

Interpreting End-to-End Deep Learning Models for Speech Source Localization Using Layer-wise Relevance Propagation

Luca Comanducci, Fabio Antonacci, Augusto Sarti

Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB), Politecnico di Milano

Piazza Leonardo Da Vinci 32, 20133 Milan, Italy

name.surname@polimi.it

Abstract—Deep learning models are widely applied in the signal processing community, yet their inner working procedure is often treated as a black box. In this paper, we investigate the use of eXplainable Artificial Intelligence (XAI) techniques to analyze learning-based end-to-end speech source localization models. We consider the Layer-wise Relevance Propagation (LRP) technique, which aims to determine which parts of the input are more important for the output prediction. Using LRP we analyze two state-of-the-art models, of differing architectural complexity that map audio signals acquired by the microphones to the cartesian coordinates of the source. Specifically, we inspect the relevance associated with the input features of the two models and discover that both networks seem to denoise and de-reverberate the microphone signals in order to compute more accurate statistical correlations between them and consequently localize the sources. To further demonstrate this fact, we estimate the Time-Difference of Arrivals (TDoAs) via the Generalized Cross Correlation with Phase Transform (GCC-PHAT) using both microphone signals and relevance signals extracted from the two networks and show that through the latter we obtain more accurate time-delay estimation results.

Index Terms—explainable artificial intelligence, speech source localization, multichannel array processing

I. INTRODUCTION

The use of deep learning techniques has become ubiquitous in most fields of science, due to their ability to obtain more accurate results than the majority of the previous state-of-the-art methodologies. Following this trend, there has been an increasing adoption of deep learning techniques also in the acoustic signal processing field [1], [2], to tackle problems such as speech source localization [3]–[5], DoA and TDoA estimation [6], [7], speech separation [8] or 3D audio [9].

The main drawback of deep learning techniques is their behavior as black-boxes, which gives the user very little control or understanding of the physical meaning, if any, of the nonlinear mappings learned by the network. This has motivated the surge of the eXplainable Artificial Intelligence (XAI) field, whose aim is to develop methods for interpreting, explaining, or visualizing features learned by deep learning models [10]. The application of XAI techniques has become common practice in the wider machine learning community, but it is still lacking in most acoustic signal processing fields. XAI has been widely adopted in image processing and computer vision [11]–[14], where the task is easier due to the

presence of visual cues. A similar reasoning can be applied in the audio domain when considering frequency-based representations as input to the network, such as spectrograms [15]–[17] or Head-Related Transfer Functions (HRTF) [18]. The adoption of visualization techniques in the case of raw audio, instead, is not straightforward. An early attempt was made in [19], where a gradient-based approach was applied to study networks created for tasks such as speaker recognition and phone classification.

We consider the Layer-wise Relevance Propagation (LRP) technique [20], an explainability technique that propagates the prediction of the network backward to the input, allowing to visualize which elements of the input data have been more important in determining the output. In [21] LRP was used in order to extensively analyze the input Ambisonics acoustic intensity features to the source localization network for several noise/reverberation conditions, but it has never been applied to end-to-end localization networks, i.e. mapping raw audio signals to cartesian coordinates.

In this paper, we analyze two deep-learning architectures for end-to-end speech source localization. The one proposed in [3], in the following *LocCNN*, whose core structure consists of a simple sequence of 1D-convolutional layers, and the much more complex [4], based on a *SampleCNN* [22] architecture, which contains residual Squeeze-and-Excitation blocks [23]. We first inspect the raw audio features and the corresponding relevance signals. We then convert them into frequency-based representations, discovering that the networks act as a “gate” function, denoting as more relevant, samples corresponding to the onset of the speech signals received by the microphones. These observations suggest that the networks discard information related to the intelligible content of the speech, to focus only on the temporal information, as also suggested by listening to the relevance signals. To find out if this is true we inspect the GCC-PHAT [24] and discover that the one computed using the relevance signals contains fewer spurious peaks than the one obtained via the microphone signals. We perform an experiment by estimating the TDoA over the whole test dataset using microphone and relevance signals. Results suggest that both networks learn to denoise and de-reverberate the signals in order to better correlate them and consequently estimate the source position. The rest of the paper is organized

as follows. In Sec. II we present LRP and how it was applied to the LocCNN and SampleCNN models, described in Sec. III along with the simulated scenarios. In Sec. IV we present the experiments and results. Finally, Sec. V concludes the paper.

