interpretations for SONYC-UST on our companion website.¹

The third dataset OpenMIC-2018, offers challenges under unique scenarios. Unlike SONYC-UST while it does not face issue of noise in data, it faces the hurdle of a strong overlap between instruments. This is because their onsets are often aligned by beats of the musical piece. This increases difficulty of filtering the signal of interest. Even with the greater complexity, the interpretations in many cases are able to emphasize the class of interest. Classes with relatively unique sounds such as ‘Bass’ or ‘Mallet-percussion’ are very well extracted. String like instruments including Violin and Guitar are also generally emphasized well.

Coherence of interpretations. We visualize interpretations generated on the test set for SONYC-UST and OpenMIC-2018 by clustering relevance vectors. Specifically, we compute the vector $r_{c,x} \in \mathbb{R}^K$ which contains relevances of all components in prediction for class c for sample x . The relevance vectors are collected for each test sample x and its predicted class c . We then apply a t-SNE [65] transformation to 2D for visualization. This is shown in Fig. 5. Each point is labeled/colored according to the class for which we generate the interpretation. Interpretations for any single class are coherent and similar to each other. This is to some extent a positive consequence of global weight matrix in Θ . Moreover, globally it can be observed that classes like ‘Machinery-impact’ and ‘Powered-Saw’ have similar relevances which are to some extent close to ‘Engine’. This is to be expected as these classes are acoustically similar. ‘Dog’ and ‘Music’ are also close in this space, likely due to the often periodic nature of barks or beats. The visualization for OpenMIC is arguably even more interesting because of larger number of classes and several inter-class relationships. Various sets of similar instruments end-up as clusters in proximity of each other. The examples include ‘Cello-Violin’, ‘Drums-Cymbals’, ‘Clarinet-Flute’, ‘Ukulele-Mandolin-Banjo’, ‘Trombone-Trumpet-Saxophone’. Moreover, the meaningfulness of clustering also extends to higher-level of grouping. For example, the data is partitioned so as the string-based, wind-based, or percussion instruments are close to each other within their respective groups. This indicates that the interpreter’s representations of what constitutes sound of an instrument aligns to some extent to human

ConvBlocks	\mathcal{L}_{NMF}	\mathcal{L}_{FID}	top-1
B3	0.104	1.788	53.0
B6	0.118	1.698	57.8
B2+B3	0.093	1.966	51.8
B5+B6	0.103	1.572	61.5
B4+B5+B6	0.079	1.546	65.5
Input	0.102	2.384	34.5

TABLE VI: Ablation study for hidden layers: loss values on ESC50 (fold 1) test set for different subsets of hidden layers. Current choice indicated in bold.

understanding.

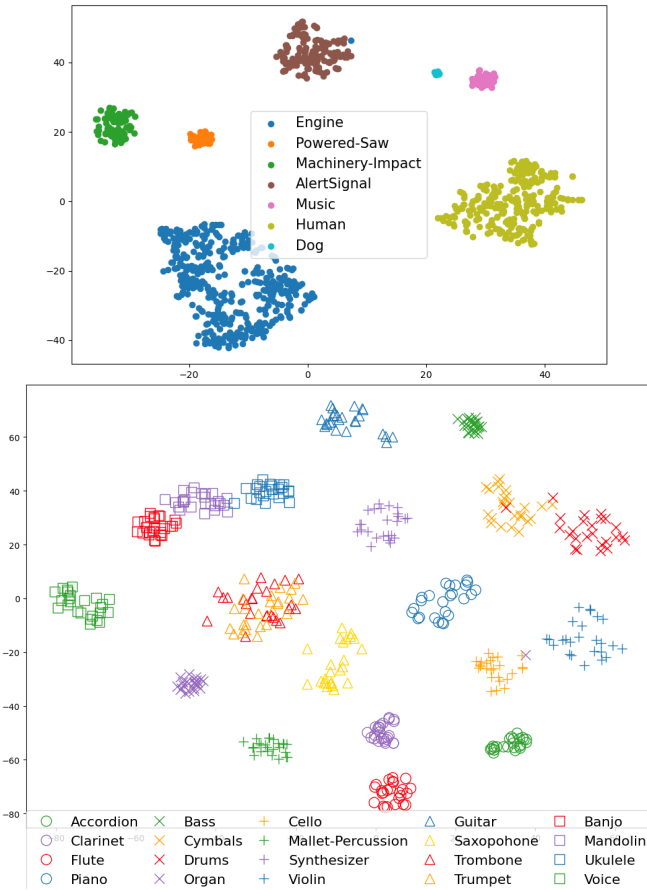


Fig. 5: Visualized relevances (following a t-SNE transformation) of generated interpretations on test sets of SONYC-UST (top) and OpenMIC-2018 (bottom), colour-coded according to interpreted class. For clarity in case of OpenMIC, we show up to random 25 interpretations of a class.

C. Ablation studies

Tab. VI and Tab. VII present ablation studies for loss hyperparameters and choice of hidden layers. The values in bold indicate our current choices for post-hoc interpretation. The metrics and loss values given here are for a single run. Selecting the **hidden layers** of the classifier that should be accessed by the interpreter is an important choice. At first

α	β	\mathcal{L}_{NMF}	\mathcal{L}_{FID}	macro-AUPRC
10.0	0.8	0.028	0.386	0.900
10.0	8.0	0.048	0.386	0.879
10.0	0.08	0.028	0.388	0.876
1.0	0.8	0.045	0.375	0.921
100.0	0.8	0.027	0.445	0.612

TABLE VII: Ablation study for loss hyperparameters: loss values on SONYC-UST test set for different weights of loss functions. Current choice indicated in bold.

glance, this model selection task might appear to be computationally too expensive as total possible choices is exponential in number of hidden layers. However, practical considerations can heavily reduce the search space. An upper bound to the number of layers could be set depending upon the desired size of interpreter. In our experiments throughout the paper, we limited ourselves to at most 3 layers. Crucially, layers close to the output are more favourable, for multiple reasons. They generally result in better fidelity and inherently tie the interpreter much closer to the output of classifier. Moreover, the latter layers are also expected to capture higher level features. We illustrate how selecting different subsets of hidden layers affects optimization of our fidelity and reconstruction loss by doing an ablation study. It’s results are reported in table VI. The classifier consists of 6 major convolutional blocks (B1–B6).

Loss weights. We illustrate the effect of varying loss weights on optimization in table VII. Too high emphasis on \mathcal{L}_{NMF} , that is, high α can hurt the fidelity of interpreter while a high β (sparsity loss) can result in poorer reconstruction. Importantly, there is a good range of values wherein the system can be regarded as operating reasonably.

Number of components. Choosing K , also known as order estimation, is typically data and application dependent. It controls the granularity of the discovered audio spectral patterns. Determining the optimal value has been a long standing problem within the NMF community [66]. Our choice for this parameter was guided by three main factors:

- Choices made previously in literature for similar pre-learning of \mathbf{W} [47], who demonstrated reasonable acoustic scene classification results with a dictionary size of $K = 128$. We used this as a reference to guide our choice.
- Dataset specific details which include number of classes, samples for each class, variability of recordings etc. For eg. acoustic variability of ESC-50 (larger number of classes), prompted us to use a dictionary of larger size compared to SONYC-UST. We use highest number of components for OpenMIC, which has largest dataset size among the three and reasonably high acoustic variability.
- When tracking loss values for different K , we observed a plateauing effect for larger dictionary sizes as illustrated in Fig. 6 for ESC-50 and OpenMIC-2018. In case of OpenMIC, this effect is prominent for reconstruction loss \mathcal{L}_{NMF} . The fidelity remains high even for small K .

Architectural modification for OpenMIC. Given the complexity of music data in OpenMIC-2018, we allowed for

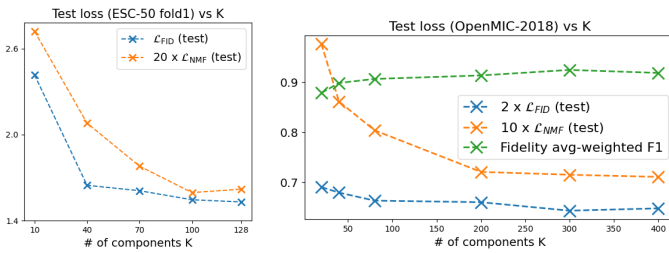


Fig. 6: Ablation study for number of components. Loss values on test data for ESC-50 and OpenMIC-2018.

modification in architecture of Θ as mentioned in Sec. III-D. We analyze its effects in supplement discussion available on our website.²