II. LAYER-WISE RELEVANCE PROPAGATION

The LRP technique works by redistributing relevance values associated with neurons in upper layers backward to the input. This procedure is subject to a conservation property, meaning that the relevance received by a given layer must be redistributed to the lower layer in an equal amount [25]. The propagation of the relevance score R_k at a given layer to a generic neuron j from an immediately lower layer can be obtained by applying the rule

$$R_j = \sum_k \frac{z_{jk}}{\sum_j z_{jk}} R_k, \quad (1)$$

where R_j corresponds to the relevance score at neuron j due to neuron k . The quantity z_{jk} models how much neuron j has contributed to making neuron k relevant, while the denominator enforces the conservation property. The way in which z_{jk} is computed is denoted as *rule* and differs depending on the type of layers contained in the network.

We will now present the rules used in order to analyze the LocCNN and SampleCNN models. Specifically, for the input convolutional layer of both networks, we use the w^2 -rule, chosen since it was proposed for real-valued unbounded inputs [26], as is the case of raw audio signals. Following the literature [27], we apply to the 1D convolutional layers the γ rule, which favors the effect of positive against negative contributions and is therefore well-suited to the output of ReLU and Sigmoid activations. We then apply the ϵ rule [20] to the fully connected layers. We disregard the contribution due to the dropout layer and to the activations since these are all monotonically increasing [21]. Following [28] we use a custom canonizer to handle the skip connections in SampleCNN and we apply the signal-takes-it-all rule [29] to handle gating functions in the squeeze-and-excitation boxes.

III. SPEECH SOURCE LOCALIZATION MODELS AND ENVIRONMENTAL SETUP

In this section we describe how source localization is formulated in the LocCNN and SampleCNN models, provide a brief overview of their architectures, and describe how the application environment was simulated.




Fig. 1. Manipulation of relevant input features results using the random — —, amplitude and LRP — strategies.




Fig. 2. Signal at a microphone placed in $[0.575 \text{ m}, 7 \text{ m}, 1.2 \text{ m}]^T$ (a), and corresponding relevance using LocCNN (b) and SampleCNN (c) for a source placed in $[1.53 \text{ m}, 5.71 \text{ m}, 1.15 \text{ m}]^T$.

A. Signal Model and Problem Formulation

Let us consider an array of M microphones and a source located in $\mathbf{s} = [s_x, s_y, s_z]^T$ emitting a signal $s[n]$, where n denotes the discrete time index. The signal acquired by the m -th microphone can be described as

$$x_m[n] = h_m[n] * s[n] + e_m[n], \quad (2)$$

where $*$ denotes convolution, h_m is the Room Impulse Response (RIR) between \mathbf{m} and \mathbf{s} and e_m an additive noise term. If we denote $\mathbf{x}_m = [x_m[0], x_m[1], \dots, x_m[N-1]]^T$ the vector containing the signal acquired by the m -th microphone, where a N -sample windowing has been applied and we stack them into the matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M] \in \mathbb{R}^{N \times M}$, then the source localization task performed by LOC-CNN can be defined as finding the function f such that $\hat{\mathbf{s}} = f(\mathbf{X})$, where $\hat{\mathbf{s}}$ is an estimate of the location of \mathbf{s} in cartesian coordinates.

B. Network Architectures

The LocCNN architecture consists of 5 consecutive 1D convolutional blocks, alternated with MaxPooling. The SampleCNN architecture is instead more complex and consists of 5 residual squeeze-and-excitation boxes. We slightly modified it by adding ReLU and sigmoid activations which showed better performances and were prescribed in the paper where the SampleCNN architecture was first proposed [22]. We also removed the input normalization layer. To compare it with LocCNN we modified its output layer using a fully connected layer with 3 neurons since the network was originally proposed for a multi-source scenario. Full details available in [3] and [4].

C. Setup

We consider a $3.6 \text{ m} \times 8.2 \text{ m} \times 2.4 \text{ m}$ room where we deploy a uniform linear array (ULA) composed of $M = 16$ microphones, with inter-microphone distance 0.16 m at a height of 1.2 m . Externally positioned with respect to the array we position 3940 sources on a $1.75 \text{ m} \times 2 \text{ m}$ grid along the x-y axes (width-length) and vary their z (height) coordinate by randomly sampling it from a uniform distribution in the interval of $[1 \text{ m}, 1.5 \text{ m}]$. We split the dataset using 3152, 788, and 288 samples for train, validation, and test datasets, respectively. The test sources are placed on a $0.5 \text{ m} \times 0.5 \text{ m}$ grid placed in the middle of the whole sources grid. We consider speech signals extracted from the CMU ARCTIC database [30], synchronously sampled at $F_s = 16 \text{ kHz}$, using a different recording for each source and simulate the propagation to the microphones using the `gpuRIR` library [31]. The environmental conditions of the room are varied considering the