D. By-design Interpretation

The performance of all systems is given in Tab. VIII. We compute the same metrics as used to evaluate the classifiers for each dataset. Mean performance along with variance across 3 runs is reported. Pairwise system significance results are available in supplements on our website². We make the following key observations:

- Among the **interpretable neural networks for audio**, $L2I_{BD}$ w/ Θ_{ATT} clearly outperforms APNet. The size of the models plays an important role in this. $L2I$ learns with the help of a network architecture that feeds it with higher quality representations for prediction compared to architecture in APNet. It is generally able to sustain a comparable performance w.r.t the base network $BASE-f$ while imposing an interpretable structure for final prediction model.
- **Comparison with NMF baselines.** While TDL-NMF performs better than unsupervised-NMF, $L2I$ variants are noticeably better than both. This highlights a unique advantage of combining NMF representations with deep neural network representations, wherein, the NMF structure leads to interpretability and using deep networks as source provides higher prediction performance compared to directly using NMF activations generated from input.
- We also validate our design of training procedure for by-design interpretable network $g(x)$, by comparing it with the two variants of proposed system, $L2I$ -PostHoc and $L2I_{BD}$ -NoPred. The performance of $L2I_{BD}$ w/ Θ_{ATT} compared to $L2I$ -Posthoc highlights that $g(x)$ tends to perform better when hidden layers of f are trained jointly with interpreter. $L2I_{BD}$ -NoPred performs the worst among the three, emphasizing the benefits of updating the hidden layers of f with classification loss imposed on $f(x)$ rather than on $g(x)$.

VI. CONCLUSION

We have presented a framework to tackle both post-hoc and by-design for audio classification networks. To this end, a novel interpreter is designed with the key idea of using an NMF-inspired regularizer. This enables listenable concept-based interpretations. We motivate listenability as an important

attribute for audio interpretability. Efficacy of the proposed framework is established through extensive qualitative and quantitative experimentation. In particular, we quantitatively evaluate both post-hoc and by-design interpretations on three popular datasets pertaining to audio event and music instrument recognition tasks. We perform a user-study to confirm usefulness of our interpretations. In addition, through a visualization of the generated interpretations, we show that they are coherent across samples from different classes and cluster in a fashion that aligns well with human understanding of sound. Further works concern the extension of this framework to other machine learning audio-based tasks.

Limitations and future work: From the point of view of both interpretability and applicability our system faces certain limitations. Firstly, the current system’s application is more suited for traditional supervised classification tasks wherein multiple audio sources constitute an acoustic scene and identifying importance of a given source for the decision is a useful information as interpretation. In this regard, it could be interesting to extend the given system to arbitrary audio processing tasks. Directions such as extending the dictionary learning or interpreter learning to few-shot data scenarios are also highly relevant and interesting. Finally, it can be promising to explore the explicit representation of perceptually grounded audio concepts. For instance, representing “shrill sounds” via a single component is a hard task. Instead, it might be possible to represent it by a small set of components. However, this ability needs to be explored further. Representing complex perceptual grounded concepts can make the interpretations highly accessible and unlock the ability to generate effective global interpretations for a model.

ACKNOWLEDGEMENTS

This work has been supported by the research chair Data Science & Artificial Intelligence for Digitalized Industry and Services of Télécom Paris. This work also benefited from the assistance of French National Research Agency (ANR) under reference ANR-20-CE23-0028 (LIMPID project).

REFERENCES

- [1] P.-J. Kindermans, S. Hooker, J. Adebayo, M. Alber, K. T. Schütt, S. Dähne, D. Erhan, and B. Kim, “The (un) reliability of saliency methods,” in *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer, 2019, pp. 267–280.
- [2] D. Alvarez-Melis and T. S. Jaakkola, “On the robustness of interpretability methods,” *arXiv preprint arXiv:1806.08049*, 2018.
- [3] B. Dimanov, U. Bhatt, M. Jamnik, and A. Weller, *You shouldn’t trust me: Learning models which conceal unfairness from multiple explanation methods*. IOS Press, 2020.
- [4] A. Fisher, C. Rudin, and F. Dominici, “All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously,” *J. Mach. Learn. Res.*, vol. 20, no. 177, pp. 1–81, 2019.
- [5] C. Rudin, C. Chen, Z. Chen, H. Huang, L. Semenova, and C. Zhong, “Interpretable machine learning: Fundamental principles and 10 grand challenges,” *CoRR*, vol. abs/2103.11251, 2021.
- [6] C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nature machine intelligence*, vol. 1, no. 5, pp. 206–215, 2019.
- [7] A. S. Bregman, *Auditory scene analysis: The perceptual organization of sound*. MIT press, 1994.
- [8] T. D. Griffiths and J. D. Warren, “What is an auditory object?” *Nature Reviews Neuroscience*, vol. 5, no. 11, pp. 887–892, 2004.

System	ESC-50 (in %)	SONYC-UST	OpenMIC-2018
	accuracy	macro-AUPRC	avg-weighted-F1
L2I _{BD} w/ Θ_{ATT}	70.1 ± 1.5	0.581 + 0.008	0.825 ± 0.005
APNet [13]	63.6 ± 1.7	0.422 ± 0.012	0.563 ± 0.025
Unsupervised-NMF	39.4 ± 2.3	0.373 ± 0.006	0.659 ± 0.018
TD-NMF [47]	46.7 ± 2.7	0.431 ± 0.018	0.699 ± 0.012
L2I-Posthoc	65.4 ± 3.4	0.567 ± 0.007	0.825 ± 0.003
L2I _{BD} -NoPred	64.4 ± 1.1	0.563 ± 0.004	0.746 ± 0.006
FLINT [29]	75.3 ± 3.6	0.556 ± 0.008	0.827 ± 0.002
BASE- f	82.5	0.601	0.831

TABLE VIII: Classification performance for by-design interpretation. The evaluated systems include our proposed by-design interpretable network, denoted as L2I_{BD} w/ Θ_{ATT} , its variant with modified loss function (L2I_{BD}-NoPred), interpreter trained for post-hoc interpretation (L2I-Posthoc), classification models based on traditional NMF (unsupervised NMF and TD-NMF) and audio prototypical network APNet. The base classification network used for post-hoc interpretation (BASE- f) and FLINT are used as references for high performance networks not suitable for audio interpretation.