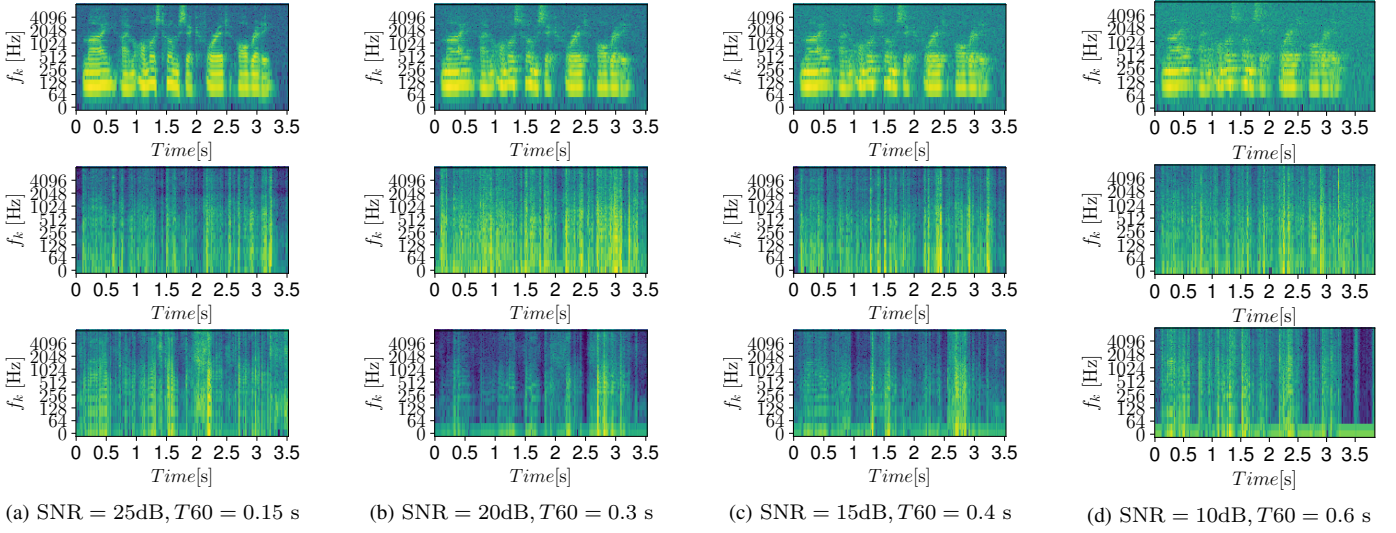


Fig. 3. STFT of the signal measured at microphone placed in $[0.57 \text{ m}, 7 \text{ m}, 1.2 \text{ m}]^T$ from a source placed in $[1.48 \text{ m}, 5.37 \text{ m}, 1.33 \text{ m}]^T$. Top row: microphone signals. Middle row: LocCNN relevance signals. Bottom row: SampleCNN relevance signals. Sub-captions denote corresponding environmental conditions.




Fig. 4. GCC-PHATs between two microphones placed in $[1.47 \text{ m}, 7 \text{ m}, 1.2 \text{ m}]^T$ and $[1.92 \text{ m}, 7 \text{ m}, 1.2 \text{ m}]^T$. Source in $[1.48 \text{ m}, 5.43 \text{ m}, 1.23 \text{ m}]^T$. Top row: microphone signals. Middle row: LocCNN relevance signals. Bottom row: SampleCNN relevance signals. ℓ denotes the time lag in samples. Sub-captions indicate the environmental conditions.

following reverberation times (T60) [0.15 s, 0.3 s, 0.4 s, 0.6 s] and Additive White Gaussian Noise (AWGN) with a Signal-to-noise ratio (SNR) of [10 dB, 15 dB, 20 dB, 25 dB]. The inputs to both localization techniques were computed by applying a 320 ms rectangular window to the microphone signals.

D. Network Training

We trained both models with a batch size of 100, following [3] and the Mean Squared Error (MSE) as a loss function. We used the Adam optimizer with a learning rate of 0.001 for LocCNN and 0.01 for SampleCNN, which were halved after 100 consecutive epochs with no validation loss improvement. We set 1000 as the maximum number of epochs and ended the training when no validation loss improvement was present for 200 consecutive epochs.

IV. EXPERIMENTS

In this section, we present results that aim to explain the learning process of LocCNN and SampleCNN. Both networks were implemented in PyTorch, while LRP via Zennit [32]. Localization error in terms of Mean Absolute Error (MAE) is comprised between 0.088 m, in the simplest scenario and 0.11 m for the most challenging one, using both LocCNN and SampleCNN. Results are computed over all windowed signals for each test source, for a total of 2911 examples. The code, and model weights as well as additional results are publicly available at <https://lucacoma.github.io/XAISrcLoc/>.

A. Manipulation of relevant input features

Following [15], we first perform a sanity check aimed at assessing the reliance of the LocCNN and SampleCNN

TABLE I
TDOA ESTIMATION, PROBABILITY OF ANOMALOUS ESTIMATES FOR DIFFERENT MICROPHONE SPACINGS.

SNR	T60	$d = 0.15$ m			$d = 0.45$ m			$d = 0.75$ m		
		Signal	LocCNN	SampleCNN	Signal	LocCNN	SampleCNN	Signal	LocCNN	SampleCNN
25 dB	0.15 s	2.64 %	0.0 %	0.29 %	3.0 %	3.0 %	5.0 %	3.25 %	21.05 %	14.87 %
20 dB	0.3 s	6.39 %	0.17 %	0.21 %	8.0 %	4.0 %	3.0 %	11.41 %	23.43 %	14.84 %
15 dB	0.4 s	13.0 %	10.0 %	3.0 %	13.0 %	10.0 %	3.0 %	20.69 %	21.46 %	17.06 %
10 dB	0.6 s	15.98 %	1.04 %	0.22 %	27.0 %	7 %	3.0 %	34.57 %	17.96 %	23.89 %

models on features deemed as relevant by LRP, before delving into further analyses. We apply three different strategies to manipulate a fraction of the input signals by zeroing out chosen samples. We use the *random* strategy as a baseline, where we put samples to zero following a uniform distribution. Then we apply the *amplitude* strategy, where samples are zeroed out according to the absolute amplitude value (i.e. first zeroing samples corresponding to higher absolute amplitude values of the microphone signal), finally, we consider the *LRP* strategy, where samples are zeroed according to the maximal relevance values indicated by LRP. The idea behind this procedure is that, if the network is actually relying on samples deemed relevant by LRP, then the performance should deteriorate more for smaller fractions of manipulation in the case of the LRP strategy, with respect to the others.

We present in Fig. 1 manipulation results averaged over the 4 environmental conditions taken into account, by removing samples considering percentages going from 0% up to 70%. It is clear that using the LRP strategy the MAE increases more steadily with respect to other strategies and performance decrease becomes comparable to the amplitude strategy only when zeroing 60% of samples, this is reasonable, since considering such short windowed signals, at that point, the audio would almost contain only noise. These results show that the samples indicated as relevant by LRP are also relevant to the network and allow us to proceed in further in our study.

B. Input features visualization

We continue our study by visually inspecting the input provided to the network and the relative relevance signal. We show in Fig. 2(a) one window of the signal acquired by a microphone and in Fig. 2(b) Fig. 2(c) the corresponding relevances obtained via LocCNN and SampleCNN, respectively. As expected, from a simple visual inspection of the 1D signals it is hardly possible to draw any conclusions.

In order to understand if it is possible to derive some visual explanations from the time-frequency representation for the input signal we compute the Short Time Fourier Transform (STFT) of the whole whole speech signal emitted by the source, obtained by concatenating the different windows in which it was divided before inputting it into the network. We using 512 frequency points, Hann window of length 512 samples, and hop size of 128 samples. We show the STFT for a single source and microphone in Fig. 3 for all considered environmental conditions computed using the microphone (top

row) and LocCNN (middle row) and SampleCNN relevance (bottom row) signals. From the inspection of the STFT we can speculate that the relevance does not preserve the phonemes of the input speech signal¹, but instead simply the pauses between the emission of such phonemes, i.e. corresponding to the temporal contour, suggesting the reasoning that, when performing single source speech localization, the network focuses more on the time at which the signals are emitted than on the speech content. This behavior seems consistent for all noise/reverberation combinations considered and is more evident for the SampleCNN model.