- [9] B. J. Dyson and C. Alain, "Representation of concurrent acoustic objects in primary auditory cortex." *The Journal of the Acoustical Society of America*, vol. 115, no. 1, pp. 280–288, 2004.
- [10] B. W. Schuller, T. Virtanen, M. Riveiro, G. Rizos, J. Han, A. Mesaros, and K. Drossos, "Towards sonification in multimodal and user-friendly explainable artificial intelligence," in *ICMI*, 2021, pp. 788–792.
- [11] D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Advances in Neural Information Processing Systems*, T. Leen, T. Dietterich, and V. Tresp, Eds., vol. 13. MIT Press, 2001.
- [12] J. Parekh, S. Parekh, P. Mozharovskiy, F. d'Alché Buc, and G. Richard, "Listen to interpret: Post-hoc interpretability for audio networks with nm," in *NeurIPS*, 2022.
- [13] P. Zinemanas, M. Rocamora, M. Miron, F. Font, and X. Serra, "An interpretable deep learning model for automatic sound classification," *Electronics*, vol. 10, no. 7, p. 850, 2021.
- [14] E. Humphrey, S. Durand, and B. McFee, "Openmic-2018: An open dataset for multiple instrument recognition," in *ISMIR*, 2018, pp. 438–444.
- [15] 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 international conference on knowledge discovery and data mining*. ACM, 2016, pp. 1135–1144.
- [16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017, pp. 4765–4774.
- [17] H. Lakkaraju, E. Kamar, R. Caruana, and J. Leskovec, "Faithful and customizable explanations of black box models," in *AAAI/ACM Conference on AI, Ethics, and Society*, 2019, pp. 131–138.
- [18] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," *arXiv preprint arXiv:1412.6806*, 2014.
- [19] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, "Smoothgrad: removing noise by adding noise," *arXiv preprint arXiv:1706.03825*, 2017.
- [20] K. Schulz, L. Sixt, F. Tombari, and T. Landgraf, "Restricting the flow: Information bottlenecks for attribution," in *ICLR*, 2019.
- [21] G. Montavon, W. Samek, and K.-R. Müller, "Methods for interpreting and understanding deep neural networks," *Digital Signal Processing*, vol. 73, pp. 1–15, 2018.
- [22] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *ICCV*, 2017, pp. 618–626.
- [23] D. Alvarez-Melis and T. S. Jaakkola, "A causal framework for explaining the predictions of black-box sequence-to-sequence models," *arXiv preprint arXiv:1707.01943*, 2017.
- [24] R. Kusters, Y. Kim, M. Collery, C. d. S. Marie, and S. Gupta, "Differentiable rule induction with learned relational features," *arXiv preprint arXiv:2201.06515*, 2022.
- [25] P. Jacob, É. Zablocki, H. Ben-Younes, M. Chen, P. Pérez, and M. Cord, "SteeX: steering counterfactual explanations with semantics," in *ECCV*. Springer, 2022, pp. 387–403.
- [26] L. A. Hendricks, Z. Akata, M. Rohrbach, J. Donahue, B. Schiele, and T. Darrell, "Generating visual explanations," in *ECCV*. Springer, 2016, pp. 3–19.
- [27] O. Li, H. Liu, C. Chen, and C. Rudin, "Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions," in *AAAI Conference on Artificial Intelligence*, vol. 32, 2018.
- [28] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin, and J. K. Su, "This looks like that: deep learning for interpretable image recognition," in *NeurIPS*, 2019, pp. 8928–8939.
- [29] J. Parekh, P. Mozharovskiy, and F. d'Alché Buc, "A framework to learn with interpretation," in *NeurIPS*, 2021.
- [30] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, and R. Sayres, "Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav)," *arXiv preprint arXiv:1711.11279*, 2017.
- [31] A. Ghorbani, J. Wexler, J. Y. Zou, and B. Kim, "Towards automatic concept-based explanations," in *NeurIPS*, 2019, pp. 9277–9286.
- [32] C.-K. Yeh, B. Kim, S. O. Arik, C.-L. Li, P. Ravikummar, and T. Pfister, "On concept-based explanations in deep neural networks," *arXiv preprint arXiv:1910.07969*, 2019.
- [33] S. Becker, M. Ackermann, S. Lapuschkin, K.-R. Müller, and W. Samek, "Interpreting and explaining deep neural networks for classification of audio signals," *arXiv preprint arXiv:1807.03418*, 2018.
- [34] M. Won, S. Chun, and X. Serra, "Toward interpretable music tagging with self-attention," *arXiv preprint arXiv:1906.04972*, 2019.
- [35] H. Muckenhirn, V. Abrol, M. Magimai-Doss, and S. Marcel, "Understanding and visualizing raw waveform-based CNNs." in *Interspeech*, 2019, pp. 2345–2349.
- [36] S. Mishra, B. L. Sturm, and S. Dixon, "Local interpretable model-agnostic explanations for music content analysis." in *ISMIR*, 2017, pp. 537–543.
- [37] S. Mishra, E. Benetos, B. L. Sturm, and S. Dixon, "Reliable local explanations for machine listening," in *IJCNN*, 2020, pp. 1–8.
- [38] V. Haunschmid, E. Manilow, and G. Widmer, "audiolime: Listenable explanations using source separation," *arXiv preprint arXiv:2008.00582*, 2020.
- [39] S. Chowdhury, V. Praher, and G. Widmer, "Tracing back music emotion predictions to sound sources and intuitive perceptual qualities," *arXiv preprint arXiv:2106.07787*, 2021.
- [40] F. Foscarin, K. Hoedt, V. Praher, A. Flexer, and G. Widmer, "Concept-based techniques for" musicologist-friendly" explanations in a deep music classifier," *ISMIR*, 2022.
- [41] P. Smaragdis, "Non-negative matrix factor deconvolution; extraction of multiple sound sources from monophonic inputs," in *ICA*. Springer, 2004, pp. 494–499.
- [42] K. W. Wilson, B. Raj, P. Smaragdis, and A. Divakaran, "Speech denoising using nonnegative matrix factorization with priors," in *ICASSP*. IEEE, 2008, pp. 4029–4032.
- [43] J. Le Roux, H. Kameoka, N. Ono, A. De Cheveigne, and S. Sagayama, "Computational auditory induction as a missing-data model-fitting prob-