In fact, all these considerations are expected. From a physical understanding of sound propagation, there is no inherent reason why individual samples should be more relevant than others. The information related to the source-sensor geometry can instead be extracted by analyzing the correlations between individual signals acquired by different microphones. To understand if this is the case also for the considered models, we compute the Generalized Cross Correlation with Phase Transform (GCC-PHAT) [24] considering two microphones spaced by 0.45 m as shown in Fig 4 using the concatenated microphone signals (top row), and the concatenated relevance signals corresponding to LocCNN (middle row) and SampleCNN (bottom row) for all considered environmental conditions. From a simple visual inspection, it seems that the GCC-PHAT obtained via the relevance signals contains fewer spurious peaks, enabling to estimate more easily the TDoA between the sensors and then localizing the sources. This becomes more noticeable as the reverberation and noise conditions worsen. But is this actually what the network is learning? In order to try to answer this question we perform a further experiment assessing the performance of the TDoA estimation.

C. TDoA estimation

We consider three couples of microphones, centered in the middle of the array, with a spacing of $d = 0.15$ m, $d = 0.3$ m, $d = 0.75$ m, respectively. For each couple, we compute the TDoA as the maximum peak of the GCC-PHAT, for all the considered test sources and time windows, for a total of 2911 examples for each spacing. We quantify the TDoA estimation performance in terms of the probability of anomalous estimates P_a [7]. Following [33] we consider as anomalous estimates the ones exceeding half of the Signal

¹Listening examples available in the accompanying website.

Correlation Time, defined as the width of the main lobe of the autocorrelation function at -3 dB, for each of the considered speech signals. We compute the GCC-PHAT using both the signals arriving at the microphones and the corresponding relevance. Results are shown in Table I. The results show that the TDoA estimation performance worsens in all cases with the worsening of the environmental conditions and the increase of the microphone spacing. However, the error increases using the measured signals is more steep with respect to the one obtained using the signals obtained via LRP. The relevance signals indicate which part of the input signals are deemed important by the networks to estimate the source location. This means that both models learn from the microphone signal information that enables to better estimate of the TDoA, suggesting, that the networks leverage on statistical correlation between the microphone signals to estimate the source position.

V. CONCLUSION

In this paper, we studied the learning process of two end-to-end deep learning models, differing in complexity, for speech source localization, mapping raw multichannel audio signals to cartesian positions. We applied the Layer-wise Relevance Propagation technique to understand which parts of the input are more important in determining the predicted output. The performed analyses suggest that the networks learn to denoise and de-reverberate the signals in order to better correlate them and consequently estimate the TDoA. We believe that the obtained results should contribute to motivating the adoption of XAI techniques as a standard practice in deep learning applications for acoustic signal processing problems.