- lem with bregman divergence,” *Speech Communication*, vol. 53, no. 5, pp. 658–676, 2011.
- [44] C. Dittmar and D. Gärtner, “Real-time transcription and separation of drum recordings based on nmf decomposition.” in *DAFx*, 2014, pp. 187–194.
- [45] N. Bertin, R. Badeau, and G. Richard, “Blind signal decompositions for automatic transcription of polyphonic music: Nmf and k-svd on the benchmark,” in *ICASSP*, vol. 1. IEEE, 2007, pp. 1–65.
- [46] C. Févotte, E. Vincent, and A. Ozerov, “Single-channel audio source separation with nmf: divergences, constraints and algorithms,” *Audio Source Separation*, pp. 1–24, 2018.
- [47] V. Bisot, R. Serizel, S. Essid, and G. Richard, “Feature learning with matrix factorization applied to acoustic scene classification,” *IEEE/ACM TASLP*, vol. 25, no. 6, pp. 1216–1229, 2017.
- [48] R. Gao, R. Feris, and K. Grauman, “Learning to separate object sounds by watching unlabeled video,” in *ECCV*, 2018, pp. 35–53.
- [49] S. Parekh, A. Ozerov, S. Essid, N. Q. Duong, P. Pérez, and G. Richard, “Identify, locate and separate: Audio-visual object extraction in large video collections using weak supervision,” in *WASPAA*. IEEE, 2019, pp. 268–272.
- [50] J. Le Roux, J. R. Hershey, and F. Wening, “Deep nmf for speech separation,” in *ICASSP*. IEEE, 2015, pp. 66–70.
- [51] S. Wisdom, T. Powers, J. Pitton, and L. Atlas, “Deep recurrent nmf for speech separation by unfolding iterative thresholding,” in *WASPAA*. IEEE, 2017, pp. 254–258.
- [52] J. Goldstein, “Auditory nonlinearity,” *The Journal of the Acoustical Society of America*, vol. 41, no. 3, pp. 676–699, 1967.
- [53] J. Le Roux, F. J. Wening, 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.
- [54] M. Ilse, J. Tomczak, and M. Welling, “Attention-based deep multiple instance learning,” in *ICML*. PMLR, 2018, pp. 2127–2136.
- [55] K. J. Piczak, “Esc: Dataset for environmental sound classification,” in *ACM Multimedia*, 2015, pp. 1015–1018.
- [56] M. Cartwright, A. E. M. Mendez, J. Cramer, V. Lostanlen, G. Dove, H.-H. Wu, J. Salamon, O. Nov, and J. Bello, “SONYC urban sound tagging (SONYC-UST): A multilabel dataset from an urban acoustic sensor network,” in *DCASE*, October 2019, pp. 35–39.
- [57] A. Kumar, M. Khadkevich, and C. Fügen, “Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes,” in *ICASSP*. IEEE, 2018, pp. 326–330.
- [58] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *ICASSP*. IEEE, 2017, pp. 776–780.
- [59] A. Arnault and N. Riche, “CRNNs for urban sound tagging with spatiotemporal context,” *DCASE2020 Challenge, Tech. Rep.*, October 2020.
- [60] K. Koutini, H. Eghbal-Zadeh, V. Haunschmid, P. Primus, S. Chowdhury, and G. Widmer, “Receptive-field regularized cnns for music classification and tagging,” *arXiv preprint arXiv:2007.13503*, 2020.
- [61] M. Cartwright, J. Cramer, A. E. M. Mendez, Y. Wang, H.-H. Wu, V. Lostanlen, M. Fuentes, G. Dove, C. Mydlarz, J. Salamon *et al.*, “Sonymc-ust-v2: An urban sound tagging dataset with spatiotemporal context,” *arXiv preprint arXiv:2009.05188*, 2020.
- [62] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [63] S. Bang, P. Xie, H. Lee, W. Wu, and E. Xing, “Explaining a black-box by using a deep variational information bottleneck approach,” in *AAAI Conference on Artificial Intelligence*, vol. 35, no. 13, 2021, pp. 11 396–11 404.
- [64] D. Alvarez-Melis and T. Jaakkola, “Towards robust interpretability with self-explaining neural networks,” in *NeurIPS*, 2018, pp. 7775–7784.
- [65] M. LJPvd and G. Hinton, “Visualizing high-dimensional data using t-sne,” *J Mach Learn Res*, vol. 9, no. 2579–2605, p. 9, 2008.
- [66] V. Y. Tan and C. Févotte, “Automatic relevance determination in nonnegative matrix factorization with the/spl beta/-divergence,” *IEEE TPAMI*, vol. 35, no. 7, pp. 1592–1605, 2012.