REFERENCES

- [1] M. Cobos, J. Ahrens, K. Kowalczyk, and A. Politis, "An overview of machine learning and other data-based methods for spatial audio capture, processing, and reproduction," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2022, no. 1, pp. 1–21, 2022.
- [2] M. J. Bianco, P. Gerstoft, J. Traer, E. Ozanich, M. A. Roch, S. Gannot, and C.-A. Deledalle, "Machine learning in acoustics: Theory and applications," *JASA*, vol. 146, no. 5, pp. 3590–3628, 2019.
- [3] J. M. Vera-Diaz, D. Pizarro, and J. Macias-Guarasa, "Towards end-to-end acoustic localization using deep learning: From audio signals to source position coordinates," *Sensors*, vol. 18, no. 10, p. 3418, 2018.
- [4] H. Sundar, W. Wang, M. Sun, and C. Wang, "Raw waveform based end-to-end deep convolutional network for spatial localization of multiple acoustic sources," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 4642–4646.
- [5] L. Comanducci, F. Borra, P. Bestagini, F. Antonacci, S. Tubaro, and A. Sarti, "Source localization using distributed microphones in reverberant environments based on deep learning and ray space transform," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 28, pp. 2238–2251, 2020.
- [6] S. Chakrabarty and E. A. Habets, "Multi-speaker doa estimation using deep convolutional networks trained with noise signals," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 8–21, 2019.
- [7] L. Comanducci, M. Cobos, F. Antonacci, and A. Sarti, "Time difference of arrival estimation from frequency-sliding generalized cross-correlations using convolutional neural networks," in *Proc. ICASSP*. IEEE, 2020, pp. 4945–4949.
- [8] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [9] P. Morgado, N. Nvasconcelos, T. Langlois, and O. Wang, "Self-supervised generation of spatial audio for 360 video," *NeurIPS*, vol. 31, 2018.
- [10] W. Samek and K.-R. Müller, "Towards explainable artificial intelligence," in *Explainable AI: interpreting, explaining and visualizing deep learning*. Springer, 2019, pp. 5–22.
- [11] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European conference on computer vision*. Springer, 2014, pp. 818–833.
- [12] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *ICLR 2014*, 2014.
- [13] A. Mahendran and A. Vedaldi, "Understanding deep image representations by inverting them," in *IEEE CVPR*, 2015.
- [14] A. Dosovitskiy and T. Brox, "Inverting visual representations with convolutional networks," in *IEEE CVPR*, 2016.
- [15] S. Becker, J. Vielhaben, M. Ackermann, K.-R. Müller, S. Lapuschkin, and W. Samek, "Audiomnist: Exploring explainable artificial intelligence for audio analysis on a simple benchmark," *Journal of the Franklin Institute*, vol. 361, no. 1, pp. 418–428, 2024.
- [16] T. Pellegrini and S. Mouysset, "Inferring phonemic classes from cnn activation maps using clustering techniques," in *Annual conference Interspeech*, 2016, pp. pp–1290.
- [17] L. Comanducci, P. Bestagini, M. Tagliasacchi, A. Sarti, and S. Tubaro, "Reconstructing speech from cnn embeddings," *IEEE Signal Processing Letters*, vol. 28, pp. 952–956, 2021.
- [18] E. Thuillier, H. Gamper, and I. J. Tashev, "Spatial audio feature discovery with convolutional neural networks," in *IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*. IEEE, 2018, pp. 6797–6801.
- [19] H. Muckenhirn, V. Abrol, M. Magimai-Doss, and S. Marcel, "Understanding and visualizing raw waveform-based cnns," in *INTERSPEECH*, 2019, pp. 2345–2349.
- [20] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation," *PloS one*, vol. 10, no. 7, p. e0130140, 2015.
- [21] L. Perotin, R. Serizel, E. Vincent, and A. Guérin, "Cnn-based multiple doa estimation using acoustic intensity features for ambisonics recordings," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 22–33, 2019.
- [22] J. Lee, J. Park, K. L. Kim, and J. Nam, "Samplecnn: End-to-end deep convolutional neural networks using very small filters for music classification," *Applied Sciences*, vol. 8, no. 1, p. 150, 2018.
- [23] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [24] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," *IEEE Trans. Acoust., Speech, and Signal Process.*, vol. 24, no. 4, pp. 320–327, 1976.
- [25] W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K.-R. Müller, *Explainable AI: interpreting, explaining and visualizing deep learning*. Springer Nature, 2019, vol. 11700.
- [26] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, "Explaining nonlinear classification decisions with deep Taylor decomposition," *Pattern recognition*, vol. 65, pp. 211–222, 2017.
- [27] H. Cho, E. K. Lee, and I. S. Choi, "Layer-wise relevance propagation of interactionnet explains protein–ligand interactions at the atom level," *Scientific reports*, vol. 10, no. 1, p. 21155, 2020.
- [28] F. Pahde, G. Ü. Yolcu, A. Binder, W. Samek, and S. Lapuschkin, "Optimizing explanations by network canonization and hyperparameter search," in *IEEE CVPR*, 2023, pp. 3818–3827.
- [29] L. Arras, F. Horn, G. Montavon, K.-R. Müller, and W. Samek, "What is relevant in a text document?": An interpretable machine learning approach," *PloS one*, vol. 12, no. 8, p. e0181142, 2017.
- [30] J. Kominek and A. W. Black, "The cmu arctic speech databases," in *Fifth ISCA workshop on speech synthesis*, 2004.
- [31] D. Diaz-Guerra, A. Miguel, and J. R. Beltran, "gpur: A python library for room impulse response simulation with gpu acceleration," *Multimedia Tools and Applications*, vol. 80, pp. 5653–5671, 2021.
- [32] C. J. Anders, D. Neumann, W. Samek, K.-R. Müller, and S. Lapuschkin, "Software for dataset-wide xai: From local explanations to global insights with Zennit, CoRelAy, and ViRelAy," *CoRR*, vol. abs/2106.13200, 2021.
- [33] J. Ianniello, "Time delay estimation via cross-correlation in the presence of large estimation errors," *IEEE Trans. Acoust., Speech, and Signal Process.*, vol. 30, no. 6, pp. 998–1003, 1982.